



Building a Satellite Image Classification Model with Residual Neural Network

Kayode Sherifdeen and Samon Daniel

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

July 11, 2024

Building a Satellite Image Classification Model with Residual Neural Network

Kayode Sheriffdeen, Samon Daniel

10th 07,2024

Abstract

Satellite image classification plays a crucial role in various fields such as agriculture, urban planning, and environmental monitoring. Accurate classification of satellite images helps in extracting valuable information and making informed decisions. In recent years, deep learning models, particularly Residual Neural Networks (ResNet), have shown remarkable performance in image classification tasks. This abstract presents an overview of building a satellite image classification model using a ResNet architecture.

The process begins with data preparation, including gathering a dataset of satellite images and preprocessing the data through resizing, cropping, and normalization. The dataset is then divided into training, validation, and testing sets to facilitate model development and evaluation.

The Residual Neural Network is constructed by defining its architecture, which consists of convolutional layers with residual blocks, pooling layers, dense layers, and an output layer. The model is compiled with an appropriate loss function and optimizer. The training process involves setting up parameters such as batch size, number of epochs, and learning rate. The model is trained on the training set, and the training progress is monitored using evaluation metrics and visualizations.

After training, the model is evaluated by predicting classes on the validation set and assessing performance metrics such as accuracy, precision, and recall. Fine-tuning of the model can be performed based on the evaluation results to optimize performance.

Once the model is trained, it can be deployed and tested on unseen satellite images. The trained model is saved for future use, and it can be loaded to predict classes on new images. The model's performance is assessed on the testing set to measure its effectiveness in classifying satellite images.

In conclusion, this abstract presents the key steps involved in building a satellite image classification model using a Residual Neural Network. The combination of deep learning techniques and ResNet architecture provides a powerful approach for accurate classification of satellite images, enabling better decision-making in various applications.

Introduction:

Satellite image classification is a vital task in various domains, including environmental monitoring, urban planning, and agricultural analysis. It involves categorizing satellite images into different classes based on their content, enabling the extraction of valuable insights and aiding decision-making processes. With the advancements in deep learning, specifically the emergence of Residual Neural Networks (ResNet), satellite image classification has achieved new heights of accuracy and efficiency.

In this introduction, we will provide an overview of building a satellite image classification model using a Residual Neural Network. We will discuss the significance of accurate classification and the role of ResNet in enhancing classification performance.

Accurate classification of satellite images holds immense importance in numerous applications. In environmental monitoring, it helps in detecting and tracking changes in land cover, identifying deforestation areas, and monitoring the health of ecosystems. Urban planning benefits from satellite image classification by facilitating the identification of land use patterns, transportation networks, and infrastructure planning. Furthermore, in agriculture, classifying satellite images aids in crop monitoring, disease detection, and optimizing resource allocation.

Traditional methods of satellite image classification relied on handcrafted features and shallow machine learning algorithms. However, these approaches often struggled to capture the complex patterns and hierarchical structures present in satellite images. The advent of deep learning, particularly Residual Neural Networks, revolutionized the field by effectively leveraging the power of deep neural networks to extract intricate features and improve classification accuracy.

Residual Neural Networks, introduced by Kaiming He et al. in 2015, address the problem of vanishing gradients, which can hinder the training of very deep neural networks. The key innovation of ResNet lies in the introduction of residual blocks that allow networks to learn residual mappings instead of explicitly learning the

underlying mapping. By utilizing skip connections, ResNet enables the flow of information across layers, facilitating the training of much deeper networks. This architectural design has proven to be highly effective in image classification tasks, including satellite image classification.

In the subsequent sections, we will delve into the process of building a satellite image classification model with a Residual Neural Network. We will cover data preparation, model construction, training, evaluation, and deployment stages. By following these steps, we can harness the power of deep learning and ResNet to develop an accurate and robust satellite image classification model.

Overall, the integration of Residual Neural Networks into satellite image classification empowers us to unlock the full potential of satellite imagery for various applications. With the ability to capture complex patterns and hierarchical structures, ResNet-based models offer unprecedented accuracy and reliability, contributing to informed decision-making and advancing our understanding of the Earth's dynamic systems.

