

Tutorials March 8, 2023 7 min read

How to evaluate and load a PyTorch model with Giskard?

This tutorial teaches you how to upload a PyTorch model (built from scratch or pre-trained) to Giskard, and identify potential errors and biases.



Favour Kelvin



Table of contents

- 🔥 Training and testing a PyTorch model
- ⏪ Before loading your PyTorch model
- 🏃 Build a PyTorch model and analyse it with Giskard

Note: The API used by this tutorial is deprecated. To use the new API, please refer to the official documentation:

<https://docs.giskard.ai>

🔥 Training and testing a PyTorch model

Have you ever spent hours training a PyTorch model, only to realize you made a mistake? Or have you ever gotten decent results from your model, but you are not sure if it is because you built the model correctly or just because deep learning is so powerful that even a flawed architecture can produce acceptable outcomes? AI is not perfect. This trade-off is something to be aware of, especially when considering the limitations of AI and its suitability for various problems. However, AI has

🇫🇷 Upload a PyTorch pre-trained model

✅ Conclusion

Want to subscribe to our monthly newsletter?

Enter your email address

Subscribe



previously considered almost impossible.

Training a PyTorch model can be a time-consuming process, so it is frustrating when you realize you made a mistake in the code. It is also hard to know if your model is truly effective, or if it is just luck that it's producing good results. To ensure your model will work well on new data, it is important to test it and compare the accuracy and loss to the values you saw at the end of training. If there's a significant difference, it could mean the model is overfitted to the training data and may not perform well on unseen data.

Giskard is a tool designed to address some of the challenges of working with AI. It allows you to quickly test your model to make sure there are no biases and errors in your model. In this tutorial, we will guide you through how



tuned pretrained model to Giskard for analysis to find edge cases and bugs.

⏪ Before loading your PyTorch model

You'll need to have [Git](#) and [Docker](#) installed. To get started, follow these installation [instructions](#) to install the Giskard Python library.

🏃 Build a PyTorch model and analyse it with Giskard

If you have a model which was built from scratch, there is a chance that it contains performance issues and edge cases. Giskard helps to detect problems related to model performance, robustness, discriminative behaviour, etc.

Below are the steps to build, scan and upload a PyTorch model, which based




```
1 import time
2
3 import torch
4 import numpy as np
5 import pandas as pd
6 from torch import nn
7 from torchtext.datasets import AG_NEWS
8 from torch.utils.data import DataLoader
9 from sklearn.metrics import accuracy_score
10 from torchtext.data.utils import
    get_tokenizer
11 from torch.utils.data.dataset import
    random_split
12 from torchtext.vocab import
    build_vocab_from_iterator
13 from torchtext.data.functional import
    to_map_style_dataset
14
15
16 # Define constants.
17 DEVICE = torch.device("cpu")
18
19 TARGET_MAP = {0: "World", 1: "Sports", 2:
    "Business", 3: "Sci/Tech"}
20 TARGET_COLUMN_NAME = "label"
21 FEATURE_COLUMN_NAME = "text"
22
23 LOADERS_BATCH_SIZE = 64
24
25
26 # Fetch data.
27 train_data, test_data = AG_NEWS()
28
29 # Simple English tokenizer provided by
    torchtext.
30 tokenizer = get_tokenizer("basic_english")
31
32 # Build a vocabulary from all the tokens we
    can find in the train data.
33 vocab =
    build_vocab_from_iterator((tokenizer(text)
```



```
38     return vocab(tokenizer(raw_text))
39
40
41 def preprocess_label(raw_label):
42     return int(raw_label) - 1
43
44
45 def collate_fn(batch):
46     label_list, text_list, offsets = [], [],
47     [0]
48
49     for _label, _text in batch:
50
51         label_list.append(preprocess_label(_label))
52         processed_text =
53         torch.tensor(preprocess_text(_text),
54                     dtype=torch.int64)
55         text_list.append(processed_text)
56
57     offsets.append(processed_text.size(0))
58
59     label_list = torch.tensor(label_list,
60                             dtype=torch.int64)
61     offsets =
62     torch.tensor(offsets[:-1]).cumsum(dim=0)
63     text_list = torch.cat(text_list)
64
65     return label_list.to(DEVICE),
66           text_list.to(DEVICE), offsets.to(DEVICE)
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161 # Create the datasets.
162 train_dataset =
163     to_map_style_dataset(train_data)
164 test_dataset =
165     to_map_style_dataset(test_data)
166
167
168 # We further divide the training data into a
169 train and validation split.
170 train_split, valid_split =
171     random_split(train_dataset, [0.95, 0.05])
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
```



```
71 test_dataloader = DataLoader(test_dataset,
    batch_size=LOADERS_BATCH_SIZE, shuffle=True,
    collate_fn=collate_fn)
```

[view raw](#)
giskard_newspaper_classification_2_01_data_preparation.py
hosted with ❤️ by GitHub

