Environmental Conservation

cambridge.org/enc

Can artificial intelligence support Bactrian camel conservation? Testing machine learning on aerial imagery in Mongolia's Gobi Desert

Chris McCarthy[1], Simon Phillips[2], Troy Sternberg[3,4], Adiya Yadamsuren[5], Battogtokh Nasanbat[6], Kyle Shaney[7], Buho Hoshino[8] and Erdenebuyan Enkhjargal[9]

[1]Zanvyl Krieger School of Arts & Sciences, Johns Hopkins University, Baltimore, MD, USA; [2]Centre for Biocultural Diversity, School of Anthropology and Conservation, University of Kent, Canterbury, UK; [3]School of Geography, University of Oxford, Oxford, UK; [4]CEI Centre for International Studies ISCTE – University Institute Lisbon, Avenida das Forças Armadas, Lisbon, Portugal; [5]Institute of Zoology, Mongolia, Ulaanbaatar, Mongolia; [6]Institute of Biology, Mongolian Academy of Sciences, Ulaanbaatar, Mongolia; [7]Department of Biology, Health, and the Environment, The University of Texas at San Antonio, San Antonio, TX, USA; [8]Lab of Environmental Remote Sensing, Department of Environmental Sciences, College of Agriculture, Food and Environment Sciences, Rakuno Gakuen University, Hokkaido, Japan and [9]Graduate School of Global Studies, Doshisha University, Kyoto, Japan

Research Paper

Cite this article: McCarthy C et al. (2025) Can artificial intelligence support Bactrian camel conservation? Testing machine learning on aerial imagery in Mongolia's Gobi Desert. *Environmental Conservation* **52**: 149–156. doi: 10.1017/S0376892925100118

Received: 23 February 2025
Revised: 2 June 2025
Accepted: 2 June 2025
First published online: 18 July 2025

Keywords:
Artificial intelligence; camels; conservation monitoring; deep learning; desert ecosystems; drone technology; Gobi Desert; machine learning; wildlife surveillance; YOLOv8

Corresponding author:
Chris McCarthy; Email: cmccar27@jh.edu

Summary

Monitoring wildlife populations in vast, remote landscapes poses significant challenges for conservation and management, particularly when studying elusive species that range across inaccessible terrain. Traditional survey methods often prove impractical or insufficient in such environments, necessitating innovative technological solutions. This study evaluates the effectiveness of deep learning for automated Bactrian camel detection in drone imagery across the complex desert terrain of the Gobi Desert of Mongolia. Using YOLOv8 and a dataset of 1479 high-resolution drone-captured images of Bactrian camels, we developed and validated an automated detection system. Our model demonstrated strong detection performance with high precision and recall values across different environmental conditions. Scale-aware analysis revealed distinct performance patterns between medium- and small-scale detections, informing optimal drone flight parameters. The system maintained consistent processing efficiency across various batch sizes while preserving detection quality. These findings advance conservation monitoring capabilities for Bactrian camels and other wildlife in remote ecosystems, providing wildlife managers with an efficient tool to track population dynamics and inform conservation strategies in expansive, difficult-to-access habitats.

Introduction

The Gobi Desert, spanning southern Mongolia and northern China, hosts diverse wildlife, including significant populations of Bactrian camels (Hare 2008). As one of the world's largest desert ecosystems, the Gobi Desert presents unique challenges for wildlife monitoring and conservation (Payne et al. 2020). Traditional wildlife observation methods in such vast, often inaccessible terrains include ground-based vehicle surveys, observation posts at water sources, camera traps and satellite collar tracking, all of which are time-consuming, resource-intensive and prone to human error (Linchant et al. 2015). Recently, the integration of unmanned aerial vehicles and artificial intelligence (AI) has opened new avenues for efficient and accurate wildlife monitoring in remote habitats (Hodgson et al. 2018).

Bactrian camels exist as both wild (*Camelus bactrianus ferus*) and domestic (*Camelus bactrianus*) populations in these landscapes. The wild Bactrian camel is Critically Endangered, with fewer than 1000 individuals remaining (Hare 2008). Conservation efforts focus on protecting key habitats and water resources, establishing protected areas and reducing human–wildlife conflict (Reading et al. 2001, Yadamsuren et al. 2012). However, effective conservation is hampered by the Gobi Desert's vastness and the camels' wide-ranging behaviour, which make traditional field-based surveys with limited personnel challenging (Kaczensky et al. 2014).

Drone technology has revolutionized wildlife monitoring, offering non-invasive, quick surveys with minimal disturbance (Hodgson et al. 2016). However, the volume of data from drone surveys poses a new challenge: efficient and accurate image analysis (Weinstein 2018). Machine learning, especially deep learning models for object detection, shows promise in addressing this challenge. These algorithms can rapidly process vast amounts of visual data, potentially surpassing human capabilities in speed and accuracy (Norouzzadeh et al. 2018, Tabak et al. 2019).

