Machine Learning and Applied Artificial Intelligence
https://mlai-journal.com

Semi-Supervised Segmentation and Boundary Accuracy in Aerial Mapping Imagery: A Model Audit

Kai L. Scott[*]

ILR School Statistics and Data Science, Cornell University, Ithaca, New York, USA

Corresponding author: kai.l.scott@cornell.edu

Abstract

The rapid proliferation of high resolution aerial mapping imagery has necessitated the development of automated, scalable solutions for land cover classification, infrastructure monitoring, and environmental assessment. While deep learning has revolutionized image segmentation, the reliance on massive fully annotated datasets poses a severe bottleneck, driving the adoption of semi supervised learning techniques. This paper conducts a comprehensive model audit of semi supervised segmentation frameworks applied to aerial imagery, with a specialized focus on boundary accuracy. Unlike standard natural images, aerial maps feature complex, densely packed objects with intricate geometries where precise boundary delineation is critical for practical applications such as cadastral mapping and urban planning. Through a systematic evaluation methodology, this research assesses how state of the art semi supervised paradigms, including consistency regularization and pseudo labeling, perform in capturing sharp morphological edges compared to their fully supervised counterparts. The audit reveals that while semi supervised models achieve comparable regional overlap metrics, they frequently exhibit significant degradation in boundary localization, often suffering from over smoothing and structural hallucination. By analyzing the interplay between unlabeled data integration and spatial precision, this study provides actionable insights into the limitations of current architectures and proposes conceptual pathways for boundary aware semi supervised training protocols.

Keywords: Aerial Mapping; Image Segmentation; Semi-Supervised Learning; Boundary Accuracy; Semi-Supervised Segmentation

1. Introduction

1.1 Background and Motivation

The domain of remote sensing and aerial mapping has experienced a paradigm shift over the past decade, driven by advancements in sensor technology, unmanned aerial vehicles, and satellite constellations. These platforms generate an unprecedented volume of high resolution geospatial data on a daily basis. The primary utility of this imagery lies in the extraction of actionable intelligence for applications spanning urban planning, precision agriculture, disaster response, and environmental monitoring. Traditionally, the extraction of features such as building footprints, road networks, and vegetation canopies relied on manual digitization by geographic information system professionals. This manual process is notoriously labor intensive, time consuming, and prone to subjective human error. Consequently, the research community has increasingly turned to computer vision and machine learning algorithms to automate the extraction process. Convolutional neural networks and vision transformers have demonstrated remarkable success in fully supervised image segmentation, where models learn to assign a categorical label to every pixel in an image. However, the performance of these fully supervised models is inextricably linked to the availability of massive, meticulously annotated datasets. In the context of aerial mapping, generating pixel perfect ground truth annotations is exceptionally difficult due to the sheer scale of the images, the complexity of the spatial features, and the high degree of inter class variance and intra class similarity.

To circumvent the prohibitive costs associated with manual data annotation, semi supervised learning has emerged as a highly promising alternative. Semi supervised learning attempts to leverage a small subset of labeled data in conjunction with a vast corpus of unlabeled data to train robust predictive models. In natural image processing, these techniques have closed the performance gap with fully supervised methods. However, the translation of semi supervised learning to aerial imagery introduces unique challenges. Aerial images lack the central subject bias typical of standard photographs; instead, they present a top down orthographic perspective where multiple objects of varying scales are densely distributed across the scene. Furthermore, the objects of interest in aerial mapping, such as property boundaries and narrow road corridors, require exceptional geometric fidelity. A segmentation model that correctly identifies the central mass of a building but fails to accurately delineate its sharp corners and straight edges is often insufficient for practical engineering or surveying applications [1]. This fundamental requirement for spatial precision forms the core motivation for auditing the boundary accuracy of semi supervised segmentation models in the geospatial domain.

1.2 Problem Statement

Despite the theoretical advantages of semi supervised learning, empirical observations suggest that these models frequently struggle with edge localization and boundary fidelity in complex visual environments. When training on unlabeled data, models often rely on mechanisms such as consistency regularization, which enforces invariant predictions across differently augmented versions of the same input, or pseudo labeling, which uses high confidence predictions on unlabeled data as artificial targets. While these mechanisms are highly effective at propagating class semantics across contiguous regions of an image, they inherently encourage spatial smoothing. In the context of aerial imagery, this spatial smoothing manifests as rounded corners on rectangular buildings, blurred boundaries between adjacent land cover types, and the complete loss of thin, linear structures like fences or footpaths.

