A Multimodal Deep Learning Framework for Locating Nomadic Pastoralists to Strengthen Public Health Outreach

Anonymous Authors[1]

Abstract

Nomadic pastoralists are systematically underrepresented in the planning of health services and frequently missed by health campaigns due to their mobility. Previous studies have developed novel geospatial methods to address these challenges but rely on manual techniques that are too time and resource-intensive to scale on a national or regional level. To address this gap, we developed a computer vision-based approach to automatically locate active nomadic pastoralist settlements from satellite imagery. We curated labeled datasets of satellite images capturing approximately 1,000 historically active settlements in the Omo Valley of Ethiopia to train and evaluate deep learning models, studying their robustness to low spatial resolutions and limits in labeled training data. We deployed our best model on a region spanning 5,400 square kilometers in the Omo Valley, resulting in the identification of historical settlements with a 270-fold reduction in manual review.

1. Introduction

Nomadic pastoralists migrate over large areas of remote terrain to support herds of livestock, which makes them susceptible to systematic underrepresentation in demographic surveys and census-reliant health campaigns (Randall, 2015). Systematic underrepresentation of mobile populations can lead to biased national statistics and underfunding of pastoral regions. Furthermore, underrepresentation of subgroups within pastoral populations, such as those that are most remote or mobile, can lead to imprecise and ineffective policy decisions by health officials working in these regions (Wild et al., 2019). The implications of such bias are particularly significant considering the impact of climate change, food insecurity, conflict, and infectious diseases among nomadic pastoralist populations in numerous regions of Sub-Saharan Africa (Barnes et al., 2017). Previous studies have developed geospatial and remote sensing techniques to address these difficulties, including an approach capable of generating representative sampling strategies among nomadic populations using remote sensing data that was piloted and validated among the Nyangatom of Ethiopia's South Omo Valley (Wild et al., 2019). However, the pilot of this methodology relied on manual enumeration that limited scalability. In this study, we addressed these challenges by leveraging machine learning to automatically detect nomadic settlements from remotely sensed imagery.

In the last decade, the availability of remote sensing data and satellite imagery has increased substantially due to advancements in satellite missions and technology. The influx of this data has significantly expanded the scope and granularity of earth observation, providing unique opportunities to perform comprehensive mapping of local landmarks in a variety of spatial and temporal contexts (Li et al., 2018). In recent years, deep learning-based computer vision models have been applied widely to systematically address population-level issues by using remotely sensed imagery. Recent approaches have leveraged the wide breadth of temporal, spectral, and spatial data available to improve mapping efforts by making them increasingly scalable, robust, and accurate across diverse geographical settings (Li et al., 2018; 2022; Leonita et al., 2018; Levin & Duke, 2012).

Here, we developed a novel computer vision model for the automatic localization of active nomadic pastoralist settlements from satellite imagery and evaluated the scalability of this method to a level of implementation that is compatible with national health campaigns. We showed that leveraging auxiliary settlement data such as roadway and waterway proximity can substantially improve model precision. Furthermore, we demonstrated that these strategies could augment performance in the face of limited training data, which is commonly observed in downstream model applications. To evaluate the robustness of our approach to regions with diverse geographical characteristics, we deployed our best model on pastoral regions in the South Omo Valley, Ethiopia. This approach holds potential to improve inclusion of underrepresented nomadic pastoralist popula-

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

Table 1. Composition of datasets used for model training and evaluation. "Continental" refers to randomly sampled geographical locations across continental Africa, excluding the country of interest. "National" refers to randomly sampled geographical locations in the country of interest. Satellite images containing active settlements were designated with "Positive" labels and all other images were designated with "Negative" labels.

TYPE	LABEL	TRAIN	VALIDATION	TEST
CONTINENTAL	NEG	6906	1482	1469
NATIONAL	NEG	1371	280	306
INACTIVE	NEG	1432	320	310
ACTIVE	POS	480	114	101

tions in health data and services.

