

Remote Sensing Cannabis Cultivation in Four Northern California Centers of Production with Faster-RCNN Object Detection

Seth LaRosa

Department of Environmental Science, Policy, and Management
University of California, Berkeley

ABSTRACT

Cannabis cultivation has long been present in Northern California, yet unlike conventional agricultural commodities, its spatial extent remains largely unknown. Proposition 64 (2016) legalized recreational cannabis while granting California's 58 counties authority to permit or prohibit cultivation, creating varied regulatory environments across the state. The introduction of this long-prohibited crop into legal markets – and the persistence of its unauthorized cultivation – creates the need for baseline spatial data to inform environmental policy and establish effective regulatory frameworks, yet detection of cannabis cultivation has been historically difficult. We applied Faster R-CNN deep learning models to 60cm NAIP aerial imagery across four Northern California “production centers” (i.e. enduring, dense locales of cultivation prior to legalization) from 2016-2022 to detect greenhouse and outdoor cultivation sites. Models achieved a combined 78% classification accuracy, identifying 13,554 unique sites across 20,306 site-year observations. We quantified cultivation area change and characterized biophysical context including water access, road proximity, protected areas, and terrain, which varied dramatically across regions. Total cultivation area in these production-rich locales increased 96% from 344 to 674 hectares, with greenhouse expansion accounting for 58% of total growth. Greenhouse cultivation expanded 4.5 fold from 55 to 248 hectares, with 3,231 conversions from outdoor to greenhouse infrastructure. The locale with a total ban on cannabis prohibition grew faster (+185% sites) with greenhouse share rising from 5% to 59%, while regions with permit programs grew slower (+37% sites). Most cultivation (92%) occurred on parcels lacking cultivation permits. This study provides the first machine learning application for systematic cannabis detection and establishes methodology for monitoring land use change in agricultural sectors lacking conventional data systems.

Keywords: cannabis agriculture, machine learning, RCNN, object detection, land use policy, Proposition 64, agricultural transformation

INTRODUCTION

Global cannabis markets are expanding. An estimated 209 million people use cannabis worldwide, making it the most widely consumed controlled substance (Zheng et al., 2021]. Legal cannabis markets reached \$35 billion globally in 2023 with projections of \$102 billion by 2030 (Grand View Research, 2024]. Twenty-four U.S. states have legalized adult-use cannabis. Canada, Uruguay, Mexico, Germany, and Thailand have established national legal frameworks. This expansion represents a fundamental shift in drug policy

affecting agricultural landscapes, regulatory systems, and environmental management across multiple continents.

California occupies a unique position as both an agricultural leader and one of the first states to implement recreational cannabis legalization. The state produces over 400 agricultural commodities valued at over \$60 billion annually (California Department of Food and Agriculture, 2024], this figure excludes cannabis entirely despite cultivation representing an estimated \$12.9 billion in production value (Department of Cannabis Control, 2024). Cannabis cultivation began expanding in Northern California during the 1960s counterculture movement. Medical legalization under Proposition 215 (1996) accelerated this growth (Polson, 2021]. By 2016, cannabis represented a large but essentially unmapped agricultural sector. Proposition 64's passage that year established the Adult Use of Marijuana Act (State of California, 2016; State of California Legislature, 2017], creating California as a critical test case for how legalization frameworks function in different jurisdictions.

Prohibition excluded cannabis from conventional agricultural monitoring. Federal illegality meant no census data, no USDA surveys, no systematic land use tracking. This created, at a micro-geographic scale, 'centers of exception' i.e. established production regions operating outside formal regulatory frameworks during prohibition times. These centers emerged as hotbeds of cultivation activity, but their actual scale and extent remained shrouded in conjecture and a lack of data. The scapegoating of environmental harm attributed to cannabis cultivation (and specifically to these locales) was generally predicated on anecdote rather than empirical measurement, due to the inability of research to access and assess cultivation sites (Carah et al 2015; Gianotti et al 2017).

Environmental concerns emerged as key justifications for regulatory oversight. Studies documented water diversions threatening salmonids (Bauer et al., 2015), deforestation and habitat fragmentation (Wang et al., 2017), pesticide contamination (California Department of Fish and Wildlife, 2019), and sediment loading to waterways (California State Water Resources Control Board, 2017). These impacts received intense scrutiny despite cannabis cultivation occupying far smaller acreage than timber extraction or conventional agriculture. The disproportionate attention reflected legal stigma and policy urgency rather than purely scientific assessment of relative environmental harm (Dillis et al., 2021a; Dillis et al., 2021b; Polson, 2019).

Proposition 64 preserved local control mechanisms, granting California's 58 counties authority to permit, regulate, or prohibit cannabis cultivation within their jurisdictions (State of California, 2016; State of California Legislature, 2017]. Local control gave municipalities authority which created a patchwork of regulatory environments; some counties established permitting programs while others enacted total bans. The resulting heterogeneity generated a natural quasi-experiment for examining how cultivation practices respond to regulatory policy. Assessing these policy effects requires basic spatial data that has not been collected due to cannabis' prohibited past.

Remote Sensing of Cannabis Sativa

Remote sensing offered potential solutions but faced technical challenges. Early spectral approaches sought unique absorption signatures for Cannabis sativa. Daughtry and Walthall (1998] identified distinctive reflectance patterns at 550nm, 670nm, and 720nm

wavelengths under controlled greenhouse conditions. Subsequent field testing revealed these signatures proved insufficiently stable across varied environmental conditions and growth stages to enable reliable discrimination (Walthall et al., 2006; Bicakli et al., 2022]. Cannabis spectral characteristics overlap extensively with other herbaceous vegetation. Additionally, cultivation can occur across complex mixed canopy environments, further limiting the efficacy of spectral detection.

Foundational spatial quantification of cannabis cultivation in Northern California relied on manual digitization. Butsic and Brenner (2016) established baseline data for Humboldt County through visual detection, requiring intensive human labor reviewing satellite imagery. Butsic et al. (2018) extended this approach across the Emerald Triangle, documenting cultivation as environmental threat. Manual methods provided accurate site identification but remained labor-intensive and difficult to scale across large geographies or repeated temporal observations. Visual detection remains valid and highly accurate but does not meet the spatial and temporal monitoring demands of this widespread land use.

