

Proceeding Paper

Improving Remote Sensing Classification with Transfer Learning: Exploring the Impact of Heterogenous Transfer Learning †

ROUBA Maria^{1, *}, LARABI Mohammed Lamine² and ATTAF Dalila²

¹ National Higher School of Geodesic Sciences and Space Techniques, Algerian Spatial Agency; rou-bamaria645@gmail.com

² Algerian Spatial Agency; arabi.mohamed.lamine@gmail.com; dalila.ataf@gmail.com

* Correspondence: roubamaria645@gmail.com

† Presented at the title, place, and date.

Abstract: Deep Learning (DL) has become increasingly popular in recent years, with researchers and businesses alike successfully applying it to a wide range of tasks. However, one challenge that DL faces in certain domains, such as Remote Sensing (RS), is the difficulty of creating large, well-annotated training datasets. This is due to the high cost of acquiring and labeling RS data. This challenge significantly limits the development of DL in RS. RS data can come from multiple sources, such as satellites, airplanes, and drones, and use different sensor technologies. Training DL models on data from one source may not produce the same accuracy on data from other sources, even if they cover the same region. Transfer learning (TL) can help to address this challenge by relaxing the requirement for large training datasets. Specifically, TL allows us to train a model on data from one source and then adapt it to data from another source, even with less training data. This makes TL a promising approach for solving both the problem of multisource adaptation and the problem of insufficient training data in the target domain. This paper evaluates the homogenous and heterogeneous TL approach that addresses model transfer across different domains. Transfer gain is measured through specific statistical metrics such as precision, kappa, recall, and F1-score, and a positive gain is empirically shown in the vast majority of cases. The proposed method is evaluated on the challenging task of Multispectral RS image (MSI) classification due to the complexity and variety of natural scenes. This work is examined in terms of its social, economic, and environmental consequences. Additionally, potential future directions for research and the achievement of established goals are explored.

Keywords: Transfer Learning, Heterogenous Learning, Homogenous Learning, Multispectral Remote Sensing, Sentinel-2, Landsat-9

Citation: To be added by editorial staff during production.

Academic Editor: Firstname Last-name

Published: date



Copyright: © 2023 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Most machine learning (ML) methods employed nowadays assume that training and test data are from the same feature space and distribution [1]. This implies that when the data distribution changes, models trained from scratch with newly collected data are required. However, there are situations where obtaining new data for training new models can be financially or logistically challenging, particularly in environments where data collection demands substantial computational power or financial resources. Hence, the ability to leverage existing knowledge becomes highly valuable, offering a means to circumvent the expensive endeavor of reconstructing a model from scratch and amassing sufficient new data to create a dependable system. This is where the concept of transfer learning (TL) becomes indispensable.

TL is a concept that allows information learned in one setting to be used in another, or at least part of it. This can improve the training process by avoiding the need to develop

new models from scratch and instead adapting previously built models to a new setting. TL has proven to be a powerful technique with successful results in many different environments. However, there are several challenges that need to be considered before deciding whether to use TL.

The classification of multispectral remote sensing (RS) images (MSI) is a critical and important topic that has been extensively studied. It is essential for a variety of applications, including land cover mapping, vegetation monitoring, urban analysis, resource management, and decision-making. The complexity and variety of natural scenes pose significant challenges to MSI classification. This is due to the unique characteristics of RS data, such as spectral variability, spatial heterogeneity, and temporal dynamics. These characteristics make it difficult to distinguish between different land cover classes, and collecting data to train deep learning (DL) models can be expensive or even unfeasible. Therefore, using homogeneous and heterogeneous TL to reduce costs and labor time and enhance classification accuracies would be extremely useful. TL has shown promise in improving the performance of MSI classification. In recent years, a number of studies have used TL to achieve state-of-the-art results on a variety of MSI classification tasks. However, selecting the right source domain and task is essential for successful TL in RS. The performance of the target model is significantly influenced by the source model.

The importance of selecting the right source dataset and task for TL (TL) in multispectral RS image (MSI) classification for both homogeneous and heterogeneous TL was investigated using Sentinel-2 (S2) and Landsat-9 (L9) MSI imagery as heterogeneous domains. Homogeneous TL was evaluated for S2 to S2 and L9 to L9, and heterogeneous TL was evaluated for S2 to L9 and L9 to S2. A variety of metrics were used for evaluation, including precision, kappa, recall, and F1-score. The results showed the importance of TL and suggested using homogeneous TL whenever possible.