2. Methods

2.1. Dataset

A comprehensive dataset of active settlements was derived from the Omo Valley of Ethiopia by using geographic bounds supplied by the ArcGIS hub and with feedback from experts at the International Livestock Research Institute (ILRI). Although satellite images in the Omo Valley of Ethiopia differ in characteristics such as soil composition, terrain, occupied footprint, and settlement structure, common recognizable features can be analyzed to determine whether an image contains an active settlement. For example, active settlements are generally surrounded by darker circular enclosures, more distinctly defined in appearance, and feature the appearance of village huts and a lack of vegetation growth due to livestock grazing. In contrast, inactive settlements are characterized by burn marks and vegetation overgrowth.

2.2. Model Experiments

2.2.1. MODEL ARCHITECTURE

We trained deep convolutional neural networks (CNNs) to perform a binary classification task as a proxy for our overall mapping task. Specifically, we trained models with the objective of identifying whether a satellite image contains an active nomadic pastoralist settlement. We tested prominent CNN and Transformer-based architectures including Vision Transformers (Khan et al., 2022), EfficientNet (Tan & Le, 2019), DenseNet (Huang et al., 2017), ResNeXt (Xie et al., 2017), ResNet (He et al., 2016), and HRNet (Wang et al., 2020) with a focus on understanding the effect of macroscopic model architecture choices on performance . In addition, we experimented with a wide array of training hyperparameters, including learning rate, optimizer choice, and regularization while keeping our architecture fixed. We observed that training an EfficientNet-B6 model with an Adam optimizer, batch size of 16, and a learning rate of 0.001 yielded an optimal performance within 40 epochs.

2.2.2. EVALUATION METRICS

We used multiple metrics in the assessment of our model performance but focused on precision at high recall thresholds (i.e. 0.95) due to the prioritization of maximizing coverage of settlements over the sheer correctness of our predictions. We measured 1) the AUPRC to understand the progression of precision at various recall thresholds and 2) precision under high recall thresholds to assess the preciseness of our predictions while maximizing settlement coverage. We additionally recorded the top F1 score, which represents the harmonic mean between precision and recall in our evaluation.

2.2.3. MULTIMODAL ABLATION DESIGN

We leveraged publicly available infrastructure data on roadways and waterways to develop a multi-modal learning approach to improve our model's precision at high recall thresholds, particularly when faced with limits in labeled training data. We specifically leveraged the observation that active settlements tend to be located closer to roadways and waterways and furthermore, that the distributions of infrastructure proximities between active settlements and other points differ substantially. To capture these differences, we relied on a Gaussian discriminant analysis (GDA) procedure (Hastie & Tibshirani, 1996) to construct auxiliary features encoding information on distances to nearby infrastructure, such as roadways and waterways. Specifically, we sampled subsets of our dataset separated by class to compute infrastructure proximity metrics and subsequently fitted multivariate Gaussian distributions to each class. We then sampled probabilities from these distributions to form our auxiliary features. We experimented with three different fusion strategies. $Aux_{(1)}$ refers to adding auxiliary metric features to class logits $\hat{y} = f_\phi(E_\psi(\vec{x}))$. $Aux_{(2)}$ refers to appending auxiliary features to global-average-pooled, low-dimensional feature map embeddings $\vec{z} = E_\psi(\vec{x})$. $Aux_{(1,2)}$ refers to simultaneously employing both strategies (Fig. 4).

2.3. Model Deployment

We deployed our best model on satellite images covering 5,400 square kilometers in the South Omo Valley region of Ethiopia at resolutions of 0.5 and 3.0 m / pixel. Predicted probabilities were converted into binary predictions by using a threshold that obtained 0.95 recall of known locations. Positively predicted, adjacent images were merged by computing the centroid of their spatial union. A full manual review of predicted active settlements was performed to assess the overall precision of our model deployment.

Table 2. Model performances across different spatial resolution settings that capture the resolution bounds of prevalent satellite products. Resolution is measured by meters per pixel.