Recent studies demonstrate the potential of combining drone imagery with machine learning for wildlife monitoring in various ecosystems (Hodgson et al. 2018, Duporge et al. 2021). While

© The Author(s), 2025. Published by Cambridge University Press on behalf of Foundation for Environmental Conservation. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

these approaches have shown promise in other environments, their application to desert ecosystems, particularly for wildlife detection in the complex terrain of the Gobi Desert, remains largely unexplored. Similar techniques have been applied to other large mammals (Kellenberger et al. 2018, McCarthy et al. 2024), but the Gobi Desert's varied desert terrain and the specific challenges of camel detection warrant dedicated research. Several factors make camel detection in the Gobi Desert particularly challenging: the camels' coloration often blends with the desert background, making visual distinction difficult, and the terrain requires robust detection algorithms capable of performing across different backdrops (Kaczensky et al. 2014). These challenges are compounded by the need to process high-resolution imagery efficiently while maintaining detection accuracy across different scales and environmental conditions.

This study addresses these challenges by evaluating the effectiveness of YOLOv8, a state-of-the-art object detection model, for automated camel detection in drone imagery across the Gobi Desert. Using a comprehensive dataset of 1479 high-resolution drone-captured images of Bactrian camels, we assess the model's performance across different desert landscapes with the explicit goal of enhancing current conservation practices. We test whether automated detection can complement traditional survey methods to develop more robust population monitoring protocols. By establishing the technical feasibility and practical limitations of this approach, we aim to provide conservation managers with evidence-based guidance for integrating drone surveys and machine learning into existing monitoring frameworks. This integration is particularly important for the endangered wild Bactrian camel population, for which improved survey efficiency could significantly enhance conservation outcomes.

Methodology

Study area and data collection

This study on the Mongolian Gobi Desert encompasses four *aimags* (provinces): Dornogovi, Ömnögovi, Bayankhongor and Govi-Altai (Fig. 1). The area features sandy and rocky deserts, oases and sparse vegetation. Local herders and conservation teams from the Mongolia Wild Camel Foundation helped identify key Bactrian camel locations across these regions, providing valuable knowledge about typical movement patterns and gathering areas that informed our survey site selection and flight planning.

Data collection occurred between June and July 2024, during peak summer camel activity, when the animals regularly visit water sources. Our field methodology employed a targeted, expertise-based approach rather than systematic transects. Knowledgeable local guides helped direct us to known camel locations across the region (Fig. 1). Upon locating camel groups, we conducted drone flights to collect imagery, with flight sessions varying in duration (typically 15–40 min) based on several factors: the number of camels present in each group, the opportunity to capture images under diverse lighting conditions (such as when passing clouds created shadows) and drone battery limitations. Our field team conducted drone surveys while maintaining a safe distance of 500–1000 m from the camels to minimize disturbance. Images were captured across various daylight hours and weather conditions, including early morning, midday and late afternoon, and under different sky conditions ranging from clear to overcast. This variation in lighting and atmospheric conditions provided a more comprehensive dataset, helping us to develop a more robust model capable of performing under real-world monitoring scenarios. While we collected data across these various environmental conditions, our analytical approach focused on establishing overall model performance rather than comparing effectiveness across specific terrains or lighting scenarios, as this preliminary study aimed to validate the fundamental capability of YOLOv8 for camel detection in the Gobi Desert environment.

We captured 1479 high-resolution aerial images using a single DJI Mini Pro 4 drone equipped with a 1/1.3-inch CMOS sensor (48 MP, f/1.7 aperture, 24-mm equivalent focal length). The drone operated at 100 m above ground level to minimize disturbance to the camels while achieving a ground sampling distance of 2.5 cm/pixel. All images were processed to a standardized 4096 × 2304 pixel resolution. Flight paths were designed to cover representative samples of Gobi Desert terrain, including dunes, gravel plains and areas near water sources.

The study area encompasses desert landscapes typical of the Gobi Desert region (Fig. 2). These include, in approximate percentages of our dataset, sparsely vegetated sandy dunes (31%), gravel plains with desert shrubs (28%), semi-stabilized terrain with saxaul (*Haloxylon ammodendron*) vegetation (24%) and bare rocky outcrops (17%). These terrain types influence image contrast and feature recognition for camel detection.

The DJI Mini Pro 4 was selected as it represents an accessible and cost-effective platform for wildlife monitoring, offering the necessary performance capabilities while remaining within the budget constraints typical of conservation projects. Despite being a consumer-grade drone, its high-resolution camera and stable flight characteristics at an altitude of 100 m proved sufficient for reliable camel detection.

We utilized *LabelImg* software (Tzutalin 2015) for object detection annotation, choosing this tool for its offline capabilities and direct compatibility with YOLO format annotations. Following a standardized protocol, all 1479 images were annotated using bounding boxes with a single class designation of 'camel', focusing on fundamental detection capability rather than demographic classification. The protocol specified that boxes should encompass the entire visible body of each camel, excluding shadows. For partially visible camels, boxes were drawn around the visible portion if at least 50% of the animal was visible. In cases of overlapping camels, individual boxes were drawn for each animal where distinct boundaries could be determined. When camels were too closely clustered to distinguish individuals, the entire group was marked with a single bounding box and noted in our documentation.

