Improving Landsat predictions of rangeland fractional cover with multitask learning and uncertainty

Brady W. Allred[1,2], Brandon T. Bestelmeyer[3], Chad S. Boyd[4], Christopher Brown[5], Kirk W. Davies[4], Lisa M. Ellsworth[6], Tyler A. Erickson[5], Samuel D. Fuhlendorf[7], Timothy V. Griffiths[8], Vincent Jansen[9], Matthew O. Jones[2], Jason Karl[9], Jeremy D. Maestas[10], Jonathan J. Maynard[11], Sarah E. McCord[3], David E. Naugle[1], Heath D. Starns[12], Dirac Twidwell[13], Daniel R. Uden[13]

[1]W.A. Franke College of Forestry and Conservation, University of Montana, Missoula, MT, 59812, USA

[2]Numerical Terradynamic Simulation Group, University of Montana, Missoula, MT 59812, USA

[3]Jornada Experimental Range, USDA Agricultural Research Service, Las Cruces, NM 88003, USA

[4]Eastern Oregon Agricultural Research Center, USDA Agricultural Research Service, Burns, OR 97720 USA

[5]Google, Inc., Mountain View, CA 94043, USA

[6]Fisheries and Wildlife Department, Oregon State University, Corvallis, Oregon 97331 USA

[7]Natural Resource Ecology and Management, Oklahoma State University, Stillwater, OK 74078, USA

[8]USDA Natural Resources Conservation Service, Landscape Initiatives Team, Bozeman, MT 59715, USA

[9]Department of Forest, Rangeland, and Fire Sciences, University of Idaho, Moscow, ID 83844, USA

[10]USDA Natural Resources Conservation Service, West National Technology Support Center, Portland, OR 97232, USA

[11]Sustainability Innovation Lab at Colorado, University of Colorado at Boulder, Boulder, CO 80309, USA

[12]Texas AgriLife Research, Texas A&M University, Sonora, TX 76950 USA

[13]Department of Agronomy and Horticulture, University of Nebraska–Lincoln, Lincoln, NE 68588, USA

Abstract

1. Operational satellite remote sensing products are transforming rangeland management and science. Advancements in computation, data storage, and processing have removed barriers that previously blocked or hindered the development and use of remote sensing products. When combined with local data and knowledge, remote sensing products can inform decision making at multiple scales.

2. We used temporal convolutional networks to produce a fractional cover product that spans western United States rangelands. We trained the model with 52,012 on-the-ground vegetation plots to simultaneously predict fractional cover for annual forbs and grasses, perennial forbs and grasses, shrubs, trees, litter, and bare ground. To assist interpretation and to provide a measure of prediction confidence, we also produced spatially-explicit, pixel-level estimates of uncertainty. We evaluated the model with 5,780 on-the-ground vegetation plots removed from the training data.

3. Model evaluation averaged 6.3% mean absolute error and 9.6% root mean squared error. Model performance increased across all functional groups compared to the previously produced fractional product.

4. The advancements achieved with the new rangeland fractional cover product expand the management toolbox with improved predictions of fractional cover and pixel-level uncertainty. The new product is available on the Rangeland Analysis Platform (https://rangelands.app/), an interactive web application that tracks rangeland vegetation through time. This product is intended to be used alongside local on-the-ground data, expert knowledge, land use history, scientific literature, and other sources of information when making interpretations. When being used to inform decision-making, remotely sensed products should be evaluated and utilized according to the context of the decision and not be used in isolation.

1 Introduction

The ability to monitor rangeland vegetation and to quantify changes in cover with satellite remote sensing is revolutionary to the rangeland management discipline. Whereas on-the-ground data collection and monitoring is constrained logistically, satellite remote sensing scales easily, measuring 100% of the landscape across space and through time. Satellite measurements are modeled to predict rangeland indicators, providing key information for land managers and practitioners. Chief among these indicators is vegetation cover at species or functional group levels. Historically, these indicators were categorical or thematic, and occurred at local, regional, or national levels (Homer et al., 2015). More recently, fractional cover is used to preserve the inherent complexity and heterogeneity of the landscape, as it estimates the proportion of an area covered by vegetation or land cover type (Boyte, Wylie, & Major, 2016; Collins et al., 2015; Xian, Homer, Rigge, Shi, & Meyer, 2015; Zhang, Okin, & Zhou, 2019). Fractional cover predictions, combined with local data and knowledge, can inform decision making at multiple scales, providing land managers flexibility that is largely absent with categorical classifications.

Fractional cover products derived from machine learning algorithms are widely available for United States rangelands (Jones et al., 2018; Zhang et al., 2019; Rigge et al., 2020). On-the-ground data is correlated to remotely sensed measurements using regression tree approaches, with models developed and predictions performed individually for each desired component. Although robust, univariate regression trees do not capitalize on the ability to learn from shared representation among dependent variables. That is, relationships among functional groups will not be learned and therefore may not be reflected in the final prediction. A learned multitask model, however, will examine all output variables together and learn from inherent interactions and relationships present in the data, improving learning efficiency and accuracy of predictions (Caruana, 1997). Variables that covary will be reflected as such in the model, e.g., functional groups or species that are mutually exclusive or inversely related.