RESOLUTION	F1	AUPRC	PRECISION@95
10.0 M/PIXEL	0.828	0.901	0.554
8.0 M/PIXEL	0.845	0.912	0.627
5.0 M/PIXEL	0.900	0.952	0.734
3.0 M/PIXEL	0.939	0.972	0.840
2.0 M/PIXEL	0.960	0.976	0.972
1.0 M/PIXEL	0.943	0.973	0.938
0.5 M/PIXEL	0.957	0.984	0.968

3. Results

3.1. Spatial Resolution Ablation

We observed that models trained with images at higher spatial resolutions outperformed those trained on images of lower resolutions (Table 2). Specifically, models trained at lower spatial resolution settings, including 5.0 m/pixel and 8.0 m/pixel achieved a precision at 95% recall of 0.734 and 0.627, respectively. Models trained with 3.0 m/pixel resolution imagery performed comparably to those of 0.5 m/pixel resolution imagery, displaying moderate drops in AUPRC and precision at 95% recall of approximately 0.01 and 0.13, respectively.

3.2. Auxiliary Data Fusion Ablation

We performed experiments to investigate the potential for publicly available auxiliary data to improve model precision at high recall thresholds. We observed that top-performing fusion approaches that separately leveraged waterway and roadway proximity data led to improvements in the precision at 95% recall by >0.13 and >0.12, respectively, relative to a non-fusion baseline. A GDA-based fusion approach that jointly incorporated waterway and roadway features outperformed our non-fusion baseline in precision at 95% recall by >0.03 (Table 3). In evaluating the direct contributions of GDA, we found that modeling distance features with our GDA approach led to gains in precision at 95% recall compared to fusing distance features directly.

3.3. Improving Model Robustness in Low Data Regimes

To address challenges posed by limited training data, we tested models that leveraged a GDA-based fusion approach with different classes of auxiliary data, including nearest waterway distance, nearest roadway distance, and a combination of both distance features. We observed that models trained with top-performing fusion approaches outperformed a baseline approach. Specifically, our top-performing waterway fusion model outperformed our baseline model across all low data regimes, with improvements of >0.10, >0.28, >0.40, and >0.29 in precision at 95% recall relative to training sets containing 25, 50, 75, and 100 active settlement examples, respectively. Our top-performing roadway fusion model outperformed our baseline model by >0.25 in precision at 95% recall when trained on 100 active settlement examples and performed comparably in all other settings. Similarly, the top-performing roadway-waterway fusion model performed comparably to our baseline model in all low data regimes (Table 1).

Table 3. Model performance comparisons across GDA-based fusion approaches leveraging different classes of auxiliary data.

DATA	STRATEGY	F1	AUPRC	PRECISION@95
	BASELINE	0.925	0.964	0.774
WATER	$Aux_{(1)}$	0.921	0.968	0.799
	$Aux_{(2)}$	0.935	0.968	0.913
	$Aux_{(1,2)}$	0.937	0.965	0.780
ROAD	$Aux_{(1)}$	0.932	0.964	0.900
	$Aux_{(2)}$	0.932	0.965	0.810
	$Aux_{(1,2)}$	0.931	0.967	0.881
BOTH	$Aux_{(1)}$	0.926	0.960	0.809
	$Aux_{(2)}$	0.934	0.968	0.771
	$Aux_{(1,2)}$	0.938	0.966	0.751

Figure 1. Graphical comparisons of model performance among top GDA-based fusion models in low data regimes relative to a non-fusion baseline model. Model performance in this graph is measured by precision at 95% recall. The GDA-based waterway fusion model outperformed its fusion counterpart models and its non-fusion baseline.

3.4. Model Deployment

We deployed our best model on the full Omo Valley target region spanning 5,400 square kilometers at image resolutions of 0.5 m/pixel and 3.0 m/pixel (Fig. 3). We subsequently performed a full manual review of active settlement predictions. Under image resolutions of 0.5 m/pixel and 3.0 m/pixel, our model achieved a precision at 95% recall of

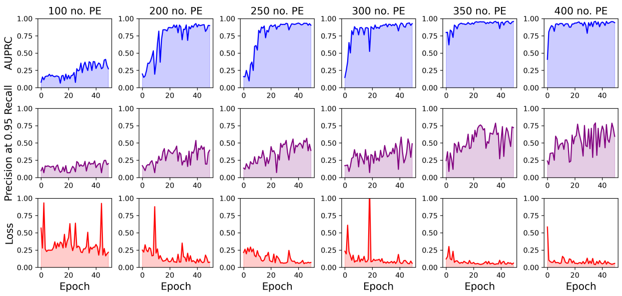

Figure 2. Graphical comparisons of model training progressions under exposure to different numbers of active settlement examples (no. PE) during training. Overall, training stability improves with higher numbers of active settlements in the training dataset.