Importance of accurate classification for various applications

Accurate classification of satellite images holds immense importance across various applications due to its ability to extract valuable insights and support decision-making processes. Here are some key domains where accurate classification plays a crucial role:

Environmental Monitoring: Accurate classification of satellite images aids in monitoring and assessing changes in land cover, vegetation patterns, and natural resources. It enables the identification of deforestation areas, urban expansion, water bodies, and ecological changes. This information is vital for environmental conservation, land management, and understanding the impact of human activities on ecosystems.

Urban Planning: Satellite image classification provides valuable information for urban planning and development. It helps in identifying land use patterns, transportation networks, infrastructure planning, and monitoring urban growth. Accurate classification assists in assessing urban sprawl, identifying areas for potential development, and optimizing resource allocation for sustainable urban management.

Agriculture and Crop Monitoring: Satellite image classification is essential for precision agriculture and crop monitoring. It enables the identification of crop types, growth stages, and disease detection. Accurate classification helps farmers optimize resource allocation, assess crop health, and make informed decisions regarding

irrigation, fertilization, and pest control. It contributes to improving crop yield, reducing environmental impact, and ensuring food security.

Disaster Management: During natural disasters like floods, fires, or earthquakes, accurate classification of satellite images provides critical information for emergency response and disaster management. It helps in assessing the extent of damage, identifying affected areas, and facilitating rescue operations. Timely and accurate classification supports decision-making to allocate resources effectively and plan recovery efforts.

Climate Change Analysis: Satellite image classification assists in studying the impact of climate change on the Earth's surface. It helps in monitoring glacier retreat, coastal erosion, vegetation dynamics, and changes in sea ice extent. Accurate classification enables the quantification of environmental changes over time, contributing to climate modeling, understanding climate patterns, and formulating mitigation strategies.

Natural Resource Management: Accurate classification of satellite images aids in managing natural resources such as forests, water bodies, and mineral deposits. It assists in forest inventory, biodiversity assessment, water resource management, and mineral exploration. Accurate classification provides crucial information for sustainable resource management, conservation planning, and balancing economic development with environmental preservation.

Defense and Security: In defense and security applications, satellite image classification plays a vital role in identifying and monitoring critical infrastructure, military installations, and border surveillance. Accurate classification helps in detecting changes, tracking movements, and identifying potential security threats. It supports situational awareness, intelligence gathering, and strategic decision-making.

Accurate classification of satellite images empowers decision-makers, researchers, and policymakers with valuable insights into our changing world. By leveraging advanced classification techniques, we can unlock the full potential of satellite imagery, leading to improved environmental management, sustainable development, and informed decision-making across a wide range of applications.

Data Preparation

Data preparation is crucial in building a satellite image classification model with a Residual Neural Network (ResNet). Properly preparing the dataset ensures that the model receives clean, normalized, and representative data for effective training and evaluation. Here are the key components of data preparation:

Gathering the Satellite Image Dataset:

Identify the specific satellite imagery dataset suitable for your classification task. Consider factors such as resolution, spectral bands, and coverage area.

Ensure the dataset aligns with your classification objectives and contains representative samples of each class you want to classify.

Preprocessing the Dataset:

Resize and Crop Images: Normalize the size of images in the dataset to a consistent resolution suitable for the model. Consider computational constraints and the level of detail required for classification. Crop the images if necessary to remove irrelevant or empty areas.

Normalize Pixel Values: Normalize the pixel values of the images to a common scale (e.g., between 0 and 1) to ensure consistent input for the model.

Handle Class Imbalance: If your dataset exhibits a significant class imbalance (unequal representation of classes), consider applying techniques such as oversampling, undersampling, or data augmentation to address the imbalance and improve model performance.

Splitting the Dataset:

Divide the dataset into training, validation, and testing sets. The training set is used to train the model, the validation set is used to tune hyperparameters and evaluate model performance during training, and the testing set is used to assess the final model's performance.

Ensure that the data split maintains a representative distribution of classes across the sets. Common splits include 70-15-15 or 80-10-10 for training-validation-testing, respectively.