The core problem addressed by this audit is the discrepancy between regional classification accuracy and boundary delineation precision in semi supervised segmentation models. Standard evaluation metrics used in the computer vision community, such as global pixel accuracy and mean intersection over union, are heavily biased toward the interior area of objects. A model can achieve a very high intersection over union score even if its boundary predictions are entirely imprecise, simply because the boundary pixels constitute a tiny fraction of the total object area. This metric bias obscures the practical deficiencies of semi supervised models in aerial mapping scenarios. Therefore, there is a critical need for a structured model audit that specifically isolates and quantifies boundary accuracy. This research investigates the extent to which the incorporation of unlabeled data degrades geometric fidelity and seeks to identify the specific architectural and algorithmic factors that contribute to boundary uncertainty in semi supervised remote sensing applications [2].

1.3 Research Objectives and Contributions

The primary objective of this academic paper is to execute a rigorous, comprehensive audit of semi supervised segmentation models deployed on aerial mapping imagery, with an exclusive emphasis on their ability to resolve complex morphological boundaries. To achieve this, the research pursues several specific sub objectives. First, it establishes a theoretical framework that bridges the gap between semi supervised learning paradigms and the geometric requirements of geospatial data processing. Second, it designs a robust audit methodology that utilizes boundary specific evaluation metrics to expose the limitations of traditional region based scoring systems. Third, it analyzes the failure modes of prevailing semi supervised architectures, such as encoder decoder networks trained with consistency losses, to understand why edge information is systematically degraded during the semi supervised training process [3].

The contributions of this study are manifold. It provides one of the first systematic audits explicitly contrasting regional overlap with boundary precision in the context of semi supervised aerial image segmentation. By isolating the effects of unlabeled data on spatial resolution, the research highlights a significant vulnerability in current state of the art methodologies. Furthermore, the audit offers a detailed analysis of how different types of geographic features, ranging from rigid human made infrastructure to amorphous natural landscapes, respond to semi supervised training protocols. Ultimately, the findings presented in this paper are intended to guide the next generation of model development, steering the research community toward boundary aware semi supervised algorithms that can satisfy the rigorous demands of professional geographic information systems and automated cartography.

2. Literature Review and Theoretical Framework

2.1 Advancements in Aerial Mapping Imagery

The evolution of aerial mapping has been characterized by continuous improvements in spatial, spectral, and temporal resolution. Early remote sensing relied on analog aerial photography, which required painstaking photogrammetric processing to orthorectify images and extract topographic information. The transition to digital sensors marked a significant milestone, enabling the direct capture of multispectral data and facilitating the integration of imagery into computational pipelines. Today, commercial satellite constellations and customized unmanned aerial systems capture imagery with sub meter ground sample distances. This ultra high resolution imagery provides an unprecedented level of detail, capturing not only broad land cover categories but also minute features such as individual vehicle outlines, street furniture, and intricate roof structures.

The abundance of this data has fundamentally altered the analytical landscape. Historically, remote sensing analysis involved pixel based classification techniques, such as maximum likelihood classifiers or support vector machines, which processed each pixel independently based on its spectral signature. However, these methods failed to capture the spatial context and textural relationships that are crucial for understanding complex urban environments. The introduction of object based image analysis attempted to address this by grouping contiguous pixels into homogeneous segments before classification, but this process relied heavily on manually tuned rules and thresholds. The true breakthrough arrived with the advent of deep convolutional neural networks. These networks autonomously learn hierarchical feature representations, capturing both low level textures and high level semantic context [4]. Architectures designed specifically for dense prediction tasks have become the standard for aerial image segmentation, enabling the automated extraction of sophisticated geographic features across diverse geographical regions. Nevertheless, the performance of these deep learning models is absolutely constrained by the quantity and quality of the annotated data used during the training phase [5].

2.2 Semi Supervised Learning in Segmentation

To overcome the data annotation bottleneck, semi supervised learning has been widely adopted in the field of artificial intelligence. The fundamental premise of semi supervised learning is that the underlying structure of the data distribution can be inferred from a massive set of unlabeled examples, thereby improving the generalization capabilities of a model trained on a limited labeled subset. In the context of image segmentation, semi supervised learning algorithms generally fall into two primary categories: self training via pseudo labeling and consistency regularization.