0.71 and 0.61, respectively. Furthermore, we obtained a 270-fold search space reduction, reducing our number of active settlement candidates from 300,000 to ≈1,100 for manual review.

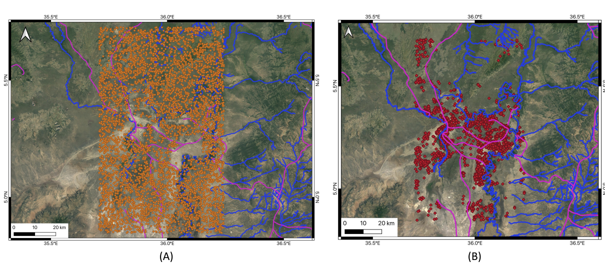

Figure 3. Maps of the Omo Valley deployment region in Ethiopia. Active settlements (detections in red) tend to be located near roadways and waterways compared to a random sample of points in the area of interest (orange).

4. Discussion

Our results demonstrate the potential for computer vision approaches to perform efficient and accurate localization of active nomadic pastoralist settlements from satellite imagery. Initial model experiments provided critical insights into the effect of model design choices and satellite image parameters on performance. We leveraged these insights to develop methods that use publicly available settlement auxiliary information in model training, achieving considerable performance improvements in low data regimes. We then deployed our best model over an extensive area of interest in the Omo Valley of Ethiopia, validating the scalability of our approach and its practical use in health and census campaigns. Collectively, our results demonstrate promising evidence of the potential for our methodology's application in public health evaluations and health campaigns to increase inclusion and equity for underrepresented nomadic pastoralist populations.

Our experiments leveraging publicly available infrastructure proximity data reveal that auxiliary information about active settlements can be successfully leveraged to improve model performance, boosting precision at 95% recall by as much as >0.13. The observation that improvements are observed after leveraging waterway and roadway proximity data, separately, indicate that the observed effect is potentially generalizable to other types of auxiliary information for which the behavior of associated distributions is substantially different among classes of interest. There was no clear dominance of a single fusion strategy in our experiments, as some fusion strategies achieved higher performances relative to our model baseline under different classes of auxiliary data. The effectiveness of our GDA-based fusion approaches may offer important insights for other machine learning mapping studies that seek to leverage differences in associated auxiliary information to improve the precision of their vision-based predictions.

The successful deployment of our models over a 5,400 square-kilometer region in Omo Valley, Ethiopia demonstrated the real-world effectiveness of our approach in performing a timely, automatic mapping of active settlements. Modest differences in performance between models trained on images obtained at 1.0 and 3.0 m / pixel resolution suggest that large-scale deployments can be successfully carried out with satellite images obtained at lower spatial resolutions of 3.0 m / pixel. Overall, we found that deploying our model resulted in a substantial 270-fold reduction of the search space for active settlements, reducing our number of active settlement candidates from 300,000 to 1,100 for manual review. This search space reduction gives credence to the expanded flexibility and feasibility that our approach offers for analyzing large areas of interest that would be impractical to analyze manually. Crucially, our approach addresses scalability concerns noted in several related studies (Weibel et al., 2008; Jean-Richard et al., 2015; Himelein et al., 2014) and with further development, shows promise in being implemented at a level consistent with national health campaigns.

5. Conclusion

In this study, we developed a computer vision-based approach for the localization of active nomadic pastoralist settlements from satellite imagery. We highlighted key considerations that are important to the integration and development of these models in health campaigns and demographic surveillance.

