

High-resolution seismic reservoir monitoring with multitask and transfer learning

Ahmed M. Ahmed & Ilya Tsvankin

*Center for Wave Phenomena and Dept. of Geophysics, Colorado School of Mines, Golden CO 80401
email ahmedmohamedahmed@mines.edu*

ABSTRACT

High-resolution real-time monitoring of reservoir changes is essential during CO₂ injection or hydrocarbon production. This study leverages convolutional neural networks (CNNs) that employ multitask (MTL) and transfer (TL) learning to accurately predict relevant reservoir parameters from seismic data. CNNs are initially trained to estimate the P-wave velocity from 2D multicomponent seismic data and then are fine-tuned through TL to obtain the S-wave velocity, density, and saturation. This methodology is applied to a synthetic CO₂ sequestration model based on California's Kimberlina storage reservoir. When using MTL, CNNs are trained simultaneously on multiple related tasks by taking advantage of their commonalities. After pretraining the model on a 2D line, it can be fine-tuned to predict the reservoir parameters from the data in the crossline direction. Our work addresses the challenge of training-data scarcity, promotes efficient use of computational resources in reservoir monitoring, and helps increase the accuracy of real-time monitoring of the fluid movement inside the reservoir.

Key words: CO₂ monitoring, time-lapse, multitask-learning, transfer-learning

1 INTRODUCTION

Monitoring the time-lapse changes in subsurface reservoirs is crucial for effectively managing oil and gas production and CO₂ injection for carbon capture and storage (CCS). These changes, particularly the fluid movement and saturation levels, directly impact operational strategies and environmental safety. Elastic time-lapse full-waveform inversion (FWI) of seismic data is capable of estimating the parameters of realistic heterogeneous anisotropic models (Queißer and Singh, 2013; Liu and Tsvankin, 2021). However, FWI faces such challenges as parameter crosstalk and trade-offs, which are partially mitigated by choosing an optimal model parameterization (Mardan et al., 2022; Pan et al., 2018).

The variation in fluid saturation within a reservoir leads to changes in the elastic parameters (the P- and S-wave velocities and density), with density being the most difficult parameter to estimate in time-lapse FWI (Pan et al., 2018). Evaluating fluid saturation from time-lapse seismic data involves complex and often error-prone processes (Ajo-Franklin et al., 2013; Vasco et al., 2019). Finally, conventional (physics-based) inversion methods, such as FWI, are computationally intensive and may not be suitable for efficient continuous long-term seismic monitoring.

Here, we employ recent advances in computing science and machine learning to solve complex nonlinear inverse problems that arise in reservoir monitoring. Deep-learning (DL) algorithms use a data-driven input-output relationship that circumvents some of the FWI constraints. Without relying solely on physical and geologic information (Biswas et al., 2019; Zhang and Gao, 2022; Liu et al., 2023; Wang et al., 2023b; Wu et al., 2019; Liu et al., 2020). Among commonly used ML techniques are convolutional neural networks (CNNs), recurrent neural networks (RNNs), and transformer methodologies. For example, Um et al. (2022, 2024) present a DL network for monitoring CO₂ saturation. Feng et al. (2023) introduce a benchmark data set for predicting the P- and S-wave velocities using CNNs. Long short-term memory (LSTM) and CNNs are employed by Simon et al. (2023) for velocity model-building.

Wang et al. (2023a) combine CNNs (U-Net autoencoder - autoencoder) and transformers in seismic interpretation for fault detection. Application of transformers in geophysics is still nascent compared to CNNs, partly because transformers are historically

more aligned with Natural Language Processing (NLP) algorithms (Maurício et al., 2023). If CNNs are trained with transfer learning (TL), they may perform better than pure transformer-based networks, especially for relatively small data sets (Kornblith et al., 2019; Maurício et al., 2023).

Incorporating multitasking into CNNs is a promising approach for predicting multiple reservoir parameters. Multitasking with multiple encoders and decoder CNN architectures is a powerful tool for capturing complex geospatial relationships by simultaneously processing multiple seismic attributes. The encoder compresses the input data into a compact representation for the inversion tasks. The multiple decoders handle several tasks simultaneously by learning various properties of seismic data. However, these architectures require extensive and diverse training data sets to capture the variability in geologic formations. They also incur higher computational costs due to their complex structure.

Transfer learning (TL) leverages models pretrained on similar data sets, which enables the application of the learned features to new (but related) tasks. TL enhances model generalizability and accelerates the training process, which helps overcome data scarcity and/or lack of computational resources (Simon et al., 2023). The effectiveness of TL depends on the relevance of the features learned from the source domain for the target domain, which makes selection of appropriate pretrained model particularly important. The papers by Tajbakhsh et al. (2016), Becherer et al. (2019), and Kumar et al. (2022) highlight the value of fine-tuning in optimizing neural-network models and a delicate balance in adjusting pretrained models to new tasks. Although geologic and geophysical model-driven CNNs are potentially capable of estimating the elastic parameters (Li et al., 2022), they are limited by the memory and data requirements, particularly for 3D data sets (Li et al., 2023). The results also depend on the quality of the input data, and CNNs require extensive computational resources.

This study aims to leverage the strengths of CNNs with TL and MTL to efficiently predict relevant reservoir parameters from seismic data. We propose a novel approach in which CNNs, initially trained to estimate the P-wave velocity from multicomponent seismic data, are fine-tuned through TL to obtain other medium parameters. We start by discussing multitasking (MTL), transfer learning (TL), and the corresponding neural network architectures. Then we present a methodology for preparing training data using the available physics information and prior knowledge. The network’s performance is validated using noise-contaminated data generated for the realistic synthetic model of CO₂ injection into Kimberlina reservoir in California. Finally, we discuss the impact of the available prior information and of the quality of the simulated training data on the accuracy of the reconstructed time-lapse parameter variations.

2 METHODOLOGY

2.1 Multitask learning (MTL)

Applying deep learning to seismic inversion allows CNN algorithms to recognize patterns and infer relationships in seismic data for parameter predictions. MTL is designed to enhance model generalization by leveraging domain-specific information across related tasks. It involves parallel training on a shared representation that includes feature extraction from the encoders. After obtaining these shared features, the network branches into multiple decoders, each responsible for a separate task. The outputs from these decoders are evaluated, and the cumulative losses are used in the back-propagation to further train the network. This feedback loop allows the network to develop shared representations that are beneficial for all tasks (Figure 1).