Machine learning model development, training and optimization

The dataset was split into training (1035 images, 70%), validation (222 images, 15%) and test sets (222 images, 15%), ensuring a representative distribution of camel instances across all sets. We employed YOLOv8 (Jocher et al. 2023) for camel detection, utilizing a configuration optimized for processing aerial imagery at 4096 × 2304 resolution. Data augmentation techniques included rotation (±15°), scaling (±20%), translation (±20%), horizontal flipping (probability 0.5) and HSV colour space adjustments (hue: ±0.015, saturation: ±0.7, value: ±0.4) to enhance model robustness. Training was conducted on Google Colab Pro+ using an NVIDIA A100-SXM4-40GB GPU. The model was configured with a batch size of 1 and standardized input dimensions of 4096 × 2304 pixels using an AdamW optimizer with an initial learning rate of 0.01,

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press

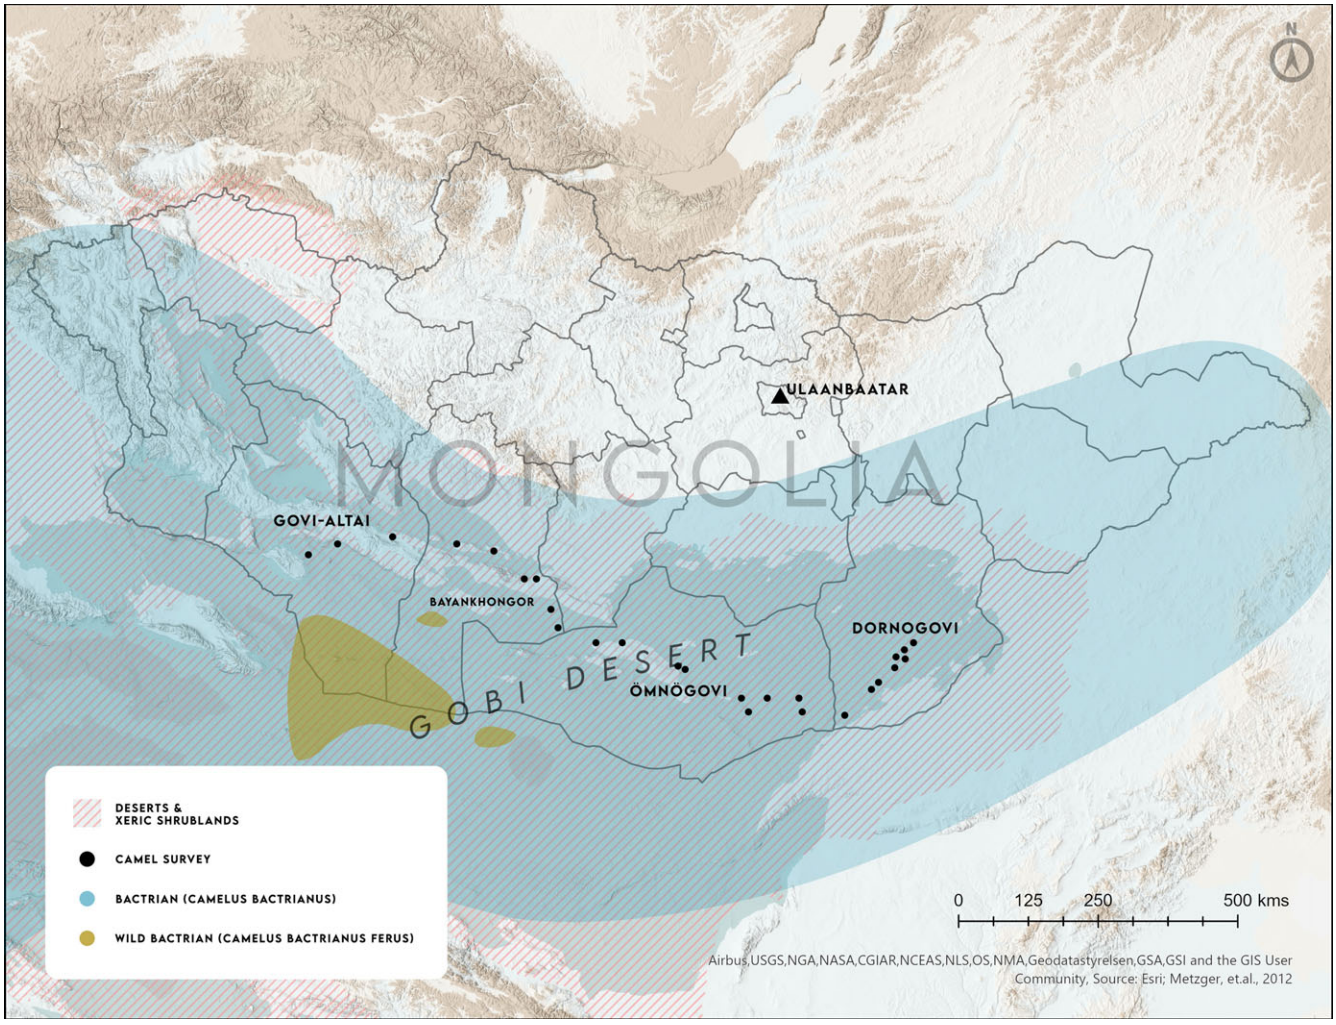


Figure 1. Map of the study area in the Mongolian Gobi Desert. The map shows the boundaries of Dornogovi, Ömnögovi, Bayankhongor and Govi-Altai *aimags* (provinces), with black dots indicating camel survey locations.

final learning rate of 0.001, momentum of 0.937 and weight decay of 0.0005. Training proceeded for 50 epochs with a cosine learning rate schedule and 3 warmup epochs. We employed automatic mixed precision training and RAM caching to optimize memory usage, with two worker processes for data loading.