This paper is organized as follows. Section 2 delves into the fundamental concepts of TF, including domain and task definitions, and differentiates between homogeneous and heterogeneous TL. Section 3: Methodology outlines the meticulous data preparation process, dataset details, and the architecture of our neural network models used in the experiments. Section 4: Experiments is the heart of our paper, where we present the results and analyses of two distinct TF experiments: Homogeneous TL and Heterogeneous TL. This paper is conducted in Section 5.

2. Deep Transfer Learning

In TF, knowledge is transferred from one domain to another to improve the training process, either in terms of model performance or training speed. TF can be used to address the scarcity and cost of collecting training data. When discussing TF, it is necessary to define some concepts. According to [1], a domain D is defined as a collection of data with a shared feature space X and marginal probability distribution $P(X)$, which can be represented as $D = \{X, P(X)\}$. Similarly, a task T is defined as a set of data with a shared target space Y and objective predictive model M , which can be represented as $T = \{Y, M\}$. TF can be defined as the process of reducing the cost of learning a predictive model in a target domain DT by leveraging knowledge from a source domain DS and a learning task TS , where the source and target domains may have different feature spaces X and/or target spaces Y . This leads to the definition of two fundamental types of TF: homogeneous and heterogeneous.

2.1.A. Homogenous Transfer Learning

Homogeneous (TL) is described as a situation where the feature spaces of data in both the source and target domains match precisely ($X_S = X_T$), the corresponding outcome spaces are identical ($Y_S = Y_T$), and the dimensions of these spaces are also equivalent ($d_S = d_T$). [4]. This alignment simplifies the task of transfer learning (TL), as it establishes a seamless match between the data attributes and the target objectives. When transitioning

knowledge from one domain to another, the transferred model simply requires specialization in tackling the particular task within the target domain. This process can be accomplished through fine-tuning the model, wherein the model is trained using data from the target domain while maintaining the weights acquired during training in the source domain. This is more efficient than randomly initializing and optimizing the model from scratch, as the model already has some knowledge from the previous task. Homogeneous (TF) has proven to be a successful technique in bolstering model performance and expediting training. To illustrate, in reference [5], the authors show that knowledge transfer can be used to improve the performance of image classification on the ImageNet dataset [6].

2.2.B. Heterogeneous Transfer Learning

Heterogeneous (TL) is distinguished by the presence of dissimilar feature spaces and/or target spaces between the source and target domains, as depicted in (see Figure 1). In other words, $X_S \neq X_T$ and/or $Y_S \neq Y_T$. This signifies that the domains may lack shared features, and the dimensions of these features may also differ. Therefore, heterogeneous TF is more challenging than homogeneous TF, as it necessitates bridging the gap between disparate features and their respective quantities. While knowledge transfer remains a possibility in heterogeneous TF, it becomes more challenging due to the need to translate valuable information originally represented in terms of the source domain into an appropriate format for the target domain. Heterogeneous TF is not always possible or advisable, as it can be more difficult to implement and may not lead to the same performance improvements as homogeneous TF.

When the domains, features, or tasks exhibit substantial dissimilarity, a phenomenon known as negative transfer learning (TL) can occur. Negative TL arises when the process of transferring knowledge from the source domain to the target domain has an adverse impact on the model's performance within the target domain. In such cases, it is often more prudent to opt for training a model from the ground up, utilizing advanced techniques like data augmentation, active/meta learning, and others, rather than attempting to transfer a model from the source domain. The different types of TL methods can be distinguished based on their feature space, the difference between the domains, and the tasks that the predictive models are intended to perform.