We describe a new rangeland fractional cover product that spans the western United States. We build upon previous advancements (Jones et al., 2018) by 1) utilizing a learned multitask approach to model the dynamic interactions of functional groups; and 2) generating pixel-level estimates of prediction uncertainty. We produce the fractional cover product annually for the Landsat 5, 7, and 8 satellite time periods (1984-2019+) at a moderate resolution of 30 m. It is made available for analysis, download, and visualization through the Rangeland Analysis Platform (https://rangelands.app/).

2 Materials and Methods

2.1 Data

2.1.1. Rangeland Analysis Platform - fractional cover datasets

Jones et al. (2018) (hereafter referred to as fractional cover version 1.0) described the initial model and product released on the Rangeland Analysis Platform in 2018. The new model and subsequent product described in this paper (hereafter referred to as fractional cover version 2.0) supersedes the initial version.

2.1.2 Vegetation field data

We used vegetation field data collected by the Bureau of Land Management Assessment, Inventory, and Monitoring and Landscape Monitoring Framework, and the Natural Resources Conservation Service National Resources Inventory programs. Combining these data resulted in 57,792 field plots collected from 2004 to 2018 across western US rangelands of mixed ownership, private and public. We followed methods outlined in Jones et al. (2018), with species being aggregated into the following functional groups: annual forbs and grasses, perennial forbs and grasses, shrubs, trees, litter, and bare ground. We randomly divided the vegetation field data into training (90%, 52,012 field plots) and validation (10%, 5,780 field plots) datasets.

2.1.3 Landsat imagery

We used Landsat 5 TM, 7 ETM+, and 8 OLI surface reflectance products for predictors of fractional cover. We masked Landsat pixels identified as clouds, cloud shadow, snow, and saturated surface reflectance to calculate 64-day means throughout a given year for surface reflectance bands 2-7. The 64-day periods resulted in six measurements per year, with start dates occurring on day of year 001, 065, 129, 193, 257, and 321. To supplement surface reflectance measurements, we calculated normalized difference vegetation index (NDVI) and normalized burn ratio two (NBR2) for each 64-day period. These indices represent specific vegetation domains and have been successful in modeling rangeland fractional cover (Jones et al., 2018). We reprojected and bilinearly resampled all Landsat imagery to a geographic coordinate system of approximately 30 m resolution.

2.2 Model

We used Landsat surface reflectance measurements, vegetation indices, and spatial location (XY coordinates) as covariates to predict rangeland fractional cover. To generate a multivariate response, we used a temporal convolutional network to learn and predict cover for each functional group simultaneously. We used a temporal convolution (i.e., one dimensional convolution) as features associated with each vegetation field data plot varied sequentially through time, but not space. Temporal convolutions work well for satellite time series

4

classification (Pelletier, Webb, & Petitjean, 2019; Zhong, Hu, & Zhou, 2019) and may also outperform standard recurrent neural networks such as long short-term memory (Bai, Zico Kolter, & Koltun, 2018).

We combined temporal convolutions with dropout, pooling layers, and fully connected layers (Figure 1). We used an Adam optimizer with a learning rate of 0.0001, a batch size of 32, a convolutional kernel width of three, and a dropout rate of 20% (Srivastava, Hinton, Krizhevsky, Sutskever, & Salakhutdinov, 2014); the number of filters increased from 32 to 128 over three layers, and the dilation rate increased from one to four. We utilized average pooling with a pooling size of 12 to reduce temporal sequences to a single value. We performed convolutions on Landsat surface reflectances and vegetation indices separately due to the differing domains they represent, concatenating layers prior to a fully connected layer (Figure 1). The final layer contained six units, corresponding to the six functional groups. To produce uncertainty estimates, we implemented dropout during prediction (Gal & Ghahramani, 2015), utilizing a 10% dropout rate before the fully connected layer. We repeated predictions four times, averaged results to obtain the predictive output, and calculated variance to estimate uncertainty (Gal & Ghahramani, 2015).

