Potential of Domain Adaptation in Machine Learning in Ecology and Hydrology to Improve Model Extrapolability

Haiyang Shi[1,2]

[1] Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

[2] The National Key Laboratory of Ecological Security and Sustainable Development in Arid Region, Chinese Academy of Sciences, Urumqi, 830011, China

Correspondence: Haiyang Shi (hiayang@illinois.edu)

Abstract

Due to the heterogeneity of the global distribution of ecological and hydrological ground-truth observations, machine learning models can have limited adaptability when applied to unknown locations, which is referred to as weak extrapolability. Domain adaptation techniques have been widely used in machine learning domains such as image classification, which can improve the model generalization ability by adjusting the difference or inconsistency of the domain distribution between the training and test sets. However, this approach has rarely been used explicitly in machine learning models in ecology and hydrology at the global scale, although these models have often been questioned due to geographic extrapolability issues. This paper briefly describes the shortcomings of current machine learning models of ecology and hydrology in terms of the global representativeness of the distribution of observations and the resulting limitations of the lack of extrapolability and suggests that future related modelling efforts should consider the use of domain adaptation techniques to improve extrapolability.

Main

In classical machine learning, it is assumed that the training and test sets are drawn from the same distribution. However, this assumption may not always hold true in real-world applications in fields like ecology and hydrology, which involve geographical variability. In the construction of big data models in ecology and hydrology, when models are applied to specific locations, there can be significant differences (dissimilarities) in the domain distributions of the test and training sets, potentially leading to decreased model performance (Shi et al., 2023). Domain Adaptation learning (Farahani et al., 2021), a type of transfer learning, aims to enhance the model's self-adaptive ability for the prediction set by adjusting the deviation between the domains of the training and test sets. Domain adaptation has been widely applied in fields like computer vision (Farahani et al., 2021), but is rarely used in big data models in ecology and hydrology. This indicates that previous models have overlooked the significant differences in domain distributions between test and training sets implied by geographical variability (including differences in climate, hydrology, underlying surfaces, etc.), highlighting an urgent need for improvements to substantially increase the models' generalizability.

Domain adaptation learning was initially proposed in the field of machine learning (Ben-David et al., 2006) and has since been widely applied in areas such as computer vision (Long et al., 2015). This approach arises when training data does not accurately reflect the distribution of test data, potentially leading to performance degradation when the trained model is applied to the test data. To address this issue, domain adaptation was introduced (Ben-David et al., 2006). In this context, the training and test sets are referred to as the source domain and target domain, respectively. Domain adaptation seeks to identify and adjust common underlying factors between the source and target domains to reduce mismatches in the feature space between the domains (Ben-David et al., 2006). Existing domain adaptation methods can be broadly categorized into shallow and deep architectures. Shallow domain adaptation methods (Gopalan et al., 2011; Pan et al., 2010) primarily use instance-based and feature-based techniques to align domain distributions. This includes minimizing the distance between domains (common distance metrics include Wasserstein distance, Kullback-Leibler divergence, etc.), and aligning the covariance of source and target domain data (correlation alignment methods) (Sun et al., 2016). On the other hand, deep domain adaptation methods (Ganin & Lempitsky, 2015; Long et al., 2015) employ neural networks. These methods often use convolutional networks, autoencoders, or adversarial networks to reduce differences between domains. Such research has demonstrated higher accuracy in a series of image classification experiments (Ganin & Lempitsky, 2015; Liang et al., 2020; Long et al., 2015) and effectively improves adaptability in situations with significant domain variations.

In the field of geospatial big data research, while some studies in remote sensing image classification (Matasci et al., 2015; Tuia et al., 2016) have begun to explore domain adaptation learning, its application in constructing big data models for hydrology and ecology, which are