2. Wrap a dataset with Giskard

We need to wrap a dataset with `giskard.Dataset`. This wrapper allows to perform model scanning on a given data.

```
1 from giskard import Dataset
2
3
4 # Prepare data to wrap.
5 raw_data = pd.DataFrame({TARGET_COLUMN_NAME:
    TARGET_MAP[label_id - 1], FEATURE_COLUMN_NAME:
    text} for label_id, text in test_data)
6
7 # Wrap it with Giskard.
8 wrapped_data = Dataset(raw_data, name="Test
    Dataset", target="label", column_types=
    {FEATURE_COLUMN_NAME: "text",
    TARGET_COLUMN_NAME: "category"})
```

[view raw](#)
giskard_newspaper_classification_05_dataset_wrapping.py
hosted with ❤️ by GitHub

3. Build your PyTorch model

```
1 # Define model.
2 class TextClassificationModel(nn.Module):
```



```
6         self.fc = nn.Linear(embed_dim,
7                               num_class)
8         self.init_weights()
9         def init_weights(self):
10            init_range = 0.5
11            self.embedding.weight.data.uniform_(-
12            init_range, init_range)
13            self.fc.weight.data.uniform_(-
14            init_range, init_range)
15            self.fc.bias.data.zero_()
16
17            def forward(self, text, offsets):
18                embedded = self.embedding(text,
19                offsets)
20                return
21                self.fc(embedded).softmax(axis=-1)
22
23 model =
24 TextClassificationModel(vocab_size=len(vocab)
25 , embed_dim=64, num_class=4).to(DEVICE)
26
27 # Train and evaluate model.
28 criterion = torch.nn.CrossEntropyLoss()
29 optimizer =
30 torch.optim.SGD(model.parameters(), lr=5)
31 scheduler =
32 torch.optim.lr_scheduler.StepLR(optimizer, 1,
33 gamma=0.1)
34
35 def train_epoch(dataloader):
36     model.train()
37
38     train_accuracy = total_count = 0
39     for label, text, offset in dataloader:
40         optimizer.zero_grad()
41         predicted_label = model(text, offset)
42         loss = criterion(predicted_label,
43 label)
```



```
    predicted_label = torch.argmax(
label).sum().item()
41     total_count += label.size(0)
42
43     return train_accuracy / total_count
44
45
46 def validation_epoch(dataloader):
47     model.eval()
48
49     validation_accuracy = total_count = 0
50     with torch.no_grad():
51         for label, text, offsets in
dataloader:
52             predicted_label = model(text,
offsets)
53             validation_accuracy +=
(predicted_label.argmax(1) ==
label).sum().item()
54             total_count += label.size(0)
55
56     return validation_accuracy / total_count
57
58
59 total_accuracy = None
60 for epoch in range(1, 3):
61     start_time = time.perf_counter()
62
63     train_epoch(train_dataloader)
64     accu_val =
validation_epoch(valid_dataloader)
65
66     if total_accuracy is not None and
total_accuracy > accu_val:
67         scheduler.step()
68     else:
69         total_accuracy = accu_val
70
71     print("-" * 65)
72     print(f"| end of epoch {epoch: .3f} |
time: {time.perf_counter() - start_time
:5.2f}s | valid accuracy {accu_val:8.3f} ")
```



```
18 print(f'Test accuracy {test_accuracy:.5f}')
```

[view raw](#)

[giskard_newspaper_classification_04_model_training.py](#)
hosted with ❤️ by GitHub

4. Define a prediction function

In order to encapsulate the prediction logic, we need to put it inside a "prediction_function", which will be used to infer the predictions during the analysis with Giskard.



```
1 def infer_predictions(_model:
  torch.nn.Module, _data_loader: DataLoader) ->
  np.ndarray:
2     _model.eval()
3     pred = list()
4
5     for _, text, offsets in _data_loader:
6         with torch.no_grad():
7             probs = model(text,
  offsets).cpu().detach().numpy()
8
9             pred.append(probs)
10
11     pred = np.concatenate(pred, axis=0)
12     return pred
13
14
15 def prediction_function(df: pd.DataFrame) ->
  np.ndarray:
16     # Placeholder for label.
17     if df.shape[1] == 1:
18         df.insert(0, TARGET_COLUMN_NAME,
  np.zeros(len(df)))
19
```

```
collate_fn=collate_fn)
23     predictions = infer_predictions(model,
24     dataloader)
24     return predictions
```

[view raw](#)
giskard_newspaper_classification_06_prediction_function.py
hosted with ❤️ by GitHub



5. Wrap your model with Giskard

Same as with wrapping the dataset, we need to wrap the model to perform an analysis using the Giskard framework. This is done via the `giskard.Model` wrapper.

```
1  from giskard import Model
2
3
4  # Wrap model with Giskard.
5  wrapped_model =
6  Model(model=prediction_function,
7         name="SimpleNewsClassificationModel",
8         feature_names=["text"],
9         model_type="classification",
10        classification_labels=list(TARGET_MAP.values(
11        )))
12 # Validate wrapped model.
13 wrapped_test_metric =
14 accuracy_score(wrapped_data.df[TARGET_COLUMN_
15 NAME],
```