Recent machine learning advances showed improved detection capabilities. Sujud et al. (Sujud et al., 2021) achieved high classification accuracy using random forest classifiers applied to PlanetScope satellite imagery (3.7m resolution) in illegal Lebanese cannabis parcels with government-provided ground truth data. Irawadi et al. (Irawadi et al., 2025) used Sentinel-2 satellite data (10m resolution) with multitemporal phenology statistics to detect small illicit plantations in Indonesian forested mountains, using law enforcement eradication records as training data. Both approaches relied on spectral or phenological characteristics of cannabis plants distinguishing them from surrounding vegetation. While these studies yielded highly precise results in their respective study areas, we sought methods not dependent on plant spectral signatures or growth stage patterns, given the variable conditions, diverse terrain and infrastructure types found in Northern California cannabis production.

Ferreira et al. (2019) demonstrated CNN ensemble approaches for detecting cannabis cultivation in Brazil using IKONOS satellite imagery (3.2m resolution), combining near-infrared spectral analysis with false-color representations to achieve 92% model accuracy. Their work showed promise for infrastructure-focused detection but operated at medium resolution capturing mature outdoor plantations. High-resolution aerial imagery combined with object detection models targeting cultivation infrastructure; greenhouses, hoop houses, cleared plots, offered detection signals independent of crop spectral characteristics or phenological timing.

This infrastructure-focused approach confirmed the possibility of detecting cannabis cultivation sites based on the presence of typical cultivation landscape features like greenhouses, hoop houses, access roads, water tanks and irrigation infrastructure -- all spatial organization patterns detectable from aerial imagery. These physical features remain visible regardless of plant phenology or spectral variation. Faster R-CNN models perform well at detecting such structured objects in complex scenes. The approach effectively transcends spectral limitations by recognizing that cultivation infrastructure, not crop reflectance, provides a more temporally reliable detection signal.

The introduction of a "new" agricultural commodity in California's agricultural economy created opportunity to examine regulation, environmental monitoring, and land use planning in this emerging sector. Understanding this transformation requires developing

detection methodology that effectively captures the extent of California's cannabis landscape and the changes taking place in these spaces.

We address three objectives:

- 1) *Validate Faster R-CNN performance for cannabis cultivation detection across varied terrain using publicly available imagery.*
- 2) *Quantify temporal changes in cultivation extent and the prevalence of greenhouse versus outdoor production from 2016 through 2022.*
- 3) *Assess whether transformation patterns differed across regulatory contexts and if local policy influenced cultivation during legalization's implementation.*

METHODS

Study Areas

We analyzed four established cannabis production centers in Northern California, anonymized as Regions A through D to address privacy and endangerment concerns. These regions were selected as "centers of production" -- areas with documented cultivation history predating Proposition 64's passage. Production centers are not typical of cannabis cultivation, per se, but rather, they provide cannabis-concentrated catchments that are ideal for refinement of remote sensing methodologies. Study areas were defined using 2-kilometer buffers around each production center to capture cultivation activity while maintaining spatial independence between regions.

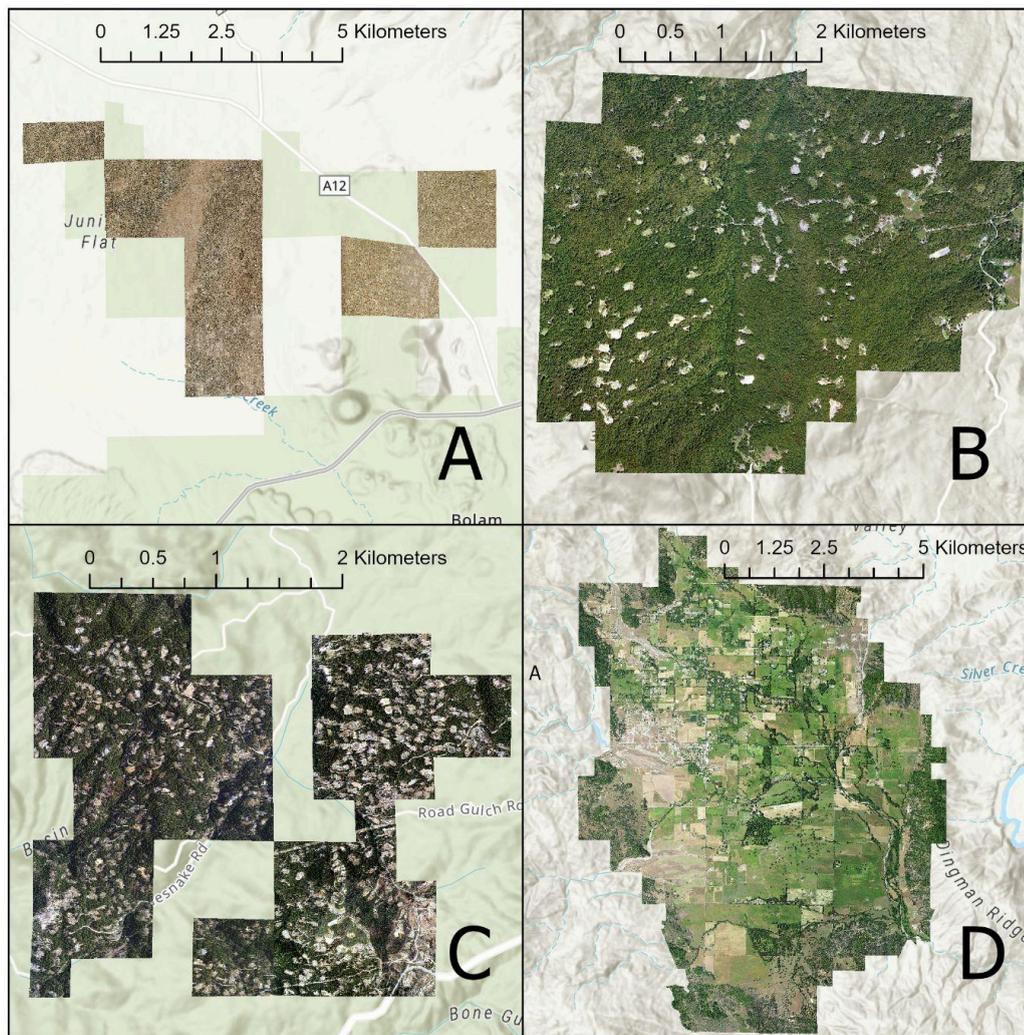


Figure 1. Study area locations in Northern California. Four cannabis production centers (Regions A-D, anonymized) selected as established cultivation areas with documented pre-legalization activity. Regions span diverse biophysical conditions including arid valley floor and rural subdivisions (Region A), forested mountainous terrain (Regions B and C), and mixed agricultural valley with surrounding mountains (Region D). Study area boundaries defined using 2-kilometer buffers around production centers. Three regions operated under county permit programs following Proposition 64 passage, while one region (Region A) enacted total cannabis prohibition.

