Computer Vision Workflows to Operationalize Multiple Physical Carrying Capacity Indicators in Linear Recreation Corridors[⋆,⋆⋆]

Jordan W. Smith[a,*], Chase C. Lamborn[a]

[a]Institute of Outdoor Recreation and Tourism, Department of Environment and Society, S. J. & Jessie E. Quinney College of Agriculture and Natural Resources, Utah State University, 5215 Old Main Hill, Logan, 84322, Utah, USA

ARTICLE INFO

Keywords:
computer vision
semantic segmentation
object detection
carrying capacity
recreation corridors
parking occupancy
river recreation

Abstract

Linear recreation corridors are increasingly managed using carrying capacity frameworks, yet implementation is often constrained by limited and transparent measurement of physical conditions that structure where and when use concentrates. This study develops and demonstrates computer vision workflows that operationalizes multiple physical carrying capacity indicators for a heavily used river recreation corridor in Idaho, USA. The first workflow combines high-resolution aerial imagery with semantic segmentation to map corridor-scale parking opportunity; the second uses fixed on-site cameras with object detection to quantify temporal patterns in parking occupancy and on-river watercraft use. In both workflows, model performance improved substantially through iterative training, resulting in reliable parking supply estimation as well as parking area and on-river use. The resulting indicators revealed strong spatial heterogeneity in potential parking supply, pronounced temporal clustering in observed parking use, and reach-specific peaks in on-river use during high-demand periods. Together, these outputs provide a replicable and transparent basis for diagnosing when and where physical constraints are most likely to drive congestion within a recreation corridor. The approach supports indicator-based capacity planning by enabling trigger-based management during predictable peak windows, prioritizing infrastructure investments at constrained access areas, and facilitating coordination among agencies with shared management responsibilities.

1. Introduction

Carrying capacity remains central to sustainable outdoor recreation and nature-based tourism management because it provides an applied basis for aligning visitor use with acceptable resource and experiential conditions (Hammitt et al., 2015). In linear recreation corridors—river segments with multiple access points, trail networks, and scenic byways—capacity pressures are often amplified by rapid growth in participation, highly seasonal demand, and tight coupling between transportation access, staging infrastructure, and movement through the corridor (Fisher & Krutilla, 1972). These settings are therefore high-conflict and high-consequence: congestion, informal parking, safety risks, and encounter-related crowding often emerge as system behaviors rather than as isolated "site problems."

At the same time, corridor capacity problems are also a classic applied machine learning (ML) challenge: translating high-volume, heterogeneous imagery into replicable, decision-ready indicators under real-world constraints (limited labels, variable lighting and occlusion, and the need for defensible outputs). Solutions to these challenges are timely because many management domains outside the traditional ML core increasingly depend on computer vision (CV) systems that can be deployed, maintained, and trusted in operational settings—not merely demonstrated in benchmark conditions. In corridor contexts, managers and partner agencies need spatially explicit, frequently updated measures of physical constraints—where vehicles can plausibly park, when parking fills, and how use concentrates along the corridor—yet these quantities remain difficult to measure reliably with conventional monitoring approaches.

Contemporary carrying capacity practice has evolved away from identifying a single "maximum number" toward indicator–based approaches that define desired conditions, select measurable indicators, establish thresholds, and link monitoring to management actions (Interagency Visitor Use Management Council, 2016). However, implementation frequently stalls at a practical bottleneck: measurement. In corridors especially, managers often lack timely, spatially explicit data on the physical constraints that govern where

[⋆]This research was funded jointly by the USDA Forest Service, Boise National Forest and the Institute of Outdoor Recreation and Tourism at Utah State University. The authors would like to thank Eli Plehn and Maizy Troxell for their work in setting up and maintaining cameras throughout the study. The authors would also like to thank Van Smith for his assistance in manually labeling cars, trucks, and boats during the training process; "count it!" A previous version of this paper was presented at the 2025 River Management Society Annual Meeting in Ashland, Oregon.

[⋆⋆]High-resolution aerial imagery used in this study is publicly available from the US Geological Survey. Derived spatial indicator outputs and aggregated time-series summaries supporting the findings of this study are available from the corresponding author upon reasonable request. Raw camera imagery and full-resolution detection outputs are not publicly available due to privacy, site security, and data management considerations. The analytical workflow and source code are proprietary and are not publicly available; access may be provided under a licensing or research collaboration agreement subject to institutional approval.

[*]Corresponding author

jordan.smith@usu.edu (J.W. Smith)

https://www.linkedin.com/in/jordan-smith-38351b35/ (J.W. Smith); chase.lamborn@usu.edu (C.C. Lamborn)

ORCID(s): 0000-0001-7036-4887 (J.W. Smith); 0000-0003-0854-6270 (C.C. Lamborn)

and when use concentrates (e.g., parking opportunity and occupancy, access area turnover, and on-corridor use). Traditional monitoring tools—manual counts, periodic observations, and simple automated counters—can be costly to scale and may provide limited information beyond raw "passes," with known errors such as group undercounting and false triggers (Lindsey et al., 2018; Watson et al., 2000; Ziesler & Pettebone, 2018).

A multi-phasic view of recreation provides useful context for corridor planning by highlighting that constraints can originate in travel and access and still shape on-site experiences (Clawson & Knetsch, 1966; Clawson, 1963; Hammitt, 1980; Stewart, 1998). However, the practical monitoring bottleneck in most corridors is the on-site portion of the system: staging infrastructure and on-corridor conditions that can be directly observed, quantified, and linked to operations. This manuscript therefore focuses on developing scalable, image-based methods that quantify on-site physical constraints while retaining the multi-phasic framing as a conceptual bridge to broader corridor governance.

CV paired with deep learning offers a practical pathway to address monitoring bottlenecks by automating classification and counting from imagery, enabling multi-class monitoring and richer temporal signatures of corridor demand (Fennell et al., 2022; Mitterwallner et al., 2024). Yet, CV has been used sparingly in recreation and protected area management, and almost never to operationalize multiple physical carrying capacity indicators in corridor contexts where parking, access, and on-corridor use interact. Technically, the gap is not simply model availability; it is the lack of end-to-end, replicable workflows that (i) reduce labeling burden, (ii) integrate segmentation of aerial imagery and ground-based detection, and (iii) convert predictions into interpretable indicators that can be paired with standards and management actions.

Accordingly, this manuscript develops and demonstrates applied CV workflows that operationalize multiple physical carrying capacity indicators for a heavily used river recreation corridor in central Idaho (USA). We integrate (i) high-resolution aerial imagery with semantic segmentation to quantify spatial variation in parking opportunity, and (ii) in-field cameras with object detection to estimate temporal patterns in parking occupancy and on-corridor use density. To make corridor-scale mapping feasible without exhaustive hand digitizing, we use a foundation segmentation model to generate candidate objects that are then manually reviewed and attributed for supervised training. We evaluate model behavior under field-relevant constraints and demonstrate how the resulting indicators can support trigger-based interventions, infrastructure prioritization, and cross-agency coordination within standard visitor use management frameworks.

2. Problem Formulation and Indicator Framework

2.1. Measuring Capacity in Linear Recreation Corridors

Linear recreation corridors are commonly managed as if the recreation experience begins at a parking area or launch point and ends upon exit. In contrast, the multi-phasic tradition in outdoor recreation research conceptualizes recreation as a "package" comprised of five phases—anticipation, travel-to, on-site, travel-back, and recollection—each contributing to perceived quality and behavioral responses (Clawson & Knetsch, 1966; Clawson, 1963; Hammitt, 1980; Stewart, 1998). While the five-phase model has been critiqued for excessive linearity and porous boundaries, it provides a useful organizing lens for corridor planning because it emphasizes that constraints can arise in different parts of a trip and still affect the on-site experience that managers are typically tasked to protect.

For this study, the multi-phasic framing is used primarily to motivate why corridor monitoring should distinguish among physical bottlenecks and why indicators should be spatially explicit. Our empirical focus is narrower: we operationalize indicators that characterize the on-site portion of the corridor system, specifically (i) staging infrastructure (parking opportunity and parking occupancy) and (ii) in-corridor conditions (on-river use). These touchpoints are where management actions are most frequently deployed, where measurable physical constraints are most likely to cause congestion, and where imagery-based monitoring is most feasible at scale.

Expectations formed before arrival can shape tolerance for congestion and judgments about acceptability (Heywood, 2002). Travel-to and travel-back phases can introduce roadway throughput limits, queuing, and safety exposure that spill over into entry experiences and redistribute arrivals among access nodes (Aas & Onstad, 2013; Arnberger & Brandenburg, 2007; Sidder & Hall, 2024). Post-trip recollection can influence repeat visitation and displacement (Cole & Hall, 2012). We do not measure these phases directly here; instead, we treat them as contextual mechanisms that motivate why on-site indicators should be interpretable and transferable across jurisdictions and time periods, and why errors that inflate or suppress on-site constraints (e.g., erroneous estimates of parking opportunity) can meaningfully affect downstream planning.

2.2. Capacity Indicators Logic and the Role of Computer Vision

Modern carrying capacity practice is operationalized through desired conditions, indicators, standards/thresholds, and management actions linked to monitoring and adaptive response (Interagency Visitor Use Management Council, 2016). This logic reframes capacity from a single "maximum number" to a defensible decision process: managers specify what conditions must be maintained, identify measurable indicators of those conditions, define standards that distinguish acceptable from unacceptable change, and select tools to prevent or correct deviations.

This paper adopts that logic but narrows scope to physical capacity indicators, defined here as the infrastructural and operational throughput potential of a corridor system: how much staging space exists, how it is used through time, and how use manifests along the corridor itself. In corridor settings, physical bottlenecks are multiple and distributed—so a single proxy (e.g., "number of parking spaces" or "total passes") cannot represent the system state through time and space. The methodological implication is that indicator-based planning depends on measurement methods that are scalable and repeatable under field constraints.

Computer vision becomes relevant because it provides a practical mechanism for generating such indicators from heterogeneous imagery sources. The recreation and protected area literature documents growing feasibility of using motion-triggered cameras for visitor monitoring (Lupp et al., 2021), including video-based approaches that combine detection and tracking to improve counting under occlusion and group travel (Moreno et al., 2025). More broadly, object detection has become a standard approach for processing camera-trap and trail-camera imagery at scale, including open models and evaluation studies that demonstrate reduced manual processing with acceptable accuracy under many conditions (Fennell et al., 2022; Guidosse et al., 2025; Mitterwallner et al., 2024; Staab et al., 2021). At the same time, applied CV in public-land contexts faces persistent challenges: domain shift across seasons and lighting (glare, canopy shadow, snow/ice, night imagery), occlusion and clustering, and the need for operationally lightweight privacy safeguards that support adoption and legitimacy (Lupp et al., 2021; Mitterwallner et al., 2024; Wilkins et al., 2022).

Our applied ML contribution is to integrate these considerations into an end-to-end workflow that links model outputs to interpretable, corridor-scale indicators. Overhead imagery is used to map relatively stable "supply" features (parking opportunity) via semantic segmentation, while fixed cameras are used to estimate dynamic "demand" signals (parking occupancy and on-river concurrency) via object detection. The workflow is designed to be label-efficient and operationally reproducible, using iterative refinement (active learning and hard-example mining) to improve performance on rare targets and difficult conditions, and producing outputs that can be summarized consistently across sites and time bins for different planning and management needs.