characterized by high geographic variability, remains limited. In many mainstream data products of hydrological and ecological geographic variables based on machine learning, all available observational data are used to build a unified prediction model, which is then directly applied on a regional or global scale (Hassani et al., 2020; Jung et al., 2020; Li et al., 2023). However, the training sets for these models are often densely distributed in economically developed regions like the United States and Europe (Figure 1), but are sparse in many underdeveloped areas and regions with minimal human footprint, such as Africa, Central Asia, and Siberia. The challenges in data sharing also mean that the availability of training set data is not maximized. As a result, current machine learning models in the fields of ecology and hydrology are more likely to learn patterns specific to regions with dense observational networks, such as the United States and Europe, and are insufficiently trained on areas with scarce observations. Therefore, the extrapolative potential of these models has not been fully recognized. This highlights a significant gap in the predictive capabilities of these models across different geographic regions, indicating a need for models that can adapt to diverse geographical conditions and data distributions.
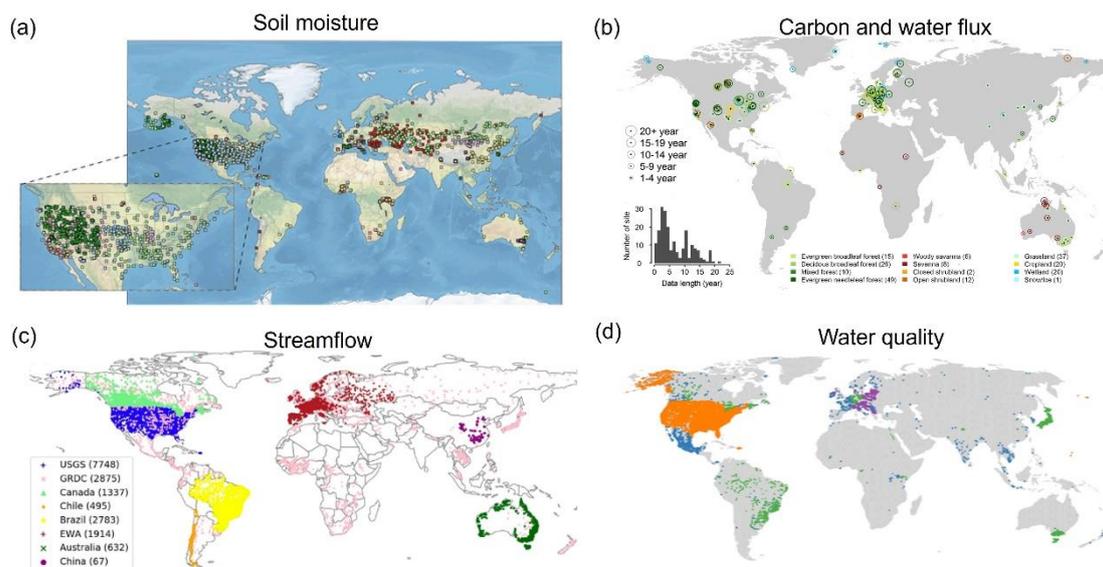
Figure 1: Distribution of observational data available for training global ecological and hydrological machine learning models. (a) Network of soil moisture observation stations, modified from the ref. (Dorigo et al., 2021), (b) Network of carbon and water flux stations, modified from the ref. (Pastorello et al., 2020), (c) Network of runoff gauge stations, modified from the ref. (Lin et al., 2019), (d) Network of water quality (total suspended solids) stations, modified from the ref. (Virro et al., 2021).

Many models in geospatial research employ methods like "leave-one-out cross-validation" (e.g., in the ET product of FLUXCOM (Jung et al., 2019)) or random cross-validation (such as the commonly used k-fold cross-validation (Shi et al., 2022a, 2022b)) to assess model accuracy. However, these approaches often overlook the significant impact that differences in domain

distribution (or dissimilarity) between the left-out site or fold and the training set can have on the precision of predictions. These models typically don't assess and mitigate the differences between the source domain (training set) and the target domain (prediction set), leading to potentially unreliable predictions in locations where there is a substantial difference from the training data's domain. A few studies have focused on "model extrapolation" related to domain adaptability (Jung et al., 2020; Meyer & Pebesma, 2021), attempting to establish a relationship between model accuracy and the dissimilarity between the test set and the training set in leave-one-out cross-validation (Meyer & Pebesma, 2021). This approach aims to provide a geographical distribution of predictability. For instance, in ET (Evapotranspiration) simulation predictions, Shi et al. (Shi et al., 2023) compared the accuracy of models trained with all data against those trained with data from specific land cover types. They found that in some sites, models trained with all data were less accurate than those trained with specific land cover types. This suggests that quantifying the distance between the prediction set and the training set could be crucial in understanding these differences. However, these studies have not explicitly employed domain adaptation methods to optimize the domain distribution of the training set data. By doing so, they could potentially improve the adaptability of the models to different test sets, significantly enhance the accuracy of the models' predictions, and produce high-precision data on a global scale. There's a growing recognition of the need for such domain adaptation in ecological and hydrological modeling, especially considering the diverse and often sparse data distributions in these fields.