The regions provided variation in regulatory response to Proposition 64. Three regions implemented county-level permitting programs following legalization, though program access and limitations differed across counties. One region enacted total cannabis prohibition, banning all commercial cultivation within county jurisdiction (Getz et al., 2024). All regions contained established cultivation infrastructure prior to 2016, making them comparable in terms of existing production capacity.

Regions spanned diverse biophysical conditions including elevation ranges from valley floors to mountainous terrain exceeding 1,500 meters, climates from Mediterranean to semi-arid to montane, and vegetation ranging from oak woodland to mixed conifer forests. Terrain complexity varied from relatively flat agricultural valleys to steep watersheds. This

environmental heterogeneity provided rigorous testing conditions to evaluate model performance.

Image Acquisition and Processing

We obtained National Agriculture Imagery Program (NAIP) aerial imagery at 60cm spatial resolution for years 2016, 2018, 2020, and 2022. NAIP acquires four-band imagery (red, green, blue, near-infrared) during peak growing season, typically June through September in California. This temporal coverage captured pre-legalization baseline (2016), implementation period (2018-2020), and post-legalization conditions (2022). The 60cm resolution enabled discrimination of individual greenhouse structures and outdoor cultivation plots.

Imagery preprocessing applied standard radiometric and geometric corrections. We mosaicked individual NAIP tiles covering each region using ArcGIS Pro. No atmospheric correction was applied as detection focused on spatial patterns and infrastructure rather than spectral values. We converted imagery to 8-bit unsigned integer format and created three-band RGB composites for model input, as Faster R-CNN architectures expect standard image formats rather than multispectral data cubes.

Deep Learning Model Development

We developed separate Faster R-CNN models for greenhouse and outdoor cultivation detection using ArcGIS Pro's deep learning framework. Models employed ResNet-50 convolutional neural network backbones pre-trained on ImageNet, providing robust feature extraction for complex scene understanding. Training data consisted of 1,411 manually digitized greenhouse polygons and 570 outdoor cultivation polygons distributed across all study regions and years. Training samples captured variation in structure types (greenhouse, hoop house, mixed), sizes (small personal plots to large commercial operations), and environmental contexts (forested, agricultural, mixed development).

We generated training chips using ArcGIS Pro's Export Training Data for Deep Learning tool with RCNN_Masks format. Chip size was set to 512×512 pixels with 25% overlap between adjacent chips to ensure complete coverage of training features. We applied standard data augmentation including horizontal and vertical flips and 90-degree rotations to increase effective training sample size and model robustness. Models trained for 50 epochs with learning rate 0.0001, batch size 4, and validation percentage 10%. Learning rate proved critical for model performance - initial testing at 0.001 produced significantly worse results, while 0.0001 enabled convergence to optimal detection accuracy.

Model deployment applied trained models to complete NAIP imagery for each region-year combination, generating polygon features representing detected cultivation sites with associated confidence scores. Initial detection outputs contained false positives requiring systematic quality control. We implemented area-based filtering (20 to 5,000 square meters) and minimum confidence threshold (0.55) based on iterative manual validation of detection subsets. These filters achieved ~83% precision while maintaining acceptable recall rates. The area range excluded obviously spurious detections while retaining legitimate cultivation sites ranging from personal grows to commercial operations.

Training data consisted exclusively of positive examples without negative samples teaching models to reject non-cultivation features. The outdoor model occasionally misclassified scrubby or brushy landscapes as cultivation sites, mistaking natural vegetation patterns for the distinct clearing and planting configurations of outdoor grows. This confusion likely contributed to the outdoor model's lower precision score (AP 0.635) compared to greenhouse detection (AP 0.712). Ferreira et al. (2019) demonstrated that incorporating negative training samples e.g. natural vegetation, non-cannabis clearings, and similar landscape features, improves model discrimination and reduces such false positives, a refinement that would strengthen future implementations.



Figure 2. ML detection results from site C in 2022, showing the RCNN model's ability to identify outdoor cultivation sites (purple overlays) across the forested landscape.

Site Tracking and Conversion Analysis

We tracked individual cultivation sites across years using spatial overlap analysis. Detection polygons in successive years were assessed for spatial intersection. When polygons from consecutive years exhibited overlap area equal to or exceeding 50% of the smaller polygon's area, they received identical site identifiers, indicating persistence at the same location. This overlap threshold accommodated minor variations in polygon boundaries between NAIP acquisitions while preventing false linkages between distinct nearby sites. The 50% criterion ensured that tracked sites represented genuine persistence rather than incidental proximity of separate operations. This tracking generated 13,554 unique site identifiers across 20,306 total site-year observations.

Physical conversion from outdoor to greenhouse cultivation represented substantial infrastructure investment indicating long-term commitment to site development. We identified conversions through spatial intersection analysis combined with the same overlap criteria. Sites detected as outdoor cultivation in any year from 2016-2020 that exhibited 50% or greater spatial overlap with greenhouse cultivation polygons detected in 2022 received conversion classification. This conservative definition required both outdoor infrastructure

presence and greenhouse infrastructure replacement at the same spatial location, documenting actual physical transformation rather than simple abandonment and new site development.

Environmental Characterization

We integrated detection outputs with publicly available environmental datasets to characterize biophysical and spatial features of cultivation sites. National Hydrography Dataset (NHD) flowlines enabled assessment of proximity to surface water. TIGER/Line road networks documented road accessibility. California Protected Areas Database (CPAD) identified cultivation proximity to conservation lands. USGS 3DEP elevation data generated slope derivatives.

Spatial proximity was calculated using ArcGIS Pro Generate Near Table with planar distance measurement in Web Mercator projection. For river proximity (NHD flowlines), we established thresholds at 50m (immediate riparian zone), 100m (close proximity), and 500m (watershed-scale concern). Road proximity (TIGER/Line data) used thresholds at 20m, 50m, 100m, and 500m. Protected area analysis combined spatial intersection identifying sites on conservation land with proximity thresholds at 100m and 500m from boundaries. Terrain slope was calculated from digital elevation models as percent rise, with sites classified as high erosion risk if mean slope exceeded 30% or if 50% or more of site area occurred on slopes greater than or equal to 30%.

These proximity analyses characterize the potential to assess environmental exposure rather than quantifying actual impacts. A site near a stream may or may not divert water. A site on steep slopes may or may not generate erosion. The characterization establishes baseline spatial relationships enabling future targeted assessment, but environmental outcomes depend on site-specific management practices.