3. Methods

3.1. Study Area and Management Context

This study focuses on the Payette River recreation corridor in Idaho (USA), a high-demand river corridor where physical capacity constraints translate directly into management concerns (Figure 1). The Main Payette near Banks is a nationally recognized whitewater run with concentrated summer boating use and pronounced weekend peaks, supported by multiple formal access sites and frequent informal pullouts along a highway corridor (American Whitewater, 2026). Access capacity is constrained at key put in and take out areas; Banks and Beehive Bend are widely described as difficult to park at during peak season, contributing to spillover parking, safety exposure, and localized conflict. These pressures are amplified by transportation bottlenecks at the SH-55 / Banks–Lowman Road (SH-17) intersection, where the Idaho Transportation Department documents seasonal traffic increases and has implemented interim controls to address summer weekend congestion (Idaho Transportation Department, 2026). Management authority is multi-jurisdictional: the USDA Forest Service and Bureau of Land Management oversee several access and parking sites and concessionaire operations, while state transportation and local law enforcement manage roadway operations and safety enforcement (Bureau of Land Management, 2026). This combination of high demand, constrained access, and fragmented authority makes the Payette corridor a suitable testbed for developing CV methods that yield transferrable, corridor-scale indicators.

3.2. Quantifying Parking Opportunity

We operationalized parking opportunity as the spatial extent of surface types that, based on high-resolution aerial imagery and basic terrain constraints, represent plausible locations where vehicles could park within the corridor. The workflow comprises two core stages: (i) label development using SAM-assisted candidate polygons with manual attribution and (ii) supervised semantic segmentation to map parking-related classes across the corridor. This is followed by a standardized indicator derivation step that summarizes mapped parking opportunity (Figure 2).

3.2.1. Data Sources and Preprocessing

Model inputs were derived from high-resolution (0.6 m) four-band NAIP imagery (U.S. Geological Survey, 2026) accessed via Google Earth Engine (Gorelick et al., 2017). Imagery values were normalized to a common range for training and inference. The corridor analysis area was defined by buffering primary highways by 500 m and intersecting that buffer with a 500 m buffer around the river, yielding a road–river intersection polygon that constrains both labeling and prediction to areas most relevant to access and parking behavior. To reduce implausible parking predictions on steep terrain, we applied a binary slope constraint retaining areas $\leq$ 5% slope, derived from a 1 m digital elevation model from the USGS 3DEP product suite (U.S. Geological Survey, 2024). This constraint was applied twice: first to screen candidate-label generation to low-slope tiles, and later to mask predicted class labels outside low-slope terrain during inference.

3.2.2. Label Development

We generated candidate objects using SAM (ViT-H) (Kirillov et al., 2023); these objects were later manually

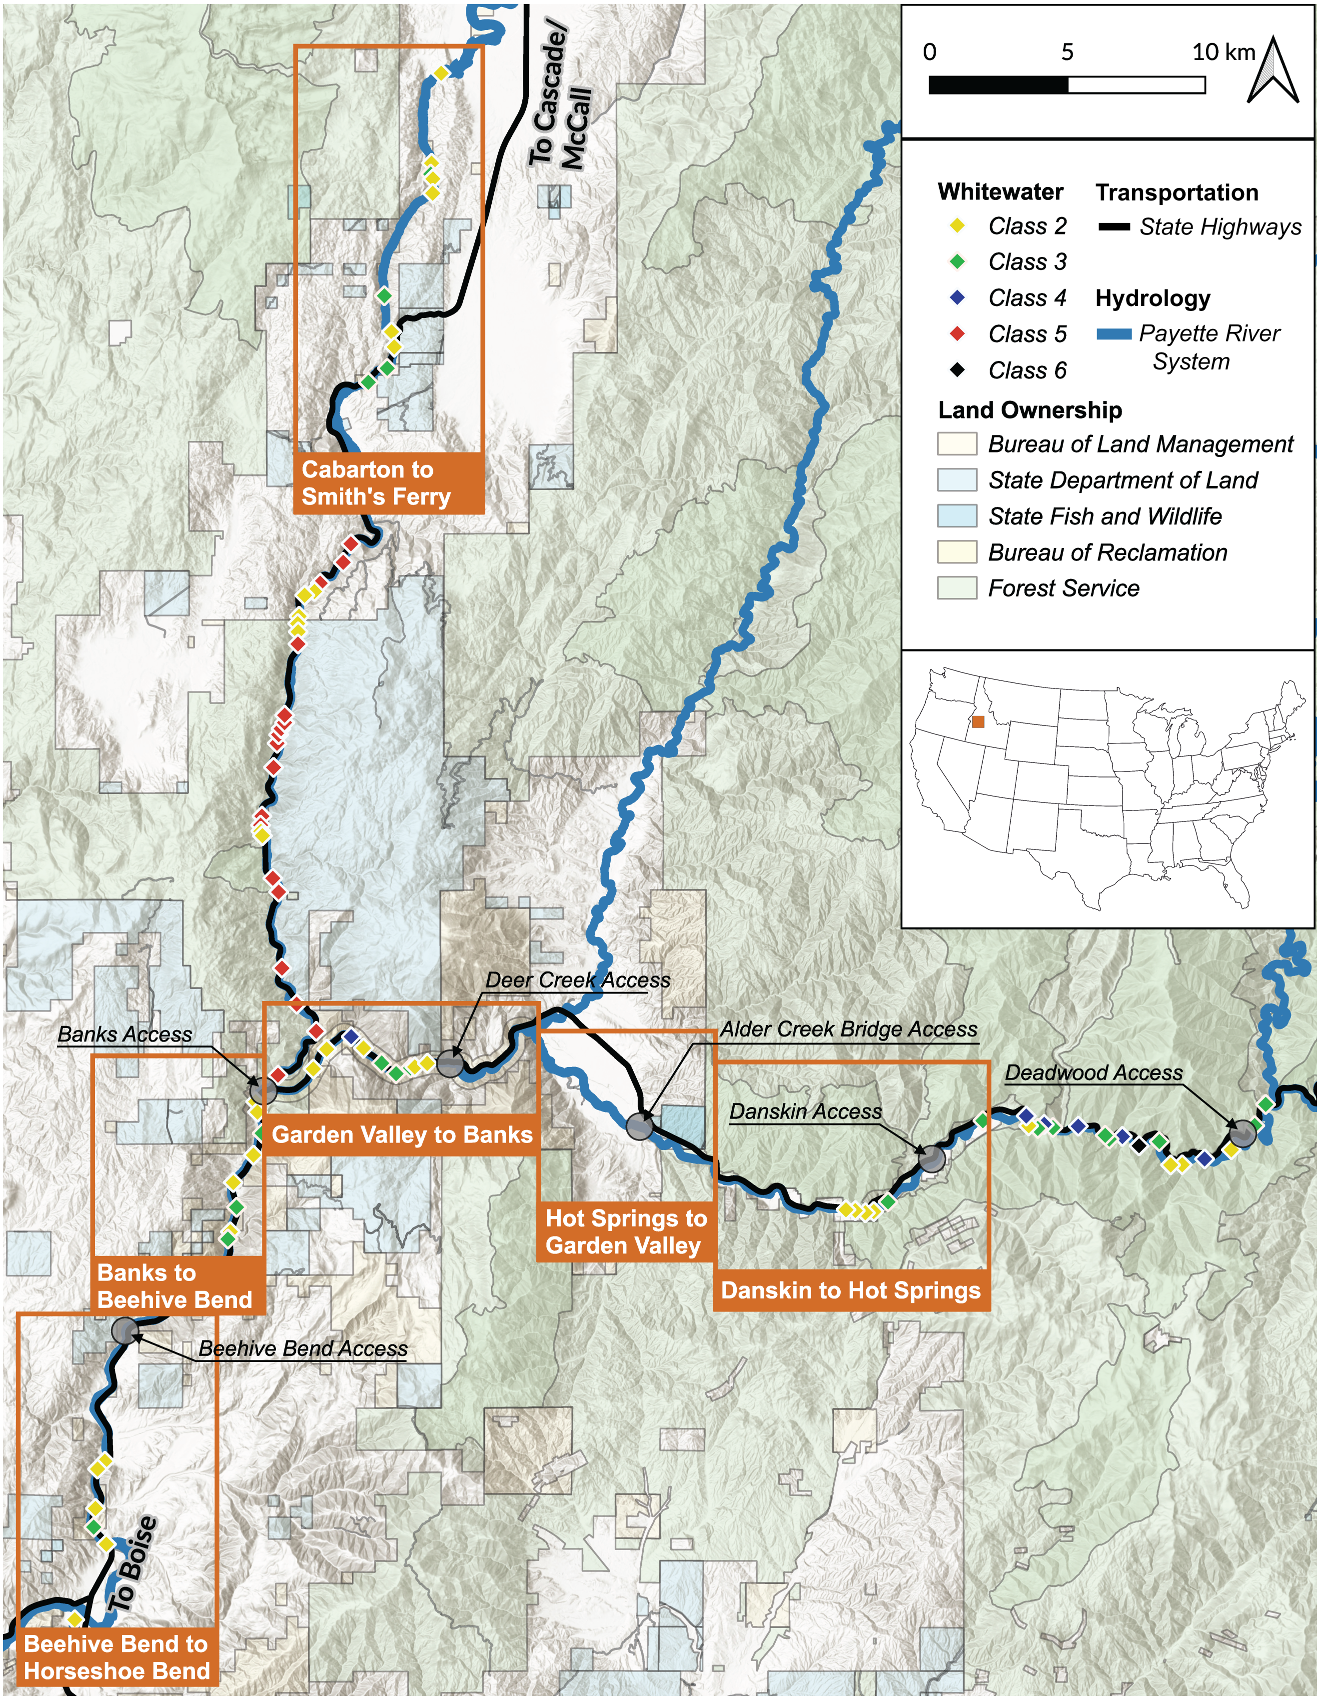

Figure 1: Study Area. The workflows were developed and tested within the Payette River recreation corridor in central Idaho (USA). The corridor is well known for its world class whitewater rafting and kayaking opportunities.