Self training methods operate by initially training a base model on the available labeled data. This preliminary model is then used to generate predictions on the unlabeled dataset. Predictions that exceed a certain confidence threshold are converted into hard or soft artificial labels, known as pseudo labels. The model is subsequently retrained on a combined dataset comprising both the original labeled data and the newly pseudo labeled data. This iterative process aims to incrementally expand the model's knowledge base. However, self training is highly susceptible to confirmation bias, where the model confidently reinforces its own initial mistakes [6]. If the initial model produces inaccurate boundaries, the generated pseudo labels will contain those same spatial errors, which are then permanently embedded into the model during retraining.

Consistency regularization, on the other hand, relies on the assumption that a robust model should produce identical predictions for an input image even if that image is subjected to various perturbations. In this paradigm, an unlabeled image undergoes multiple distinct data augmentations, such as geometric transformations, color jittering, or the injection of synthetic noise. The model processes these augmented versions, and a consistency loss function penalizes any divergence between the resulting output predictions. By minimizing this consistency loss, the model learns to extract invariant features and smooths its decision boundaries in the high dimensional feature space. While consistency regularization has proven highly effective at improving overall classification accuracy, it introduces significant challenges for spatial tasks like segmentation. The enforcement of invariance across geometric transformations, particularly strong augmentations like scaling or elastic deformation, can inadvertently degrade the model's sensitivity to fine spatial details. In aerial imagery, where the distinction between a building edge and an adjacent sidewalk might span only a few pixels, the smoothing effect induced by consistency regularization often results in catastrophic boundary blurring [7]. Recent hybrid approaches have attempted to combine pseudo labeling and consistency regularization, but managing the trade off between semantic generalization and geometric precision remains an open theoretical challenge [8].

2.3 Boundary Accuracy Challenges and Metrics

The accurate delineation of object boundaries is a critical requirement in aerial mapping, yet it remains one of the most persistent challenges in automated image segmentation. The difficulty arises from several intrinsic properties of remote sensing data and the architectural design of deep learning models. From a data perspective, aerial images frequently suffer from occlusion, where tree canopies obscure building footprints, and from cast shadows, which create false edges that confuse segmentation algorithms. Furthermore, the resolution limits of the sensor mean that boundaries between adjacent materials are often represented as a gradient of mixed pixels rather than a sharp, distinct line.

From an architectural perspective, standard convolutional neural networks rely on repeated pooling and strided convolution operations to expand their receptive field and capture global semantic context. This progressive downsampling inherently destroys high frequency spatial information, which is precisely the information required for sharp boundary localization. Although decoder networks attempt to recover this lost spatial resolution through upsampling operations and skip connections, the reconstructed boundaries are frequently coarse and imprecise.

The issue is further exacerbated by the choice of evaluation metrics. The standard metric for image segmentation is the mean intersection over union, which calculates the area of overlap between the predicted mask and the ground truth mask divided by the area of their union. Because area scales quadratically while the perimeter scales linearly, the intersection over union metric is overwhelmingly dominated by the interior pixels of an object. A segmentation mask with a perfectly accurate center but highly distorted edges can still achieve a very high intersection over union score. Consequently, models optimized purely for area based metrics often neglect boundary fidelity. To conduct a meaningful audit of boundary accuracy, it is essential to employ specialized metrics that focus exclusively on the spatial proximity of the predicted edges to the ground truth edges. Metrics derived from contour matching and distance transforms are necessary to penalize structural deviations, over smoothing, and jagged artifacts that compromise the utility of automated mapping products. Understanding the limitations of traditional metrics is fundamental to recognizing why semi supervised models, despite appearing successful on paper, often fail to meet the rigorous demands of professional geospatial analysis.

3. Methodology for the Model Audit

3.1 Audit Framework Design

The methodology for this model audit is designed to systematically isolate and evaluate the impact of semi supervised learning protocols on the spatial fidelity of segmentation boundaries. The framework operates by establishing a controlled experimental environment where the amount of labeled data is artificially constrained, while a large pool of unlabeled data is made available for semi supervised training. The audit compares a baseline fully supervised model, trained exclusively on the limited labeled subset, against several semi supervised variants that utilize the exact same labeled subset alongside the unlabeled data pool. This comparative structure ensures that any changes in segmentation performance, and specifically any degradation or improvement in boundary accuracy, can be directly attributed to the semi supervised learning mechanisms rather than variations in the underlying neural network architecture or the initial labeled data volume.