We evaluated model performance by calculating mean absolute error (MAE), root mean square error (RMSE), residual standard error, (RSE), and the coefficient of determination (r^2) of the validation dataset. We compare evaluation metrics to fractional cover version 1.0. We developed the model using the Keras library within Tensorflow and performed all image processing and predictions in Google Earth Engine (Gorelick et al., 2017) and Google Cloud AI Platform (AI Platform, 2020), respectively.
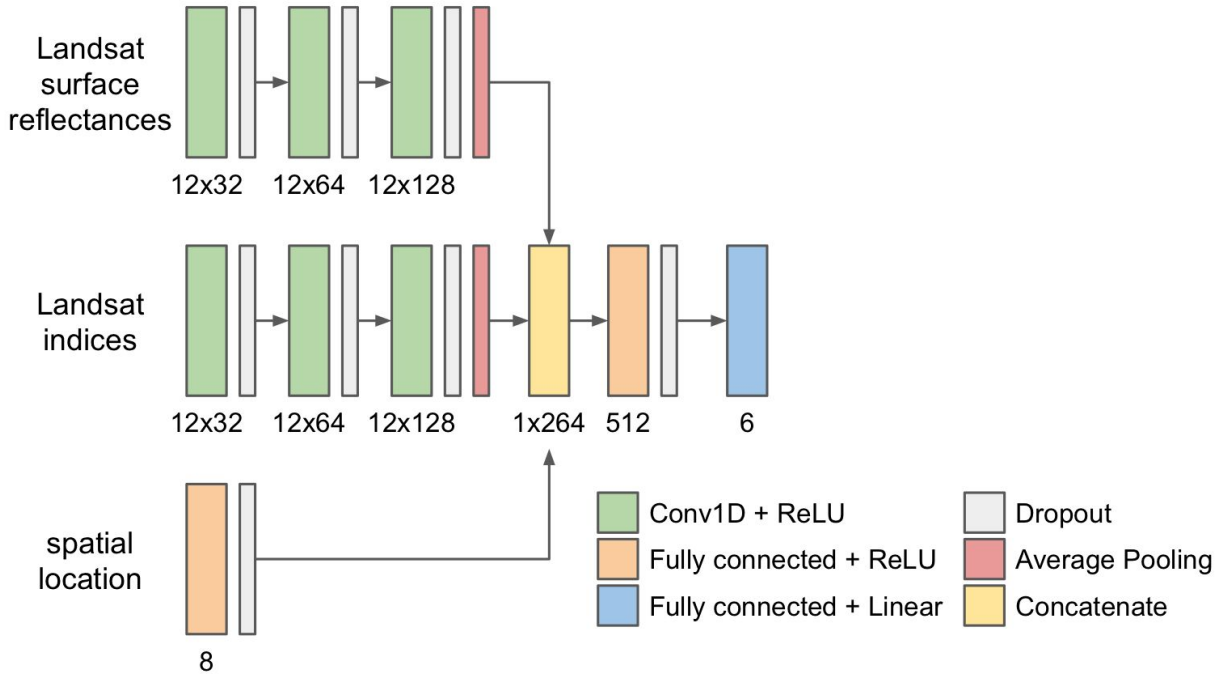
Figure 1. Model architecture to predict cover of rangeland functional groups. Inputs include Landsat surface reflectances, Landsat vegetation indices, and spatial location. Landsat surface reflectances and indices are temporal sequences across twelve 64-day timesteps. The final layer outputs percent cover of the six rangeland functional groups.

3 Results

3.1 Model evaluation

Model results and evaluation metrics suggest strong relationships between predicted and on-the-ground measurements. Evaluation metrics of the validation dataset averaged 6.3 and 9.6% (MAE and RMSE, respectively) across rangeland functional groups. Residual standard errors of predicted and on-the-ground measurements varied from 4.6 to 12.7% among functional groups (Table 1). Coefficient of determination values ranged from 0.57 to 0.77 for most functional groups (Table 1). On average, the model slightly over predicted lower values and under predicted higher values (Figure 2). Model performance increased compared to fractional cover version 1.0 (Jones et al., 2018; Table 2), and is comparable to other US rangeland fractional products available over disparate geographies (Zhang et al., 2019; Rigge et al., 2020).

Table 1. Model evaluation metrics (mean absolute error, MAE; root mean square error, RMSE; residual standard error, RSE; and coefficient of determination, r^2) calculated using the respective validation dataset for fractional cover versions 1.0 and 2.0.