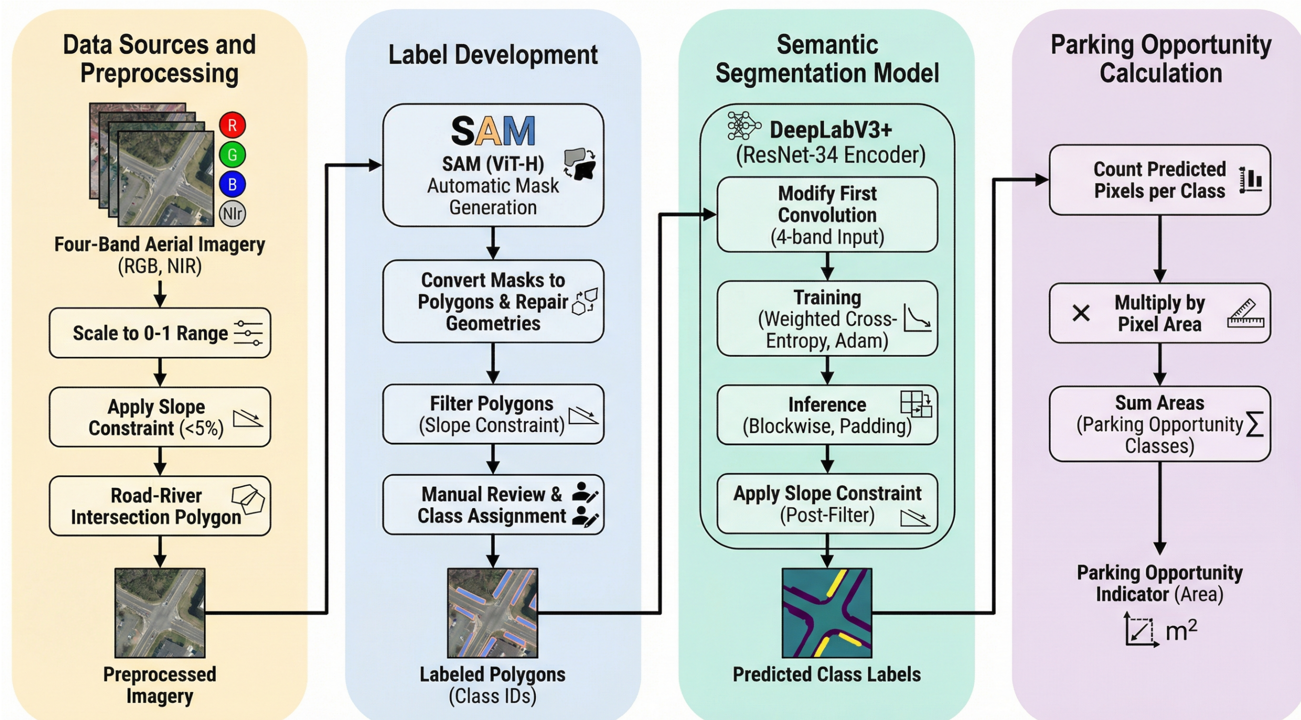

Figure 2: Workflow for Quantifying Parking Opportunity. In stage 1, NAIP aerial imagery is queried and exported from Google Earth Engine, road vectors are retrieved from the U.S. Census TIGER/Line datasets, and the river area is obtained from the NHD dataset. Major highways and the rivers are intersected and buffered to define the corridor. A USGS D3EP 1-meter digital elevation model is clipped and used to identify candidate "flat" areas ($\leq$ 5% slope). These constrained polygons then inform a two-stage image analysis workflow. In stage 2, the Segment Anything Model (SAM (ViT-H)) generates terrain-constrained object proposals from tiled NAIP imagery, which are exported as GIS-editable vectors for manual labeling. In stage 3, labeled polygons are rasterized into training masks, paired 512×512 image–mask chips are created, and a DeepLabV3+ semantic segmentation model is trained to predict imagery classes across the corridor. Finally, windowed inference produces a classified GeoTIFF, with predictions explicitly masked outside low-slope ($\leq$ 5%) areas, and outputs are polygonized and summarized (e.g., per-class area) for visualization and comparison to manual labels.

classified (see Hauser et al. (2025) for a review of the evolution of efficient labeling strategies for geospatial data). SAM outputs pixel masks that were converted to vector polygons and filtered using the slope constraint in two steps: imagery tiles not intersecting low-slope terrain were skipped, and extracted polygons were retained only if they intersected low-slope areas. Each retained polygon was assigned an integer `class_id` field initialized as background (`0`). These polygons served as label proposals for manual review, during which geometries were edited (e.g., accept/reject decisions, merge/split operations, and repairs) and non-zero class labels were assigned (Wang et al., 2024). Only polygons with `class_id = 0` were used as labeled classes for training; all remaining areas were treated as background).

3.2.3. Semantic Segmentation Model

Semantic segmentation was performed using DeepLabV3+ (Chen et al., 2018) with a ResNet-34 encoder (He et al., 2016) implemented in a standard segmentation library. The encoder was initialized with ImageNet weights trained on three-channel RGB imagery. To support four-band NAIP inputs (RGB + near-infrared), we modified the first convolutional layer to accept four channels by copying pretrained weights for the RGB channels and initializing the near-infrared channel from the mean of the RGB kernels. This approach preserves the benefits of transfer learning while adapting the model to the additional spectral band commonly used in land-surface classification.

Training used a weighted multi-class cross-entropy objective to address pixel-level class imbalance (background dominating the landscape), with class weights specified in the implementation. Optimization used Adam (learning rate 10-4), batch size 8, and a fixed 15-epoch schedule. To improve robustness to orientation and illumination differences, the final iteration of the model was trained using on-the-fly augmentations (random flips, 90° rotations, affine shift/scale/rotate transforms, and brightness/contrast perturbations) (Sharma & Gosain, 2025)

Inference was performed blockwise using the imagery's internal tiling (Q. Liu et al., 2026). To accommodate downsampling constraints, blocks were padded to the nearest multiple of 16 using replicate padding and cropped back to their original size after prediction. The slope constraint was then re-applied so any predicted labels outside low-slope terrain were reset to background.

3.2.4. Calculating Parking Opportunity

Parking opportunity was computed as class-specific (paved parking and dirt/gravel parking) areas within the corridor. For each mapped class, the number of predicted pixels was multiplied by pixel area (derived from imagery resolution) to compute area by class. Total parking opportunity was computed by summing areas across the subset of classes designated as "parking opportunity" in the labeling schema.

3.3. Quantifying Parking Opportunity and On-River Capacity

We quantified corridor recreation demand using YOLO26 applied to a large image corpus collected by motion-triggered cameras deployed at both parking areas and river viewpoints (Figure 3). The corridor is treated as two interacting observational domains: parking areas, which provide a proxy for access pressure and arrival intensity, and the river corridor itself, which provides a proxy for on-water activity. The image corpus is standardized and cleaned prior to modeling (stable filenames, manifest completeness, deduplication, and daytime filtering). All subsequent steps operate via an authoritative manifest that provides file paths and timestamps and enables reproducible sampling, labeling, training, and full-corpus inference.

3.3.1. Baseline Inference and Corridor-Wide Active Learning (Loop 1)

The workflow begins with an initial, sample-based inference pass using a COCO-pretrained detector adapted to a two-class ontology: vehicle and watercraft. A deterministic stratified sample is drawn to represent both parking and river cameras and to span temporal variation, stratified by site and hour-of-day. For each sampled image, the model produces bounding boxes, class labels, and confidence scores, which are stored alongside per-image summary metrics such as class counts and total detections. These outputs are written to a structured table suitable for downstream selection and interactive review.

From this inference output, an active learning selector constructs the first labeling batch (500 images). Selection is designed to increase information gain rather than maximize random coverage (Agarwal et al., 2020). Images are prioritized when predictions are uncertain (e.g., near a decision boundary), while also ensuring diversity across sites, camera viewpoints, and time-of-day. The selector also includes hard negatives likely to drive false positives and a subset of predicted-empty images to stabilize specificity.

Manual labeling is conducted in an interactive environment (Voxel51, (Moore & Corso, 2026)) where model predictions are displayed and corrected. Annotators create ground-truth bounding boxes and assign the two target classes, yielding a vetted training dataset aligned with the project's operational definitions. After labeling, automated quality assurance checks validate label integrity and bounding box geometry (presence of expected fields, plausible distributions, finite coordinates, and well-formed boxes). Datasets that pass QA advance to training.

A deterministic train/validation split is then assigned across labeled samples in a single bulk operation to ensure reproducibility. Because watercraft are rare, the split procedure can enforce a minimum number of watercraft-positive images in the validation set to support stable evaluation. The labeled dataset is exported into a standard YOLO training layout, converting bounding boxes to YOLO label files, writing a dataset configuration file, and producing a first fine-tuned detector using a fully parameterized training configuration (image size, batch size, device selection, caching strategy, and augmentation recipe).

3.3.2. River-Focused Active Learning (Loop 2)

A second active learning loop specializes the detector for river imagery, reflecting the scientific priority of measuring on-water recreation and the distinct visual conditions of river scenes (glare, occlusion, complex backgrounds). A second inference pass is run on a river-only candidate pool using deterministic stratified sampling across river sites and times of day. Predictions and per-image summary metrics are recorded in the same structured format as Loop 1.

A river-focused selection procedure then constructs a second labeling batch (1,000 images). While it retains the same objectives—uncertainty, diversity, hard negatives, and zero-case coverage—stratification is tuned for river scenes where rare targets and difficult lighting can produce different error modes than parking lots. The selected batch is labeled in the same interactive environment, again using displayed predictions to accelerate annotation and reduce missed objects. These labels are used to train a second, more fine-tuned model that incorporates river-focused supervision while retaining corridor-wide learning from Loop 1.

3.3.3. Watercraft Mining to Address Extreme Case Imbalance

Even with two active learning loops, watercraft remain sparse, so the workflow adds a targeted mining phase to increase positive training examples (L. Liu et al., 2026; Oksuz et al., 2021). The operating confidence threshold for watercraft is calibrated by sweeping candidate confidence values on validation data and computing precision, recall, and F1 under a fixed intersection-over-union (IoU) matching rule; this yields a threshold that balances false positives and false negatives.

Using the calibrated threshold and the updated detector, large-scale inference is then executed across the full river corpus in chunks for stability. Images are ranked by a mining score designed to prioritize confident and meaningful watercraft detections, producing a candidate list for rapid review. High-ranked candidates can be staged for annotation and incorporated into training through an oversampling mechanism that increases model exposure to watercraft examples, mitigating imbalance-driven underperformance.