Data preparation aims to create a clean, balanced, and representative dataset for training and evaluating the ResNet model. In addition to these steps, it is essential to preprocess the data further to match the input requirements of the ResNet architecture, such as converting images to the appropriate format (e.g., RGB or grayscale) and organizing the data into the required directory structure.

By carefully preparing and preprocessing the dataset, you can ensure that the ResNet model receives high-quality input data, leading to more accurate and reliable classification results.

Preprocessing the dataset

Preprocessing the dataset is a crucial step in building a satellite image classification model with a Residual Neural Network (ResNet). It involves transforming and preparing the data to ensure it is in a suitable format and quality for training the model. Here are the key preprocessing steps for the dataset:

Image Resizing and Cropping:

Resize the satellite images to a consistent resolution appropriate for the ResNet model. Choose a resolution that balances computational resources and the level of detail required for classification.

Crop the images if necessary to remove irrelevant or empty areas that do not contribute to the classification task. Focus on retaining the relevant features within the image while reducing unnecessary background noise.

Normalization of Pixel Values:

Normalize the pixel values of the images to a common scale to ensure consistent input for the ResNet model. This normalization step helps in stabilizing training and improving convergence.

One common approach is to scale the pixel values between 0 and 1 by dividing each pixel value by the maximum pixel value (e.g., 255 for 8-bit images). Alternatively, you can use other normalization techniques such as z-score normalization.

Handling Missing Data:

Check for and handle any missing or corrupted data in the dataset. Missing data can be problematic during training and may lead to biased or inaccurate results.

If an image or a portion of an image is missing, consider either removing the affected sample or applying image inpainting techniques to fill in the missing regions.

Data Augmentation:

Apply data augmentation techniques to increase the diversity of the training dataset and improve the model's ability to generalize. Data augmentation artificially creates new training samples by applying transformations such as rotations, translations, flips, and brightness adjustments to the original images.

Augmentation helps in mitigating overfitting and enhances the model's robustness by exposing it to a broader range of variations in the data.

Class Balancing:

Address class imbalance in the dataset if present. Class imbalance occurs when some classes have significantly fewer samples compared to others. This can lead to biased training and poor performance on underrepresented classes.

Apply techniques such as oversampling (e.g., duplicating minority samples) or undersampling (e.g., randomly removing samples from the majority class) to balance the class distribution. Alternatively, you can use more advanced methods like Synthetic Minority Over-sampling Technique (SMOTE) to generate synthetic samples.

These preprocessing steps ensure that the dataset is in a suitable format and quality for effective training and evaluation of the ResNet model. It is important to maintain consistency in preprocessing steps across the training, validation, and testing sets to ensure fair evaluation and reliable performance metrics.

Remember to document the preprocessing steps applied to the dataset, as this information is crucial for reproducibility and understanding the impact of preprocessing choices on the model's performance.

Splitting the dataset into training, validation, and testing sets

Splitting the dataset into training, validation, and testing sets is a crucial step in building a satellite image classification model with a Residual Neural Network (ResNet). This division allows for model training, hyperparameter tuning, and evaluation. Here's how you can split the dataset:

Shuffle the Dataset:

Before splitting, shuffle the dataset to ensure that the samples are randomly ordered. This helps in avoiding any potential bias that might be present in the original ordering of the dataset.

Define the Split Ratio:

Determine the ratio in which you want to divide the dataset into training, validation, and testing sets. Common split ratios include 70-15-15 or 80-10-10 for training-validation-testing, respectively. However, the split ratio can vary depending on the size of the dataset and the specific requirements of your project.

Split the Dataset:

Allocate a percentage of the dataset for training, validation, and testing based on the defined split ratio. One approach is to use the `train_test_split` function available in various machine learning libraries, which randomly splits the dataset into two parts: one for training and the other for validation and testing combined.

After splitting the dataset into training and validation/testing, further divide the validation/testing portion into separate validation and testing sets. The validation set is used for hyperparameter tuning and model evaluation during training, while the testing set remains untouched until the final evaluation of the trained model.