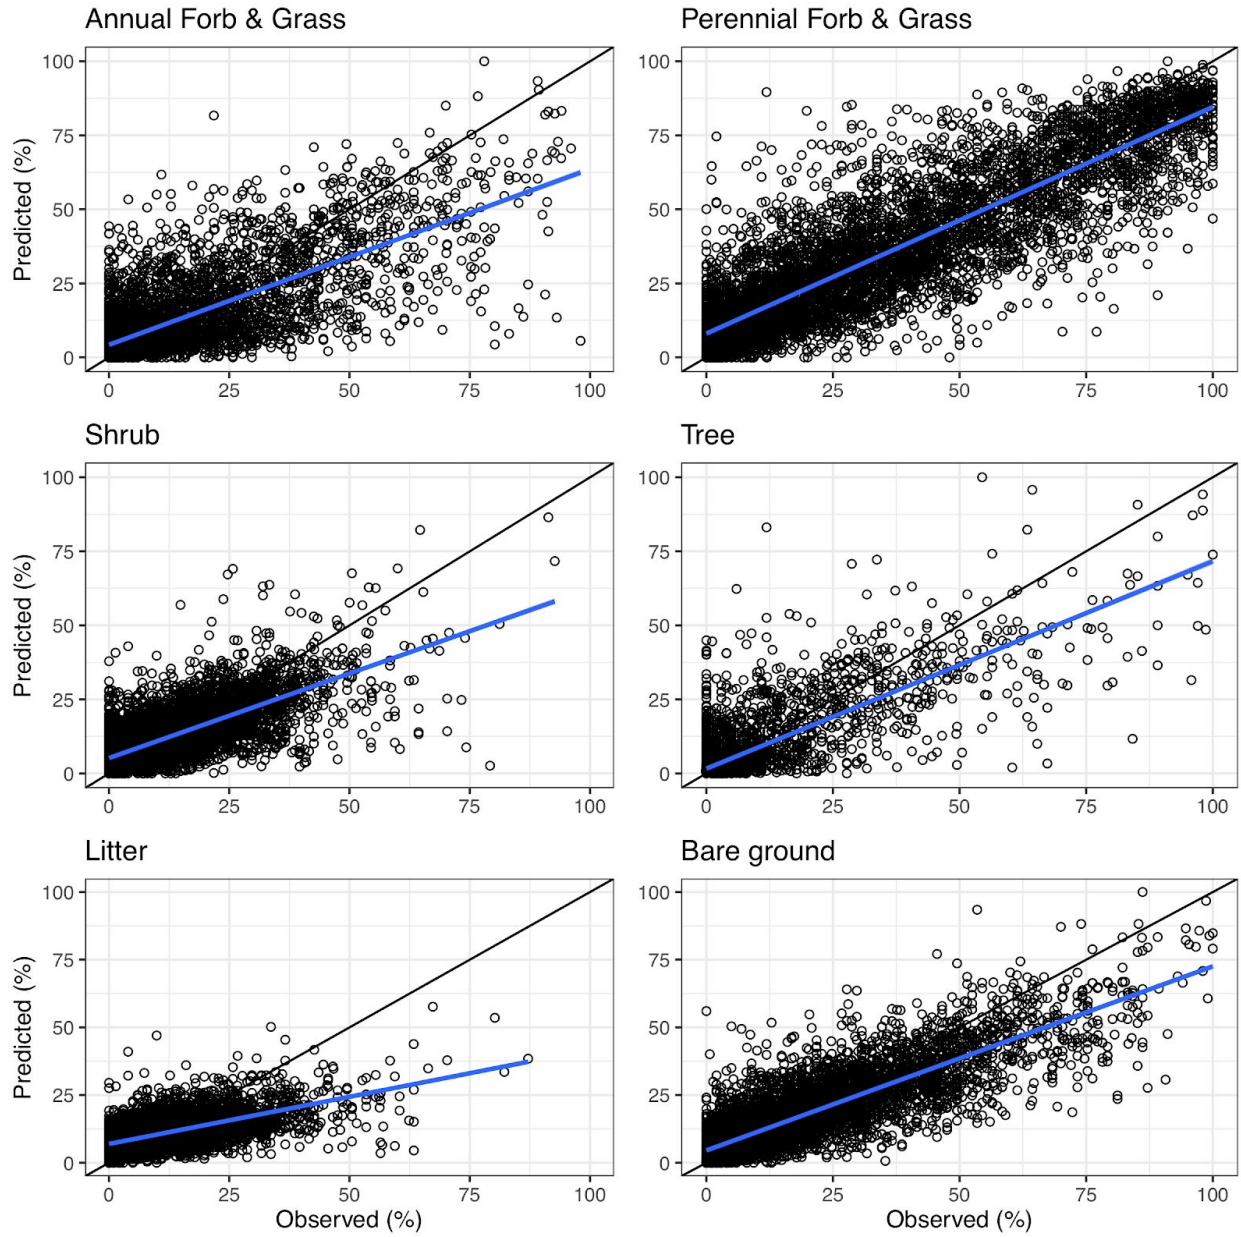|	Annual Forb & Grass	Perennial Forb & Grass	Shrub	Tree	Litter	Bare ground	Average
fractional cover version 2.0 (this paper)							
MAE	7.0	10.3	5.8	2.8	5.7	6.7	6.3
RMSE	11.0	14.0	8.3	6.8	7.9	9.8	9.6
RSE	8.8	12.7	6.6	5.9	4.6	7.9	-
r^2	0.58	0.77	0.57	0.65	0.37	0.73	-
fractional cover version 1.0 (Jones et al., 2018)							
MAE	7.8	11.1	6.9	4.7	-	7.3	7.56
RMSE	11.8	14.9	9.9	8.5	-	10.6	11.14
RSE	-	-	-	-	-	-	-
r^2	0.43	0.71	0.43	0.52	-	0.71	-

Figure 2. Model predictions of fractional cover relative to observed on-the-ground measurements for the validation dataset, separated by rangeland functional group. Diagonal black line represents a 1:1 relationship; blue line is the linear fit between predicted and observed values. Coefficient of determination (r^2) and residual standard error (RSE) are reported in Table 1.