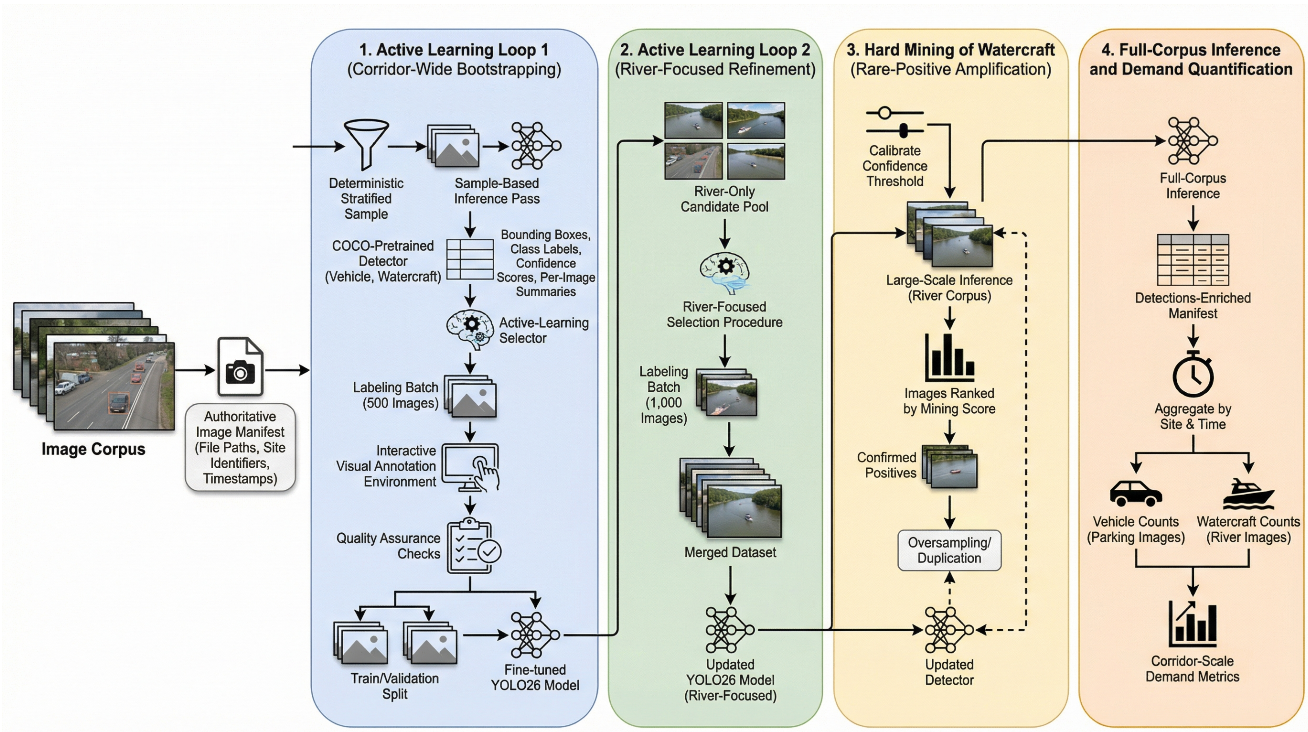

Figure 3: Computer Vision–Enabled Active Learning Workflow for Quantifying Recreation Demand. Schematic of the end-to-end workflow used to estimate recreation demand from a pre-cleaned time-stamped image corpus spanning parking and river camera domains. The process begins with sample-based baseline inference to generate preliminary detections and per-image metrics, followed by an uncertainty- and diversity-aware selection of an initial corridor-wide labeling batch ($n = 500$). Human annotation in an interactive visual environment produces ground-truth detections that are quality-checked, deterministically split into training and validation subsets with rare-class coverage, exported to YOLO format, and used to fine-tune the detector. A second active learning loop repeats these steps on a river-only stratified labeling batch ($n = 1{,}000$) to specialize performance for in-river scenes. To mitigate severe watercraft class imbalance, a calibrated confidence-threshold sweep informs corpus-scale mining of likely watercraft positives, which are oversampled during a final fine-tuning stage. The workflow concludes with full-corpus inference using class-specific thresholds, yielding detections-enriched outputs that can be aggregated into spatial–temporal indicators of access demand (vehicles in parking areas) and on-water use (watercraft on the river).

3.3.4. Full-Corpus Inference and Demand Quantification Outputs

The workflow concludes with full-corpus inference over the cleaned manifest. Distinct class-specific confidence thresholds are applied for vehicles and watercraft to reflect different base rates and error costs. The output is a detections-enriched manifest containing per-image class counts and full detection records.

These outputs enable demand quantification by aggregating vehicle detections at parking cameras as an access-demand proxy and watercraft detections at river cameras as an on-water use proxy. Because each image retains site and timestamp context, detections are summarized into spatial and temporal metrics (by camera, site, reach, hour-of-day, and day), enabling corridor-scale analysis of recreation pressure, use patterns, and change over time.

4. Results

4.1. Parking Opportunity Indicators

4.1.1. Overview of Parking Opportunity Workflow Outputs

"Parking opportunity" represents the mapped extent of surface types that appear plausibly usable for vehicle parking within the analysis area, with predictions constrained to low-slope terrain to reduce implausible classifications on steep hillsides. The workflow produced three primary outputs. First, it generated a classified raster surface map that assigns each pixel to one of the modeled land-use classes, including two parking-related classes (paved parking and dirt/gravel parking). Second, it produced a polygon layer used to support interpretation and quality review of mapped parking opportunity. Third, it yielded area-based indicators summarized by reach and for the corridor overall. The focal indicator is the estimated area of parkable surface, which can be interpreted as a spatially explicit measure of parking "supply" relative to corridor access demands.

4.1.2. Assessment of Label Quality and Training Data Representativeness

Label development began with automated candidate-object generation, followed by manual attribution. The candidate generation step produced 12,942 objects after terrain filtering; however, manual review revealed that most candidate objects were trees, underscoring the importance of human screening and class assignment when applying general-purpose segmentation tools to corridor infrastructure features. After manual classification, 817 polygons were retained with non-background labels and used to build the training mask. Notably, the manual editing process did not involve rejecting candidate polygons outright, while merges were common across chips, indicating that the dominant labeling task was consolidating fragmented objects into coherent surface features rather than dividing over-generalized polygons.

The resulting labeled set was imbalanced across classes. Of the 817 labeled polygons, 16 represented paved parking and 50 represented dirt/gravel parking, while 321 were roofs, 227 were water, and 203 were road. The comparatively sparse labeling of parking surfaces motivated two methodological adjustments that are directly relevant to interpreting the results: first, the semantic segmentation model was configured to penalize misclassification of the parking classes substantially more than other classes (`Model V2: Weighted Loss`), and second, data augmentation was applied during training to increase the effective diversity of training examples and reduce overfitting to a small number of parking instances (`Model V3: Augmented`).

4.1.3. Segmentation Model Performance

Model performance was assessed using training dynamics, class-wise IoU, and qualitative visual validation of predictions against labeled masks. Training loss declined rapidly over the 15-epoch run (Figure 4, Panel A). This learning curve indicates that the trained model fit the labeled data well under the specified weighting and augmentation regime and performed notably better than an earlier, unweighted and non-augmented configuration.

Across classes, mean IoU reached 0.64 for the final model. Class-specific IoU values were 0.85 for background, 0.60 for paved parking, 0.56 for dirt/gravel parking, 0.44 for roofs, 0.76 for water, and 0.79 for roads (Figure 4, Panel A). These values exceeded those from the initial version of the model that did not apply additional penalization for parking misclassification or use augmentation. Visual validation examples comparing aerial imagery chips, groundtruth masks, and predictions are presented in Figure 5, which provides an interpretive check on whether pixel-wise agreement corresponds to meaningful corridor features.

From the standpoint of capacity indicators, the error mode with the greatest practical consequence is systematic confusion between non-parking hardscape and parking surfaces. The most common observed error was the model interpreting segments of the highway as paved parking. This matters because false positives of paved parking inflate estimated opportunity in precisely the places where parking is operationally constrained or unsafe.

4.1.4. Parking Opportunity Indicators for the Payette River Corridor

Parking opportunity indicators were generated by converting predicted pixels in each parking class into mapped area and summarizing totals for the corridor and for individual reaches relevant to managers. Figure 6 shows corridor-wide (left) and reach-specific (right) totals. These totals provide a first-order estimate of the physical space potentially available for parking across the corridor.

Spatial variation in parking opportunity is highly uneven along the corridor. The reach-level summaries (Figure 6) provide the primary basis for comparing where supply is likely to be most constrained and where additional monitoring or management attention may be warranted. Interpreting these patterns requires caution: mapped opportunity is not equivalent to authorized or safe parking, and the corridor's effective capacity will depend on site design, access control, and how demand concentrates during peak periods. Nonetheless, the reach-level patterning offers a defensible, replicable way to characterize how physical parking "supply" is distributed relative to the corridor's access areas and travel path.

4.2. Parking Opportunity Indicators

4.2.1. Workflow Outputs and Indicator Definition

Parking occupancy was quantified from fixed cameras positioned to view parking areas. For each image, the detector produced bounding boxes and a vehicle class label with an associated confidence score. We converted these detections to a per-image vehicle count, then aggregated counts to standardized time bins (e.g., hourly) by parking area to generate a consistent, camera-based proxy for access pressure through time. Because camera viewsheds and lot geometries differ by area, the indicator is interpreted as an occupancy proxy (vehicle counts within the camera's field of view), rather than as a direct estimate of percent occupancy unless explicitly calibrated to area-specific stall capacity (not reported here).

4.2.2. Model Assessment

Model performance for vehicles improved across iterations as the workflow progressed from the initial baseline run to two active learning loops and the final model (Figure 4, Panel B). Across iterations, the vehicle class achieved consistently strong detection performance and remained comparatively stable relative to watercraft, reflecting the larger number of vehicle instances available for learning. In the final iteration, vehicle detection performance remained high (P = 0.856, R = 0.733; $mAP50$ = 0.865; $mAP50\text{–}95$ = 0.584), supporting use of vehicle counts as a reliable indicator for comparing area-level pressure over time. The most consequential errors for the occupancy indicator are missed vehicles under partial occlusion (under-counting)

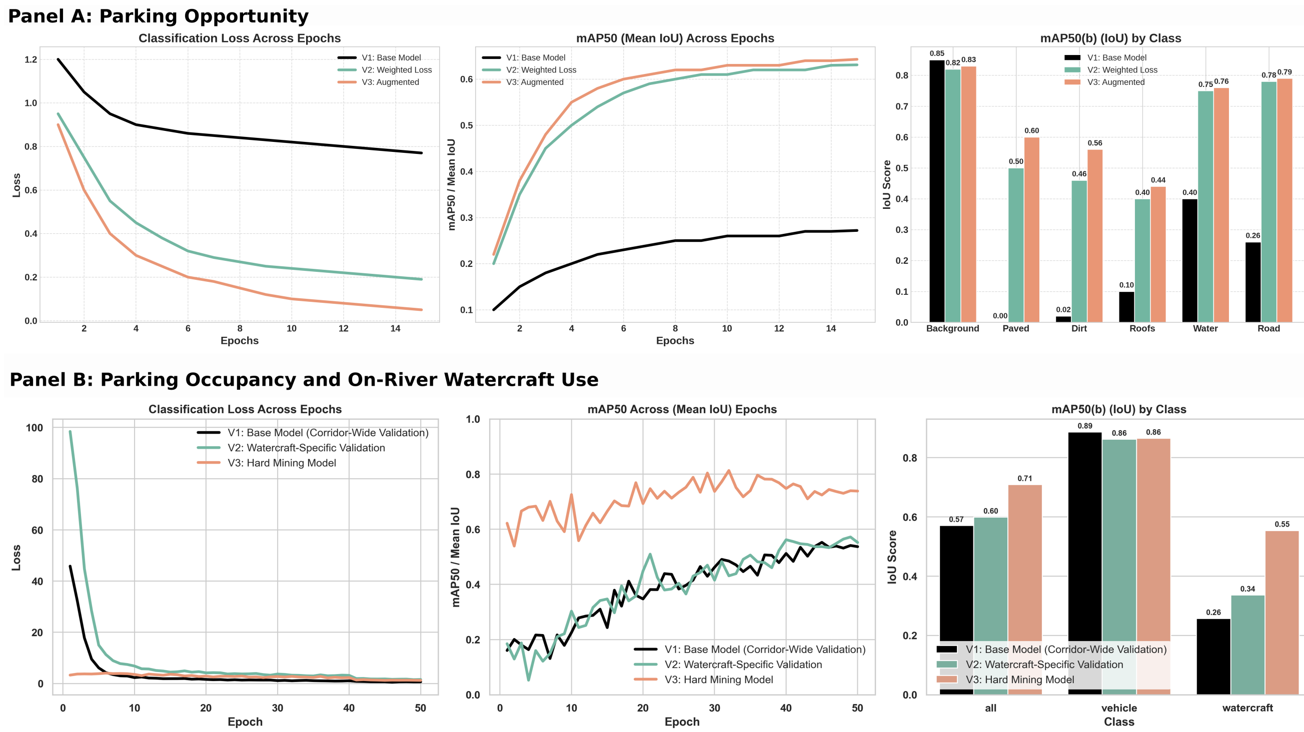

Figure 4: Training Dynamics and Class-Wise Segmentation Performance for the Parking Opportunity Model. (A) Cross-entropy training loss and class-wise IoU for the DeepLabV3+ semantic segmentation model over 15 epochs under the base model, class-weighted, and augmented training configurations. IoU is computed as the ratio of the area of overlap between prediction and reference labels to the area of their union; higher values indicate better pixel-level agreement.**(B)** Cross-entropy training loss and class-wise IoU for the YOLO26 object detection model over 50 epochs under the base model, watercraft-specific training, and hard mining configurations.

and spurious detections on visually similar background features (over-counting); both directly translate into bias in per-image counts.