Evaluation metrics and performance analysis

For evaluation, we employed standard object detection metrics including precision, recall, F1-score and mean Average Precision (mAP) at Intersection over Union (IoU) thresholds ranging from 0.5 to 0.95. Model predictions were saved with confidence scores and in YOLO-compatible text format for detailed analysis. The Scale-Aware Performance Analysis evaluated detection characteristics across different sizes, categorizing detections based on their pixel area into small ($<64 \times 64$ pixels) and medium-sized (64×64–128×128 pixels) instances. Model inference speed was evaluated on the NVIDIA A100 GPU configuration, measuring processing time per image, including pre-processing, inference and post-processing steps. The test dataset was processed through the model using a batch size of 1 to assess both detection accuracy and computational efficiency under consistent conditions.

Practical application assessment

Our assessment framework focused on two key aspects: detection accuracy and processing efficiency. Detection accuracy was measured through the comprehensive metrics suite, capturing true positives, false positives and missed detections. Processing efficiency was evaluated through detailed timing analysis of each pipeline stage (pre-processing, inference and post-processing) and GPU memory utilization during operation, informing real-world deployment requirements.

Results

Model performance in desert environments

Our implementation of YOLOv8-large successfully detected Bactrian camels across the Gobi Desert terrain and achieved strong performance metrics. During the training process (Fig. S1), the model showed systematic improvement, with training losses decreasing steadily across all components: box loss (2.2 to 1.2), classification loss (2.0 to 0.5) and Distribution Focal Loss (1.6 to 1.0). Validation metrics followed similar patterns, indicating robust learning without overfitting.

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press

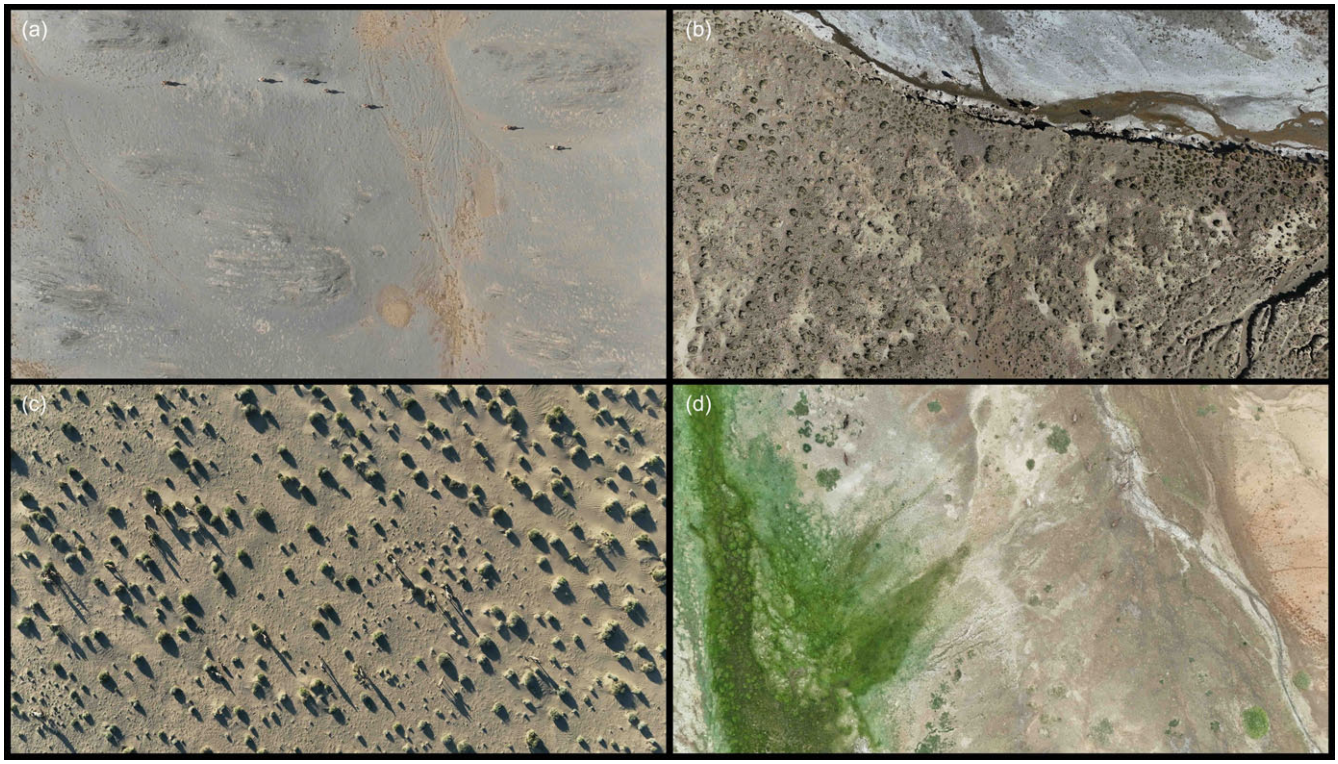


Figure 2. Representative landscapes from the study area in the Mongolian Gobi Desert captured at 100 m altitude with 4096 × 2304 resolution (2.5 cm/pixel ground sampling distance). (a) Barren rocky terrain characteristic of upland areas; (b) ephemeral riverbed terrain in arid landscape; (c) semi-stabilized terrain with saxaul (*Haloxylon ammodendron*) vegetation; and (d) oasis habitat with concentrated green vegetation. These diverse landscapes represent the range of backgrounds against which camel detection must function, allowing for clear identification of individual animals against varying desert backdrops.