The MTL architecture, conceptualized by Caruana (1997), fosters inter-task learning with different tasks acting as mutual regularizers, which increases prediction accuracy (Vafaeikia et al., 2020; Zhang and Yang, 2022). This architecture is further analyzed in recent studies by Li et al. (2023), where MTL has been applied to seismic inversion tasks, demonstrating its effectiveness in predicting reservoir properties. The main goal of MTL is to obtain the mapping function $f : X \rightarrow Y$ that connects the input domain $\mathbf{x}(i) \in X$ to the output domain $\mathbf{y}(i) \in Y$ across n input-output pairs $\{[\mathbf{x}(1), \mathbf{y}(1)], [\mathbf{x}(2), \mathbf{y}(2)], \dots, [\mathbf{x}(n), \mathbf{y}(n)]\}$. This mapping seeks to minimize the loss across all tasks without increasing the model complexity by using regularization:

$$\min_W \sum_{t=1}^T \sum_{i=1}^{n_t} \left\{ \mathcal{L}[w_t, x_t(i), y_t(i)] \right\} + \lambda \Omega(W), \quad (1)$$

where $t = 1, 2, \dots, T$ is a task with the corresponding data set $[\mathbf{x}_t(i), \mathbf{y}_t(i)]$, n_t is the number of training samples for the task t , w_t are the model parameters for t , \mathcal{L} is the total loss for the model parameters w_t , $W = \{w_t\}_{t=1}^T$ is the combined set of model parameters for all related tasks, and λ is the weighting factor for the regularization term Ω .

Our implementation utilizes a CNN-based autoencoder-decoder framework to extract shared features among tasks and translate them into the pertinent elastic parameters. Despite the challenges posed by cross-task interference and the need for a large network size to accommodate multitask learning, this approach has proved effective by leveraging multiple prediction tasks.

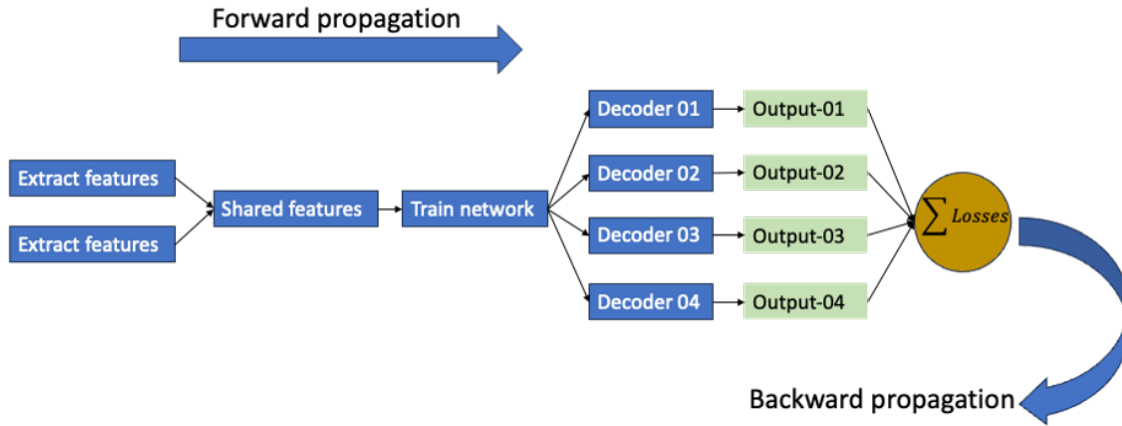


Figure 1. Flowchart of a multitask learning (MTL) CNN architecture.




Figure 2. Transfer-learning workflow includes sequential steps of freezing layers, fine-tuning the non-frozen layers to adapt to new data, training the model with the new data, and evaluating the model's performance for the new task.

2.2 Transfer learning (TL)

Our project is focused on monitoring the CO₂ movement in the reservoir by employing a neural network trained with seismic data to predict the velocities, density, and saturation. CNN models, however, tend to have a large size, which necessitates using substantial memory and storage. Also, training these models is computationally expensive due to the complexity of the data and the intricacies of the simulation. Finally, fine-tuning the models requires a careful optimization of numerous hyperparameters.

To mitigate these challenges, we employ transfer learning, which enables us to utilize pretrained models that can be adapted to a specific task with relatively minor adjustments. The proposed methodology involves training the CNN on a well-sampled parameter and then fine-tuning it for other parameters in a different domain (Ruder, 2019). This approach mitigates the problems caused by a limited sample size for specific parameters and can produce accurate results without requiring an extensive data set for fine-tuning.

Pretraining on a sufficiently large data set before transferring knowledge to a new task significantly increases the accuracy and learning capabilities of the network compared to training models from scratch (Zhao et al., 2024), which enhances both performance and generalization (Figure 2). TL can mitigate the need for new data collection by transferring knowledge from a domain similar to the target (Weiss et al., 2016; Zhuang et al., 2020). Depending on the domain variability, TL can be categorized as homogeneous (same feature space) or heterogeneous (different feature spaces). For example, Simon et al. (2023) apply TL to velocity model-building by starting with horizontally layered media before training the network on more complex 2D models with dipping structures. TL implementation involves feature extraction or fine-tuning, both aimed at reducing computational demands (Ruder, 2019). We demonstrate the effectiveness of this approach by training our model using data acquired on a 2D line (inline) and then applying the pretrained model to fine-tune the parameters on the orthogonal line (crossline). This approach facilitates hyperparameter optimization. Unlike MTL, which balances the focus across multiple tasks, TL is designed to prioritize the target task.