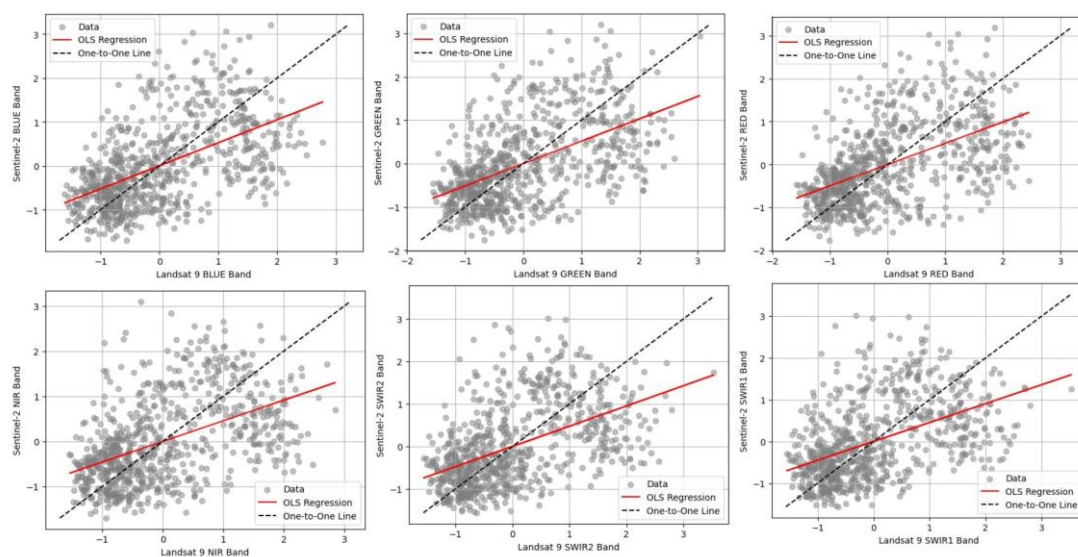


Figure 1. Comparing surface reflectance values from S₂ and L₉ data over the same geographic area. Common bands such as Blue, Green, Red, NIR, SWIR-1, and SWIR-2 are typically selected in both domains. This uniform band selection ensures that the spectral

information used for comparison is consistent across different data sources. In all plots, x-axis represents the L_9 values, and y-axis is the S_2 values.

3. Methodology

3.1 Source and Target Data Preparation

The first step in TL is to pre-process the data and create the appropriate training and validating splits. S_2 and L_9 imageries are downloaded and saved. Then, only the bands (features) that are considered most relevant to the task are kept. These features, which are briefly described in Table 1, are the same for both domains. Once the final datasets are created, they are split into training and test subsets. A 70:30 data split is used, meaning that 70% of the data samples are used for training and 30% are used for testing. Finally, the training and test sets are scaled. At this point, the data is fully processed and the training and test datasets can be used to train and validate the predictive models. The entire data processing stage is performed on both the source and target domain data independently.

Table 1. Selected MSI Bands.

	Sentinel-2		Landsat-9	
	Band Resolution (m)	CV (nm)	Band Resolution (m)	CV (nm)
Blue	10m	490nm	30m	452nm
Green	10m	560nm	30m	561nm
Red	10m	665nm	30m	665nm
NIR	10m	842nm	30m	865nm
SWIR1	20m	1610nm	30m	1609nm
SWIR2	20m	2190nm	30m	2200nm

As discussed in [8,16,20], the key to achieving successful TL and avoiding negative effects is to discover and exploit shared underlying structures between DS ($X; Y$) and DT ($X; Y$).

3.2 Dataset Description

Our dataset comprises two subsets derived from satellite images acquired by S_2 and L_9 , encompassing seven spectral bands (blue, green, red, near-infrared, shortwave infrared 1, shortwave infrared 2), along with calculated indices: NDVI (Normalized Difference Vegetation Index) and NDWI (Normalized Difference Water Index).

The S_2 data was collected in 2023 using a single raster image, along with a cloud mask to detect and mask cloudy areas. In contrast, the L_9 data was gathered over three years, from 2021 to 2023, using a series of five raster images, and it also incorporates a cloud mask for the identification and exclusion of cloudy regions. Each subset contains three CSV files, each with 10,000 samples, resulting in a total of 30,000 samples. The first file focuses on distinguishing olive-bearing land from other land types, featuring two classes: Olive and Non-Olive. The second file targets the differentiation between palm vegetation and other land cover types, comprising Palm and Non-Palm classes. Lastly, the third file aims to detect the presence of buildings within land sections, with classes labeled as Building and Non-Building. This dataset offers a valuable resource for land classification tasks, including land cover mapping, urban planning, and agriculture monitoring, harnessing the richness of spectral bands and computed indices to facilitate accurate and efficient land feature classification in various applications.