The model's final performance exceeded our initial expectations (Table 1). The high mAP indicates that the model could reliably detect camels even in challenging desert conditions, where animals often blend in with the terrain. The balanced precision and recall values demonstrate the model's ability to avoid both false detections and missed animals.

Field application success rates

The confusion matrix analysis (Fig. S2) reveals the practical effectiveness of our approach. From 2251 actual camel instances in the test set, the model successfully identified 2120 (94.2% true positive rate) while generating only 406 false positives. Perhaps more importantly for wildlife monitoring applications, the model missed just 131 camels, suggesting its reliability for population surveys.

Impact of detection scale

The model's performance varied notably with the size of the detected camels in the images, with significant differences across 3108 total detections (Table 2). Medium-sized detections (64 × 64–128 × 128 pixels) showed markedly higher confidence scores (0.657 ± 0.121) compared to small-scale detections (<64 × 64 pixels, 0.548 ± 0.146), indicating optimal drone flight heights for future surveys. This scale-dependent performance helps establish practical guidelines for field applications.

Table 1. Core performance metrics of the YOLOv8 model. mAP50 represents the mean Average Precision at the 50% Intersection over Union (IoU) threshold; mAP50–95 is the mean Average Precision averaged across IoU thresholds from 50% to 95%; Precision indicates the proportion of detections that are correct; Recall shows the proportion of actual camels that are detected; and F1-score is the harmonic mean of precision and recall. For all metrics, higher values indicate better performance.

Metric	Value
mAP50	0.907
mAP50–95	0.450
Precision	0.876
Recall	0.882
F1-score	0.879

Detection confidence analysis

Analysis of 596 detections across 50 test images revealed a mean confidence score of 0.609 ± 0.143. Confidence scores ranged from 0.250 to 0.811, demonstrating varying levels of detection certainty. The distribution of confidence scores showed that 198 instances (33.2%) were high-confidence detections, with scores above 0.7, while 289 instances (48.5%) achieved medium confidence scores between 0.5 and 0.7. The remaining 109 instances (18.3%) were low-confidence detections, with scores below 0.5. This distribution pattern indicates that the majority of detections (81.7%) achieved medium to high confidence scores, suggesting reliable detection performance across the test set.

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press

Table 2. Scale-based detection performance, comparing metrics between camels appearing as small objects (<64 × 64 pixels, typically more distant from the drone) versus medium-sized objects (64 × 64–128 × 128 pixels, typically closer to the drone) in the images. No large-scale detections (>128 × 128 pixels) were present in our dataset given the flight altitude. Confidence values represent the model's certainty in its detections, not overall accuracy, which explains why these values are lower than the overall precision/recall metrics.

Scale category	Count	Percentage	Mean confidence	SD	Minimum confidence	Maximum confidence
Small (<64 × 64)	1378	44.34%	0.548	±0.146	0.250	0.784
Medium (64 × 64–128 × 128)	1730	55.66%	0.657	±0.121	0.250	0.811

Computational performance and performance across terrain types

The processing efficiency evaluation on our NVIDIA A100-SXM4-40GB GPU, running *CUDA* 12.4 and *PyTorch* 2.5.1, demonstrated consistent performance across batch sizes (Table S1). The model, comprising 43.6M parameters, achieved efficient scaling as batch size increased from 1 to 16, improving per-image processing time while maintaining detection quality.

The model demonstrated performance across representative terrain types characteristic of the Gobi Desert (Fig. 3). The examples illustrate detection capabilities in saxaul shrub landscapes, where scattered vegetation creates complex backgrounds; in vegetated and dry terrain, demonstrating adaptability to mixed ground cover; along stream corridors with patches of green vegetation; and in barren rocky landscapes with minimal vegetation. These environments represent the range of detection challenges encountered in the Gobi Desert ecosystem.

Discussion

Our study demonstrates the effectiveness of drone-based Bactrian camel monitoring in the Gobi Desert, achieved through the successful implementation of deep learning technology. Our approach using YOLOv8 achieved strong detection performance, a level comparable to wildlife detection systems in other challenging environments (Kellenberger et al. 2018, McCarthy et al. 2024).

The balanced precision and recall values, combined with the confusion matrix analysis (Fig. S2) showing 2120 true positives and only 131 false negatives from 2251 instances, indicate robust detection capabilities that address key challenges identified in traditional wildlife monitoring approaches (Linchant et al. 2015). This performance level suggests that automated drone-based surveys could reduce the resource intensiveness and human error associated with conventional Bactrian camel population monitoring methods in the Gobi Desert (Kaczensky et al. 2014).

Scale-aware analysis revealed distinct performance patterns across different detection scales. Medium-sized detections constituted the majority of total detections with higher confidence scores, while small-scale detections showed lower confidence scores (Table 2). This scale-dependent performance aligns with findings from similar wildlife monitoring studies (Hodgson et al. 2018) and suggests that optimization of drone flight altitudes and imaging parameters is needed for improved detection reliability.

The high input resolution proved crucial for achieving these detection rates, particularly for smaller and distant camels. The 4096 × 2304 resolution ensured sufficient pixel density (2.5 cm/pixel) to maintain detection confidence even for small targets while remaining computationally feasible, as demonstrated by our processing efficiency analysis.