4 Discussion

We provide next generation predictions of annual, fractional cover of rangeland functional groups by implementing a multitask learning approach across the western US (Figures 3 and 4). We improved upon our previous efforts (Jones et al., 2018) by 1) utilizing a neural network that models the dynamic interactions of functional groups; 2) reducing errors and improving model fit; and 3) providing spatially-explicit, pixel-level estimates of uncertainty alongside predictions. These efforts represent one of the larger applications of deep learning and remote sensing, providing billions of predictions across a large region and long time period. The results not only provide a better product, but also begin to provide guidance to the end user on application and utilization.
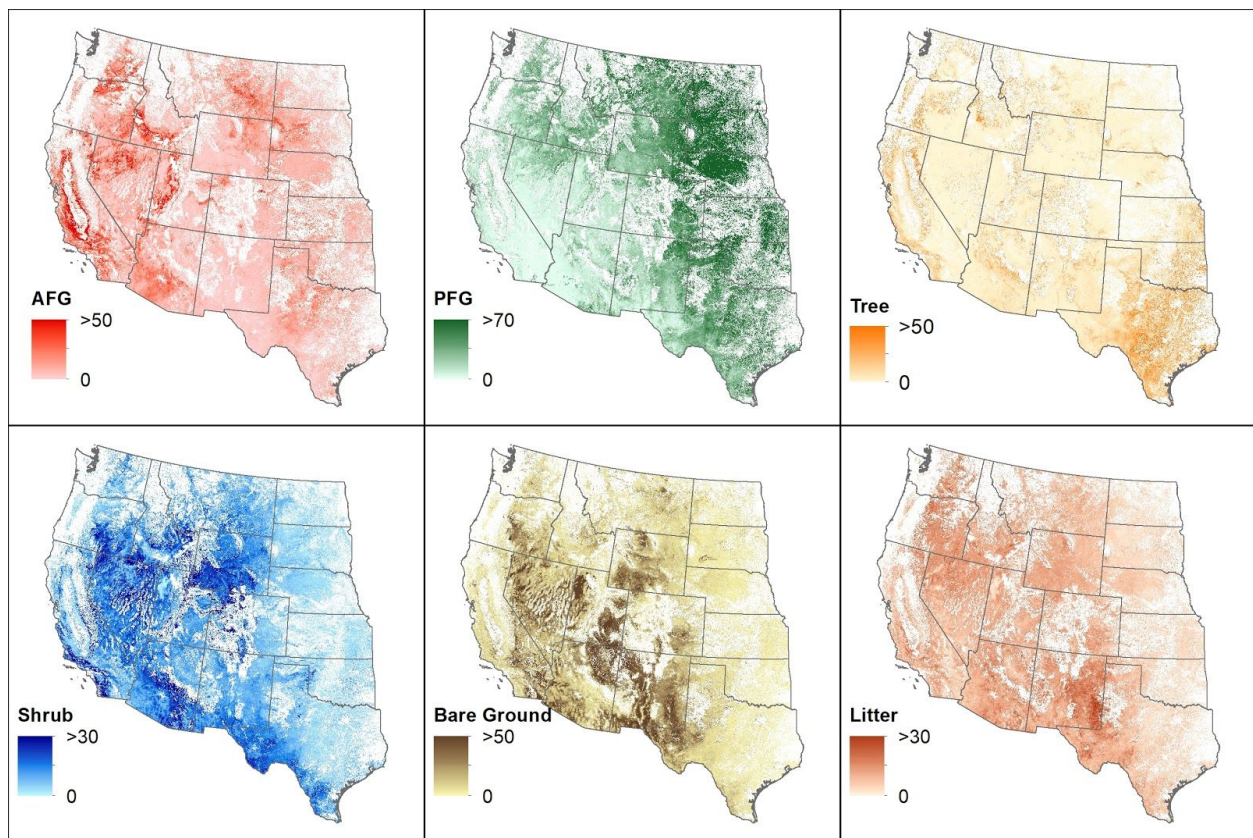

Figure 3. Fractional cover predictions of annual forbs and grasses (AFG), perennial forbs and grasses (PFG), shrubs, trees, litter, and bare ground for 2019. White areas are non-rangeland as identified by Reeves and Mitchell (2011).
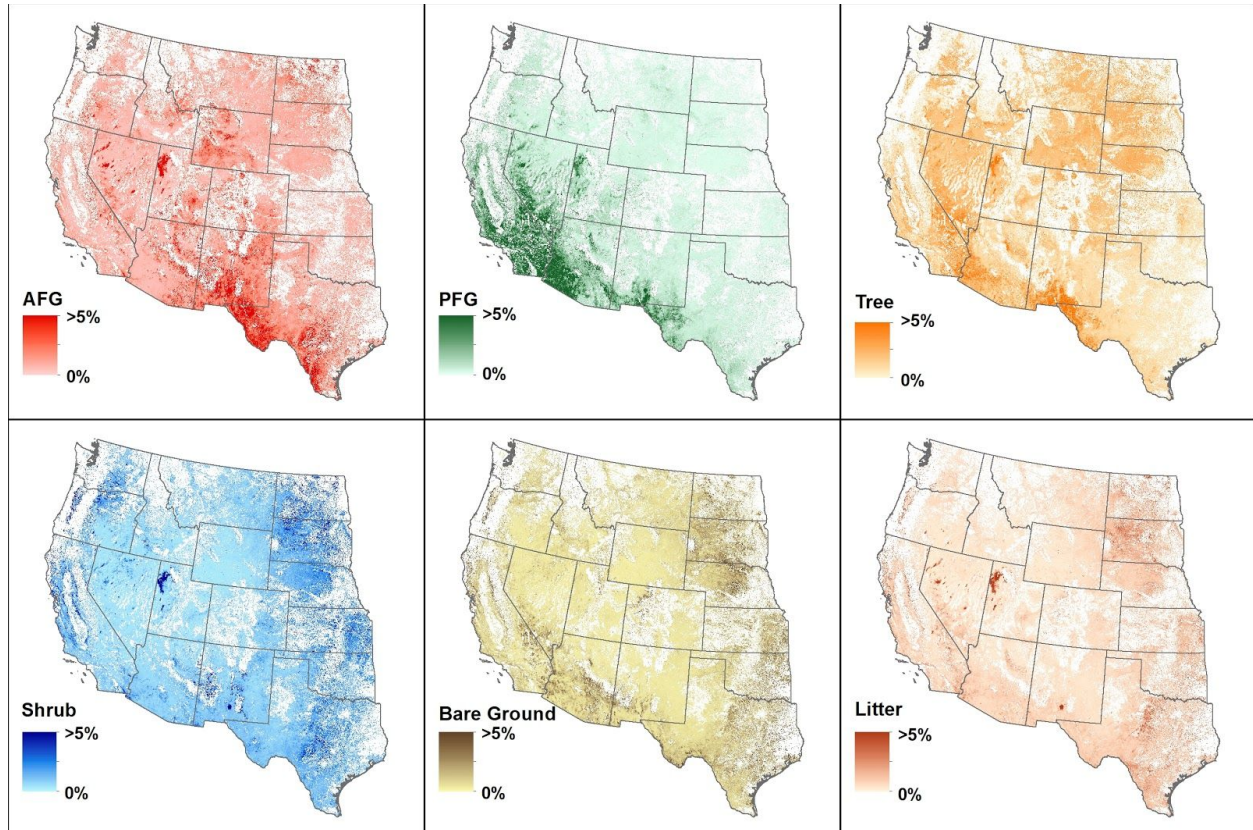
Figure 4. Fractional cover uncertainty of annual forbs and grasses (AFG), perennial forbs and grasses (PFG), shrubs, trees, litter, and bare ground for 2019. White areas are non-rangeland as identified by Reeves and Mitchell (2011).