6. Impact Statement

Our approach provides a strong framework for the integration of computational remote sensing in large-scale, demographic health campaigns for nomadic pastoralists.

References

Barnes, A. N., Davaasuren, A., Baasandagva, U., and Gray, G. C. A systematic review of zoonotic enteric parasitic diseases among nomadic and pastoral people. *PloS one*, 12(11):e0188809, 2017.

Bauer, K. On the politics and the possibilities of participatory mapping and gis: using spatial technologies to study common property and land use change among pastoralists in central tibet. *cultural geographies*, 16(2):229–252, 2009.

Chambers, R. Participatory mapping and geographic information systems: whose map? who is empowered and who disempowered? who gains and who loses? *The Electronic Journal of Information Systems in Developing Countries*, 25(1):1–11, 2006.

Drusch, M., Del Bello, U., Carlier, S., Colin, O., Fernandez, V., Gascon, F., Hoersch, B., Isola, C., Laberinti, P., Martimort, P., et al. Sentinel-2: Esa's optical high-resolution mission for gmes operational services. *Remote sensing of Environment*, 120:25–36, 2012.

Gammino, V. M., Diaz, M. R., Pallas, S. W., Greenleaf, A. R., and Kurnit, M. R. Health services uptake among nomadic pastoralist populations in africa: A systematic review of the literature. *PLoS neglected tropical diseases*, 14(7):e0008474, 2020.

Greter, H., Jean-Richard, V., Crump, L., Schelling, E., Zinsstag, J., Bechir, M., Alfaroukh, I. O., and Bonfoh, B. The benefits of'one health'for pastoralists in africa: proceedings. *Onderstepoort Journal of Veterinary Research*, 81(2):1–3, 2014.

Hastie, T. and Tibshirani, R. Discriminant analysis by gaussian mixtures. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):155–176, 1996.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Himelein, K., Eckman, S., and Murray, S. Sampling nomads: a new technique for remote, hard-to-reach, and mobile populations. *Journal of official statistics*, 30(2):191–213, 2014.

Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

Jean-Richard, V., Crump, L., Abicho, A. A., Abakar, A. A., Mahamat Ii, A., Bechir, M., Eckert, S., Engesser, M., Schelling, E., Zinsstag, J., et al. Estimating population and livestock density of mobile pastoralists and sedentary settlements in the south-eastern lake chad area. *Geospatial health*, 10(1), 2015.

Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F. S., and Shah, M. Transformers in vision: A survey. *ACM computing surveys (CSUR)*, 54(10s):1–41, 2022.

Leonita, G., Kuffer, M., Sliuzas, R., and Persello, C. Machine learning-based slum mapping in support of slum upgrading programs: The case of bandung city, indonesia. *Remote sensing*, 10(10):1522, 2018.

Levin, N. and Duke, Y. High spatial resolution night-time light images for demographic and socio-economic studies. *Remote Sensing of Environment*, 119:1–10, 2012.

Li, H., Zech, J., Hong, D., Ghamisi, P., Schultz, M., and Zipf, A. Leveraging openstreetmap and multimodal remote sensing data with joint deep learning for wastewater treatment plants detection. *International Journal of Applied Earth Observation and Geoinformation*, 110: 102804, 2022.

Li, Y., Zhang, H., Xue, X., Jiang, Y., and Shen, Q. Deep learning for remote sensing image classification: A survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(6):e1264, 2018.

Masek, J. G., Wulder, M. A., Markham, B., McCorkel, J., Crawford, C. J., Storey, J., and Jenstrom, D. T. Landsat 9: Empowering open science and applications through continuity. *Remote Sensing of Environment*, 248:111968, 2020.

Randall, S. Where have all the nomads gone? fifty years of statistical and demographic invisibilities of african mobile pastoralists. *Pastoralism*, 5:1–22, 2015.

Robinson, L. W., Flintan, F. E., Kasyoka, S., Ng'ang'a, I. N., Otieno, K., and Sircely, J. A. Participatory rangeland management toolkit for kenya. 2020.