Permit Enrollment Analysis

We assessed regulatory compliance by analyzing cultivation occurring on parcels enrolled in California Department of Cannabis Control (DCC) state licensing databases. The DCC parcel dataset (2022 snapshot) included both fully licensed and provisional license parcels. We spatially joined this dataset to 2022 detections, classifying sites as permitted when their centroids fell within enrolled parcel boundaries. Analysis was limited to 2022 detections.

Regulatory Context Comparison

We conducted quasi-experimental comparison examining cultivation transformation patterns across regulatory regimes. Analysis grouped regions into permit program category (n=3) versus total ban category (n=1). We aggregated detection data to region-year level, calculating total site counts, cultivation area (hectares), and greenhouse share (percentage) for each region and year. Comparative analysis examined growth trajectories and greenhouse adoption patterns from 2016 baseline through 2022.

Accuracy Assessment

Model performance was evaluated using multiple metrics. Classification accuracy reached 78% across validation samples. End-to-end Mean Detection Score (combination of

precision, recall, and spatial overlap) achieved 0.71. The greenhouse model achieved Average Precision (AP) of 0.712 while the outdoor model achieved AP of 0.635. Manual validation on randomly selected detection samples confirmed 85% precision for greenhouse detection and 78% precision for outdoor detection after quality control filtering.

RESULTS

Detection Performance

Faster R-CNN models achieved robust detection performance across the 4 study regions diverse terrain. Models successfully identified Greenhouse and Outdoor cultivation sites in forested environments, agricultural valleys, arid rangeland, and mixed development contexts. Performance remained consistent across years despite variations in imagery acquisition timing and phenology. The infrastructure-focused detection approach proved effective regardless of environmental conditions or crop growth stage.

Cultivation Extent and Transformation

Detection identified 13,554 unique cultivation sites across 20,306 site-year observations from 2016-2022. Total cultivation area increased 96% from 344 hectares in 2016 to 674 hectares in 2022 (Table 1). Site counts grew from 3,398 in 2016 to 6,513 in 2022. The most rapid expansion occurred during 2018-2020, with cultivation area increasing 69% during this two-year period during the initial implementation of recreational cannabis policies. These phenomena may reflect localized activity rather than overall regional trends; Dillis et al. (in progress) CannaVision modeling indicated cultivation may in fact exhibit opposing trends at larger scales with the Emerald Triangle and other legacy counties exhibiting some of the most consistent decreases in cultivation

Cannabis Cultivation Transformation in Northern California, 2016-2022

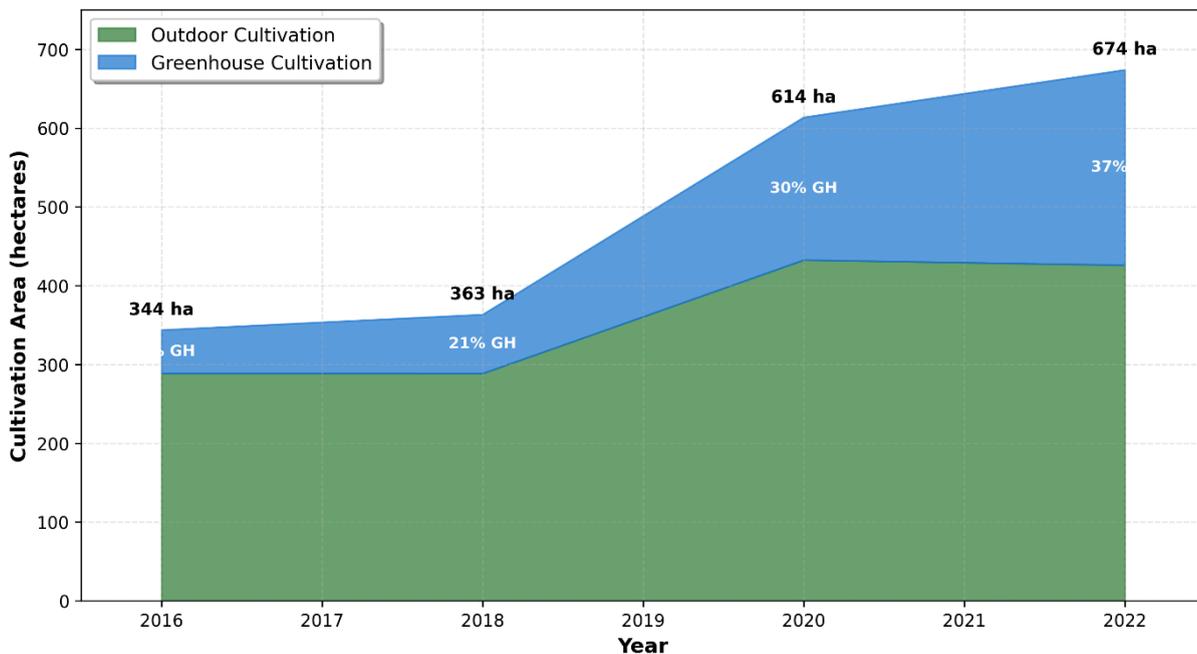


Figure 3. Cultivation area transformation by type, 2016-2022. Total cannabis cultivation area increased 96% from 344 hectares in 2016 to 674 hectares in 2022. Greenhouse cultivation expanded 4.5-fold from 55 hectares (16% of total area) in 2016 to 248 hectares (37% of total area) in 2022. Outdoor cultivation increased 48% from 289 hectares to 426 hectares over the same period. The most rapid expansion occurred during 2018-2020 with total area increasing 69% during this two-year period.

Greenhouse cultivation expanded dramatically while outdoor cultivation remained relatively stable. Greenhouse area increased 349% from 55 hectares (16.1% of total) in 2016 to 248 hectares (36.8% of total) in 2022. Outdoor cultivation area grew modestly with a 48% increase from 289 hectares to 426 hectares. However total share of outdoor cultivation declined from 83.9% to 63.2% of total cultivation area. Greenhouse share of total cultivation area more than doubled during the six-year period. This transformation represented a fundamental shift in cultivation approaches rather than simple expansion of existing practices.

We identified 3,231 sites that converted from outdoor to greenhouse cultivation between 2016 and 2022. These sites showed outdoor cultivation in at least one year during 2016-2020, then exhibited greenhouse infrastructure at the same location in 2022. These conversions represented permanent infrastructure development requiring construction materials, labor, and capital investment. Getis-Ord G_i^* hot spot analysis on conversion locations identified 200 statistically significant spatial clusters, providing evidence for peer effects and neighborhood diffusion of greenhouse adoption decisions.