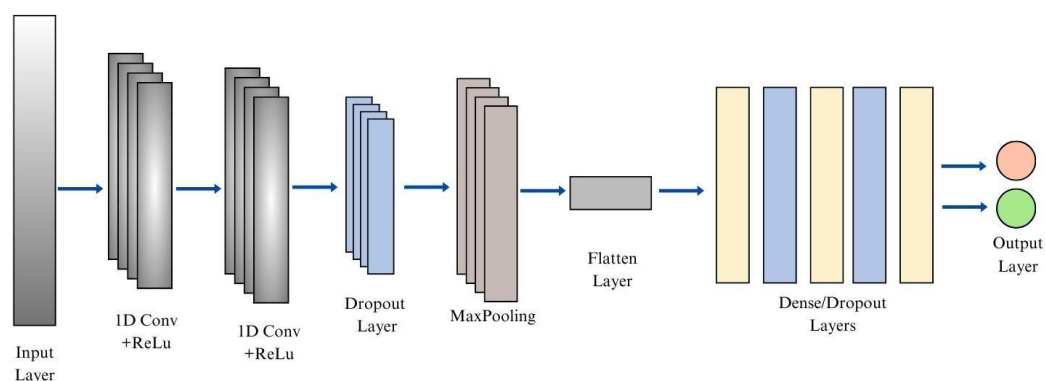


Figure 2. Convolutional Neural-network model architecture. (Find the datasets and the code used in this study on the GitHub repository: https://github.com/mrouba/TL_RS).

3.3 Model Training

After processing the data samples and constructing train and test datasets for both the source and target domains, the next step involves the creation and training of a neural network specifically designed for the source domain. However, to gain a comprehensive understanding of how the proposed solution operates, it is necessary to explain the neural network architectures selected to accomplish this task.

3.4 Neural Network Architectures

1D Convolutional Neural Networks (CNNs) is used here as shown in Figure 2. It consists of two consecutive 1D convolutional layers, followed by one Dropout and one MaxPooling layers, then, features are flattened to be injected in classical classifiers; Multiple Layer Perceptron (MLP).

3.5 Source model training process

CNN weights are randomly initialized via Xavier initialization [18]. All layers use rectified linear unit (Relu) as the activation function and Adam [19] as the optimizer with learning rate 0.001. Training was run for 100 epochs with batch size of 32.

3.6 Target model training (Fine-tuning)

TL idea is based on the conclusions discussed in [5], which state that neural networks have a tendency to acquire more general features in their initial layers and progressively more specific features in the later layers of their architecture. Consequently, even when the features received by the network differ, as long as there is a sufficient degree of similarity between the source domain (DS) and the target domain (DT), the neural network should have the capacity to adapt and leverage this shared information to successfully perform the assigned task. It is generally not advisable to remove and randomly initialize the weights of the input layer when aiming to maximize the transfer of information across domains.

The information contained in the first layer of a neural network can be pivotal for the effective functioning of subsequent layers. Altering or removing this information may lead to a significant decrease in the network's performance. In TL, the final classification layer of a pretrained CNN (SoftMax) is reconfigured to meet the new classification task, and the remaining layers are frozen during training and later used as feature extractors. However, the last two layers are unfrozen and trained to learn the new classification task. Once the TL is performed and the transferred model is capable of processing the input features of the target domain, the performance of the transferred model is evaluated. To do this, an exact copy of the transferred model is created, but with reinitialized weights. This allows the results obtained by the transferred model to be compared to the results that a model with the same architecture trained from scratch would produce.

4. Experiments

This project has been developed using Python and TensorFlow as the DL framework. The developed solution was coded in Jupyter notebook, and executed with CPU Intel Core i7-9700k with 16GB of RAM and GPU NVIDIA 1080-Ti.

4.1. MSI Homogenous TF

Sentinel-2A and Landsat-9 data were leveraged as the target and source domains to enhance the model's capabilities for classifying olive trees. First, a model was trained on Sentinel-2A and Landsat-9 data with a palm classification task. This foundational training enabled the model to grasp the intricacies of identifying palm trees in satellite imagery. Then, the model was fine-tuned on Sentinel-2A and Landsat-9 data with an olive classification task. Fine-tuning the model with data from a source domain with remarkable similarity to the target domain was found to be strategic. By employing a source domain that mirrored the characteristics of the target domain, the fine-tuning process was highly effective, enabling the model to adapt quickly to the nuances of the olive classification task and resulting in a substantial boost in accuracy.

Impressive accuracy was achieved when our model for olive tree classification was fine-tuned using a source model trained for building classification. This experiment broadened our perspective on the capabilities of machine learning models by showing that the traditional boundaries between source and target tasks can sometimes be more fluid than expected. This prompted us to rethink our preconceived notions about domain dissimilarity in machine learning. The ability of our model to fine-tune effectively suggests that there may be latent similarities or transferable knowledge across seemingly distinct domains.