Schelling, E., Wyss, K., Diguimbaye, C., Béchir, M., Taleb, M. O., Bonfoh, B., Tanner, M., and Zinsstag, J. Towards integrated and adapted health services for nomadic pastoralists and their animals: A north–south partnership. *Handbook of transdisciplinary research*, pp. 277–291, 2008.

Sirajea, I. and Bekeleb, A. Assessment of food insecurity and coping mechanisms among pastoral households of afar national regional state: The case of chifra district, ethiopia. *Ethiopian Journal of Agricultural Sciences*, 23 (1-2):145–156, 2013.

Stavi, I., Roque de Pinho, J., Paschalidou, A. K., Adamo, S. B., Galvin, K., de Sherbinin, A., Even, T., Heaviside, C., and van der Geest, K. Food security among dryland pastoralists and agropastoralists: The climate, land-use change, and population dynamics nexus. *The Anthropocene Review*, 9(3):299–323, 2022.

Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pp. 6105–6114. PMLR, 2019.

Tugjamba, N., Walkerden, G., and Miller, F. Adapting nomadic pastoralism to climate change. *Climatic Change*, 176(4):28, 2023.

Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020.

Wario, H. T., Roba, H. G., and Kaufmann, B. Shaping the herders'"mental maps": participatory mapping with pastoralists' to understand their grazing area differentiation and characterization. *Environmental Management*, 56: 721–737, 2015.

Weibel, D., Schelling, E., Bonfoh, B., Utzinger, J., Hattendorf, J., Abdoulaye, M., Madjiade, T., Zinsstag, J., et al. Demographic and health surveillance of mobile pastoralists in chad: integration of biometric fingerprint identification into a geographical information system. *Geospatial health*, 3(1):113–124, 2008.

Wild, H., Glowacki, L., Maples, S., Mejía-Guevara, I., Krystosik, A., Bonds, M. H., Hiruy, A., LaBeaud, A. D., and Barry, M. Making pastoralists count: geospatial methods for the health surveillance of nomadic populations. *The American journal of tropical medicine and hygiene*, 101 (3):661, 2019.

Wild, H., Mendonsa, E., Trautwein, M., Edwards, J., Jowell, A., GebreGiorgis Kidanu, A., Tschopp, R., and Barry, M. Health interventions among mobile pastoralists: a systematic review to guide health service design. *Tropical Medicine & International Health*, 25(11):1332–1352, 2020.

Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.

Zinsstag, J., Schelling, E., Waltner-Toews, D., and Tanner, M. From "one medicine" to "one health" and systemic approaches to health and well-being. *Preventive veterinary medicine*, 101(3-4):148–156, 2011.

A. Spatial Resolution Ablation Study

A.1. Methodology

To characterize the feasibility of training our models with different satellite imagery products, we analyzed the impact of spatial resolution on model performance by mirroring the resolution bounds of widely accessible satellite products. We initially labeled Google satellite images occupying a 128m x 128m footprint taken at a resolution of 0.5 m/pixel, as these parameters provided sufficient granularity to distinguish important settlement features. To evaluate different resolution settings, we started with 256 x 256 pixel images taken at a spatial resolution of 0.5 m / pixel to maintain label consistency. We then scaled the images down such that the resulting dimensions rendered the images at resolutions of 1.0, 2.0, 3.0, 5.0, 8.0, and 10.0 m / pixel. These specifications matched the resolution extents of prevalent satellite products including WorldView, GeoEye-1, QuickBird, Rapideye-5, Planet Scope, Sentinel-1, and Sentinel-2. Finally, we upsampled the modified images back to 256 x 256 pixels, thus creating a standalone change to spatial resolution while maintaining consistency in all other image parameters.