Year	GH Sites	GH Area (ha)	GH %	OD Sites	OD Area (ha)	OD %	Total Sites	Total Area (ha)
2016	1,377	55.25	16.1	2,021	288.61	83.9	3,398	343.86
2018	1,632	74.84	20.6	1,919	288.59	79.4	3,551	363.43
2020	3,527	181.18	29.5	3,040	432.55	70.5	6,567	613.73
2022	4,150	247.94	36.8	2,363	426.05	63.2	6,513	673.99

Table 1. Cultivation area and site counts by year and type

Regulatory Context Patterns

Comparative regional analysis revealed divergent transformation patterns across regulatory frameworks (Table 2). Regions with permit programs experienced 37% growth in site counts from 2016 to 2022 (849 to 1,164 sites average). The region with total cannabis prohibition experienced 185% growth in site counts during the same period (763 to 2,172 sites average). Greenhouse share of total area in permit regions remained relatively stable, declining slightly from 41% to 32%. Greenhouse share in the ban region increased dramatically from 5% to 59%.

The pattern of increased cannabis production, regardless of regulatory framework, contradict expectations that permitting policies would facilitate cultivation expansion while prohibition would suppress activity. It is important to note that the pattern detected in these historic centers of production may reflect localized dynamics rather than countywide trends.

The fastest growth and most dramatic methodological transformation occurred in the region legally prohibiting commercial cannabis cultivation. Permit regions exhibited more modest growth despite regulatory pathways for legal operation. This suggests other forces (market, sociocultural) drove cultivation behavior more strongly than local bans or permitting programs. Program continuity is likely another relevant factor affecting enrollment in permitting programs within the study area. Among the three permit regions, one maintained regulatory continuity through the study period, while two experienced program suspensions and cultivation ordinance revisions, rendering compliance more onerous, or impossible, in some locales.

Participation in state licensing programs remained limited across all locals. Only 7.8% of detected cultivation sites in 2022 occurred on parcels enrolled in California Department of Cannabis Control (DCC) databases. The remaining 92.2% of detections operated on sites without documented state licensing.

Metric	Permit 2016	Permit 2018-22	Change	Ban 2016	Ban 2018-22	Change
Total Sites	849	1,164	+37%	763	2,172	+185%
Total Area (ha)	80.7	123.5	+53%	101.7	179.8	+77%
GH Share (%)	40.5	32.3	-8.2	5.2	34.7	+29.4
GH Sites	392	602	+54%	140	1,352	+866%
OD Sites	457	562	+23%	623	820	+32%

Table 2. Cultivation transformation by regulatory context

Environmental Characteristics

Environmental feature analysis revealed substantial regional variation in cultivation site characteristics (Table 3). Riparian proximity ranged from 15.8% of sites in one region to 99.9% in another. One region exhibited near-universal riparian proximity (99.7-99.9% within 500m of rivers), while another region showed minimal river proximity (15.5-15.8% within 500m), suggesting fundamentally different water access strategies -- surface water diversion versus groundwater dependence or water transport – carrying their own associated environmental impacts.

Road accessibility proved nearly universal. Overall, 98.8% of sites (20,056 of 20,306) occurred within 500m of roads, 64.4% within 100m, and only 0.6% (117 sites) beyond 500m. Two regions achieved 100% road accessibility within 500m.

Protected area proximity varied substantially by region. Direct trespass on protected land was minimal, with only 0.3% of sites (62 sites) occurring on those lands. However, 31.7% of sites occurred within 500m of protected area boundaries. One region showed extreme boundary proximity (63.6-69.2% of sites within 500m), while another region showed virtually no protected area interface (0.0-0.1%).

Terrain analysis revealed contrasting slope exposure patterns. Overall, 8.0% of sites occurred on steep terrain (30% slope or greater). One region's outdoor cultivation showed extreme slope exposure (35.0% of sites on steep terrain), while another region's cultivation occurred on nearly flat terrain (0.1% steep for both greenhouse and outdoor). The region combining near-universal riparian proximity with extreme slope exposure creates conditions where sediment transport to waterways becomes a relevant concern.

Region	River (% <500m)	Road (% <500m)	CPAD (% <500m)	Steep (%)
Region A	15.5-15.8	100.0	42.2-43.1	0.1
Region B	92.1-96.1	97.2-98.3	4.8	13.7-31.2
Region C	99.7-99.9	99.9-100.0	63.6-69.2	11.0-35.0
Region D	63.7-68.4	95.2-97.9	0.0-0.1	1.0-2.8

Table 3. Environmental characteristics by region (ranges represent variation between greenhouse and outdoor cultivation types)

Although comprehensive environmental impact assessment lies beyond this study's scope, the spatial characterization demonstrates what detection methodology enables. Integration of cultivation locations with publicly available environmental datasets reveals substantial regional variation in biophysical exposure that remained invisible during prohibition. This spatial baseline provides the essential foundation for targeted environmental monitoring, policy development, and future impact quantification studies connecting cultivation presence to measurable ecological outcomes.

DISCUSSION

Cannabis legalization created monitoring demands for an agricultural commodity that had been previously difficult to detect due to prohibition. This study demonstrates that machine learning detection enables systematic observation where spectral approaches fell short and manual methods could not adequately scale. The documented transformation shows cultivation expanded across all regulatory contexts - prohibition regions and permit regions both experienced growth in extent and greenhouse adoption, though at different rates. This pattern challenges the assumption that local control - permits vs. bans - is an effective strategy to govern cultivation location and scale. Most cultivation remains unlicensed despite legalization establishing regulatory pathways, indicating substantial implementation challenges. Detection methodology provides the measurement foundation for evidence-based policy in emerging agricultural sectors, though readily accessible technology raises governance questions about surveillance, privacy, and enforcement equity.

RCNN Validation for Cannabis Detection

Our first objective tested whether RCNN models could reliably detect cannabis cultivation across varied terrain using publicly available imagery. The 78% classification

accuracy and consistent performance across diverse geographic contexts validates infrastructure-focused detection as effective methodology. This approach succeeds because it recognizes permanent and semi-permanent structures: greenhouses, hoop houses, cleared cultivation plots rather than relying on spectral signatures or crop phenology that vary with local conditions. Automated detection trades individual site precision for comprehensive regional coverage across multiple years. This tradeoff suits landscape-scale monitoring, regional trend analysis, and assessment at geographic extents matching cultivation's actual distribution.

Detecting cannabis in legalizing markets addresses a fundamental information gap. Basic spatial data (where cultivation occurs, how much exists, how it distributes across landscapes) did not exist during prohibition. This information gap persists in many locales and is only now coming into light as cannabis policies shift across California, the United States, and globally. Jurisdictions implementing legalization lack baseline knowledge about the agricultural sector they attempt to regulate. Canada, Uruguay, expanding European markets, and U.S. states all face this challenge: attempting to manage an agricultural commodity without knowing its spatial extent or distribution patterns. Detection methodology provides the measurement foundation enabling systematic observation where spatial existing data is lacking or nonexistent.