4.2.3. Site-Level Parking Occupancy Patterns

Aggregated vehicle counts revealed pronounced spatial and temporal concentration in parking area use across the corridor. Peak occupancy proxies occurred during predictable high-demand windows, with sharp within-day pulses that correspond to typical arrival and departure cycles for river-based recreation. Figure 7 provides an hour-by-day-of-week heatmap for all sites, highlighting recurrent peak windows and providing an interpretable "use signature". Complimentarily, Figure 8 bins the hour-by-day-of week data into distinct times of day relevant to outdoor recreation and transportation managers/planners. Across the study period, the highest-use areas were the Banks and Beehive parking lots, with peak demand centered around the weekends and mid-day. In contrast, lower-demand areas (Alder Creek, Danskin, and Deadwood) showed smaller peaks, indicating lower demand. Together, these indicators provide a corridor-scale depiction of where access pressure is most likely to create bottlenecks and identify candidate areas for targeted management actions or future calibration work (e.g., converting vehicle counts to percent occupancy for specific lots).

4.3. On-River Capacity Indicators

4.3.1. Workflow Outputs and Indicator Definitions

On-river capacity was quantified from cameras positioned to observe river segments, where the detector produced watercraft bounding boxes and confidence scores. Detections were converted to per-image watercraft counts and aggregated to standardized time bins (e.g., hourly) by river site/reach. Because each camera captures only its local viewshed, these values represent a practical proxy for use (how many craft are present in view at a given time), rather than a complete accounting of corridor-wide throughput.

4.3.2. Model Assessment

Watercraft detection exhibited the largest gains across model iterations, consistent with the workflow's design to address extreme rarity and heterogeneous appearance of boats across lighting, glare, and background conditions. In the initial baseline run, the watercraft class showed very low recall and limited average precision (Figure 4, Panel B). Performance increased modestly after the first active learning loop and the second loop, reflecting incremental improvements in the model's ability to identify rare targets. The largest improvement occurred after the final iteration that incorporated hard mining of watercraft examples, where recall increased substantially alongside high precision ($P = 0.843$), yielding markedly improved *mAP50* (0.554) and *mAP50–95*

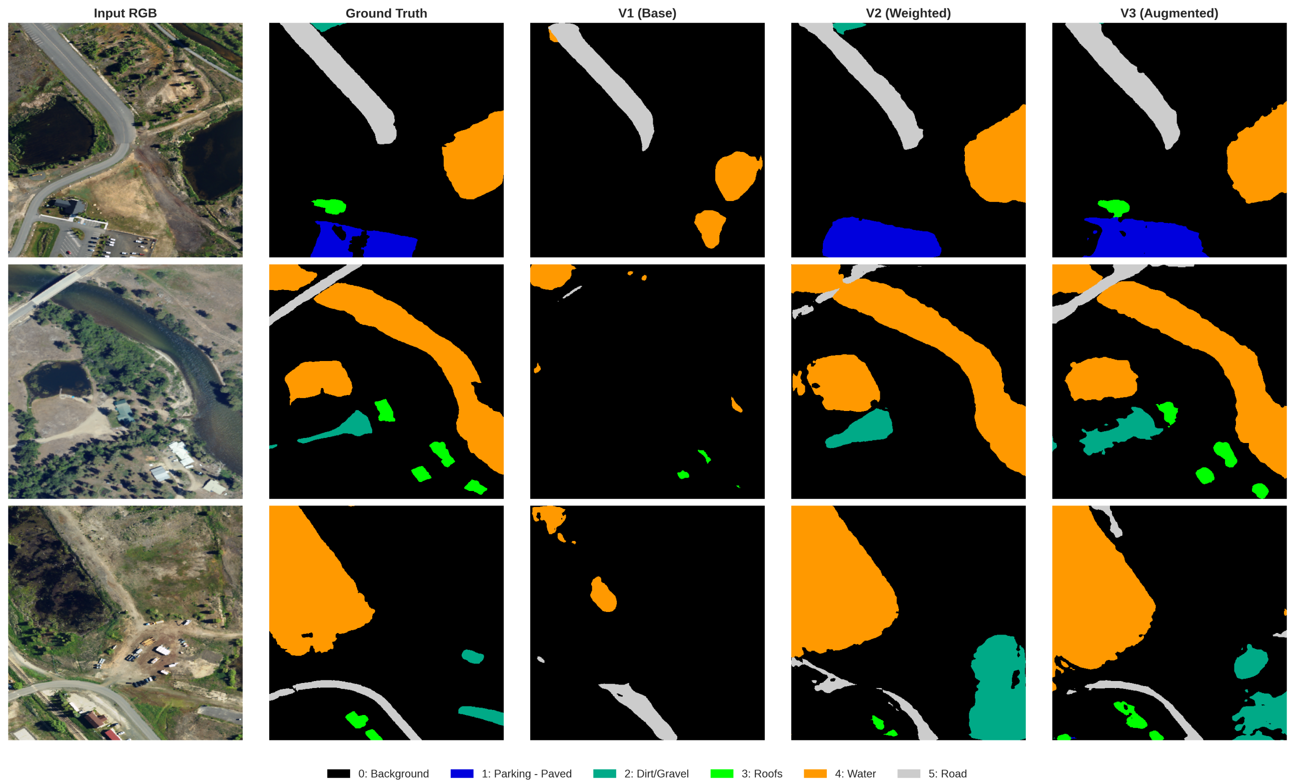

Figure 5: Visual Validation Examples for Parking Opportunity Mapping From Aerial Imagery. Example panels comparing four-band aerial imagery rendered as RGB (column 1), the manually curated reference mask derived from SAM-assisted polygons (column 2), and the DeepLabV3+ predicted class maps for the same chip locations across each iteration of the model (columns 3-5). Examples are selected to illustrate typical model behavior in corridor settings, including correct delineation of paved and dirt/gravel parking surfaces and common error modes in visually similar hardscape (e.g., misclassification of roadway segments as paved parking).

(0.442). These gains are consequential for indicator validity because low recall produces systematic underestimation of use, especially during peak conditions when craft may be small, partially occluded, or clustered.

4.3.3. Site-Level On-River Use Patterns

The resulting watercraft indicators show strong temporal clustering and clear spatial differentiation across river camera sites. On-river peaks were most pronounced during mid-morning and mid-afternoon, with the most detections on the Banks to Beehive (*Beehive* camera in Figures 7 & 8), stretch. Figure 7 shows the temporal variation in river use across sites; this illustrates both the timing of peaks and differences among reaches in peak intensity and duration. Several sites exhibited recurring high-detection pulses aligned with typical launch-to-takeout travel times and day-use patterns. Two peak use times in the mid-morning and mid-afternoon are particularly noticeable for Banks and Beehive (Figures 7 & 8), whereas other sites displayed smaller peaks or greater variability, consistent with differences in reach characteristics and use. Collectively, these site-level patterns provide a corridor-scale depiction of when and where on-water use concentrates, establishing a quantitative basis for comparing river segments and identifying peak windows for management attention.

5. Discussion

5.1. Contribution to Carrying Capacity Debates

A central theme in carrying capacity scholarship is that “capacity” is not a single number, but a management determination that depends on explicitly defined desired conditions, measurable indicators, standards/thresholds, and adaptive actions (Interagency Visitor Use Management Council, 2016). The contribution of this study is to make that logic more feasible for linear recreation corridors, where capacity constraints are distributed across space, fluctuate through time, and often fall under fragmented jurisdictional authority. By operationalizing multiple physical indicators—parking opportunity (supply), parking occupancy (demand), and on-river use (demand)—the workflow provides a transparent process that can support indicator-focused frameworks in settings where they often stall due to limited monitoring capacity (Lindsey et al., 2018; Watson et al., 2000; Ziesler & Pettebone, 2018).

Recreation corridors raise a distinct conceptual challenge for capacity debates: if the recreation experience is multi-phasic (anticipation, travel-to, on-site, travel-back, recollection), then constraints can occur in certain phases and propagate to others (Clawson & Knetsch, 1966; Clawson, 1963; Hammitt, 1980; Stewart, 1998). In the Payette

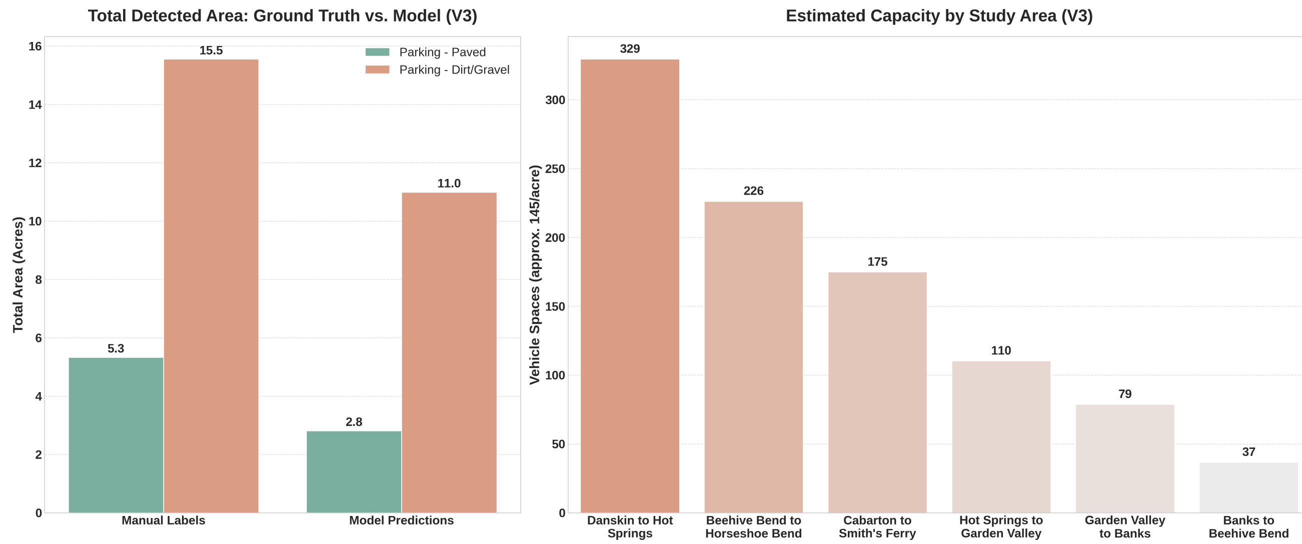

Figure 6: Corridor-wide Parking Opportunity by Surface Type for the Payette River Recreation Corridor. (Left) Total parking opportunity within the corridor analysis area, summarized as area for each parking-related class (paved parking and dirt/gravel parking). (Right) Estimated combined (paved and gravel) parking opportunity across management-relevant reaches. These corridor-wide totals represent a supply-side indicator of potential parking space and provide a baseline for comparing spatial variation and linking parking opportunity to observed occupancy and on-river use patterns in subsequent analyses.