Maintain Class Distribution:

Ensure that the distribution of different classes in the original dataset is preserved in each split. This helps in maintaining the representative nature of the data in each set and avoids introducing bias towards certain classes.

Consider using stratified sampling techniques during the splitting process to ensure that each set contains a proportional representation of different classes.

The resulting dataset splits should be disjoint, meaning that no samples should overlap between the training, validation, and testing sets. This ensures independence in evaluating the model's performance on unseen data.

It's important to note that the validation set is used for tuning hyperparameters and assessing model performance during training, while the testing set remains completely separate and is only used for the final evaluation of the trained model's performance.

By appropriately splitting the dataset, you can train the ResNet model on the training set, tune its hyperparameters using the validation set, and finally evaluate its performance on the testing set. This process helps in estimating the model's generalization capability and provides insights into its real-world performance.

Building a Residual Neural Network

To build a Residual Neural Network (ResNet) for satellite image classification, you can use popular deep learning frameworks such as TensorFlow or PyTorch. Here's a general outline of the steps involved in building a ResNet:

Import the necessary libraries:

Import the deep learning framework of your choice (e.g., TensorFlow or PyTorch) along with other required libraries for data manipulation, model building, and evaluation.

Load and preprocess the dataset:

Load the preprocessed dataset that you have prepared, including the satellite images and their corresponding labels.

Perform any additional preprocessing steps specific to your dataset, such as data augmentation or normalization.

Define the ResNet architecture:

ResNet typically consists of multiple residual blocks, which contain convolutional layers, batch normalization, activation functions (e.g., ReLU), and skip connections. Define the building blocks of the ResNet, such as the basic residual block or the bottleneck residual block, depending on the depth and complexity required by your classification task.

Stack the residual blocks to create the overall architecture of the ResNet. The number of blocks and their configurations (e.g., number of filters, stride, etc.) can be adjusted based on the complexity of the dataset and available computational resources.

Create the model:

Instantiate the ResNet model using the deep learning framework's APIs, specifying the input shape, number of classes, and other necessary parameters.

Connect the building blocks together to form the complete ResNet architecture.

Set up the training pipeline:

Define the loss function appropriate for multi-class classification, such as cross-entropy loss.

Choose an optimizer (e.g., Adam, SGD) and set the learning rate and other hyperparameters.

Configure any additional metrics you want to track during training, such as accuracy. Set up the training loop, where you forward propagate input batches through the model, calculate the loss, compute gradients, and update the model's weights using backpropagation.

Train the model:

Split the dataset into training and validation sets.

Iterate over the training set, feeding batches of data into the model and updating the weights based on the computed gradients.

Periodically evaluate the model's performance on the validation set to monitor its progress and prevent overfitting.

Continue training until the model converges or until a predefined stopping criterion is met.

Evaluate the model:

Once training is complete, evaluate the final trained model on the testing set, which was kept separate throughout the process.

Calculate relevant evaluation metrics such as accuracy, precision, recall, and F1-score to assess the model's performance.

Fine-tuning and optimization (optional):

If necessary, you can perform additional steps to optimize the ResNet model, such as adjusting hyperparameters, employing regularization techniques (e.g., dropout), or using transfer learning by initializing the model with pre-trained weights on a related task or dataset.

Remember to customize the implementation based on the specific requirements of your satellite image classification task and the deep learning framework you are using. The exact code for building a ResNet may vary depending on the framework and its APIs.

Defining the ResNet model

To define the ResNet model, I'll provide an example implementation using TensorFlow, one of the popular deep learning frameworks. In this example, I'll create a ResNet-50 model, which consists of 50 layers, including residual blocks.

The ResNet model starts with a convolutional layer, followed by batch normalization and ReLU activation. Then, it includes several residual blocks with different numbers of filters. The number of filters determines the model's capacity to capture different patterns and features.

The model ends with a global average pooling layer to aggregate spatial information and a fully connected layer with softmax activation for multi-class classification.

To use the code, you need to adjust the `input_shape` according to the shape of your input data (e.g., satellite images) and set the `num_classes` variable to the number of classes in your classification task.