The methodology extends beyond cannabis to other unauthorized or difficult-to-monitor land uses: unpermitted development, illegal mining operations, informal settlements, or any activity creating distinctive physical structures visible in aerial imagery. The demonstrated transferability across Northern California's diverse terrain suggests applicability to other contexts requiring systematic detection of activities previously undocumented or deliberately concealed.

Boom and Cultivation Transformation

Our second objective quantified temporal changes in cultivation extent and the prevalence of greenhouse versus outdoor production from 2016 through 2022. Cultivation area nearly doubled while greenhouse share increased from 16% to 37% of total area. This dual transformation reflects both overall expansion and fundamental shift in production practices. The magnitude and timing suggest responses to market development following legalization rather than gradual evolution. Physical infrastructure conversion represents permanent capital investment requiring materials, labor, and financial commitment, potentially indicating growing licitness and normalization of cultivation (Polson & Bodwitch, 2021). The spatial clustering of outdoor-to-greenhouse conversion sites suggests information diffusion through social networks or favorable site characteristics concentrating development in particular locations (Butsic et al., 2017).

Detection enables environmental characterization previously impossible during prohibition. Pioneering cannabis research in Northern California established foundational work documenting cultivation in environmental contexts, examining similar variables including slope, proximity to waterways, roads, and protected areas to assess potential risks

(Butsic et al. 2017). Integration with publicly available datasets reveals substantial regional variation in cultivation site characteristics. Some regions exhibited near-universal riparian proximity suggesting surface water dependence, while others showed minimal river access indicating groundwater reliance or water transport. Road accessibility proved nearly universal, contradicting narratives of remote wilderness cultivation (Polson 2019; Meisel, 2024). Protected area trespass remained minimal despite substantial boundary clustering in certain regions. Terrain patterns ranged from predominantly flat cultivation to concentration on steep slopes, which creates potential erosion concerns.

These proximity patterns indicate potential environmental exposure but do not quantify actual impacts. A site near streams may or may not divert water depending on extraction practices and timing. A site on steep terrain may or may not generate sediment depending on ground cover and erosion control measures. The methodological shift toward greenhouse cultivation compounds assessment challenges. Energy consumption increases exponentially across cultivation regimes: outdoor cultivation requires minimal energy inputs, greenhouse operations demand substantially more for climate control and lighting supplementation, while indoor cultivation represents the most energy-intensive extreme (Mills, 2012,2025). The documented shift toward greenhouse infrastructure indicates increased energy and material requirements compared to outdoor cultivation. Greenhouse operations concentrate resource use at discrete sites compared to dispersed outdoor cultivation. Water use patterns differ. Energy consumption increases. Environmental management frameworks developed assuming predominantly outdoor cultivation face mismatches with actual production systems, requiring policy adaptation to address greenhouse-dominated landscapes.

Regulatory Context and Policy Effectiveness

Our third objective assessed whether transformation patterns differed across regulatory contexts. The observed pattern challenges straightforward expectations about policy effects. The region with total cannabis prohibition experienced faster growth and more dramatic greenhouse adoption than regions offering permit pathways. If permit availability facilitates cultivation expansion while prohibition suppresses activity, permit regions should exhibit greater growth. The opposite pattern occurred. This suggests local regulatory frameworks exerted limited influence on cultivation behavior during legalization implementation, though multiple mechanisms could produce these patterns.

Permit program complexity and costs may have created barriers discouraging participation without effectively suppressing cultivation (Bodwitch et al 2021; Polson et al 2023). Prohibition policies sometimes spurred innovation as growers invested in permanent infrastructure anticipating future policy change (Getz et al., 2024). The ban region's greenhouse expansion could reflect such strategic positioning. Alternatively, permit regions may have possessed less available land for expansion given existing cultivation density, while the ban region had greater growth capacity. Enforcement intensity likely varied across jurisdictions. Bans may have generated less active monitoring than permit programs that required compliance verification for licensed operations. Market dynamics driving

greenhouse adoption potentially operated independently of local policy given interstate and international demand patterns.

Most cultivation occurred on unpermitted parcels regardless of regulatory context. High costs and complex requirements have been identified as barriers to compliance that may have ballooned the ranks of unlicensed cultivators (Bodwitch et al., 2021). Other research in cannabis country has documented how enforcement-first approaches failed to eliminate cultivation while generating inequitable impacts to communities (Polson & Petersen-Rockney, 2019). In our analysis, only a small fraction of detected sites participated in state licensing despite legal pathways existing in permit regions. The lack of legally compliant operators indicates substantial implementation challenges in these areas of intense historical concentration of cannabis production. This finding suggests that regulatory frameworks may have created participation barriers exceeding perceived benefits for many growers through fees, inspections, and operational requirements (Polson et al., 2024).

Policy intervention through local control represents an effort to tailor cannabis governance to diverse communities across California. In these 4 historically established production centers, cultivation patterns did not align closely with policy distinctions designed to ban or permit cannabis activity. Instead, observed land use patterns appear to have been shaped by existing cultivation practices and broader social and market conditions that operated alongside local policy. Remote sensing alone cannot adequately illuminate regulatory effects in legacy cannabis production locales, this analysis suggests that effective policy in these microgeographies depends on how new policy frameworks interact with and consider pre-existing land use regimes.

Detection Technology and Privacy Considerations

This research reveals the double-edged nature of modern remote sensing. The same information enabling better policy and land use planning becomes a surveillance mechanism with implications for privacy and population protection and harm minimization. Identifying cultivation sites using publicly available imagery and commercial software represents a shift in what constitutes “observable behavior” in ethical review of research. The technology enables deeper incursion into how people live. It's a cautionary tale for cultivators and a tool for law enforcement to target communities with sophisticated surveillance. Geospatial analysis can reveal intimate details about daily routines, occupancy patterns, and private activities without traditional warrant requirements protecting physical searches. This raises questions about boundaries between legitimate regulatory oversight and invasive surveillance. Detecting large-scale, environmentally-impactful, commercial operations may serve valid public interests. But the same technology could easily target small-scale personal cultivation or other legal but private activities. The risk of discriminatory enforcement through algorithmic bias, where certain communities face disproportionate scrutiny, demands careful legal safeguards (Kalluri et al., 2023).