2.3 Neural network architecture

We developed a CNN architecture designed for multiparameter prediction of reservoir properties by utilizing auto-encoders and auto-decoders. The starting point is open-access software InversionNet (InvNet; Wu and Lin, 2020) for P-wave velocity prediction from seismic data for a synthetic horizontally layered model. The new features of InvNet include adjustments to accommodate larger models and optimization of the activation function. We changed the model dimensions, modified the corresponding convolutional layers, and replaced rectified linear unit (LeakyRELU) activation functions by the Gaussian error linear Unit (GELU). The three CNN versions developed and used in our study are as follows:

- (i) Multitask learning network with two encoders and four decoders designed to predict model several parameters simultaneously.
- (ii) Multitask learning network with two encoders and two decoders designed to predict the P- and S-wave velocities. The transfer-learning methodology is applied to fine-tune the pre-trained model to predict the other model parameters.
- (iii) Single-task learning network with two encoders and one decoder designed to predict only the P-wave velocity, with TL employed to predict the other model parameters.

Activation functions are mathematical operations applied to the input of a neuron. Non-linearity of the model is crucial for the network to learn complex mappings between inputs and outputs, enabling it to capture intricate patterns in the data. Thus, the choice of activation functions can significantly influence the performance of trained models in seismic inversion (Bai et al., 2021; Xie et al., 2021). The GELU activation function, introduced by Hendrycks and Gimpel (2016), is smooth and differentiable, in contrast to the LeakyReLU (Figure 3). As a result, GELU has been employed in various large language models (LLM) and transfer-learning models, including BERT (Devlin et al., 2019), GPT (Brown et al., 2020), and ViT (Dosovitskiy et al., 2021).

For positive values of x in equation 2, the GELU function gradually approaches unity (Figure 3), thus allowing more of the input to pass through with increasing x . The LeakyReLU, however, passes through all positive values with no changes. For $x \approx 0$, the GELU function allows a small fraction of negative values to pass through, similar to a leakyReLU, but in a smooth manner. This helps in maintaining the gradient flow during backpropagation for small x , which enhances the learning dynamics. The probabilistic nature of GELU improves the generalization performance of the trained model. GELU avoids the "dying" ReLU problem (where neurons can become inactive and only output zero) and promotes a more diverse activation pattern across the neurons. The GELU function is represented as:

$$\text{GELU}(x) = x \Phi(x) = x \left(\frac{1}{2} \left[1 + \text{erf} \left(\frac{x}{\sqrt{2}} \right) \right] \right), \quad (2)$$

where $\Phi(x)$ denotes the cumulative distribution function (CDF) of the standard normal distribution, which effectively weighs the inputs in a nonlinear fashion based on their magnitude. This inherent probabilistic nature of GELU supports a more flexible input processing, leading to variable decision boundaries, which can be particularly beneficial in complex learning problems. To increase computational efficiency, GELU is often approximated as:

$$\text{GELU}(x) = 0.5 x \left[1 + \tanh \left(\sqrt{\frac{2}{\pi}} (x + 0.045 x^3) \right) \right]. \quad (3)$$

This approximation closely emulates the behavior of the Gaussian CDF, thus offering a practical solution for integrating GELU into neural networks. The network architecture effectively translates seismic data into subsurface parameter predictions while balancing the model complexity against computational demands. The encoder part of the network consists of eight convolutional blocks, each containing a convolutional layer, batch normalization, and GELU (Table 1). If the network inputs have a zero mean and the unit variances are not correlated, the convergence of the network speeds up. The decoder part consists of six deconvolutional blocks that produce the outputs. The input to the network includes the vertical and horizontal displacement components of the simulated shots. The loss functions of different tasks are calculated by combining the L_1 -norm and the mean-square error (MSE). The total loss function of the network is defined as:

$$L_{\text{total}} = \sum_{i=1}^k \lambda_i (l_i + \text{MSE}_i), \quad (4)$$

where k is the number of the predicted parameters, l_i is the calculated L_1 -norm for each predicted parameter, and λ_i are the scaling factors that define the contributions of each loss function to the total loss.

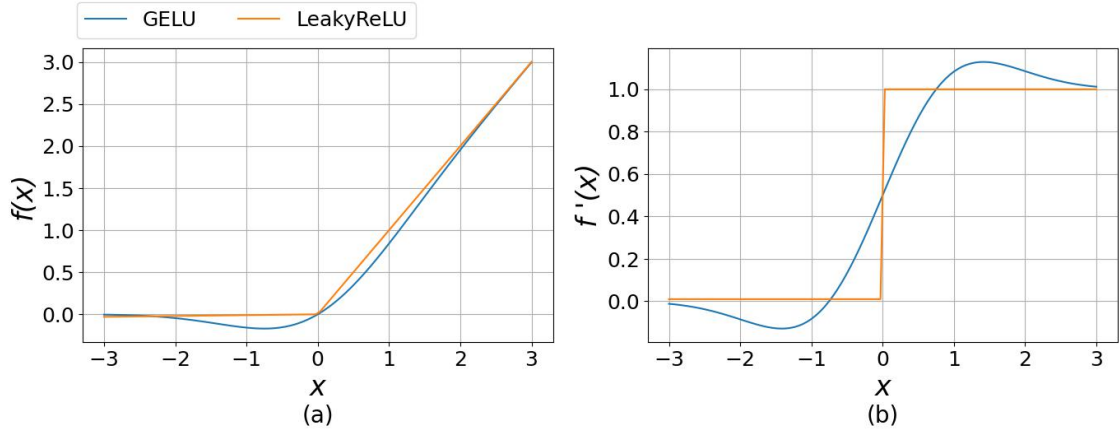


Figure 3. (a) GELU and LeakyReLU activation functions. (b) The derivatives of both functions.

Layer Type	Output shape	Kernel size	Stride	Activation
ConvBlock 1	32x500x97	7x1	2x1	GELU
ConvBlock 2	64x250x97	3x1	2x1	GELU
ConvBlock 3	64x125x97	3x1	2x1	GELU
ConvBlock 4	128x63x97	3x1	2x1	GELU
ConvBlock 5	128x32x49	3x3	2x2	GELU
ConvBlock 6	256x16x25	3x3	2x2	GELU
ConvBlock 7	256x8x13	3x3	2x2	GELU
ConvBlock 8	512x1x1	8x12	2x2	GELU
DeconvBlock 1	512x10x10	4x4	2x2	GELU
DeconvBlock 2	256x20x10	4x4	2x2	GELU
DeconvBlock 3	128x40x40	4x4	2x2	GELU
DeconvBlock 4	64x80x80	4x4	2x2	GELU
DeconvBlock 5	32x160x160	4x4	2x2	GELU
Pad	32x88x150			
DeconvBlock 6	512x10x10	3x3	1x1	Tanh

Table 1. Summary of the CNN architecture.