The fractional cover of plant functional groups and cover types in rangelands reflect numerous ecosystem processes and ecosystem services including forage production, erosion control, and wildlife habitat. Changes in one plant functional group has predictable ecological impacts on other groups and the services they provide (Uden et al., 2019). For example, woody plant encroachment into grasslands constrains herbaceous grass cover and diminishes forage production and wildlife habitat (Archer et al., 2017), whereas annual grass invasion increases fire frequencies that reduce perennial herbaceous plants and shrubs (Davies, 2008). While previous univariate modeling methods of fractional cover disregard this covariation, using a multitask model allows for the learning of fractional cover among functional groups to be done concurrently, exploiting the inherent covariation among them (Caruana, 1997). By using multitask learning the underlying relationships among rangeland functional groups are learned and incorporated into the model producing better predictions (Figure 5). Although univariate predictions can be constrained or restricted post hoc to correct or reduce such errors (Henderson, Bell, & Gregory, 2019), the goal of multitask learning is to learn and predict simultaneously. Furthermore, the shared representation of multitask learning allows for covariance dynamics and interactions to be defined by the data, eliminating the need for predetermined conditions, rules, or thresholds.
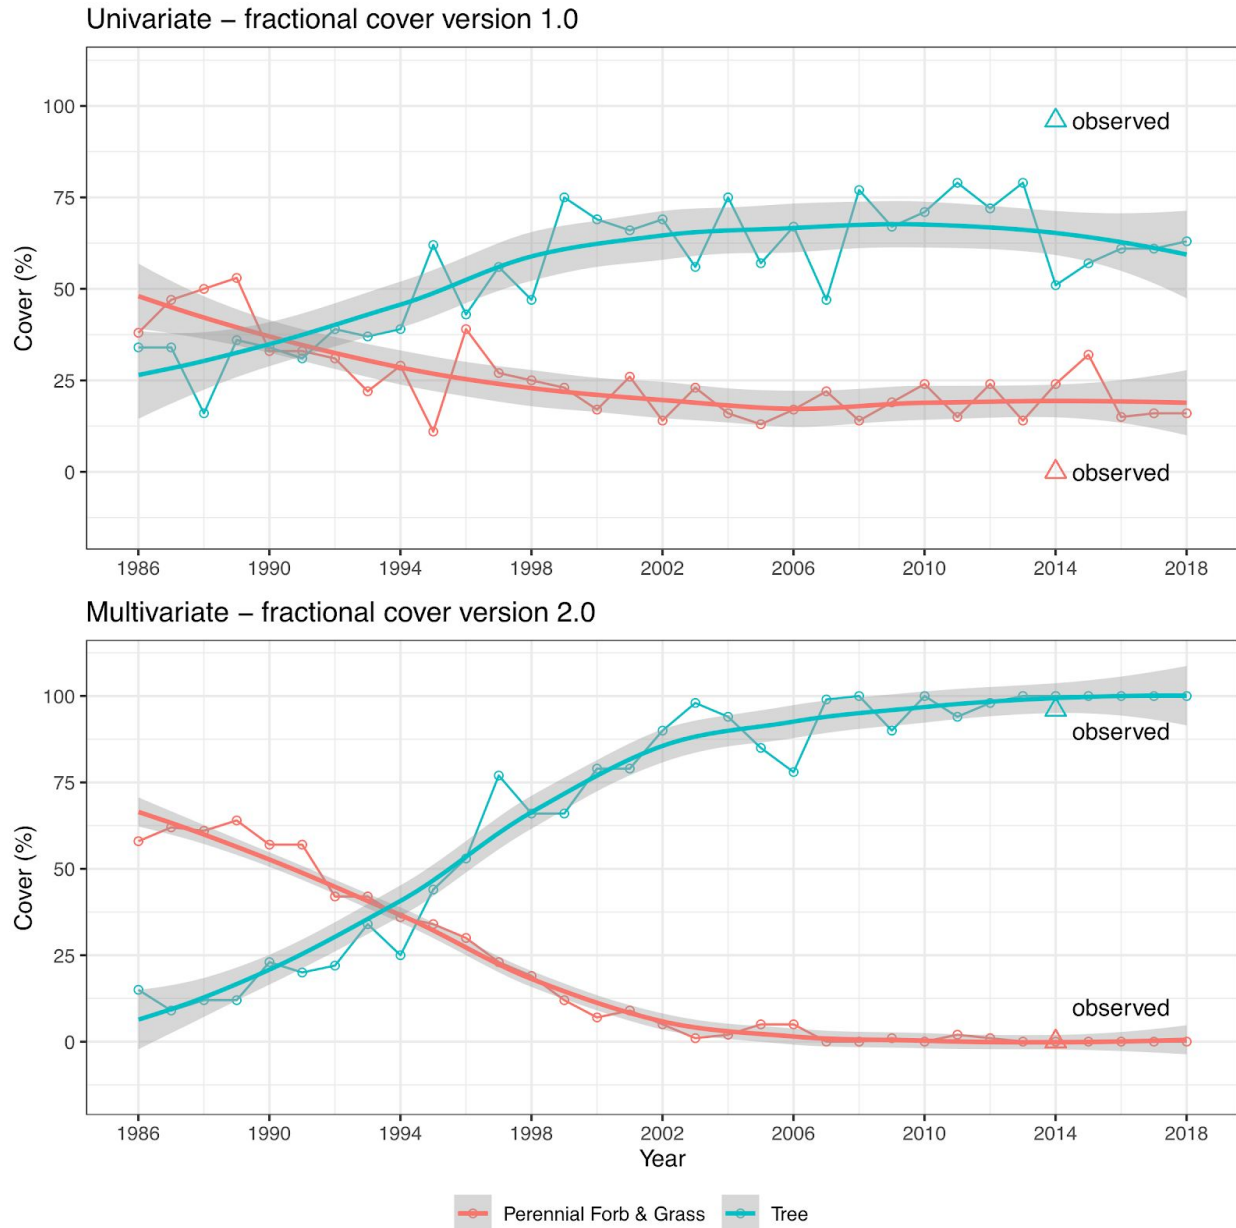
Figure 5. Univariate (top) and multivariate (bottom) predictions of perennial forb/grass and tree functional groups for a single validation plot in an area with woody encroachment. Plot data was collected in 2014 and recorded 0 and 96% cover (triangular points) for perennial forb/grass and tree, respectively. Due to shared representation, multitask models and predictions better represent functional group dynamics. Univariate predictions from fractional cover version 1.0 produced by Jones et al. (2018). Shaded lines represent locally estimated scatterplot smoothing.