The processing efficiency analysis (Table S1) demonstrates rapid processing capability (c. 63 000 images per hour), which transforms how conservation teams can monitor Bactrian camel populations across the vast Gobi Desert landscape. This efficiency enables more frequent population assessments, allowing conservationists to detect population trends, identify critical habitats and respond quickly to emerging threats. For the Critically Endangered wild Bactrian camel, this enhanced monitoring frequency could provide early warning of population declines or habitat degradation.

This approach represents a complementary method to traditional camel monitoring rather than a replacement. While traditional ground-based surveys provide valuable demographic data and behavioural observations, they are limited by accessibility constraints and observer fatigue in the vast Gobi Desert landscape. The drone-based machine learning system excels at rapidly covering extensive areas, serving as an efficient first-pass survey tool to identify camel presence and distribution patterns. Optimal monitoring strategies would probably integrate both approaches: using automated drone surveys to locate camel groups across large areas, followed by targeted traditional ground-based observations for detailed behavioural and demographic assessment. This integration would significantly improve survey efficiency while maintaining the depth of information collected. During periods of extreme weather or in particularly remote regions, the drone system may serve as the primary monitoring tool when traditional approaches are impractical or unsafe.

However, several limitations warrant consideration. The confidence score distribution (mean 0.609 ± 0.143) suggests room for improvement in detection certainty. This limitation could be particularly relevant for surveys of wild Bactrian camels, which often inhabit more remote and challenging terrain (Hare 2008). Additionally, while our model shows strong performance on the test set, further validation across different seasons and lighting conditions would be valuable for understanding the system's year-round applicability. Because we only collected images over a 2-month period during summer, the effectiveness of this method may vary during other times of year. Seasonal changes in vegetation, lighting conditions, camel behaviour and coat appearance could influence detection performance. Bactrian camels undergo significant seasonal coat changes, with thick winter coats potentially altering their visual signature compared to their summer appearance. Winter months bring additional challenges, including potential snow cover in some regions that may alter background contrasts, while spring dust storms could reduce visibility and image quality. Further research across multiple seasons would be valuable to develop a robust year-round monitoring capability and to understand how model performance might need to be calibrated for different environmental conditions.

Our findings suggest several promising directions for future research. First, integrating temporal analysis capabilities could enhance the system's ability to track camel movement patterns and habitat use. Second, exploring multi-class detection to differentiate

