User Guide for Differential Privacy Model Training Code

This code demonstrates the training of machine learning models with and without differential privacy (DP) using different classifiers on a dataset. Below is a guide to using and understanding the key components and functions in the code.

Requirements

- Python 3.x
- Required Libraries:
    - pandas
    - numpy
    - scikit-learn
    - diffprivlib

You can install the required libraries using:

```
pip install pandas numpy scikit-learn diffprivlib
```

Configuration

- **DIRECTORY**: Path to the directory containing the dataset. Default is "./DiffPrivacyDataset".
- **file_path**: Path to the dataset file, which should be a CSV file. Default is './DiffPrivacyDataset/DatasetToProcess.csv'.
- **Global Variables**:
    - **epsilonGlobalValue**: The global privacy parameter (epsilon) for differential privacy. Default is 2.5.
    - **max_iterations**: Maximum number of iterations for training with differential privacy.
    - **accuracy_threshold**: Desired accuracy threshold for the models to be considered acceptable (default is 0.85).

Key Functions

1. read_in_files()

- **Description**: Reads the dataset CSV file and returns the data as a Pandas DataFrame.
- **Usage**:

```
data = read_in_files()
```

2. Train_test_split(data)

- **Description**: Splits the dataset into features (X) and target variable (y), and further splits them into training and testing sets.

- **Usage**:

  X_train_noisy, X_train, X_test, y_train, y_test = Train_test_split(data)

3. Model Training Functions

- These functions train machine learning models using both plain and differential privacy techniques.
- The models supported are:
  - **Gaussian Naive Bayes** (GaussianNB)
  - **Decision Tree Classifier** (DecisionTreeClassifier)
  - **Random Forest Classifier** (RandomForestClassifier)
  - **Neural Network** (MLPClassifier)

Each model has two variants:

- **Plain Model** (no differential privacy)
- **Differential Privacy Model** (using diffprivlib)

4. return_accuracy()

- **Description**: Returns the accuracy of a trained model on the test set.
- **Usage**:

  accuracy = return_accuracy(model, X_test, y_test)

5. GaussianNBModels(), DecisionTreeClassifierModels(), RandomForestClassifierModels(), NeuralNetworkModels()

- **Description**: These functions are designed to train and evaluate multiple models for each classifier (plain and Differential Privacy versions). They:
  - Train the models with cross-validation.
  - Save the best models (in .pkl format) based on performance.
  - Iteratively attempt differential privacy models until the desired accuracy threshold is met.

6. Saving Models

- The trained models (both plain and DP) are saved as .pkl files.
- Files are named according to the model type, e.g., best_plain_model_GaussianNB.pkl.

Notes

- **Differential Privacy**: The models using diffprivlib are configured to add noise to the training data using a specified epsilon value. This ensures that the model training respects the privacy of the individual data points.
- **Hyperparameter Tuning**: The code uses GridSearchCV to find the best hyperparameters for each model, based on accuracy.
- **Noise Addition**: The noise is generated based on the Laplace mechanism for differential privacy, adding noise to the feature values during training to preserve privacy.