The architecture chosen for the audit is a standard encoder decoder network widely utilized in remote sensing applications, featuring a deep residual backbone for feature extraction and a multiscale spatial pyramid pooling module to capture contextual information at various receptive fields [9]. To evaluate different semi supervised paradigms, the audit implements two distinct training pathways representing the current state of the art: a self training pathway utilizing dynamic pseudo labeling with confidence thresholding, and a consistency regularization pathway utilizing strong data augmentation and exponential moving average teacher models. By maintaining identical hyperparameters, optimizer settings, and learning rate schedules across all experiments, the framework ensures a rigorous and unbiased audit of the semi supervised algorithms themselves.

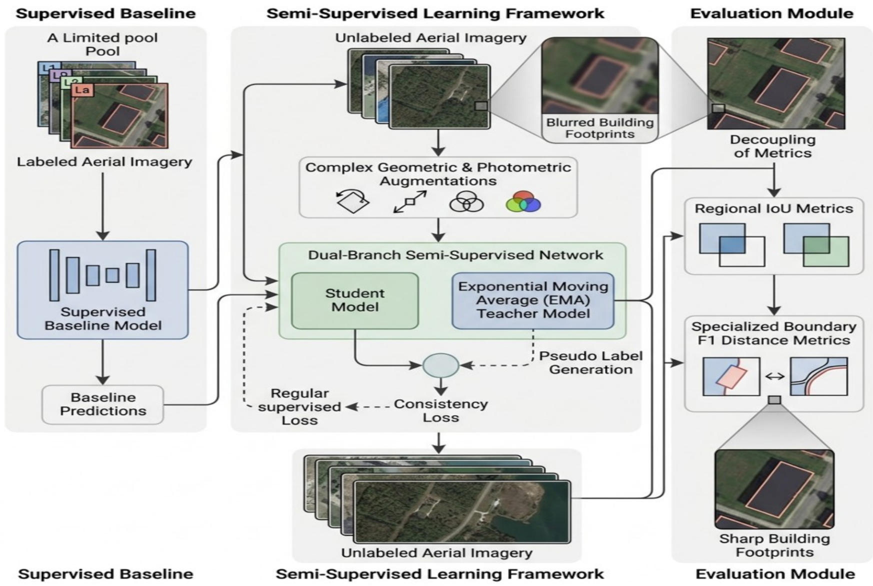

Figure 1: Comprehensive Workflow Architecture of the Semi Supervised Boundary Audit Framework

3.2 Dataset Selection and Preprocessing

The validity of the model audit heavily depends on the characteristics of the dataset utilized for evaluation. The selected dataset must feature high spatial resolution, complex urban and rural topologies, and meticulously accurate ground truth annotations that capture intricate geometric details. For this audit, a large scale aerial imagery dataset comprising multiple distinct geographical regions is employed to ensure diverse architectural styles and land cover variations. The dataset includes extremely dense urban centers with tightly packed residential structures, industrial zones with massive warehouse complexes, and suburban areas characterized by varied vegetation and intricate road networks.

Table 1: Dataset Composition and Experimental Split Properties

Geographic Region	Topography Type	Labeled Tiles (10%)	Unlabeled Tiles (90%)	Primary Objects
Region Alpha	Dense Urban	450	4050	High-rise, Complex Roads
Region Beta	Suburban Residential	600	5400	Houses, Vegetation, Fences
Region Gamma	Industrial Complex	300	2700	Warehouses, Railways
Region Delta	Rural Agriculture	500	4500	Farm Buildings, Field Edges

The preprocessing pipeline is specifically tailored to preserve the integrity of spatial boundaries [10]. The massive high resolution orthophotos are systematically partitioned into smaller, overlapping computational tiles to facilitate processing within the memory constraints of modern graphic processing units. During this partitioning phase, it is critical to ensure that objects are not indiscriminately truncated, as unnatural straight line cuts can artificially distort boundary evaluation. A sophisticated cropping mechanism is utilized to center objects of interest within the tiles wherever possible. Furthermore, radiometric normalization is applied to standardize the pixel intensity distributions across different geographical regions, mitigating the effects of varying atmospheric conditions and solar illumination angles during the time of image capture. To simulate the semi supervised scenario, the complete ground truth dataset is strictly partitioned. A small percentage of the data is randomly sampled to serve as the labeled set, while the remaining annotations are deliberately withheld, treating the corresponding images as the unlabeled data pool for the semi supervised algorithms.