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press

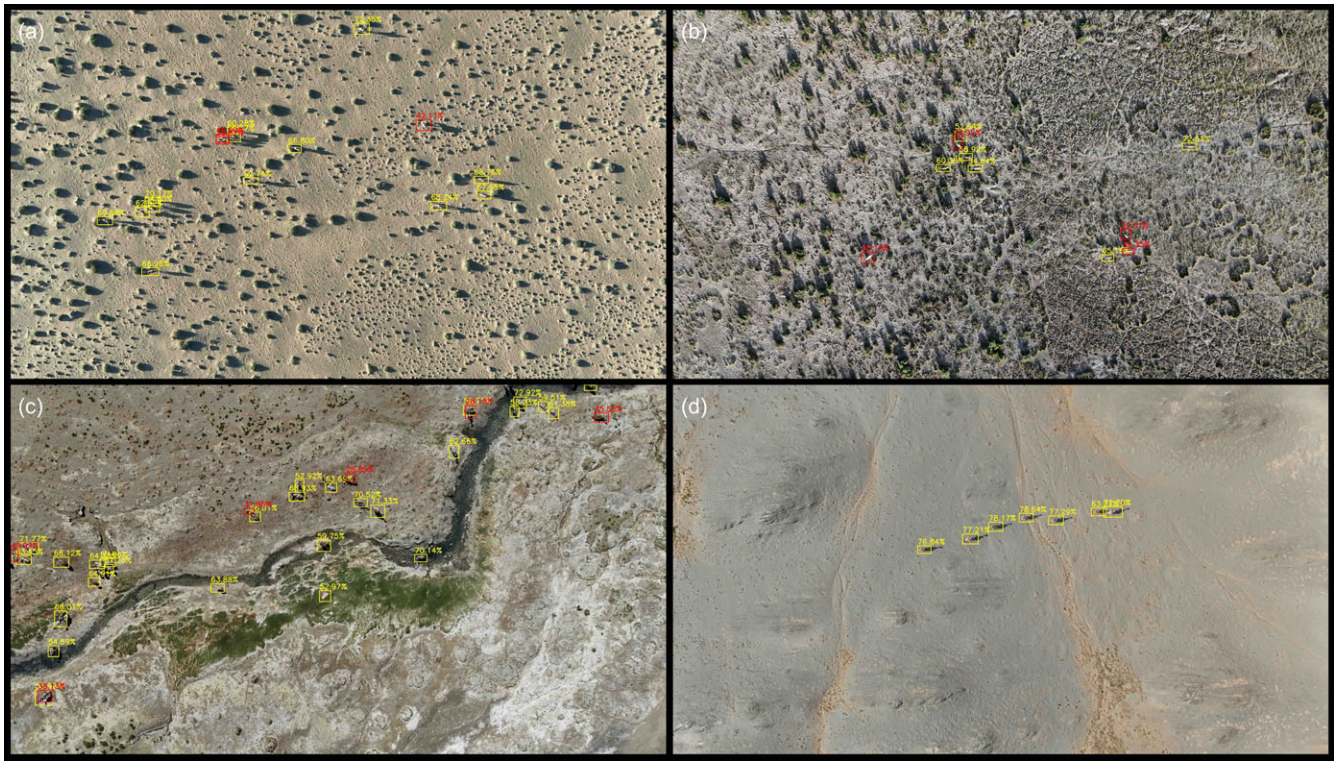


Figure 3. Representative detection examples from the test dataset across diverse Gobi Desert terrain types. (a) Camel detection in saxaul shrub landscape; (b) detection performance in vegetated and dry terrain; (c) detection along a stream corridor with green vegetation; and (d) detection in barren rocky terrain. Detection confidence is visualized through colour-coded bounding boxes, with yellow boxes indicating medium-confidence detections (0.5–0.8) and red boxes showing low-confidence detections (<0.5).

between wild and domestic Bactrian camels could provide valuable data for conservation efforts focused on the endangered wild Bactrian camel population. Finally, investigating transfer learning approaches could extend the system's utility to detecting other desert-dwelling species.

The practical implications of this work are particularly relevant for conservation efforts in the Gobi Desert. With a high true-positive rate and validated accuracy levels, the system offers a reliable tool for population monitoring, especially in areas that are difficult or dangerous to access using traditional survey methods. The demonstrated capability to process high-resolution imagery efficiently addresses one of the key challenges in drone-based wildlife surveys noted in previous research (Weinstein 2018): the balance between coverage area and detection accuracy. This achievement validates the effectiveness of the YOLOv8 architecture for wildlife detection in challenging environments, with the model's consistent performance across different terrain types demonstrating the potential of modern deep learning approaches. Our implementation advances the specific capabilities needed for reliable camel detection in desert ecosystems, providing a foundation for future automated wildlife monitoring applications in similarly challenging desert environments worldwide.

Beyond the Gobi Desert context, our approach has significant implications for wildlife monitoring practices in other arid regions. Based on our model's proven capability to distinguish camels from desert backgrounds across different terrain types (Fig. 3), the detection framework could be adapted for monitoring other camelid species in similar arid regions. For instance, it could be transferred to monitor dromedary camels in the Middle East, India and Australia's outback, where researchers face comparable challenges of vast terrains and elusive populations. While our methodology focused specifically on Bactrian camels, the scale-aware performance analysis provides insights for extending these principles to various large mammals in challenging landscapes.

The higher detection confidence observed in medium-sized objects compared to smaller objects suggests that for smaller species such as alpacas in the Andes Mountains, adjustments to drone flight altitude would be necessary to ensure adequate pixel density for reliable detection. For larger species such as Arctic caribou or other ungulates of the Gobi Desert, our current parameters may already be suitable, as indicated by the performance at our tested resolution and ground sampling distance.

Although species-specific model adaptations would be necessary, particularly regarding training datasets and class definitions, the fundamental approach of combining drone imagery with deep learning demonstrated in this study provides a versatile framework for wildlife monitoring across diverse ecosystems and species. The processing efficiency supports the practical applicability of this approach for large-scale conservation surveys. By demonstrating the feasibility of AI-powered wildlife detection in challenging desert conditions with high detection accuracy, our work contributes to the development of more efficient and accurate ecological survey methodologies, ultimately supporting better-informed conservation strategies not only for Bactrian camel populations but potentially for many other wide-ranging species in remote environments.

Conservation applications and future directions

Building on this technical validation, we propose a practical conservation monitoring framework that integrates drone surveys

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press

with AI-based image analysis and traditional methods. Key conservation applications include: (1) establishing baseline population densities across protected areas through systematic aerial surveys; (2) monitoring seasonal habitat use patterns to inform the protection of critical water sources and migration corridors; (3) rapid assessment of human–wildlife conflict zones where camel–herder interactions occur; and (4) post-disturbance population assessments following drought, disease outbreaks or extreme weather events. The validated detection performance provides the reliability needed for these conservation applications.

The next steps for such implementation include: (1) developing systematic transect designs that optimize coverage of key habitat areas, including both water sources and surrounding landscapes; (2) establishing protocols for population estimation that account for detection rates and spatial distribution patterns; and (3) exploring the integration of this technology with ground-based monitoring to create a comprehensive survey framework. While drone battery limitations restrict single-flight coverage, our approach envisions a network of strategic survey locations across the landscape, integrating knowledge of camel movement patterns to maximize detection probability. This network approach would enable conservation managers to efficiently monitor vast areas while focusing resources on critical habitats identified through initial surveys.

For distinguishing between wild and domestic camels, future research could focus on developing specialized computer vision models trained on high-resolution imagery that captures distinguishing morphological features such as differences in body conformation, coat texture and hump structure. Wild Bactrian camels typically display smaller, more symmetrical humps and a leaner overall body profile compared to their domestic counterparts, which often show more pronounced asymmetry in hump development. Capturing these distinctions would probably require lower-altitude imagery and specialized model training. This capability would be particularly valuable for conservation efforts, as accurate differentiation between wild and domestic populations is essential for assessing the true conservation status of the Critically Endangered wild Bactrian camel and for managing potential hybridization risks.