3 SYNTHETIC TEST

3.1 Data preparation

3.1.1 Forward modeling

The proposed methodology is applied to synthetic data for the model of the Kimberlina reservoir, a potential CO₂ storage site in the Southern San Joaquin Basin of California. This data set includes models corresponding to different times during CO₂ injection (from pre-injection to 20 years after its start). The available parameters include the P- and S-wave velocities, density, fluid saturation, and electrical resistivity defined on a 3D grid. Synthetic well logs of the P-wave velocity, density, and CO₂ saturation in the injection well and three monitoring wells provide additional information for monitoring CO₂ migration (Alumbaugh et al., 2023).

The velocity, density, and CO₂ saturation models are provided on a 10 m x 10 m x 10 m grid with 601 x 601 x 351 grid points. A 2D seismic line extracted at $x = 2100$ m helps map the CO₂ movement and calculate the time-lapse variations of the observed parameters. CO₂ is injected into a saline reservoir through a central well that traverses six sand-shale layers. The CO₂ movement inside the reservoir is modeled by considering buoyancy-driven migration patterns (Sigfusson et al., 2015).

Configuration	Training samples	Trained parameters per layer	Batch size	Epochs	Training duration (hrs:min:sec)
2 Enc. - 4 Dec.	12000	All	60	120	2:45
2 Enc. - 2 Dec.	12000	All	60	120	1:45
2 Enc. - 1 Dec.	12000	All	60	120	1:35
2 Enc. - 1 Dec.	2000	*see below	60	120	0:15

Table 2. Encoder-decoder configurations and their training duration for one GPU node.

*Trainable parameters are in the last layer of the encoder and all layers of the decoder.

3.1.2 Data generation

Saturation levels can be approximately estimated using empirical relationships between the P-wave velocity, porosity, and permeability validated by laboratory and field studies (Gassmann, 1951; Mavko and Mukerji, 1995). We adopt a linear fitting formula from Liu et al. (2023) to establish the velocity-saturation relationships from well-log data. This approach relies on the geologic similarity between the field of interest and those used to derive the relation between velocity and saturation. Prior knowledge of the actual baseline model is used to generate reliable training samples. The baseline velocity field, assumed to be obtained by full-waveform inversion (FWI), is perturbed only inside the reservoir to obtain the training data (labels) and ensure accurate predictions from the monitor data set. Logs from the injection and monitor wells are used to provide geologic and geophysical information that helps constrain the training samples for the inversion.

A total of 14,000 data samples were generated to model the reservoir parameters (V_P , V_S , ρ , saturation, V_P/V_S ratio, and acoustic impedance). The synthetic elastic wavefield is excited by 37 shots (point explosions) placed with an increment of 80 m at a depth of 130 m. The source signal is the Ricker wavelet with a central frequency of 10 Hz. We employ 289 receivers evenly distributed with an increment of 20 m at a depth of 140 m. A subset of 12,000 samples covering the time interval from zero to 10 years after the injection was used to train the P-wave velocity prediction model. Testing for V_P is carried out with an additional 500 samples that correspond to about 12 years post-injection. Then transfer learning is employed to fine-tune the predictions of V_S , ρ , saturation, V_P/V_S ratio, and acoustic impedance using 2,000 training samples and 500 testing samples.

3.2 Test results

3.2.1 Multitask and transfer learning

Examining the training time for various encoder-decoder configurations helps evaluate the model’s efficiency (Table 2). As expected, the time increases with the number of encoders and decoders. For 120 training epochs, the configuration with two encoders and four decoders takes the most time (approximately 2 hr 45 min using one GPU node). In contrast, the setup with a single decoder reduces the training time to about 1 hr 35 min. The targeted training of the last layer of the encoders and all decoder layers during the fine-tuning phase brings the training time down to just 15 min, thus confirming the method’s computational efficiency. Hence, focusing on fewer convolutional layers could offer an advantage for rapid convergence or in iterative training scenarios. Further analysis will explore the implications of these efficiency gains for model performance and generalizability when using a pretrained model for transfer learning.

Analyzing the learning rates and convergence patterns, particularly for the four-decoder configuration, reveals informative trends (Figure 4). The normalized loss function trajectories indicate that the velocities V_P and V_S share similar decline patterns. Therefore, their optimization paths can be aligned, whereas the density and saturation curves exhibit a more gradual decline and may require different training strategies. Figure 5 shows the reservoir parameters predicted with the MTL approach. The input to the neural network used for training included the simulated shots and the four training parameters (V_P , V_S , ρ , saturation). Our approach to optimizing the seismic inversion process involves a combination of MTL and TL. First, the neural network is trained using 12,000 samples to predict the velocity V_P (Figure 6). Then the pretrained model is employed to optimize the seismic inversion process by fine-tuning the network to predict the other medium parameters individually. During this process, it is possible to freeze and gradually unfreeze parts of the network. We emulate a field-data application where the number of training samples for each parameter is different. A smaller number of training samples is used for fine-tuning, which helps reduce computational resources without impacting the prediction accuracy. In elastic FWI, parameterizing the model in terms of the V_P/V_S ratio and acoustic impedance (IP) often reduces the crosstalk compared to using V_P , V_S and ρ . Therefore, next we perform the following tests:

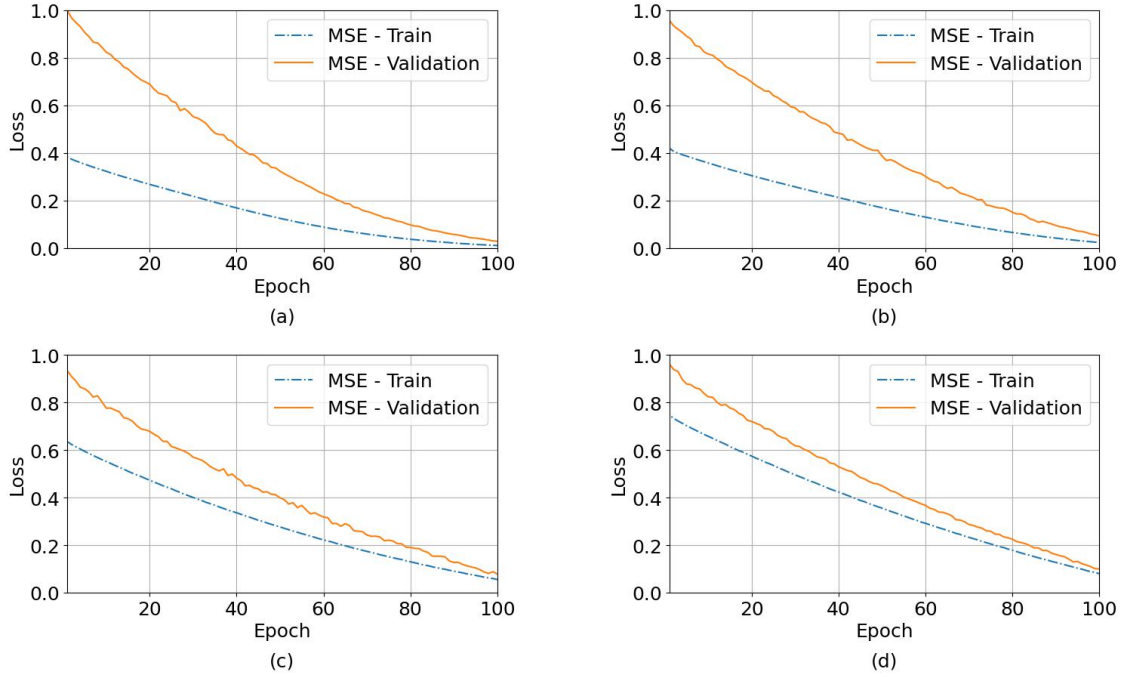


Figure 4. Normalized loss functions for different parameters with four decoders. (a) V_P , (b) V_S , (c) ρ , and (d) saturation.

Parameter	SSIM Index	Frobenius norm	l_1 -norm	l_2 -norm	MSE
V_P (noise-free)	1.03	5.22	1.54	5.22	2.01
V_P (SNR = 15%)	0.86	1.13	2.94	1.13	8.51
V_S (noise-free)	0.98	9.72	1.24	9.62	7.11
V_S (SNR = 15%)	0.95	2.03	5.54	2.03	3.22
ρ (noise-free)	1.12	5.63	8.4	5.63	2.43
ρ (SNR = 15%)	0.97	6.33	1.45	6.33	2.93
Saturation (noise-free)	1.00	0.64	8.55	0.64	3.1×10^{-5}
Saturation (SNR = 15%)	0.99	1.32	22.3	1.32	1.33×10^{-4}

Table 3. Performance metrics of the parameters predicted from the clean and noisy (the signal-to-noise ratio is 15%) data.

- (i) Use dual encoder - dual decoders to train the CNN for the V_P/V_S ratio and IP.
- (ii) Use a pretrained model to fine-tune the CNN parameters for each parameter (V_P , V_S , ρ) individually.

The results of these two tests are similar, but the second approach is more computationally efficient (Figure 8). The pretrained model also maintains prediction accuracy when the testing data are contaminated with Gaussian noise with the signal-to-noise (SNR) ratio equal to 15 (Table 3).

3.2.2 Influence of pretrained model distribution on fine-tuning efficiency

The parameters that closely match the source task’s distribution are predicted with faster convergence, which underscores the importance of the parameter distribution in model adaptation (Tajbakhsh et al., 2016; Kumar et al., 2022; Becherer et al., 2019). Thus, the range and distribution of the medium parameters strongly influence the model’s fine-tuning efficiency.

In the proposed flow, the CNNs were trained for the velocity V_P , which ranges from 1.9 km/s to 3.9 km/s, and its normalized values range approximately from 0.5 to 1.0. The histogram (Figure 11) shows a highly skewed distribution which might be a sign of outliers or non-representative samples in the data. This could potentially lead to slower convergence if the new data set does not match this distribution.