Limitations

Several limitations affect interpretation. The four study regions represent purposely selected and historically established production centers rather than a random sample of California's cultivation landscape. Patterns documented in these centers of exception may not generalize to emerging cultivation areas lacking prohibition-era production history or to regions with different regulatory environments, terrain characteristics, or market access. Geographic scope remains limited. The single control region (n=1 for ban category) prevents robust statistical inference about regulatory effects. Comparative analysis reveals patterns but cannot support causal claims about policy impacts.

Model development faced data constraints. Training samples totaled 1,411 greenhouse features but only 570 outdoor features, potentially affecting outdoor detection performance relative to greenhouse detection. We did not incorporate negative site training where models learn to reject non-cultivation features resembling target structures. Ferreira et al. (2019) demonstrated improved accuracy using negative samples. The binary classification as greenhouse or outdoor cultivation simplifies reality. Mixed operations, light deprivation structures, hoop houses, and hybrid approaches exist but receive assignment to dominant infrastructure type. Detection cannot distinguish licensed from unlicensed cultivation from imagery alone, requiring integration with parcel enrollment databases for regulatory compliance assessment.

Temporal coverage ended in 2022 due to federal NAIP program delays preventing 2024 data availability. The 60cm resolution misses very small cultivation operations below detection thresholds. Summer imagery acquisition may miss seasonal operations occurring outside peak growing season. These limitations constrain both detection completeness and temporal currency of findings.

CONCLUSION

This study demonstrates infrastructure-focused machine learning detection enables systematic cannabis cultivation monitoring across diverse environmental and regulatory contexts where spectral approaches fail and manual methods cannot scale. Detection documented substantial cultivation transformation following California's legalization, with area nearly doubling and greenhouse infrastructure becoming dominant production method. Transformation patterns suggest local regulatory frameworks exerted limited influence on cultivation behavior or environmental site selection during implementation period. Most cultivation remains unlicensed despite legal pathways existing, indicating substantial gaps between policy design and on-the-ground reality. The methodology transfers to other jurisdictions implementing cannabis legalization and potentially to other agricultural sectors lacking conventional monitoring systems. Future research could extend detection to additional California regions testing whether patterns documented in established production centers generalize to emerging cultivation areas. Critical questions remain about why local bans proved ineffective at reducing cultivation and what factors drive continued unlicensed cultivation in historical production centers, requiring methods beyond quantitative detection to more deeply understand on-the-ground cultivation patterns.

ACKNOWLEDGMENTS

This research was funded through the California Department of Cannabis Control (Grant # 93284; Licensed and Unlicensed Production Across Ban and Permit Jurisdictions). The author gratefully acknowledges support from the Land Use Change Lab at UC Berkeley's Rausser College of Natural Resources under the direction of PhD advisor Van Butsic. Undergraduate research assistants Olivia, Selena, McCree, Mia, and Gabriella provided invaluable assistance digitizing training data for model development. The Geospatial Innovation Facility at UC Berkeley provided essential technical support and access to ArcGIS Pro software. The author thanks anonymous reviewers for constructive feedback improving the manuscript.

Funding Statement

This project was funded by the California Department of Cannabis Control using monies collected through the California Cannabis Tax Fund. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the official views or policies of the State of California.

REFERENCES

- Bauer, S., Olson, J., Cockrill, A., van Hattem, M., Miller, L., Tauzer, M., & Leppig, G. (2015). Impacts of surface water diversions for marijuana cultivation on aquatic habitat in four northwestern California watersheds. *PLoS One*, 10(3), e0120016.
- Bicakli, F., Kaplan, G., & Alqasemi, A. S. (2022). Cannabis sativa L. spectral discrimination and classification using satellite imagery and machine learning. *Agriculture*, 12(6), 842.
- Bodwitch, H., Polson, M., Biber, E., Hickey, G. M., & Butsic, V. (2021). Why comply? Farmer motivations and barriers in cannabis agriculture. *Journal of Rural Studies*, 86, 155–170.
- Butsic, V., & Brenner, J. C. (2016). Cannabis (*Cannabis sativa* or *C. indica*) agriculture and the environment: a systematic, spatially-explicit survey and potential impacts. *Environmental Research Letters*, 11(4), 044023.
- Butsic, V., Lewis, D. J., Radeloff, V. C., Baumann, M., & Kuemmerle, T. (2017). Quasi-experimental methods enable stronger inferences from observational data in ecology. *Basic and Applied Ecology*, 19(1), 1–10.
- Butsic, V., Schwab, B., Baumann, M., & Brenner, J. C. (2017). Inside the Emerald Triangle: Modeling the placement and size of Cannabis production in Humboldt County, CA USA. *Ecological Economics*, 142, 70–80.
- Butsic, V., Short Gianotti, A. G., Lewis, D. J., Brenner, J. C., & Lewis, J. S. (2018). The emergence of cannabis agriculture frontiers as environmental threats. *Environmental Research Letters*, 13(12), 124017.
- California Department of Fish and Wildlife. (2019). A review of the potential impacts of cannabis cultivation on fish and wildlife resources. Report to California Fish and Game Commission. <https://nrm.dfg.ca.gov/FileHandler.ashx?DocumentID=160552>
- California Department of Food and Agriculture. (2024). California agricultural statistics review 2023-2024. Sacramento, CA. <https://www.cdfa.ca.gov/Statistics/>
- California State Water Resources Control Board. (2017). Cannabis cultivation policy: Principles and guidelines for cannabis cultivation. https://www.waterboards.ca.gov/water_issues/programs/cannabis/
- Carah, J. K., Howard, J. K., Thompson, S. E., Short Gianotti, A. G., Bauer, S. D., Carlson, S. M., Dralle, D. N., Gabriel, M. W., Hulette, L. L., Johnson, B. J., Knight, C. A., Kupferberg, S. J., Martin, S. L., Naylor, R. L., & Power, M. E. (2015). High time for conservation: adding the environment to the debate on marijuana liberalization. *Bioscience*, 65(8), 822–829.
- Daughtry, C. S. T., & Walthall, C. L. (1998). Spectral discrimination of Cannabis sativa L. leaves and canopies. *Remote Sensing of Environment*, 64(2), 192–201.

Department of Cannabis Control. (2024). California cannabis market outlook 2024 (Report prepared by ERA Economics, LLC).

<https://cannabis.ca.gov/2025/03/california-cannabis-market-outlook/>

Dillis, C., Biber, E., Bodwitch, H., Butsic, V., Carah, J., Parker-Shames, P., Polson, M., & Grantham, T. (2021). Shifting geographies of legal cannabis production in California. *Land Use Policy*, 105, 105369.