Conclusion

This study demonstrates that drone surveys combined with AI analysis can provide an effective approach for Bactrian camel monitoring in the Gobi Desert. The YOLOv8 model achieved detection performance sufficient for reliable population monitoring across varied desert terrain. By identifying optimal flight parameters through scale-aware analysis and establishing practical processing workflows, we demonstrate an efficient approach to population assessment. The integration of this technology with traditional monitoring methods offers a pathway to more comprehensive conservation strategies, which is particularly important for the Critically Endangered wild Bactrian camel. These findings contribute to the growing toolkit of technology-enhanced conservation approaches, with potential applications for monitoring other wide-ranging species in remote ecosystems. Future development of this approach, including multi-season validation and wild and domestic camel differentiation, will further strengthen conservation capacity in these challenging environments.

Supplementary material. To view supplementary material for this article, please visit https://doi.org/10.1017/S0376892925100118.

Data availability. The data that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgements. We would like to thank Mr Khasbaatar Khandmaa and the Mongolia Wild Camel Foundation for their assistance with logistics and local interviews. We thank Mr Anand Munkhuu for his support and valuable contributions to the study.

Author contributions. CM: Conceptualization, Methodology, Investigation, Fieldwork, Writing – Original Draft, Writing – Review & Editing, Software, Visualization, Funding Acquisition, Supervision, Project Administration; SP: Fieldwork, Writing – Original Draft, Writing – Review & Editing; TS: Writing – Review & Editing; AY: Conceptualization, Methodology, Investigation; BN: Writing – Review & Editing; KS: Writing – Original Draft, Writing – Review & Editing; BH: Writing – Original Draft, Writing – Review & Editing; EE: Writing – Review & Editing. All authors have read and agreed to the published version of the manuscript.

Financial support. This research was funded by the Council of American Overseas Research Centers (CAORC).

Competing interests. On behalf of all authors, the corresponding author states that there are no competing interests.

Ethical standards. Not applicable.

References

Duporge I, Isupova O, Reece S, Macdonald DW, Wang T (2021) Using very-high-resolution satellite imagery and deep learning to detect and count African elephants in heterogeneous landscapes. *Remote Sensing in Ecology and Conservation* 7: 369–381.

Hare J (2008) *Camelus ferus*. In *The IUCN Red List of Threatened Species* 2008: e.T63543A12689285 [www document]. URL https://doi.org/10.2305/IUCN.UK.2008.RLTS.T63543A12689285.en

Hodgson JC, Baylis SM, Mott R, Herrod A, Clarke RH (2016) Precision wildlife monitoring using unmanned aerial vehicles. *Scientific Reports* 6: 22574.

Hodgson JC, Mott R, Baylis SM, Pham TT, Wotherspoon S, Kilpatrick AD, et al. (2018) Drones count wildlife more accurately and precisely than humans. *Methods in Ecology and Evolution* 9: 1160–1167.

Jocher G, Chaurasia A, Qiu J (2023) YOLO by Ultralytics. *Github* [www document]. URL https://github.com/ultralytics/ultralytics

Kaczensky P, Adiya Y, von Wehrden H, Mijiddorj B, Walzer C, Güthlin D, et al. (2014) Space and habitat use by wild Bactrian camels in the Transaltai Gobi of southern Mongolia. *Biological Conservation* 169: 311–318.

Kellenberger B, Marcos D, Tuia D (2018) Detecting mammals in UAV images: best practices to address a substantially imbalanced dataset with deep learning. *Remote Sensing of Environment* 216: 139–153.

Linchant J, Lisein J, Semeki J, Lejeune P, Vermeulen C (2015) Are unmanned aircraft systems (UASs) the future of wildlife monitoring? A review of accomplishments and challenges. *Mammal Review* 45: 239–252.

McCarthy C, Banda LB, Kachamba DJ, Junior ZE, Chisambi C, Kumanga N, et al. (2024) Evaluating machine learning-based elephant recognition in complex African landscapes using drone imagery. *Environmental Research Communications* 6: 115035.

Norouzzadeh MS, Nguyen A, Kosmala M, Swanson A, Palmer MS, Packer C, et al. (2018) Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning. *Proceedings of the National Academy of Sciences of the United States of America* 115: E5716–E5725.

Payne JC, Buuveibaatar B, Bowler DE, Olson KA, Walzer C, Kaczensky P (2020) Hidden treasure of the Gobi: understanding how water limits range use of khulan in the Mongolian Gobi. *Scientific Reports* 10: 2989.

Reading RP, Henry M, Lhagvasuren B, Feh C, Kane DP, Dulamtseren S, et al. (2001) Status and distribution of khulan (*Equus hemionus*) in Mongolia. *Journal of Zoology* 254: 381–389.

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press

Tabak MA, Norouzzadeh MS, Wolfson DW, Sweeney SJ, Vercauteren KC, Snow NP, et al. (2019) Machine learning to classify animal species in camera trap images: applications in ecology. *Methods in Ecology and Evolution* 10: 585–590.

Tzutalin D (2015) tzutalin/labelImg. *Github* [www document]. URL https://github.com/HumanSignal/labelImg

Weinstein BG (2018) A computer vision for animal ecology. *Journal of Animal Ecology* 87: 533–545.

Yadamsuren A, Dulamtseren E, Reading RP (2012) *The Conservation Status and Management of Wild Camels in Mongolia*. Vienna, Austria: Verlag der Österreichischen Akademie der Wissenschaften.

https://doi.org/10.1017/S0376892925100118 Published online by Cambridge University Press