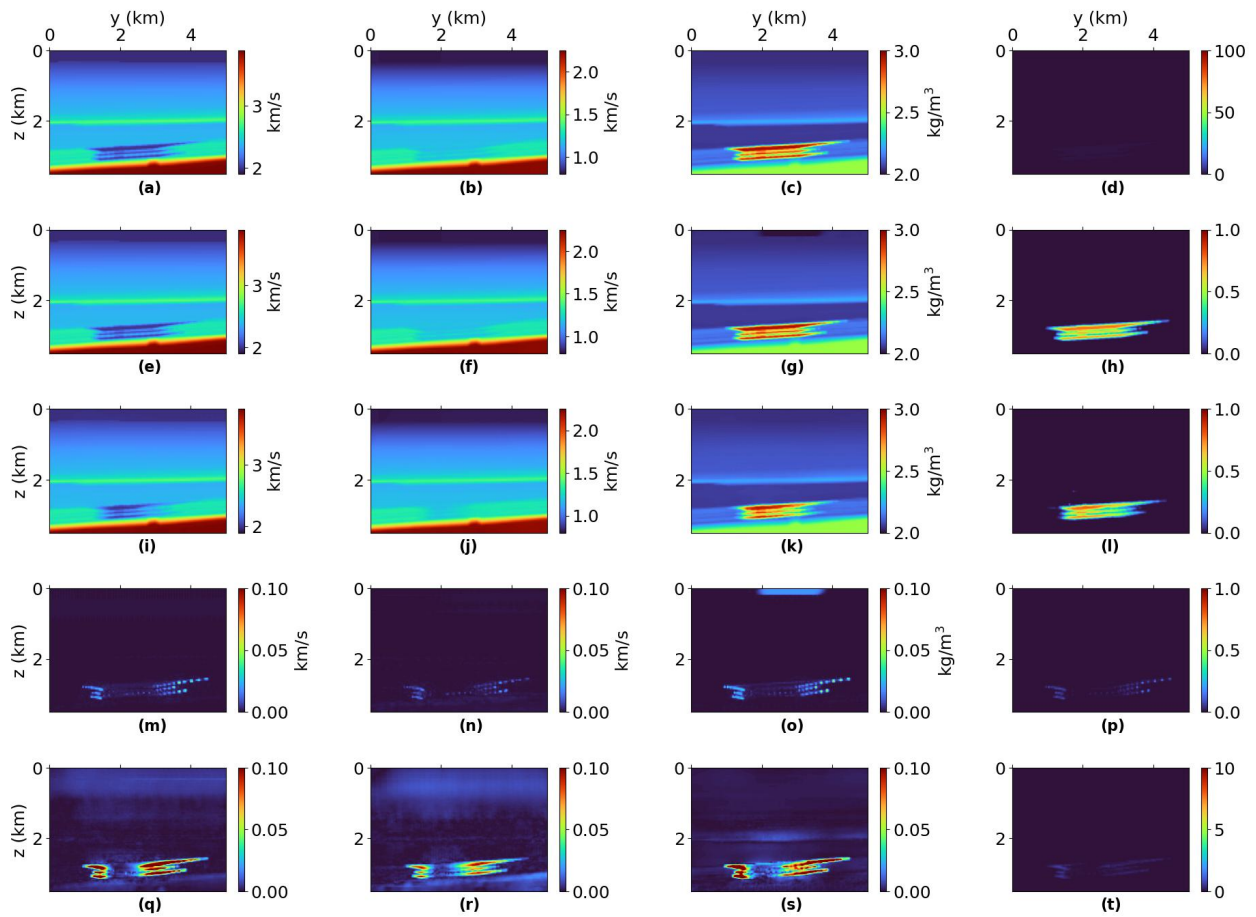


Figure 5. Predictions of the reservoir parameters. (a)-(d) The actual models of the P-wave velocity, S-wave velocity, density, and saturation, respectively. (e)-(h) The corresponding models predicted from the noise-free data. (i)-(l) The corresponding models predicted from the noisy data. (m)-(p) The normalized differences between (e) and (a), (f) and (b), (g) and (c), and (h) and (d), respectively. (q)-(t) The normalized differences between (i) and (a), (j) and (b), (k) and (c), and (l) and (d), respectively.




Figure 6. P-wave velocity field for the Kimberlina reservoir. (a) The actual model, (b) the model predicted from the noise-free data, and (c) the model predicted from the noisy data. The normalized differences: (d) between plots (b) and (a), and (e) between plots (c) and (a).




Figure 7. Reservoir parameters predicted with transfer learning. The actual models of the (a) S-wave velocity, (b) density, and (c) saturation. (d), (e), (f) The corresponding models predicted from the noise-free data. (g), (h), (i) The corresponding models predicted from the noisy data. (j), (k), (l) The normalized differences between (d) and (a), (e) and (b), and (f) and (c), respectively. (m), (n), (o) The normalized differences between (g) and (a), (h) and (b), and (i) and (c), respectively.

We fine-tuned the pretrained model to predict the velocity V_S but could not achieve satisfactory convergence. In particular, when CNN was tested on noisy data, the network did not perform well (Figure 7g). The velocity V_S ranges from 0.8 km/s to 2.25 km/s, and its normalized range is from 0.35 to 1.0, which is wider than that of V_P (Figure 11b). Because the pretraining was done with V_P , whereas V_S has a different distribution, the model requires more iterations to properly adjust its weights for predicting V_S .

Next, the network was fine-tuned to predict the V_P/V_S ratio (Figure 8e), which ranges from 1.7 to 2.47, and its normalized values range from 0.7 to 1.0 (Figure 11c). During the fine-tuning process, the model leverages its existing knowledge of V_P to achieve faster convergence for V_P/V_S .

In another test, we calculated the velocity V_S from the V_P/V_S ratio. Figure 9 compares the V_S predicted directly and calculated from the V_P/V_S ratio. The better prediction of V_P/V_S yields more accurate estimates of V_S , as discussed in more detail below. Another way to overcome the issue of slow convergence with a limited number of training samples is to remove the outliers or non-representative values. Therefore, we reduced the range of V_S from 0.8 - 2.23 km/s to 1.0-1.8 km/s prior to fine-tuning the




Figure 8. Sections of the V_P/V_S ratio and the acoustic impedance (IP). The ratio V_P/V_S : (a) the actual model, (c) the model predicted from the noise-free data, and (e) the model predicted from the noisy data. (g) and (i) The normalized differences between (c) and (a) and (e) and (a), respectively. The impedance (IP): (b) the actual model, (d) the model predicted from the noise-free data, and (f) the model predicted from the noisy data. (h) and (j) The normalized differences between (d) and (b) and (f) and (b), respectively.

pretrained model. With this more narrow distribution, the model can converge faster, as the reduced range of the input, helps focus the fine-tuning process. It takes only 120 training epochs for the model to converge and achieve an accurate V_S prediction for the noisy data (Figure 10). If the distribution of the target task is different from that of the source task, there are two possible scenarios. First, if the number of training samples is sufficient for fine-tuning, the network learns the new distribution and converges to an accurate solution. Second, if the number of samples is limited, the network requires more training epochs to converge.




Figure 9. Comparison between the directly predicted V_S models and the ones calculated from the predicted V_P/V_S ratio. (a) The model predicted from the noise-free data, (c) the model calculated from the V_P/V_S ratio predicted with the noise-free data, and (e) the normalized difference between plots (c) and (a). (b) The model predicted from the noisy data, (d) the model calculated from the V_P/V_S ratio predicted with the noisy data, and (f) the normalized difference between plots (d) and (b).




Figure 10. Sections of the shear-wave velocity. The predictions are made using a reduced range of V_S (1.0-1.8 km/s). (a) The actual model, (b) the model predicted from the noise-free data, and (c) the model predicted from the noisy data. The normalized differences: (d) between plots (b) and (a), and (e) between plots (c) and (a).