River Corridor, the mapped parking opportunity varied considerably along the corridor, and parking occupancy and on-river indicators revealed strong spatial and temporal variation. The methods provided here offer a more comprehensive understanding of capacity and utilization than governing bodies often have access to. While single proxies such as "parking spaces," "boats per day," or "trailhead passes" are often used, they are insufficient when used in isolation because constraints shift between parking, staging areas (e.g., river put-in and take-out), and the areas used for recreation (e.g., the river). In practice, this is precisely where the "capacity is not one number" claim becomes operational: multi-indicator monitoring clarifies where and when bottlenecks are occurring, enabling management responses that are proportional to the mechanism driving undesirable conditions rather than generic responses to high use. The workflow presented here provides a process where multiple indicators can be monitored simultaneously to provide spatial and temporal data for examining flow and identifying bottlenecks.

This paper also advances carrying capacity debates by reframing measurement as an applied ML problem. Capacity planning often assumes that indicators can be measured, updated, and monitored; in corridors, this assumption is frequently violated by the cost and logistical complexity of field-based monitoring (Fennell et al., 2022). The workflow demonstrates a practical pathway for producing repeatable indicators using heterogeneous visual data streams—aerial imagery for relatively stable infrastructure supply and fixed cameras for dynamic demand signals. Importantly, the value is not simply that deep learning models can classify surfaces or detect objects, but that the workflow links model outputs to interpretable indicators suitable for planning and operations. The iterative model improvements documented across both workflows further underscores that indicator production is not "one and done," but can improve as monitoring systems mature, labeling data improves, and models are updated (Alotaibi & Nassif, 2024).

5.2. Managerial and Planning Implications

The principal managerial utility of the proposed indicator set is that it supports diagnostic capacity management—identifying where access is limited and developing tailored interventions. For the Payette River corridor, geospatial data on parking opportunity provides a corridor-scale screening tool for locating segments where physical staging space is scarce or fragmented. These segments are natural candidates for targeted investment in access infrastructure, particularly where high observed occupancy suggests recurring demand pressure and where roadway conditions increase safety exposure. The opportunity layer can be used to identify candidate areas for upgrades (primarily the expansion of existing lots or the creation of dedicated roadside pullouts).

Second, the parking and on-river use indicators can support trigger-based "soft management" strategies that do not require immediate capital investment. Because both indicators provide time-resolved signatures of peak periods, managers can implement interventions that align with predictable windows of pressure. Examples include messaging about peak use and lot fullness, arrival time recommendations, and route guidance to distribute use across areas. If policy and infrastructure allow, the same triggers can support more structured strategies such as shuttle programs that reduce parking pressure or temporary restrictions during

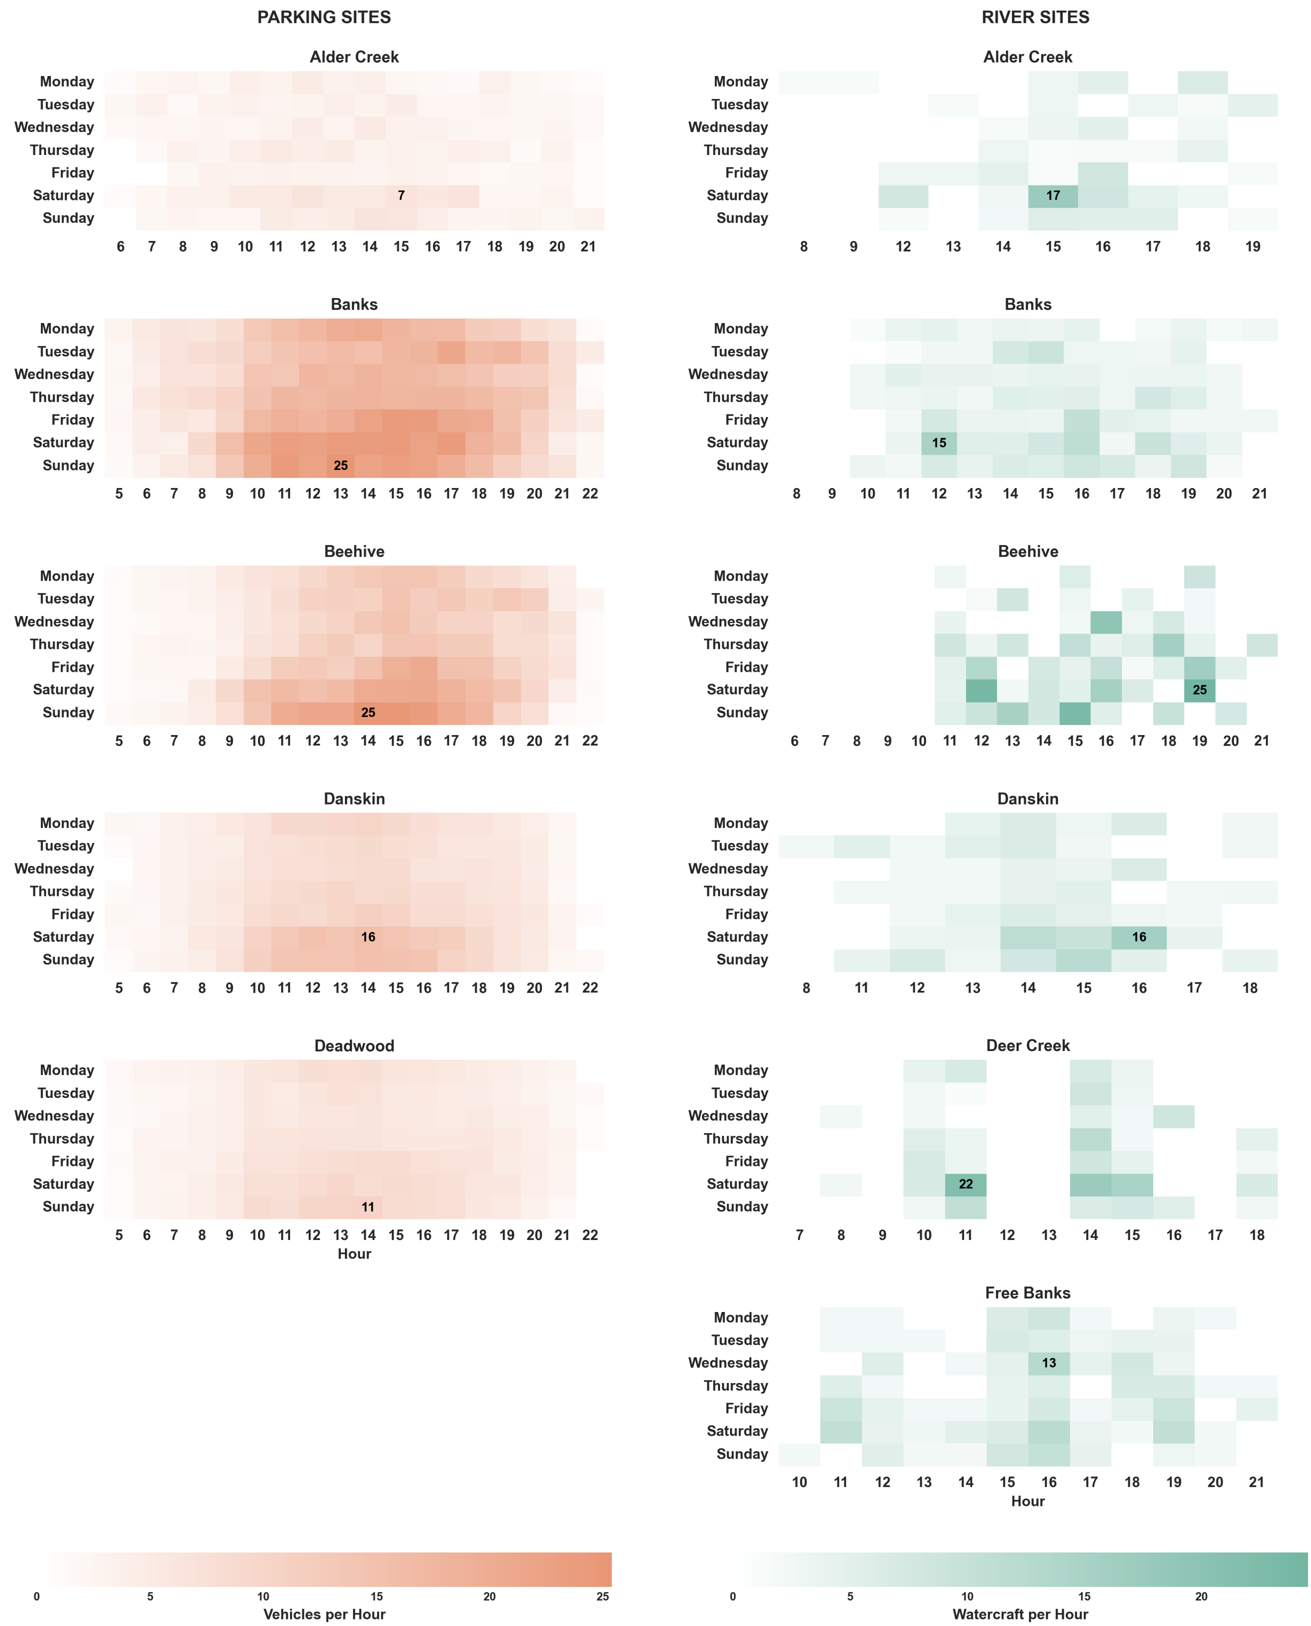

Figure 7: Heatmaps of Mean Detected Recreation Activity at Parking and River Sites. The heatmaps summarize the mean number of vehicles (unique and persistent detections) per hour as a function of hour-of-day (x-axis) and day-of-week (y-axis) for each monitored site. Parking sites (left column) represent vehicles per hour at access areas, while river sites (right column) represent on-river watercraft per hour. Each row corresponds to a site (site name shown above each panel). Color intensity indicates higher mean vehicles/watercraft per hour (note separate color scales for parking vs. river), and annotated cells highlight the peak mean value within each site panel. These patterns characterize typical temporal use profiles and help identify peak-use windows for capacity and management analyses.