This indicates that domain adaptation techniques have not yet been fully applied in global-scale ecological and hydrological machine-learning models. If implemented, these techniques could significantly reduce inconsistencies between training and prediction sets, thereby enhancing the adaptability of models to unknown locations. While there may still be challenges in finding similar training data in extreme outlier conditions, domain adaptation offers the potential to improve adaptability for many test locations.

In summary, the application of domain adaptation in the current field of hydrological and ecological geospatial big data modelling is quite scarce, and its importance has been largely overlooked. The incorporation of domain adaptation methods in future models deserves more attention and experimentation to enhance their generalizability and extrapolation capabilities. As the field continues to evolve, recognizing and addressing these gaps can lead to more robust and accurate predictive models, especially in areas with sparse data distributions.

References

Ben-David, S., Blitzer, J., Crammer, K., & Pereira, F. (2006). Analysis of representations for domain adaptation. *Advances in Neural Information Processing Systems*, *19*.

Dorigo, W., Himmelbauer, I., Aberer, D., Schremmer, L., Petrakovic, I., Zappa, L., Preimesberger, W., Xaver, A., Annor, F., Ardö, J., Baldocchi, D., Bitelli, M., Blöschl, G., Bogena, H., Brocca, L., Calvet, J.-C., Camarero, J. J., Capello, G., Choi, M., … Sabia, R. (2021). The International Soil Moisture Network: Serving Earth system science for over a decade. *Hydrology and Earth System Sciences*, *25*(11), 5749–5804. https://doi.org/10.5194/hess-25-5749-2021

Farahani, A., Voghoei, S., Rasheed, K., & Arabnia, H. R. (2021). A Brief Review of Domain Adaptation. In R. Stahlbock, G. M. Weiss, M. Abou-Nasr, C.-Y. Yang, H. R. Arabnia, & L. Deligiannidis (Eds.), *Advances in Data Science and Information Engineering* (pp. 877–894). Springer International Publishing. https://doi.org/10.1007/978-3-030-71704-9_65

Ganin, Y., & Lempitsky, V. (2015). *Unsupervised domain adaptation by backpropagation*. 1180–1189.

Gopalan, R., Li, R., & Chellappa, R. (2011). *Domain adaptation for object recognition: An unsupervised approach*. 999–1006.

Hassani, A., Azapagic, A., & Shokri, N. (2020). Predicting long-term dynamics of soil salinity and sodicity on a global scale. *Proceedings of the National Academy of Sciences*, *117*(52), 33017–33027. https://doi.org/10.1073/pnas.2013771117

Jung, M., Koirala, S., Weber, U., Ichii, K., Gans, F., Camps-Valls, G., Papale, D., Schwalm, C., Tramontana, G., & Reichstein, M. (2019). The FLUXCOM ensemble of global land-atmosphere energy fluxes. *Scientific Data*, *6*(1), Article 1. https://doi.org/10.1038/s41597-

019-0076-8

Jung, M., Schwalm, C., Migliavacca, M., Walther, S., Camps-Valls, G., Koirala, S., Anthoni, P., Besnard, S., Bodesheim, P., Carvalhais, N., Chevallier, F., Gans, F., S Goll, D., Haverd, V., Köhler, P., Ichii, K., K Jain, A., Liu, J., Lombardozzi, D., … Reichstein, M. (2020). Scaling carbon fluxes from eddy covariance sites to globe: Synthesis and evaluation of the FLUXCOM approach. *Biogeosciences*, *17*(5), 1343–1365. Scopus. https://doi.org/10.5194/bg-17-1343-2020