Figure 11. Histograms of the reservoir parameters: (a) V_P , (b) V_S , (c) V_P/V_S ratio, and (d) constrained V_S . Each parameter is normalized by its maximum value.

4 DISCUSSION

Implementing multitask learning (MTL) for simultaneous prediction of four reservoir parameters involves a complex interplay between computational resources and model optimization. The large GPU memory required for such tasks significantly limits the critical CNN parameters, such as the batch size, which in our study was constrained to 60. This limitation underscores the computational demands of MTL and highlights the need for efficient memory management strategies.

Moreover, the variable learning rates and the extended duration of training required for each parameter further emphasize the inherent complexity in the simultaneous prediction of multiple properties. In particular, Figure 5 illustrates the challenges in accurately predicting density up shallow, which is likely caused by the local optimization minima encountered during training. Clearly, it is essential to strike a balance between extending training to enhance model performance for specific parameters and the risk of overfitting other parameters.

The trial-and-error nature of neural-network training, driven by the random initialization and subsequent parameter optimization, is time consuming, especially in MTL. This issue is critical in field-data applications, where it is desirable to include additional recorded components in the training process. Transfer learning (TL) offers a promising approach for incorporating more features into the model, thereby enhancing the model’s applicability and making it more flexible for parameter predictions. The resolution and robustness of these predictions are controlled by the quality and diversity of the training samples, which should encompass reservoir geometry and thickness, well data, and other geologic and geophysical information. Such comprehensive training data sets are essential for accurately capturing the time-lapse changes in the reservoir.

In this study, we assumed that each layer is laterally homogeneous. Ongoing research focuses on adapting these methodologies for field-data applications by properly taking into account heterogeneity and anisotropy. We also plan to employ transfer learning for processing of 3D surveys using 2D prediction results.

5 CONCLUSIONS

Reservoir monitoring in hydrocarbon production and CO_2 injection can be conducted through high-resolution estimation of the elastic parameters and saturation from seismic data. We proposed an efficient data-driven approach to seismic inversion that consists of multitasking (MTL) and transfer-learning (TL) algorithms incorporated into convolutional neural networks (CNNs) that employ the GELU activation functions. The proposed methodology improves the trained model’s learning dynamics and generalization properties. The shared features of the recorded data are extracted to predict one parameter, which is followed by fine-tuning the pretrained model to predict other parameters. We addressed several key challenges in ML techniques, including mitigation of data scarcity, efficient model initialization for trainable parameters, and optimization of computational resources.

The methodology is tested on the realistic Kimberlina reservoir model to emulate field-data applications. The P-wave velocity is predicted for a 2D line (inline) using a CNN that consists of dual encoders and one decoder. Then the pretrained model is fine-tuned to estimate the S-wave velocity, density, saturation, V_P/V_S ratio, and acoustic impedance. The robustness and generalization of the CNN are validated using the input data contaminated by Gaussian noise. The proposed fine-tuning process also makes it possible to estimate the medium parameters from data acquired in the crossline direction. The performance of the proposed method depends on the availability of prior information and the accuracy of the baseline model. Ongoing work includes an extension of this methodology to 3D seismic surveys with a focus on enhancing the model adaptability for field-data applications.

6 ACKNOWLEDGEMENTS

We thank INPEX and the sponsor companies of the Center for Wave Phenomena, whose support made this research possible.

REFERENCES

- Ajo-Franklin, J., J. Peterson, J. Doetsch, and T. Daley, 2013, High-resolution characterization of a CO₂ plume using crosswell seismic tomography: Cranfield, MS, USA: *International Journal of Greenhouse Gas Control*, **18**, 497–509.
- Alumbaugh, D., E. Gasperikova, D. Crandall, M. Commer, S. Feng, W. Harbert, Y. Li, Y. Lin, and S. Samarasinghe, 2023, The Kimberlina synthetic multiphysics dataset for CO₂ monitoring investigations: *Geoscience Data Journal*.
- Bai, Y., J. Mei, A. Yuille, and C. Xie, 2021, Are transformers more robust than CNNs? (arXiv:2111.05464 [cs]).
- Becherer, N., J. Pecarina, S. Nykl, and K. Hopkinson, 2019, Improving optimization of convolutional neural networks through parameter fine-tuning: *Neural Computing and Applications*, **31**.
- Biswas, R., M. K. Sen, V. Das, and T. Mukerji, 2019, Prestack and poststack inversion using a physics-guided convolutional neural network: *Interpretation*, **7**, SE161–SE174.
- Brown, T., B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, 2020, Language models are few-shot learners: *Advances in Neural Information Processing Systems*, Curran Associates, Inc., 1877–1901.
- Caruana, 1997, Multitask learning: *Machine Learning*, **28**, 41–75.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova, 2019, Bert: Pre-training of deep bidirectional transformers for language understanding: Presented at the North American Chapter of the Association for Computational Linguistics.
- Dosovitskiy, A., L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, 2021, An image is worth 16x16 words: Transformers for image recognition at scale: Presented at the International Conference on Learning Representations.
- Feng, S., H. Wang, C. Deng, Y. Feng, Y. Liu, M. Zhu, P. Jin, Y. Chen, and Y. Lin, 2023, E-fwi: Multi-parameter benchmark datasets for elastic full waveform inversion of geophysical properties. (arXiv:2306.12386 [physics]).
- Gassmann, F., 1951, Über die elastizität poröser medien: On the elasticity of porous media, **96**, 1–23.
- Hendrycks, D., and K. Gimpel, 2016, Gaussian error linear units (GELUs): Technical report. (arXiv:1606.08415 [cs]).
- Kornblith, S., J. Shlens, and Q. V. Le, 2019, Do better ImageNet models transfer better?: 2019 IEEE/CVF Conference on computer vision and pattern recognition (CVPR), IEEE, 2656–2666.
- Kumar, A., A. Raghunathan, R. Jones, T. Ma, and P. Liang, 2022, Fine-tuning can distort pretrained features and underperform out-of-distribution: arXiv:2202.10054.
- Li, D., S. Peng, Y. Guo, Y. Lu, X. Cui, and W. Du, 2023, Reservoir multiparameter prediction method based on deep learning for CO₂ geologic storage: *Geophysics*, **88**, M1–M15.
- Li, H., J. Lin, B. Wu, J. Gao, and N. Liu, 2022, Elastic properties estimation from prestack seismic data using GGCNNs and application on tight sandstone reservoir characterization: *IEEE Trans. Geosci. Remote Sensing*, **60**, 1–21.
- Liu, N., T. He, Y. Tian, B. Wu, J. Gao, and Z. Xu, 2020, Common-azimuth seismic data fault analysis using residual UNet: *Interpretation*, **8**, SM25–SM37.
- Liu, Y., S. Feng, I. Tsvankin, D. Alumbaugh, and Y. Lin, 2023, Joint physics-based and data-driven time-lapse seismic inversion: Mitigating data scarcity: *Geophysics*, **88**, K1–K12.
- Liu, Y., and I. Tsvankin, 2021, Methodology of time-lapse elastic full-waveform inversion for VTI media: *Journal of Seismic Exploration*, **30**, 257–270.
- Mardan, A., B. Giroux, G. Fabien-Ouellet, and M. Saberi, 2022, Direct monitoring of fluid saturation using time-lapse full-waveform inversion: , 3424–3428.
- Maurício, J., I. Domingues, and J. Bernardino, 2023, Comparing vision transformers and convolutional neural networks for image classification: A literature review: *Applied Sciences*, **13**, 5521.
- Mavko, G., and T. Mukerji, 1995, Seismic pore space compressibility and Gassmann’s relation: *Geophysics*, **60**, 1743–1749.
- Pan, W., K. A. Innanen, and Y. Geng, 2018, Elastic full-waveform inversion and parametrization analysis applied to walk-away vertical seismic profile data for unconventional (heavy oil) reservoir characterization: *Geophysical Journal International*, **213**, 1934–1968.
- Queißer, M., and S. C. Singh, 2013, Full waveform inversion in the time lapse mode applied to CO₂ storage at Sleipner: *Geophysical Prospecting*, **61**, 537–555.