A.2. Discussion

Our experiments investigating the impact of spatial resolution on model training showed that overall, model performance gradually degrades at lower spatial resolution settings, which is consistent with previous studies. Specifically, moderate drops in precision at 95% recall are observed at spatial resolution settings of 3.0 m/pixel and lower. Despite these moderate drops in performance, we observed that AUPRC is reasonably maintained above 0.9 over all spatial resolution settings and that even at the lowest tested spatial resolution setting, a precision at 95% recall above 0.5 is achieved. The results of these experiments provide important insight into the spatial resolutions of satellite images that are needed to achieve adequate levels of precision at high levels of recall on our task. Moreover, the nature of performance degradation over lower spatial resolutions provides valuable information about the level of spatial granularity that is needed to distinguish important settlement features and hence, make accurate predictions. Public health researchers can leverage this information to make informed decisions on the spatial scope of their studies, adjusting for available resources and personnel.

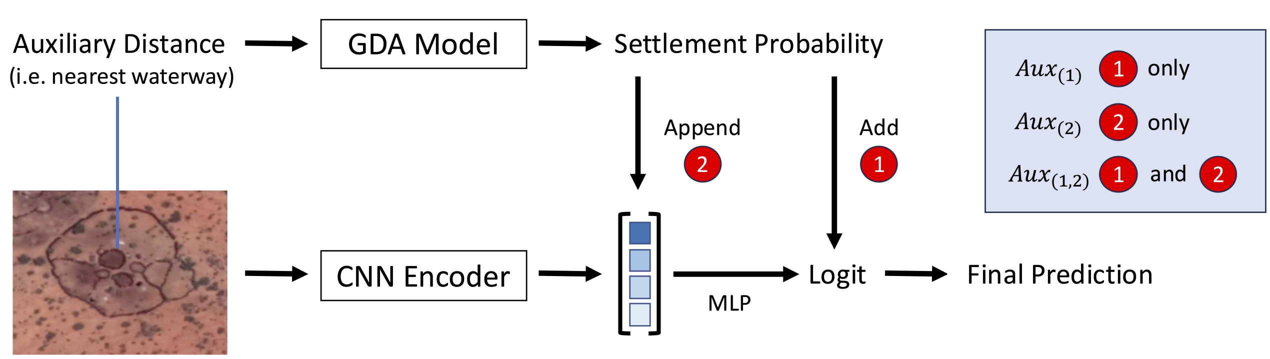

Figure 4. Visual summary of GDA-based fusion model architectures. Auxiliary distance features in our study were defined either as the distance to the nearest waterway or roadway. All displayed satellite images were sourced from the ESRI World Imagery basemap.

B. Data Regime Robustness Ablation Study

We studied the performance of our models relative to the reduction of active settlement examples in our training dataset to understand and mitigate challenges associated with low data regimes. We initially observed that the performance of our baseline model with no auxiliary data supplements saw a substantial decrease across all metrics when the number of active settlement examples in the training dataset was lowered below a count of 200, with an AUPRC and precision at 95% recall of 0.567 and 0.218, respectively. Precision at 95% recall decreased notably when the number of active settlement examples in the training dataset was lowered below a count of 350 (Fig 5). These trends were upheld in observing the training progressions of our models under different data regimes. We observed that model training became increasingly unstable with a decrease in the number of active settlement examples included during training.

Figure 5. Graphical comparisons of model performance as a function of the numbers of active settlement examples provided during training. Model performance degrades substantially when fewer than 200 examples of active settlements are included in the training dataset.

In follow-up experiments where we tested the ability for GDA-based fusion approaches to alleviate challenges posed by low data regimes, we observed that augmenting our baseline model with a waterway distance-based fusion approach led to substantial improvements to precision at 95% recall over all low data settings. These observations offer important insights for the development of computer vision models for active settlement localization, as it is conventionally difficult to obtain active settlement labels for model development. Common reasons for this trend include the low density of nomadic pastoralist settlements, sparse nature of settlement distributions, and limitations on personnel, compute power, and image acquisition. Furthermore, the success of these approaches in low data regimes suggests that strong priors on image classes can be effectively leveraged to close performance gaps presented by data limitations. This observation may be useful to adjacent machine learning mapping studies with constraints in accessible data. Future work should incorporate the domain expertise of public health officials to investigate other auxiliary priors that can be leveraged and quantify the generalizability limitations of models trained in these conditions.