Dillis, C., Butsic, V., Carah, J. K., Zipper, S. C., & Grantham, T. (2021). Cannabis farms in California rely on wells outside of regulated groundwater basins. *Environmental Research Communications*, 3(8), 081002. <https://doi.org/10.1088/2515-7620/ac1124>

Dillis, C., Polson, M., Bodwitch, H., Carah, J., Power, M. E., & Sayre, N. F. (2021). Industrializing cannabis? Socio-ecological implications of legalization and regulation in California. In J. Meisel & D. Corva (Eds.), *The Routledge Handbook of Post-Prohibition Cannabis Research* (pp. 221-230). Routledge.

Dillis, C., et al. (forthcoming 2026). [Cannavision modeling - larger scale cultivation trends].

Ferreira, A., Felipussi, S. C., Pires, R., Avila, S., Santos, G., Lambert, J., Huang, J., & Rocha, A. (2019). Eyes in the skies: A data-driven fusion approach to identifying drug crops from remote sensing images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(12), 4773–4786.

Gambardella, C., Parente, R., Ciambrone, A., & Casbarra, M. (2021). A Principal Components Analysis-Based Method for the Detection of Cannabis Plants Using Representation Data by Remote Sensing. *Data*, 6(10), 108.

Getz, C., Petersen-Rockney, M., & Polson, M. (2024). Cannabis Bans, Local Control, and the Effects and Efficacy of Proposition 64 (SSRN Scholarly Paper No. 4819697). Social Science Research Network.

Grand View Research. (2024). Legal marijuana market size, share & trends analysis report by application (medical, adult use), by product type (flower, oils and tinctures), by region, and segment forecasts, 2024-2030 (Report No. GVR-2-68038-152-8).

<https://www.grandviewresearch.com/industry-analysis/legal-marijuana-market>

Houmi, M., Mohamadi, B., & Balz, T. (2018). A hyperspectral based method to detect cannabis plantation in inaccessible areas. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42(3), 547–551.

Irawadi, D., Mauritsius, T., Kushardono, D., Budhiman, S., Diwyacitta, K., Adhitama, B. S., Ayubi, F. A., Maftukhaturrizqoh, O., & Supriyani, I. S. (2025). Smart Detection of Illicit Cannabis Plantations Using Remote Sensing Technology And Machine Learning. *Geography, Environment, Sustainability*, 130-138.

- Kalluri, P. R., Agnew, W., Cheng, M., Owens, K., Soldaini, L., & Birhane, A. (2023). The surveillance AI pipeline. arXiv preprint arXiv:2309.15084.
- Lisita, A., Sano, E. E., & Durieux, L. (2013). Identifying potential areas of Cannabis sativa plantations using object-based image analysis of SPOT-5 satellite data. *International Journal of Remote Sensing*, 34(15), 5409–5428.
- Malabadi, R. B., Mammadova, S. S., Kolkar, K. P., MR, S., Chalannavar, R. K., & Coronado, K. V. C. (2024). Cannabis sativa: A therapeutic medicinal plant-global marketing updates. *World Journal of Biology Pharmacy and Health Sciences*, 17(2), 170-183.
- Mattiuzzi, M., Bussink, C., & Bauer, T. (2014). Analysing phenological characteristics extracted from Landsat NDVI time series to identify suitable image acquisition dates for cannabis mapping in Afghanistan. *PFG Photogrammetrie Fernerkundung Geoinformation*, 5, 383–392.
- Meisel, J. (2017). Hidden in Plain Sight: Cannabis Cultivation in the Emerald Triangle. *California Geographer*, 56.
- Mills, E. (2012). The carbon footprint of indoor Cannabis production. *Energy Policy*, 46, 58–67.
- Mills, E. (2025). Energy-intensive indoor cultivation drives the cannabis industry's expanding carbon footprint. *One Earth*, 8(3).
- Polson, M. (2013). Land and law in marijuana country: Clean capital, dirty money, and the drug war's rentier nexus. *PoLAR: Political and Legal Anthropology Review*, 36(2), 215–230.
- Polson, M. (2021). 4. Cultivating Cannabis, Excepting Cannabis. In *The War on Drugs* (pp. 92-128). New York University Press.
- Polson, M., & Bodwitch, H. (2021). Prohibited commoning: Cannabis and emancipatory legalization. *Elem Sci Anth*, 9(1), 00054.
- Polson, M., & Petersen-Rockney, M. (2019). Cannabis farmers or criminals? Enforcement-first approaches fuel disparity and hinder regulation. *California Agriculture*, 73(3), 185–193.
- Polson, M., Bodwitch, H., Biber, E., Butsic, V., & Grantham, T. (2023). After legalization: Cannabis, environmental compliance, and agricultural futures. *Land Use Policy*, 126, 106531.
- Polson, M., Laudati, A., & Sayre, N. F. (2024). Transformation of Unregulated Cannabis Cultivation Under Proposition 64.
- Short Gianotti, A. G., Harrower, J., Baird, G., & Sepaniak, S. (2017). The quasi-legal challenge: Assessing and governing the environmental impacts of cannabis cultivation in the North Coastal Basin of California. *Land Use Policy*, 61, 126–134.

State of California. (2016). Proposition 64: The Control, Regulate and Tax Adult Use of Marijuana Act (AUMA). California Health and Safety Code Section 11362.1 et seq.

State of California Legislature. (2017). Assembly Bill 64: Cannabis: licensure and regulation (Medicinal and Adult-Use Cannabis Regulation and Safety Act). California Business and Professions Code Section 26000 et seq.

Sujud, L., Jaafar, H., Hassan, M. A. H., & Zurayk, R. (2021). Cannabis detection from optical and RADAR data fusion: A comparative analysis of the SMILE machine learning algorithms in Google Earth Engine. *Remote Sensing Applications: Society and Environment*, 24, 100639.

United Nations Office on Drugs and Crime. (2024). *World Drug Report 2024 (Set of 3 Booklets)*. Stylus Publishing, LLC.

Walthall, C. L., Daughtry, C. S., Vanderbilt, V., Pachepsky, L., Lydon, J., Erbe, E., Higgins, M., & Bobbe, T. (2006). Detection of illegal Cannabis cultivation using remote sensing. *Proceedings of the International Geoscience and Remote Sensing Symposium*, Denver, CO, July 31-August 4.

Wang, I. J., Brenner, J. C., & Butsic, V. (2017). Cannabis, an emerging agricultural crop, leads to deforestation and fragmentation. *Frontiers in Ecology and the Environment*, 15(9), 495-501. <https://doi.org/10.1002/fee.1634>

Zheng, Z., Fiddes, K., & Yang, L. (2021). A narrative review on environmental impacts of cannabis cultivation. *Journal of Cannabis Research*, 3(1), 35.