- Ruder, S., 2019, Neural transfer learning for natural language processing: PhD thesis, National University of Ireland, Galway.
- Sigfusson, B., S. R. Gislason, J. M. Matter, M. Stute, E. Gunnlaugsson, I. Gunnarsson, E. S. Aradottir, H. Sigurdardottir, K. Mesfin, H. A. Alfredsson, D. Wolff-Boenisch, M. T. Arnarsson, and E. H. Oelkers, 2015, Solving the carbon-dioxide buoyancy challenge: The design and field testing of a dissolved CO₂ injection system: *International Journal of Greenhouse Gas Control*, **37**, 213–219.
- Simon, J., G. Fabien-Ouellet, E. Gloaguen, and I. Khurjekar, 2023, Hierarchical transfer learning for deep learning velocity model building: *Geophysics*, **88**, R79–R93.
- Tajbakhsh, N., J. Y. Shin, S. R. Gurudu, R. T. Hurst, C. B. Kendall, M. B. Gotway, and J. Liang, 2016, Convolutional neural networks for medical image analysis: Full training or fine tuning?: *IEEE Transactions on Medical Imaging*, **35**, 1299–1312.
- Um, E. S., D. Alumbaugh, M. Commer, S. Feng, E. Gasperikova, Y. Li, Y. Lin, and S. Samarasinghe, 2024, Deep learning multi-physics network for imaging CO₂ saturation and estimating uncertainty in geological carbon storage: *Geophysical Prospecting*, **72**, 183–198.
- Um, E. S., D. Alumbaugh, Y. Lin, and S. Feng, 2022, Real-time deep-learning inversion of seismic full waveform data for CO₂ saturation and uncertainty in geological carbon storage monitoring: *Geophysical Prospecting*, 1365–2478.13197.
- Vafaeikia, P., K. Namdar, and F. Khalvati, 2020, A brief review of deep multi-task learning and auxiliary task learning. (arXiv:2007.01126 [cs, stat]).
- Vasco, D. W., M. Alfi, S. A. Hosseini, R. Zhang, T. Daley, J. B. Ajo-Franklin, and S. D. Hovorka, 2019, The seismic response to injected carbon dioxide: Comparing observations to estimates based upon fluid flow modeling: *JGR Solid Earth*, **124**, 6880–6907.
- Wang, K., T. Hu, B. Zhao, and S. Wang, 2023a, Surface-related multiple attenuation based on a self-supervised deep neural network with local wavefield characteristics: *Geophysics*, **88**, V387–V402.
- Wang, Z., J. You, W. Liu, and X. Wang, 2023b, Transformer assisted dual U-net for seismic fault detection: *Front. Earth Sci.*, **11**, 1047626.
- Weiss, K., T. M. Khoshgoftaar, and D. Wang, 2016, A survey of transfer learning: *J Big Data*, **3**, 9.
- Wu, X., L. Liang, Y. Shi, and S. Fomel, 2019, FaultSeg3D: Using synthetic data sets to train an end-to-end convolutional neural network for 3D seismic fault segmentation: *Geophysics*, **84**, IM35–IM45.
- Wu, Y., and Y. Lin, 2020, Inversionnet: An efficient and accurate data-driven full waveform inversion: *IEEE Transactions on Computational Imaging*, **6**, 419–433.
- Xie, C., M. Tan, B. Gong, A. Yuille, and Q. V. Le, 2021, Smooth adversarial training. (arXiv:2006.14536 [cs]).
- Zhang, W., and J. Gao, 2022, Deep-learning full-waveform inversion using seismic migration images: *IEEE Trans. Geosci. Remote Sensing*, **60**, 1–18.
- Zhang, Y., and Q. Yang, 2022, A survey on multi-task learning: *IEEE Trans. Knowl. Data Eng.*, **34**, 5586–5609.
- Zhao, Z., L. Alzubaidi, J. Zhang, Y. Duan, and Y. Gu, 2024, A comparison review of transfer learning and self-supervised learning: Definitions, applications, advantages and limitations: *Expert Systems with Applications*, **242**, 122807.
- Zhuang, F., Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He, 2020, A Comprehensive survey on transfer learning. (arXiv:1911.02685 [cs, stat]).