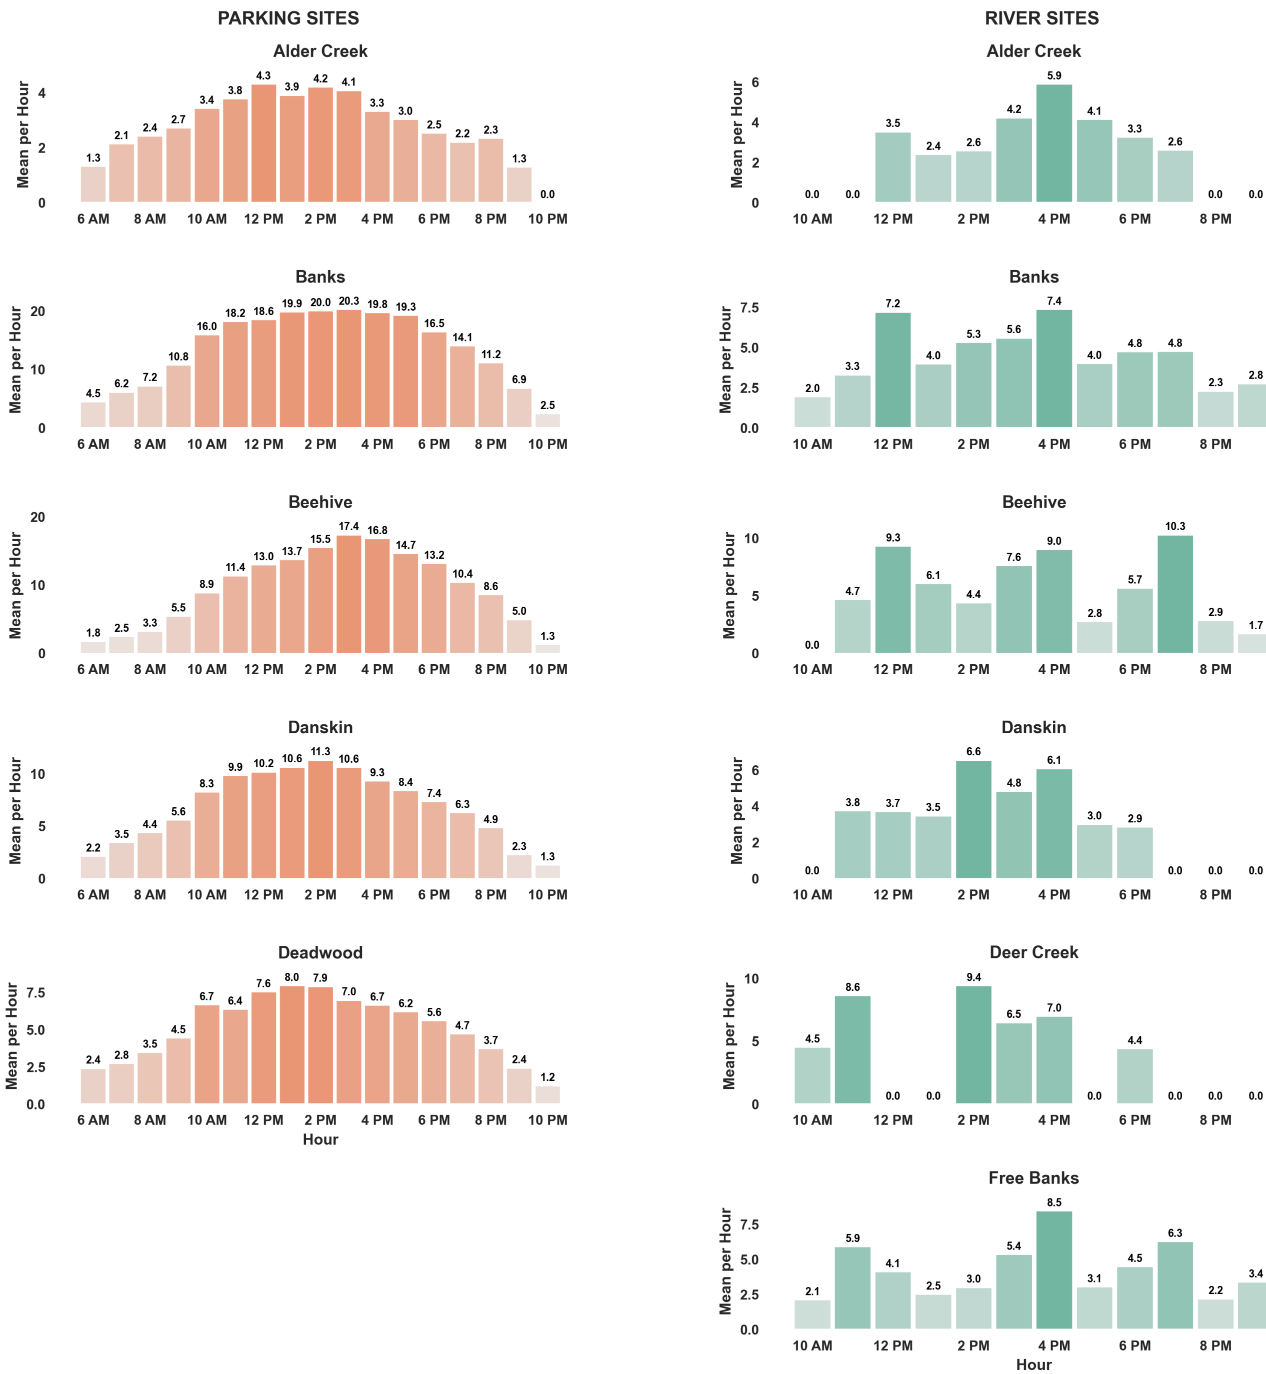

Figure 8: Mean Hourly Vehicles and Watercraft per Hour. Bars show the mean number of *unique* vehicles (parking sites) or *unique* watercraft (river sites) observed per hour. River-site values use watercraft-per-hour estimates calibrated with 67 hours of manual hand-tally counts (r = 0.35, Spearman's ρ = 0.23, R^2 = 0.12); parking-site values use model-derived vehicles-per-hour estimates. Plots are site specific and illustrate pronounced differences in both diurnal timing and overall volume among locations. Note that the y-axes are scaled independently by site to emphasize within-site temporal patterns.

peak safety risk periods (Orsi, 2015). Importantly, the indicators support more defensible decision-making because they quantify conditions rather than relying on anecdotal reports of crowding.

Third, corridor indicators are particularly valuable for cross-jurisdictional coordination, which is often the limiting factor in corridor management. The Payette River corridor involves land management agencies, transportation agencies, counties, and commercial operators. Multi-indicator monitoring provides a common empirical language for coordination: transportation agencies can interpret roadway delay and access-area pressure; land managers can interpret staging capacity and on-corridor use; outfitters can interpret peak launch windows and safety-relevant congestion. Even without formal standards, the ability to align interpretations of "peak pressure" across entities can reduce conflict and support joint planning. For example, occupancy peaks that coincide with known traffic bottlenecks provide an evidence base for coordinated traffic operations (e.g., temporary controls, signage, enforcement). Similarly, on-river use peaks that coincide with staging congestion provide an evidence base for coordinating launch-area design, safety messaging, or staggered departures.

Fourth, the indicator system supports scenario testing. Because parking opportunity produces spatially explicit "supply" estimates and camera-based indicators produce time-resolved "demand" proxies, managers can evaluate plausible interventions by adjusting indicator inputs and assessing expected impacts. For example, a proposed new access point can be evaluated by estimating parking opportunity and then modeling how redistributed arrivals might reduce pressure at current bottlenecks. A shuttle scenario can be evaluated by simulating reduced vehicle occupancy at areas and comparing expected effects on on-river use. A lot expansion scenario can be evaluated by changing the supply proxy and assessing whether occupancy peaks remain the limiting factor. While this paper does not implement a full predictive model of redistribution, the indicators provide the empirical baseline needed to support such analyses.

5.3. Transferability

The workflow is designed to transfer cleanly to other corridor types, with adjustments primarily driven by sensor availability and domain conditions. The parking capacity estimation workflow generalizes to settings where high-resolution aerial imagery is available and where parking opportunity can be inferred from surface appearance and terrain constraints. This includes other river corridors with dispersed access points, trail corridors with multiple trailheads and roadside pullouts, scenic byways with turnout networks, and coastal access corridors where parking supply is spatially fragmented. The camera-based use estimation workflow generalizes to any corridor where fixed cameras can observe staging areas and/or corridor segments, including winter trailheads, mountain bike trail systems, and river access networks.

Transferability will be strongest where three conditions hold. First, the aerial imagery domain should be sufficiently consistent with the training data or allow incremental relabeling and fine-tuning; canopy cover, strong seasonal changes (e.g., snow), and persistent shadow can alter appearance and will typically require additional labeled examples and/or seasonal models (Lupp et al., 2021; Mitterwallner et al., 2024). Second, camera placement must provide stable viewsheds; mounting constraints, and changing vegetation can impact both detection accuracy and comparability over time (Miller et al., 2017; Sunger et al., 2012). Third, legal and social contexts matter: privacy expectations and policies vary by jurisdiction, and any deployment in public settings should be designed to minimize risk and maximize legitimacy.

The model-training strategy used in the camera-based use estimation workflow—baseline inference followed by active learning and hard mining—also transfers cleanly to other rare-object monitoring problems. In many corridors, rare events or rare user types drive management concern. The demonstrated improvements in watercraft performance across iterative learning suggest that a modest amount of targeted annotation, informed by uncertainty and hard-example selection, can yield substantial gains in indicator validity (Agarwal et al., 2020; L. Liu et al., 2026; Oksuz et al., 2021).

5.4. Limitations and Ethics

This study has several limitations that affect how the indicators should be interpreted and used. First, physical capacity indicators are not equivalent to acceptable capacity. Parking opportunity, occupancy proxies, and on-river use quantify physical constraints and system pressure, but they do not define what levels of use are socially acceptable or ecologically sustainable. Standards for acceptable conditions require integration with social indicators (e.g., perceived crowding, encounter norms) and ecological indicators (e.g., bank trampling, wildlife disturbance), aligned with explicit management objectives (Interagency Visitor Use Management Council, 2016). In this sense, the indicators developed here should be viewed as enabling infrastructure for capacity frameworks, not as capacity thresholds themselves.

Second, the indicators inherit biases and uncertainties from the vision models and sensors. Segmentation of aerial imagery can confuse visually similar classes; in this study, confusion between roadway hardscape and paved parking is an example of an error mode that can inflate supply estimates in unsafe or unauthorized areas. Similarly, camera-based use detection is sensitive to occlusion, glare, low light, and clustering—conditions that can lead to undercounting or overcounting. Seasonal differences (leaf-on/leaf-off, snow, water glare) may require additional training data or season-specific models. These limitations underscore the importance of reporting model performance and using qualitative validation panels alongside quantitative metrics, particularly when indicators are used for management actions (Fennell et al., 2022; Mitterwallner et al., 2024).