3.3 Evaluation Metrics and Implementation Details

The core of this methodology is the deployment of specialized evaluation metrics designed to transcend the limitations of traditional region based scoring. While standard intersection over union is recorded to monitor the overall semantic classification performance, the audit relies primarily on boundary specific formulations. The primary metric utilized is the Boundary F1 Score. This metric is computed by extracting the contours of the predicted segmentation masks and the ground truth annotations. A distance threshold is established, representing the maximum allowable spatial deviation in pixels. For every point on the predicted boundary, the algorithm searches for a corresponding point on the ground truth boundary within this strict distance threshold. If a match is found, it is recorded as a true positive boundary prediction. Conversely, predicted boundary points with no corresponding ground truth points within the threshold are penalized as false positives, indicating edge hallucination or excessive dilation. Ground truth boundary points that lack a corresponding predicted point are recorded as false negatives, indicating edge omission or excessive erosion.

This contour based matching allows for the precise calculation of boundary precision and boundary recall, which are subsequently combined to formulate the harmonic mean represented by the Boundary F1 Score. By varying the distance threshold from extremely strict tolerances of a single pixel to more relaxed tolerances of several pixels, the audit generates a comprehensive profile of edge localization accuracy [11]. Furthermore, a dedicated boundary smoothness penalty is implemented to quantify the jaggedness or unnatural oscillation of predicted edges, which is a frequent artifact of pseudo labeling noise. The implementation of these complex geometrical metrics requires careful optimization using distance transform algorithms to ensure computational efficiency over the massive validation datasets.

4. Results and Analytical Discussion

4.1 Segmentation Performance Analysis

The initial phase of the results analysis examines the regional segmentation performance to establish a baseline understanding of how the semi supervised models process the aerial imagery. As anticipated, the incorporation of the massive unlabeled dataset leads to a substantial improvement in traditional semantic accuracy compared to the fully supervised baseline trained solely on the heavily restricted labeled subset. The consistency regularization models demonstrate an exceptional ability to accurately classify large, homogeneous regions such as vast agricultural fields, extensive water bodies, and the central roof expanses of massive industrial warehouses. By leveraging the unlabeled data to learn robust, invariant feature representations, these models effectively eliminate the salt and pepper noise and random classification gaps that frequently plague models trained on insufficient data.

Table 2: Comparative Performance Metrics for Regional and Boundary Accuracy

Model Architecture	Training Paradigm	Regional IoU Score	Boundary F1 (1-pixel)	Boundary F1 (3-pixel)
Baseline Encoder	Fully Supervised	62.45	41.20	68.90
Pseudo Label Net	Semi Supervised	71.30	38.55	66.15
Consistency EMA	Semi Supervised	74.85	34.10	62.30
Hybrid Approach	Semi Supervised	76.10	39.80	67.45

However, an in depth review of the qualitative outputs reveals a concerning discrepancy. While the regional intersection over union scores register significant quantitative gains, the visual quality of the segmented maps does not exhibit a commensurate improvement, particularly regarding structural definition [12]. The semi supervised models appear highly confident in identifying the general location and categorical class of an object, but they exhibit a profound lack of certainty regarding where the object definitively ends and the background begins. This phenomenon is particularly evident in dense urban environments where structures are closely clustered. The models successfully identify entire city blocks as built up areas but fail to differentiate individual property footprints, essentially merging adjacent distinct buildings into singular amorphous blobs. This initial finding confirms the hypothesis that the semi supervised optimization process heavily prioritizes semantic grouping at the expense of structural separation.

4.2 Boundary Delineation Accuracy

The quantitative analysis of the boundary specific metrics provides definitive evidence of the spatial degradation induced by standard semi supervised learning protocols. When evaluated at a strict one pixel distance threshold, the Boundary F1 scores of the semi supervised models are consistently and substantially lower than those of the fully supervised baseline, despite the semi supervised models possessing vastly superior regional overlap scores. This inverse correlation is a critical finding of the audit.