When developing remotely sensed products, the goal is often to maximize model performance to be as accurate as possible. Error, however, is always present and should be understood and integrated into the context of the decision to which the product is informing. Common error metrics measure the average variation of predicted model output against on-the-ground measurements. This is generally calculated by withholding a small portion (e.g., 5-20%) of the model training dataset for validation (either entirely, or in a bootstrap aggregating approach). Error metrics represent an average error for the model given the validation dataset, but do not indicate any spatial or temporal variability of error. Attempts to visualize or aggregate errors across broad regions may appear helpful, but actually do little to characterize their spatial distribution or to help judge spatial accuracy (Jones et al., 2018; Zhang et al., 2019).

The inclusion of spatially-explicit, pixel-level estimates of uncertainty elevates the utility of fractional cover version 2.0. Contrary to model error, model uncertainty provides a measure of prediction confidence, i.e., how reasonable is this prediction given the data used to build the model? If a specified location or prediction is within or close to the distribution of the model training data, uncertainty may be low and the prediction reasonable. Correspondingly, if a specified location or prediction is far from the distribution of model training data, uncertainty may be high and the prediction unreasonable. Like error, uncertainty information can be integrated into the context of the decision being made and guide product utilization, e.g., if uncertainty is high, a land manager can choose to gather more data or information, do a more detailed analysis, discuss with colleagues, etc. Although not a measure of model accuracy or error, uncertainty can be helpful in determining how to use model predictions on a case-by-case basis, as uncertainty will vary across space and time (Figure 6). While error can only be calculated using a validation dataset–which is commonly just a fraction of the size of the total amount of predictions (<<0.1% in many cases)–uncertainty can be provided for every prediction.

Figure 6. Aerial imagery (left), cover estimates (middle), and uncertainty estimates (right) for 2019 perennial forb and grass (top) and tree (bottom) cover for a small region in the southern Great Plains. Light-to-dark values represent lower-to-higher values of cover and uncertainty. Greater uncertainty for perennial forb and grass estimates are present in areas dominated by trees and vice versa.

Innovations in remotely sensed mapping of rangeland cover continue to present new opportunities to improve assessment, monitoring, and management. We provide the latest advancement to further expand the land management toolbox with improved predictions of fractional cover at a moderate resolution of 30 m, along with spatially-explicit uncertainty estimates, that can be used at such resolution or aggregated to broader scales. This product is intended to be used in combination with local on-the-ground data, expert knowledge, land use history, scientific literature, and other sources of information when making interpretations. We emphasize that when being used to inform decision-making, remotely sensed products should be evaluated and utilized according to the context of the decision and not be used in isolation. Learning how to think about and use remotely sensed data, and suitably integrate them into decision frameworks and workflows, are next steps for improving the field of rangeland monitoring.