C. Limitations and Future Work

In future work, integration of our approach with community engagement and participatory mapping efforts to provide primers on pastoralist mobility patterns should be explored. Since the early 1990s, diverse methods for participatory mapping have become commonplace in development practice (Chambers, 2006), specifically for pastoralist communities (Bauer, 2009; Robinson et al., 2020; Wario et al., 2015). Participatory mapping has the potential to offer more detailed information on pastoralist mobility patterns compared to analysis of remote sensing data alone. For example, it could assist in locating individuals, such as hired herders and family members, who are migrating with herds away from household locations. Combining our approach with participatory mapping methods and traditional ecological knowledge will require caution and keen awareness of local power dynamics, however. For example, in surveying herding destinations that the state or other powerful actors disapprove of, it will be important to carefully manage sensitive information that may lead to conflict (Bauer, 2009).

There were several limitations to this study. First, active settlement labels were designated based on inspecting high-resolution satellite imagery. Although physical aerial markers of settlement activity exist, we were not able to obtain visual evidence on the ground to validate our judgements. In future work, we aim to collaborate with field experts and local mapping authorities to obtain ground truth and rectify real-world gaps in our labeling criteria. Second, while our model substantially reduces the volume of active settlement locations that must be screened, it still necessitates a manual review of location candidates. This requirement may pose a limitation to public health efforts due to a lack of resources and personnel. Third, due to the relatively small spatial footprints of pastoralist settlements, we relied on high-resolution satellite imagery in our study, which is expensive and inaccessible on a global scale. Although we demonstrated that our approach can be feasibly applied with satellite images at spatial resolutions as low as 10.0 m/pixel, we hope to perform further studies to quantify model performance constraints at more coarse resolution settings, such as those offered by Landsat-9 at 30 m/pixel (Masek et al., 2020) and Sentinel-2 at 20 m/pixel (Drusch et al., 2012).

D. Global Health Implications

Our study has several important implications for existing global public health efforts focused on nomadic pastoralists and other mobile populations, particularly in remote regions. By addressing key bottlenecks in demographic surveys and census methodologies that have previously limited representative sampling of this population, our approach provides a more systematic and scalable way to capture data on mobile groups including nomadic pastoralists. Traditional survey techniques often rely heavily on random sampling techniques, which are inherently limited in coverage, or on manual enumeration methods that are both time and resource-intensive (Himelein et al., 2014; Schelling et al., 2008). Additionally, existing strategies require expert knowledge to be effectively deployed, further restricting their application in remote and under-resourced settings (Tugjamba et al., 2023). By leveraging deep learning for settlement identification, our method may significantly alleviate these constraints. This would not only enable demographic surveys to reach previously inaccessible locations at scale but also allow for more frequent and timely data collection. We emphasize that such approaches should be used in a context-appropriate manner and implemented in partnership with local collaborators to ensure sensitivity to local dynamics, particularly in conflict-affected settings.

Due to the cross-cutting nature of this methodological challenge across diverse global public health studies, our approach could support both research and service delivery among nomadic pastoralists across several key domains. In the context of climate change, our method could streamline existing efforts to assess how pastoralists are being affected by shifting environmental conditions and adapting to these challenges (Tugjamba et al., 2023). Similarly, given the ability to conduct full censuses of settlement locations at scale, our methodology holds potential to aid in evaluating food security dynamics, which frequently intersect both with climate stressors and conflict dynamics (Sirajea & Bekeleb, 2013; **?**; Stavi et al., 2022). From a One Health perspective, this framework could augment the study of infectious disease transmission within and between pastoralist communities, particularly in relation to zoonotic and enteric parasitic diseases, which remain a significant yet understudied risk factor. This methodology also holds potential for the design of public health campaigns as well as strategies to assess and the uptake of critical health services, ensuring that vaccinations, maternal health interventions, and disease surveillance programs more effectively reach mobile populations (Gammino et al., 2020; Wild et al., 2020). Collectively, these improvements could lead to better-informed health policy decisions for nomadic pastoralists and support integrated health frameworks such as One Health (Zinsstag et al., 2011; Greter et al., 2014).