Note that this is a simplified example, and you may need to modify the code to fit your specific requirements, such as changing the number of layers, adding regularization techniques, or adjusting the optimizer and learning rate for training.

Saving the trained model

To save a trained model in TensorFlow, you can use the `save` method provided by the `tf.keras.models.Model` class. This method allows you to save the model's architecture, weights, optimizer configuration, and training configuration.

The `save` method is called on the `resnet_model` object to save the model to the specified `model_path`. The model will be saved in the Hierarchical Data Format (HDF5) file format, which can store the model's architecture, weights, and other necessary information.

After executing this code, you should find a file named `resnet_model.h5` in your current directory, representing the saved model.

You can customize the `model_path` variable to specify the desired location and name for saving the model.

Remember to save the model after training to ensure that you can load it later for inference or further training.

Using `save_weights` will save the model's weights to the specified `weights_path`. To use the saved weights later, you'll need to recreate the model with the same architecture and load the saved weights into it.

Feel free to adjust the code based on your specific requirements, such as changing the file extension or path, or using a different format for saving the model if needed.

Conclusion

In this conversation, we covered the process of compiling, training, deploying, and testing a ResNet model in TensorFlow. Here's a summary of the key steps:

Compiling the Model: We compiled the model by specifying the loss function, optimizer, and metrics using the `compile` method. This prepares the model for training.

Training the Model: We trained the model by providing the training and validation datasets to the `fit` method. Adjustments such as batch size and the number of epochs can be made according to your requirements.

Saving the Trained Model: After training, we saved the trained model using the `save` method. This saved the model's architecture, weights, optimizer configuration, and training configuration to an HDF5 file.

Loading the Saved Model: We demonstrated how to load the saved model using the `load_model` function from `tensorflow.keras.models`. This allows you to reload the model for inference or further training.

Making Predictions: With the loaded model, we made predictions on new data by using the `predict` method. Preprocessing steps such as scaling or normalization may be required depending on the model's input requirements.

Saving and loading trained models is essential for reusing the model's learned parameters, deploying the model in production environments, or collaborating with other researchers and developers.

Remember, TensorFlow offers various tools and frameworks for serving and deploying models, such as TensorFlow Serving or TensorFlow.js, which can be explored based on your specific deployment needs.

Feel free to refer back to this conversation whenever you need guidance on compiling, training, saving, and deploying your ResNet model in TensorFlow.

References

1. Kalla, D., Smith, N., Samaah, F., & Polimetla, K. (2024). Hybrid Scalable Researcher Recommendation System Using Azure Data Lake Analytics. *Journal of Data Analysis and Information Processing*, 12(01), 76–88. <https://doi.org/10.4236/jdaip.2024.121005>
2. Docas Akinyele, J. J. Best practices for educating employees about cybersecurity in FinTech.
3. Kalla, D., Smith, N., & Samaah, F. (2023). Satellite Image Processing Using Azure Databricks and Residual Neural Network. *International Journal of Advanced Trends in Computer Applications*, 9(2), 48-55.

4. Kalla, D., Smith, N., Samaah, F., & Polimetla, K. (2021). Facial Emotion and Sentiment Detection Using Convolutional Neural Network. *Indian Journal of Artificial Intelligence Research (INDJAIR)*, 1(1), 1-13.
5. Docas Akinyele, J. J. Role of leadership in promoting cybersecurity awareness in the financial sector.
6. Kalla, D., & Kuraku, S. (2023). Phishing Website URL's Detection Using NLP and Machine Learning Techniques. *Journal on Artificial Intelligence*, 5(0), 145–162. <https://doi.org/10.32604/jai.2023.043366>
7. Daniel, S., & Olaoye, G. (2024). *Emphasize the Importance of Verifying the Legitimacy of Email Senders, Links, and Attachments Before Taking Any Action* (No. 13832). EasyChair.
8. Akinyele, D., & Daniel, S. Building a culture of cybersecurity awareness in the financial sector.
9. Kuraku, D. S., Kalla, D., Smith, N., & Samaah, F. (2023). Exploring How User Behavior Shapes Cybersecurity Awareness in the Face of Phishing Attacks. *International Journal of Computer Trends and Technology*.