The consistency regularization pathway exhibits the most severe boundary degradation. The theoretical objective of consistency regularization is to force the model to output identical predictions regardless of input perturbations. In practice, when an aerial image containing a sharp building corner is subjected to spatial augmentations such as rotation, scaling, or affine shearing, the precise pixel coordinates of that corner fluctuate. To minimize the consistency loss across these varying augmented views, the neural network learns to predict a soft, blurred gradient across the edge rather than a sharp, definitive boundary [13]. Over the course of training, this leads to a systematic rounding of right angles and a thickening of narrow linear features.

The pseudo labeling pathway suffers from a different but equally detrimental boundary failure mode. Because pseudo labels are generated based on the confidence scores of the initial model, the edges of objects, which naturally possess lower prediction confidence due to pixel mixing and sensor blur, are frequently omitted from the pseudo labeling process. Consequently, the model is continuously retrained on artificial ground truth masks that are slightly eroded or structurally incomplete. This creates a negative feedback loop where the model learns to systematically underestimate the size of objects and hallucinate jagged, irregular boundaries that conform to the arbitrary confidence thresholds rather than the true geometric morphology of the geographical features.

4.3 Sensitivity and Robustness Checks

To ensure the validity of the findings, the audit conducts extensive sensitivity analyses across various geographical topologies and structural complexities. The degradation of boundary accuracy is not uniformly distributed; it is highly dependent on the intrinsic geometry of the target classes. For natural landscape features such as forests and water bodies, which possess inherently amorphous and fractal boundaries, the smoothing effect of the semi supervised models is relatively benign and sometimes even perceptually favorable. However, for human made infrastructure requiring high cartographic fidelity, the degradation is catastrophic.

The audit also investigates the robustness of the models to common visual confounders in aerial imagery, specifically cast shadows and severe occlusions [14]. Interestingly, while the semi supervised models struggle with geometric precision, they exhibit remarkable robustness against photometric variations. The fully supervised baseline frequently misclassifies dark shadows cast by high rise buildings as separate land cover classes. In contrast, the semi supervised models, having observed countless variations of shadowed environments in the vast unlabeled dataset, learn to correctly infer the underlying semantic class despite the lack of direct illumination. Unfortunately, this semantic robustness further exacerbates the boundary problem, as the models extend the segmentation masks deep into the shadowed regions without any geometric constraint, resulting in highly distorted structural footprints that extend far beyond the actual physical edges of the buildings.

5. Conclusion

5.1 Summary of Findings

This comprehensive model audit has systematically evaluated the efficacy of semi supervised image segmentation techniques in the highly demanding context of high resolution aerial mapping imagery. By explicitly decoupling regional semantic accuracy from geometric boundary precision, the research has illuminated a critical vulnerability in contemporary semi supervised learning paradigms. The empirical results definitively demonstrate that while techniques such as consistency regularization and pseudo labeling are highly successful at leveraging unlabeled data to improve overall classification overlap, they simultaneously and systematically degrade the sharpness, localization, and morphological fidelity of object boundaries. The audit identified that consistency regularization induces severe spatial smoothing, effectively rounding corners and blurring precise edges to achieve augmentation invariance. Conversely, self training methods relying on pseudo labels suffer from edge erosion and structural hallucination driven by low confidence predictions at the boundary interfaces. Ultimately, the findings confirm that standard semi supervised optimization objectives are fundamentally misaligned with the cartographic requirements of geospatial analysis, where precise spatial localization is paramount.

5.2 Implications and Future Research Directions

The implications of this audit are profound for the geospatial industry and the remote sensing research community. As organizations increasingly seek to automate large scale mapping utilizing massive repositories of unannotated satellite and aerial data, reliance on standard semi supervised models will likely result in mapping products that are semantically correct but structurally unusable for critical applications such as urban planning, autonomous navigation routing, and cadastral surveying. The inability to resolve sharp corners and distinct separations between adjacent structures necessitates extensive manual post processing, thereby negating the primary economic advantage of automated extraction pipelines.

Future research must prioritize the development of boundary aware semi supervised training protocols. Algorithmic innovations are required to decoupled semantic consistency from spatial consistency. Potential research directions include the integration of explicit contour detection branches within the semi supervised architecture, the utilization of distance transform maps as auxiliary supervision signals, and the development of novel consistency loss functions that dynamically relax invariance constraints precisely at the high gradient regions indicative of structural edges [15]. By explicitly constraining the semi supervised learning process with geometric priors and spatial topology rules, the next generation of models can harness the massive potential of unlabeled aerial imagery without sacrificing the stringent spatial precision demanded by modern geospatial science.