Table 3. Performance Metrics for Homogeneous Transfer Learning from Different Tasks Using Landsat-9 /Sentinel-2A Data.

	Landsat-9				Sentinel-2			
	Accuracy	Precision	Recall	F1-Score	Accuracy	Precision	Recall	F1-Score
TF (Palme -> Olive)	0.98	0.98	0.98	0.98	0.95	0.95	0.95	0.95
TF (Buildings -> Olive)	0.96	0.96	0.96	0.96	0.85	0.88	0.85	0.84

The distribution of the three classes (olive, building, and palm) in Sentinel and Landsat data was visualized using t-SNE plots in Figures 3 and 4, which condensed the multi-dimensional information into a clear two-dimensional representation. This revealed how the classes were distributed and related. The olive class was notably positioned in close proximity to the building and palm classes, suggesting that it shares certain common features or characteristics with the other two classes.

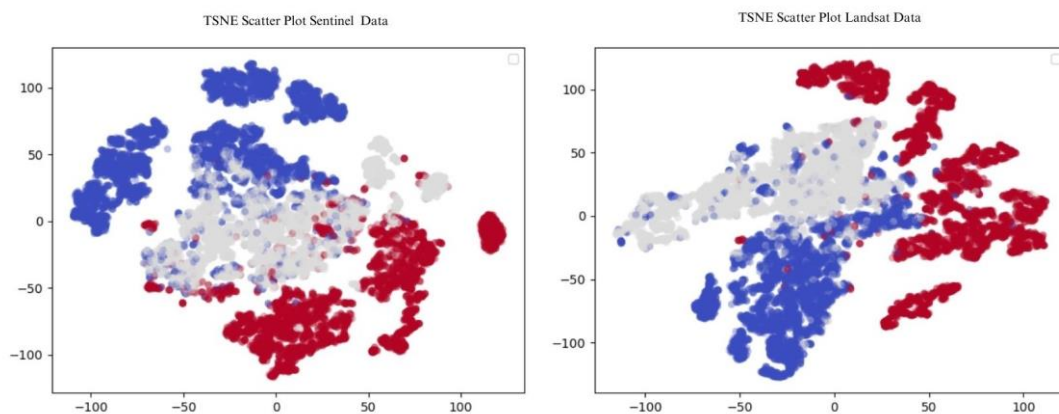


Figure 3: t-SNE Visualization for the frequency vectors of S₂ on right side and L₉ dataset on the left. Reveals Olive in Gray, Palm in Blue, and Building in Red

4.2 MSI Heterogenous TF

A model for olive tree classification was fine-tuned on S₂ data using a building/palm task classification model trained on S₂ images as the source model. The goal was to create a robust model capable of accurately identifying olive trees in L₉ data. However, the dissimilarity between the two domains (L₉ for target model and S₂ for source model) posed a significant challenge. Even though bands that were shared between the two domains were carefully selected, they exhibited a low correlation, making it difficult for the model to generalize effectively.

We undertook the task of resampling L₉ data to match the spatial resolution of S₂, which is 10 meters. It's worth noting that L₉ originally possessed a spatial resolution of 30 meters. The resampling procedure involved the utilization of cubic resampling techniques, and the outcomes were truly noteworthy. The resampled dataset exhibited a remarkable enhancement in classification accuracy when compared to the original data. This discovery underscores the critical significance of carefully selecting appropriate spatial resolution images for RS applications, given their profound influence on the performance and adaptability of machine learning models.

Table 4. Performance Metrics for HTL from Different Tasks using L₉ /S₂ Data.

	Sentinel to Landsat-9				Landsat-9 to Sentinel-2			
	Accuracy	Precision	Recall	F1-Score	Accuracy	Precision	Recall	F1-Score
TF (Palme -> Olive)	0.50	0.25	0.50	0.33	0.67	0.80	0.67	0.63
TF (Buildings -> Olive)	0.50	0.25	0.50	0.33	0.86	0.89	0.86	0.86