Li, F., Xiao, J., Chen, J., Ballantyne, A., Jin, K., Li, B., Abraha, M., & John, R. (2023). Global water use efficiency saturation due to increased vapor pressure deficit. *Science*, *381*(6658), 672–677. https://doi.org/10.1126/science.adf5041

Liang, J., Hu, D., & Feng, J. (2020). Do We Really Need to Access the Source Data? Source Hypothesis Transfer for Unsupervised Domain Adaptation. *Proceedings of the 37th International Conference on Machine Learning*, 6028–6039. https://proceedings.mlr.press/v119/liang20a.html

Lin, P., Pan, M., Beck, H. E., Yang, Y., Yamazaki, D., Frasson, R., David, C. H., Durand, M., Pavelsky, T. M., Allen, G. H., Gleason, C. J., & Wood, E. F. (2019). Global Reconstruction of Naturalized River Flows at 2.94 Million Reaches. *Water Resources Research*, *55*(8), 6499–6516. https://doi.org/10.1029/2019WR025287

Long, M., Cao, Y., Wang, J., & Jordan, M. (2015). *Learning transferable features with deep adaptation networks*. 97–105.

Matasci, G., Volpi, M., Kanevski, M., Bruzzone, L., & Tuia, D. (2015). Semisupervised transfer component analysis for domain adaptation in remote sensing image classification. *IEEE*

Transactions on Geoscience and Remote Sensing, *53*(7), 3550–3564.

Meyer, H., & Pebesma, E. (2021). Predicting into unknown space? Estimating the area of applicability of spatial prediction models. *Methods in Ecology and Evolution*, *12*(9), 1620–1633.

Pan, S. J., Tsang, I. W., Kwok, J. T., & Yang, Q. (2010). Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, *22*(2), 199–210.

Pastorello, G., Trotta, C., Canfora, E., Chu, H., Christianson, D., Cheah, Y.-W., Poindexter, C., Chen, J., Elbashandy, A., & Humphrey, M. (2020). The FLUXNET2015 dataset and the ONEFlux processing pipeline for eddy covariance data. *Scientific Data*, *7*(1), 1–27.

Shi, H., Luo, G., Hellwich, O., He, X., Xie, M., Zhang, W., Ochege, F. U., Ling, Q., Zhang, Y., Gao, R., Kurban, A., De Maeyer, P., & Van de Voorde, T. (2023). Comparing the use of all data or specific subsets for training machine learning models in hydrology: A case study of evapotranspiration prediction. *Journal of Hydrology*, *627*, 130399. https://doi.org/10.1016/j.jhydrol.2023.130399

Shi, H., Luo, G., Hellwich, O., Xie, M., Zhang, C., Zhang, Y., Wang, Y., Yuan, X., Ma, X., Zhang, W., Kurban, A., De Maeyer, P., & Van de Voorde, T. (2022a). Evaluation of water flux predictive models developed using eddy-covariance observations and machine learning: A meta-analysis. *Hydrology and Earth System Sciences*, *26*(18), 4603–4618. https://doi.org/10.5194/hess-26-4603-2022

Shi, H., Luo, G., Hellwich, O., Xie, M., Zhang, C., Zhang, Y., Wang, Y., Yuan, X., Ma, X., Zhang, W., Kurban, A., De Maeyer, P., & Van de Voorde, T. (2022b). Variability and uncertainty in flux-site-scale net ecosystem exchange simulations based on machine learning and remote

sensing: A systematic evaluation. *Biogeosciences*, *19*(16), 3739–3756. https://doi.org/10.5194/bg-19-3739-2022

Sun, B., Feng, J., & Saenko, K. (2016). Return of Frustratingly Easy Domain Adaptation. *Proceedings of the AAAI Conference on Artificial Intelligence*, *30*(1), Article 1. https://doi.org/10.1609/aaai.v30i1.10306

Tuia, D., Persello, C., & Bruzzone, L. (2016). Domain adaptation for the classification of remote sensing data: An overview of recent advances. *IEEE Geoscience and Remote Sensing Magazine*, *4*(2), 41–57.

Virro, H., Amatulli, G., Kmoch, A., Shen, L., & Uuemaa, E. (2021). GRQA: Global River Water Quality Archive. *Earth System Science Data*, *13*(12), 5483–5507. https://doi.org/10.5194/essd-13-5483-2021