Acknowledgements

Author contributions

BA, MJ, JM, DN, TG, and DT led the writing of the manuscript. BA, CB, TE, MJ, JM, SM, and DU conceived, designed, implemented, or assisted with methodology. BB, CB, KD, LE, SF, VJ, JK, and HS collected, provided, or analysed data for additional evaluation. All authors contributed to drafts and gave final approval for publication.

Data availability

Data are available on the Rangeland Analysis Platform (https://rangelands.app/) and from http://rangeland.ntsg.umt.edu/data/rap/rap-vegetation-cover/.

References

AI Platform. (2020). Retrieved 27 May 2020, from https://cloud.google.com/ai-platform

Archer, S. R., Andersen, E. M., Predick, K. I., Schwinning, S., Steidl, R. J., & Woods, S. R. (2017). Woody Plant Encroachment: Causes and Consequences. In D. D. Briske (Ed.), *Rangeland Systems: Processes, Management and Challenges* (pp. 25–84). Cham: Springer International Publishing.

Bai, S., Zico Kolter, J., & Koltun, V. (2018, March 4). *An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. arXiv [cs.LG]*. Retrieved from http://arxiv.org/abs/1803.01271

Boyte, S. P., Wylie, B. K., & Major, D. J. (2016). Cheatgrass Percent Cover Change: Comparing Recent Estimates to Climate Change − Driven Predictions in the Northern Great Basin. *Rangeland Ecology & Management*, *69*(4), 265–279.

Caruana, R. (1997). Multitask Learning. *Machine Learning*, *28*(1), 41–75.

Collins, C. D. H., Kautz, M. A., Tiller, R., Lohani-Joshi, S., Ponce-Campos, G. E., Hottenstein, J. D., & Metz, L. J. (2015). Development of an integrated multiplatform approach for assessing brush management conservation efforts in semiarid rangelands. *Journal of Applied Remote Sensing*, *9*(1), 096057.

Davies, K. W. (2008). Medusahead Dispersal and Establishment in Sagebrush Steppe Plant Communities. *Rangeland Ecology & Management*, *61*(1), 110–115.

Gal, Y., & Ghahramani, Z. (2015, June 6). *Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. arXiv [stat.ML]*. Retrieved from http://arxiv.org/abs/1506.02142

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*. doi:10.1016/j.rse.2017.06.031

Henderson, E. B., Bell, D. M., & Gregory, M. J. (2019). Vegetation mapping to support greater

sage‑grouse habitat monitoring and management: multi- or univariate approach? *Ecosphere* , *10*(8), 361.

Homer, C., Dewitz, J., Yang, L., Jin, S., Danielson, P., Xian, G., … Megown, K. (2015). Completion of the 2011 National Land Cover Database for the conterminous United States--representing a decade of land cover change information. *Photogrammetric Engineering & Remote Sensing*, *81*(5), 345–354.

Jones, M. O., Allred, B. W., Naugle, D. E., Maestas, J. D., Donnelly, P., Metz, L. J., … McIver, J. D. (2018). Innovation in rangeland monitoring: annual, 30 m, plant functional type percent cover maps for U.S. rangelands, 1984-2017. *Ecosphere* , *9*(9), e02430.

Pelletier, C., Webb, G. I., & Petitjean, F. (2019). Temporal Convolutional Neural Network for the Classification of Satellite Image Time Series. *Remote Sensing*, *11*(5), 523.

Reeves, M. C., & Mitchell, J. E. (2011). Extent of Coterminous US Rangelands: Quantifying Implications of Differing Agency Perspectives. *Rangeland Ecology & Management*, *64*(6), 585–597.

Rigge, M., Homer, C., Cleeves, L., Meyer, D. K., Bunde, B., Shi, H., … Bobo, M. (2020). Quantifying Western U.S. Rangelands as Fractional Components with Multi-Resolution Remote Sensing and In Situ Data. *Remote Sensing*, *12*(3), 412.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research: JMLR*, *15*(1), 1929–1958.

Uden, D. R., Twidwell, D., Allen, C. R., Jones, M. O., Naugle, D. E., Maestas, J. D., & Allred, B. W. (2019). Spatial Imaging and Screening for Regime Shifts. *Frontiers in Ecology and Evolution*, *7*, 407.

Xian, G., Homer, C., Rigge, M., Shi, H., & Meyer, D. (2015). Characterization of shrubland ecosystem components as continuous fields in the northwest United States. *Remote Sensing of Environment*, *168*, 286–300.

Zhang, J., Okin, G. S., & Zhou, B. (2019). Assimilating optical satellite remote sensing images and field data to predict surface indicators in the Western U.S.: Assessing error in satellite predictions based on large geographical datasets with the use of machine learning. *Remote Sensing of Environment*, *233*, 111382.

Zhong, L., Hu, L., & Zhou, H. (2019). Deep learning based multi-temporal crop classification. *Remote Sensing of Environment*, *221*, 430–443.