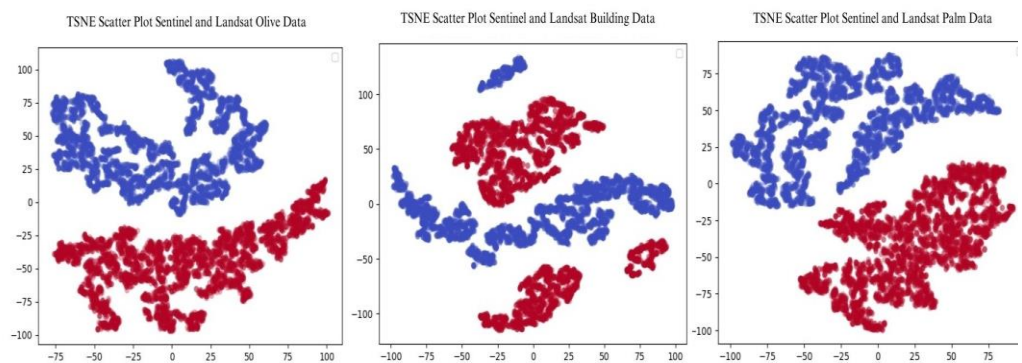


Figure 4: t-SNE Visualization for the frequency vectors of L₉ and S₂ data; S₂ in Blue, and L₉ in Red.

Table 5. Performance Metrics for HTL from Different Tasks using L₉ (resampled) /S₂ Data.

	Sentinel to Landsat-9				Landsat-9 to Sentinel-2			
	Accuracy	Precision	Recall	F1-Score	Accuracy	Precision	Recall	F1-Score
TF (Palme -> Olive)	0.51	0.69	0.51	0.35	0.97	0.98	0.97	0.97
TF (Buildings -> Olive)	0.50	0.25	0.50	0.33	0.86	0.89	0.89	0.85

According to Table 4, the utilization of L₉ imageries as the source domain for fine-tuning on S₂ data as the target domain yields notably strong performance. This remarkable outcome can be primarily attributed to the quality of the L₉ data. The high quality of

the source data significantly contributes to the effectiveness of TL process, providing a solid foundation for the model's adaptation to the target domain.

Table 5 reveals the considerable improvements resulting from the resampling of L9 data to match the spatial resolution of S2 data in the realm of HTL. Accuracy, precision, recall, and F1-Score metrics emphasize the substantial impact of spatial resolution alignment on model performance.

The quality of L9 data is a critical factor influencing the success of fine-tuning in the context of HTL. This observation holds true for both Table 4 and Table 5. When L9 data exhibits higher quality compared to S2, it significantly contributes to the effectiveness of TL process. This underscores the importance of having high-quality source data, as it provides a solid foundation for the model's adaptation to the target domain. Additionally, considering spatial resolution alignment, as demonstrated in Table 5, plays a vital role in enhancing model performance. Therefore, ensuring high-quality source data and thoughtful consideration of resolution are paramount for achieving superior performance in such transfer learning scenarios.

5. Conclusion

This study explores the concept of TL, emphasizing its vital role in domain and task considerations. It categorizes knowledge transfer into two forms: homogeneous and heterogeneous TL, illustrating how source-target domain similarities impact transfer effectiveness. Using S2 and L9 data, we explore these TL scenarios, highlighting the importance of task and source data choice for efficient transfer. Our work underscores the need for thoughtful domain and task selection to optimize TL outcomes. Additionally, our findings reveal the complexity of domain adaptation, showing that the reverse approach doesn't guarantee success. This is influenced by domain nuances, dataset characteristics, and model adaptability. Our study emphasizes the importance of strategic domain selection for effective TL.

6. Declaration of Competing Interest

This study explores the concept of Transfer Learning (TL), emphasizing its vital role in domain and task considerations. It categorizes knowledge transfer into two forms: homogeneous and heterogeneous TL, illustrating how source-target domains similarities impact TL effectiveness. Using S2 and L9 data, we explore these TL scenarios, highlighting the importance of task and source data choice for efficient TL. Our work underscores the need for thoughtful domain and task selection to optimize TL outcomes. Additionally, our findings reveal the complexity of domain adaptation, showing that the reverse approach doesn't guarantee success. This is influenced by domain nuances, dataset characteristics, and model adaptability. Our study emphasizes the importance of strategic domain selection for effective TL.