References

1. Dong, C.; Nemet, G.; Gao, X.; Barbose, G.; Sigrin, B.; O'Shaughnessy, E. Machine Learning Reduces Soft Costs for Residential Solar Photovoltaics. Sci. Rep. 2023, 13, 7213.

2. Guo, S.; Shen, Q.; Cao, H.-h.; Li, P.; Zhang, K.; Xie, H.; Yan, G.; Mao, C. Study on quality of standard decoction of Coptidis Rhizoma based on traditional decoction process. China J. Chin. Mater. Medica 2019, 44, 3985–3993.

3. Mokhtar, S.S.M.; Salimon, M.G. SMEs' Adoption of Artificial Intelligence-Chatbots for Marketing Communication: A Conceptual Framework for an Emerging Economy. In Marketing Communications and Brand Development in Emerging Markets Volume II: Insights for a Changing World; Adeola, O., Hinson, R.E., Sakkthivel, A.M., Eds.; Palgrave Macmillan: Cham, Switzerland, 2022.

4. Boubaker, O. MPPT Techniques for Photovoltaic Systems: A Systematic Review in Current Trends and Recent Advances in Artificial Intelligence. Discov. Energy 2023, 3, 9.

5. Fan, X.; Zhang, K.; Wang, S.; Qi, Y.; Dai, G.; Lu, T.; Mao, C. Discrimination between raw and ginger juice processed Fructus Gardeniae based on UHPLC-Q-TOF-MS and Heracles NEO ultra-fast gas phase electronic nose. Phytochem. Anal. 2025, 36, 377–393.

6. DUTH. 2025. Code of Conduct for Legal Professionals in the Use of AI. Deliverable D3.2. FAIR-PReSONS. Fair Predictions of Gender-Sensitive Recidivism GA No: 101160473. Available online: https://fair-presons.aegean.gr/wp-content/uploads/2025/12/FAIR-PReSONS-D3.2-EN-final.pdf (accessed on 3 February 2026).

7. Kar, A.K.; Kushwaha, A.K. Facilitators and Barriers of Artificial Intelligence Adoption in Business—Insights from Opinions Using Big Data Analytics. Inf. Syst. Front. 2023, 25, 1351–1374.

8. Al-Sukker, A.; Ross, D.; Abdel-Qader, W.; Al-Akra, M. External auditor reliance on the work of the internal audit function in Jordanian listed companies. Int. J. Audit. 2018, 22, 317–328.

9. Albayrak Unal, O.; Erkayman, B.; Usanmaz, B. Applications of artificial intelligence in inventory management: A systematic review of the literature. Arch. Comput. Methods Eng. 2023, 30, 2605–2625.

10. Gao, J.; Wang, S.; Erdogdu, F.; Marra, F.; Sarghini, F.; Chen, L. Machine learning-driven innovations in food processing: A systematic review of applications, challenges, and future developments. Trends Food Sci. Technol. 2025, 164, 105204.

11. Wei, S.; Wei, L.; Xie, B.; Li, J.; Lyu, J.; Wang, S.; Khan, M.A.; Xiao, X.; Yu, J. Characterization of volatile profile from different coriander (Coriandrum sativum L.) varieties via HS-SPME/GC–MS combined with E-nose analyzed by chemometrics. Food Chem. 2024, 457, 140128.

12. Fei, X.; Ma, L.; Wei, A.; Shi, Q. Regional quality differences in Zanthoxylum bungeanum using sensor-based analysis and climate factors. Food Chem. 2025, 490, 145135.

13. Joshi, K.D.; Han, S.; Datta, A.; Chi, L. Innovation Through Boundary Spanning: The Role of IT in Enabling Knowledge Flows Across Technological and Geographical Boundaries. Int. J. Knowl. Manag. 2017, 13, 90–110.

14. Prasshanth, C.V.; Narayanan, S.B.; Sridharan, N.V.; Vaithiyanathan, S. Fault Detection in Photovoltaic Systems Using Unmanned Aerial Vehicle-Captured Images and Rough Set Theory. Sol. Energy 2025, 290, 113348.

15. Maka, A.O.M.; Alabid, J.M. Solar Energy Technology and Its Roles in Sustainable Development. Clean Energy 2022, 6, 476–483.