Third, privacy and ethics require deliberate design. Fixed cameras in public settings raise concerns even when imagery is used for aggregate counts rather than identification. Practical safeguards include placing cameras to avoid capturing faces where possible, limiting resolution to what is needed for detection, applying masking to sensitive regions, posting signage where required or appropriate, and implementing data retention policies that minimize storage of raw imagery once detections are extracted (Lupp et al., 2021). These safeguards are not merely ethical considerations; they are also operational requirements for sustained deployment and legitimacy.

Finally, the corridor indicator system depends on assumptions about representativeness. Camera viewsheds do not capture the entire corridor, and occupancy proxies may not translate to percent occupancy without calibration. The method is therefore best used as a comparative system—tracking changes over time, contrasting sites, and identifying peaks—unless additional calibration steps are undertaken.

6. Conclusion

This manuscript provides modern measurement workflows that makes multi-indicator physical capacity and use assessment feasible. By combining semantic segmentation of aerial imagery for parking opportunity with camera-based object detection for parking occupancy and on-river use, the workflow produces spatially explicit and time-resolved indicators that are transparent and replicable. The Payette River corridor case demonstrates that these indicators can reveal where physical constraints concentrate, when peak pressure occurs, and how capacity-relevant conditions vary across access areas and corridor segments—information that is difficult to obtain consistently with conventional monitoring approaches. In applied terms, the workflows support investment targeting, trigger-based management during predictable peak windows, and cross-jurisdictional coordination by providing a common empirical basis for decision-making. More broadly, the study illustrates a transferable applied ML pattern for natural resource and recreation contexts: multi-source perception, label-efficient iterative improvement, and translation of model outputs into interpretable indicators suitable for operational planning and management.

CRediT authorship contribution statement

Jordan W. Smith: Conceived and designed the study; led all data preparation and analyses, visualizations, and developed the first draft of the manuscript. **Chase C. Lamborn:** Led the deployment of the in-field cameras; provided subsequent revisions and approved the final version.

References

Aas, Ø., & Onstad, O. (2013). Strategic and temporal substitution among anglers and white-water kayakers: The case of an urban regulated river. *Journal of Outdoor Recreation and Tourism*, *1–2*, 1–8. https://doi.org/10.1016/j.jort.2013.04.002

Agarwal, S., Arora, H., Anand, S., & Arora, C. (2020). Contextual diversity for active learning. In A. Vedaldi, H. Bischof, T. Brox, & J.-M. Frahm (Eds.), *Computer Vision – ECCV 2020* (pp. 137–153). Springer International Publishing. https://doi.org/10.1007/978-3-030-58517-4_9

Alotaibi, E., & Nassif, N. (2024). Artificial intelligence in environmental monitoring: In-depth analysis. *Discover Artificial Intelligence*, *4*(1), 84. https://doi.org/10.1007/s44163-024-00198-1

American Whitewater. (2026). Payette, Banks to Beehive Bend.

Arnberger, A., & Brandenburg, C. (2007). Past on-site experience, crowding perceptions, and use displacement of visitor groups to a peri-urban national park. *Environmental Management*, *40*(1), 34–45. https://doi.org/10.1007/s00267-004-0355-8

Bureau of Land Management. (2026). Payette River: Beehive Bend.

Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., & Adam, H. (2018). Encoder-decoder with atrous separable convolution for semantic image segmentation. *Proceedings of the European Conference on Computer Vision (ECCV)*, 801–818.

Clawson, M., & Knetsch, J. (1966). *Economics of outdoor recreation*. Johns Hopkins University Press.

Clawson, M. (1963). *Land and water for recreation*. Rand McNally.

Cole, D. N., & Hall, T. E. (2012). The effect of use density and length of stay on visitor experience in wilderness. In D. N. Cole (Ed.), *Wilderness visitor experiences: Progress in research and management* (pp. 77–95, Vol. 66). USDA Forest Service, Rocky Mountain Research Station.

Fennell, M., Beirne, C., & Burton, A. C. (2022). Use of object detection in camera trap image identification: Assessing a method to rapidly and accurately classify human and animal detections for research and application in recreation ecology. *Global Ecology and Conservation*, *35*, e02104. https://doi.org/10.1016/j.gecco.2022.e02104

Fisher, A. C., & Krutilla, J. V. (1972). Determination of optimal capacity of resource-based recreation facilities. *Natural Resources Journal*, *12*(3), 417–444.

Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., & Moore, R. (2017). Google Earth Engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, *202*, 18–27. https://doi.org/10.1016/j.rse.2017.06.031

Guidosse, Q., Breyne, J., Cioppa, A., Maréchal, K., Rubens, U., Van Droogenbroeck, M., & Dufrêne, M. (2025). Combining camera traps and artificial intelligence for monitoring visitor frequencies in natural areas: Lessons from a case study in the Belgian Ardenne. *Journal of Outdoor Recreation and Tourism*, *49*, 100856. https://doi.org/10.1016/j.jort.2025.100856

Hammitt, W. E., Cole, D. N., & Monz, C. A. (2015). *Wildland recreation: Ecology and management* (3rd). Wiley Blackwell.

Hammitt, W. E. (1980). Outdoor recreation: Is it a multi-phase experience? *Journal of Leisure Research*, *12*(2), 107–115. https://doi.org/10.1080/00222216.1980.11969430

Hauser, S., Augner, L., & Schmitt, A. (2025). Perfect labelling: A review and outlook of label optimization techniques in dynamic Earth observation. *Remote Sensing*, *17*(7). https://doi.org/10.3390/rs17071246

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.

Heywood, J. L. (2002). The cognitive and emotional components of behavior norms in outdoor recreation. *Leisure Sciences*, *24*(3-4), 271–281. https://doi.org/10.1080/01490400290050727

Idaho Transportation Department. (2026). SH-55 Banks-Lowman Intersection Improvements.

Interagency Visitor Use Management Council. (2016). *Visitor use management framework: A guide to providing sustainable outdoor recreation* (tech. rep.). Interagency Visitor Use Management Council.

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollar, P., & Girshick, R. (2023). Segment Anything. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.

Lindsey, G., Singer-Berk, L., Wilson, J. S., Oberg, E., & Hadden-Loh, T. (2018). Challenges in monitoring regional trail traffic. *Transportation Research Record*, *2672*(43), 98–109. https://doi.org/10.1177/0361198118787996

Liu, L., Liang, Y., Yan, X., Huangfu, L., Samtani, S., Yu, Z., Zhang, Y., & Zeng, D. D. (2026). Hard sample mining: A new paradigm of efficient and robust model training. *IEEE Transactions on Neural Networks and Learning Systems*, *37*(2), 526–546. https://doi.org/10.1109/TNNLS.2025.3610948

Liu, Q., Huang, T., Yang, J., & Xiang, W. (2026, February). From pixels to images: A structural survey of deep learning paradigms in remote sensing image semantic segmentation. https://doi.org/10.48550/arXiv.2505.15147

Lupp, G., Kantelberg, V., Förster, B., Honert, C., Naumann, J., Markmann, T., Pauleit, S., Lupp, G., Kantelberg, V., Förster, B., Honert, C., Naumann, J., Markmann, T., & Pauleit, S. (2021). Visitor counting and monitoring in forests using camera traps: A case study from Bavaria (Southern Germany). *Land*, *10*(7). https://doi.org/10.3390/land10070736

Miller, A. B., Leung, Y.-F., & Kays, R. (2017). Coupling visitor and wildlife monitoring in protected areas using camera traps. *Journal of Outdoor Recreation and Tourism*, *17*, 44–53. https://doi.org/10.1016/j.jort.2016.09.007

Mitterwallner, V., Peters, A., Edelhoff, H., Mathes, G., Nguyen, H., Peters, W., Heurich, M., & Steinbauer, M. J. (2024). Automated visitor and wildlife monitoring with camera traps and machine learning. *Remote Sensing in Ecology and Conservation*, *10*(2), 236–247. https://doi.org/10.1002/rse2.367

Moore, B. E., & Corso, J. J. (2026). FiftyOne.

Moreno, H., Gómez, A., & Andújar, D. (2025). Deep learning-based video analysis for visitor detection and tracking in protected areas. *Journal of Outdoor Recreation and Tourism*, *50*, 100890. https://doi.org/10.1016/j.jort.2025.100890

Oksuz, K., Cam, B. C., Kalkan, S., & Akbas, E. (2021). Imbalance problems in object detection: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *43*(10), 3388–3415. https://doi.org/10.1109/TPAMI.2020.2981890

Orsi, F. (Ed.). (2015). *Sustainable transportation in natural and protected areas*. Routledge.

Sharma, S., & Gosain, A. (2025). Addressing class imbalance in remote sensing using deep learning approaches: A systematic literature review. *Evolutionary Intelligence*, *18*(1), 23. https://doi.org/10.1007/s12065-024-01012-8

Sidder, S. A., & Hall, T. E. (2024). Intra-spatial displacement at high-use recreation destinations: Employing spatial analyses to explore visitor behavior. *Journal of Outdoor Recreation and Tourism*, *46*, 100770. https://doi.org/10.1016/j.jort.2024.100770

Staab, J., Udas, E., Mayer, M., Taubenböck, H., & Job, H. (2021). Comparing established visitor monitoring approaches with triggered trail camera images and machine learning based computer vision. *Journal of Outdoor Recreation and Tourism*, *35*, 100387. https://doi.org/10.1016/j.jort.2021.100387

Stewart, W. P. (1998). Leisure as multiphase experiences: Challenging traditions. *Journal of Leisure Research*, *30*(4), 391.

Sunger, N., Teske, S. S., Nappier, S., & Haas, C. N. (2012). Recreational use assessment of water-based activities, using time-lapse construction cameras. *Journal of Exposure Science & Environmental Epidemiology*, *22*(3), 281–290. https://doi.org/10.1038/jes.2012.4

U.S. Geological Survey. (2024). 3D Elevation Program 1-Meter resolution digital elevation model.

U.S. Geological Survey. (2026). National Agriculture Imagery Program. https://doi.org/10.5066/F7QN651G

Wang, X., Kim, H., Rahman, S., Mitra, K., & Miao, Z. (2024). Human-LLM collaborative annotation through effective verification of LLM labels. *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–21. https://doi.org/10.1145/3613904.3641960

Watson, A. E., Cole, D. N., Turner, D. L., & Reynolds, P. S. (2000). *Wilderness recreation use estimation: A handbook of methods and systems* (General Technical Report No. RMRS-GTR-56). USDA Forest Service, Rocky Mountain Research Station.

Wilkins, E. J., Van Berkel, D., Zhang, H., Dorning, M. A., Beck, S. M., & Smith, J. W. (2022). Promises and pitfalls of using computer vision to make inferences about landscape preferences: Evidence from an urban-proximate park system. *Landscape and Urban Planning*, *219*, 104315. https://doi.org/10.1016/j.landurbplan.2021.104315

Ziesler, P. S., & Pettebone, D. (2018). Counting on visitors: A review of methods and applications for the National Park Service's visitor use statistics program. *Journal of Park and Recreation Administration*, *36*(1). https://doi.org/10.18666/JPRA-2018-V36-I1-8104