References

1. Sinno Jialin Pan and Qiang Yang. "A survey on transfer learning." In: IEEE Transactions on knowledge and data engineering 22.10 (2009), pp. 1345–1359.
2. Alhammadi, M. S., et E. P. Glenn. « Detecting Date Palm Trees Health and Vegetation Greenness Change on the Eastern Coast of the United Arab Emirates Using SAVI ». International Journal of Remote Sensing 29, no 6 (mars 2008): 1745-65. <https://doi.org/10.1080/01431160701395195>.
3. Huang, Yixiang, Zhaoqing Wang, Xin Jiang, Ming Wu, Chuang Zhang, et Jun Guo. « Pointshift: Point-Wise Shift MLP for Pixel-Level Cloud Type Classification in Meteorological Satellite Imagery ». In IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium, 607-10. Kuala Lumpur, Malaysia: IEEE, 2022. <https://doi.org/10.1109/IGARSS46834.2022.9883178>.
4. Oscar Day and Taghi M Khoshgoftaar. "A survey on heterogeneous transfer learning." In: Journal of Big Data 4.1 (2017), pp. 1–42.
5. Jason Yosinski et al. "How transferable are features in deep neural networks?" In: arXiv preprint arXiv:1411.1792 (2014).
6. Olga Russakovsky et al. "ImageNet Large Scale Visual Recognition Challenge." In: International Journal of Computer Vision (IJCV) 115.3 (2015), pp. 211–252. DOI: 10.1007/s11263-015-0816-y. Kussul, Nataliia, Mykola Lavreniuk, Sergii Skakun, et Andrii Shelestov. « Deep Learning Classification of Land Cover and Crop Types Using Remote Sensing Data ». IEEE Geoscience and Remote Sensing Letters 14, no 5 (mai 2017): 778-82. <https://doi.org/10.1109/LGRS.2017.2681128>.

7. Zirui Wang et al. "Characterizing and avoiding negative transfer." In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2019, pp. 11293–11302.
8. Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. In International Conference on Machine Learning, pages 1180–1189, 2015
9. « Pires de Lima et Marfurt - 2019 - Convolutional Neural Network for Remote-Sensing Sc.pdf », s. d.
10. Pires de Lima, Rafael, et Kurt Marfurt. « Convolutional Neural Network for Remote-Sensing Scene Classification: Transfer Learning Analysis ». Remote Sensing 12, no 1 (25 décembre 2019): 86. <https://doi.org/10.3390/rs12010086>.
11. Vulova, Stenka, Alby Duarte Rocha, Fred Meier, Hamideh Nouri, Christian Schulz, et Birgit Kleinschmit. « Modeling Urban Evapotranspiration with Sentinel-2, Open Geodata, and Machine Learning in Summertime ». In IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium, 7697-7700. Kuala Lumpur, Malaysia: IEEE, 2022. <https://doi.org/10.1109/IGARSS46834.2022.9884064>.
12. Wang, Anna X., Caelin Tran, Nikhil Desai, David Lobell, et Stefano Ermon. « Deep Transfer Learning for Crop Yield Prediction with Remote Sensing Data ». In Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, 1-5. Menlo Park and San Jose CA USA: ACM, 2018. <https://doi.org/10.1145/3209811.3212707>.
13. « Wang et al. - 2018 - Deep Transfer Learning for Crop Yield Prediction w.pdf », s. d.
14. « Xie et al. - 2016 - Transfer Learning from Deep Features for Remote Se.pdf », s. d.
15. Xie, Michael, Neal Jean, Marshall Burke, David Lobell, et Stefano Ermon. « Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping ». Proceedings of the AAAI Conference on Artificial Intelligence 30, no 1 (5 mars 2016). <https://doi.org/10.1609/aaai.v30i1.9906>.
16. M. Long, Y. Cao, J. Wang, and M. I. Jordan. Learning transferable features with deep adaptation networks. In Proceedings of the 32nd International Conference on International Conference on Machine Learning, 2015.
17. Zhao, Huizhen, Fuxian Liu, Han Zhang, et Zhibing Liang. « Convolutional Neural Network Based Heterogeneous Transfer Learning for Remote-Sensing Scene Classification ». International Journal of Remote Sensing 40, no 22 (17 novembre 2019): 8506-27. <https://doi.org/10.1080/01431161.2019.1615652>.
18. Xavier Glorot and Yoshua Bengio. "Understanding the difficulty of training deep feedforward neural networks." In: Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings. 2010, pp. 249–256.
19. Diederik P Kingma and Jimmy Ba. "Adam: A method for stochastic optimization." In: arXiv preprint arXiv:1412.6980 (2014).
20. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In Computer Vision and Pattern Recognition (CVPR), volume 1, page 4, 2017