6. Scan your model with Giskard

Now that we are all set to analyse trained model, we will use `giskard.scan` to find problematic issues.

```
1 from giskard import scan
2
3
4 results = scan(wrapped_model, wrapped_data)
5 display(results)
```

`giskard_twitter_sentiment_05_model_scanning.py` [view raw](#)
hosted with ❤️ by GitHub



By calling `display(results)` one can conveniently view detected issues directly in the jupyter notebook. For example, the given model is not robust to typos in the text.

7. Generate tests based on the scanning results

the changes he applied in the new version of a model helped to resolve the detected problems:

```
1 test_suite = results.generate_test_suite("My first test suite")
2 test_suite.run()
```

[view raw](#)

`giskard_twitter_sentiment_06_test_suites_generation.py`
hosted with ❤️ by GitHub



By calling `test_suite.run()`, the user can also visualise the result of test-suite execution.

Further analysis with Giskard UI

Additionally to the [Giskard Python library](#), the Giskard server is the app that you can install either locally or on your cloud instance. It provides a convenient UI to debug tests, compare models, collaborate with other users, etc. To learn more about its functionality check the [documentation](#).

```
1 giskard server start
```

giskard_general_server_start hosted with ❤️ by [view raw](#)
GitHub

2. Start the Giskard worker

Next, we need to setup the giskard worker. This is the machine, which executes test-suites and dataset inspection.

```
1 giskard worker start -u  
http://localhost:19000/
```

giskard_general_worker_start hosted with ❤️ by [view raw](#)
GitHub

3. Upload your model and dataset to the Giskard server

Now, we are all set to upload the necessary artefacts, including dataset, trained model and generated test suite to the Giskard UI:

```
1 from giskard import GiskardClient  
2  
3  
4 # Uploading the test suite will automatically  
save the model, dataset, tests, slicing &
```



```
7
8 client = GiskardClient(
9     url="http://localhost:19000", # URL of
    your Giskard instance
10     token=token
11 )
12
13 my_project =
    client.create_project("project_id",
        "PROJECT_NAME", "Project description.")
14
15 # Upload to the current project ☒
16 test_suite.upload(client, "project_id")
```

[giskard_twitter_sentiment_07_artifacts_upload.py](#) [view raw](#)
hosted with ❤️ by GitHub



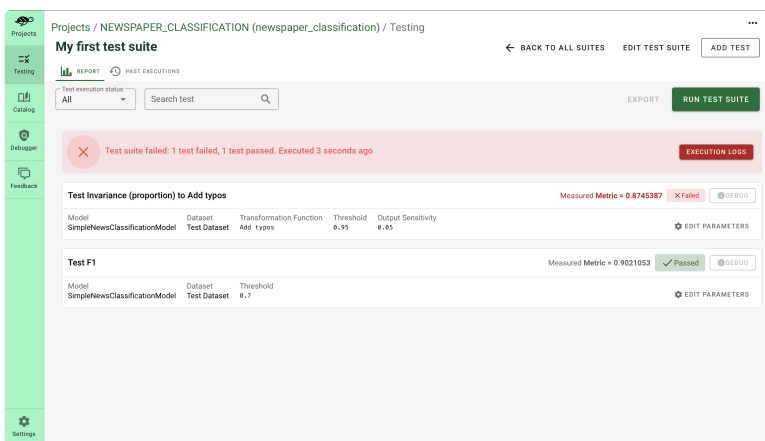
4. Debug your model and dataset

With the Debugger you can conveniently analyse each prediction. You can check, if it is correct or not, and analyse the features contribution to the prediction. You can tweak and change inputs to see how it affects output.

Debugger page. We can view the predicted probabilities of the classes, correctness of the prediction and the words that contributed the most to the output.

5. Testing

As was stated above, the Giskard UI provides a convenient way to execute test suites. We can modify the uploaded suite by adding new types of tests or create a new one. Giskard provides a wide variety of model's performance tests.



Testing page. Example of a test suite wherein one test has failed and another passed.

Fine-tuning adapts a pre-trained model to the new data without training it from scratch. After fine-tuning, you may want to check the model performance and behaviour. This can be done by uploading a fine-tuned model to the Giskard Server to perform inspection and testing routines, described in the previous section.

In this example, we are going to prepare and upload the XLM-ROBERTA model, which is fine-tuned on the SST-2 binary text classification dataset.

1. Data preparation

```
1 from enum import Enum
2
3 import torch
4 import numpy as np
5 import pandas as pd
6 import torch.nn as nn
7 from torch.optim import AdamW
8 import torchtext.transforms as T
9 import torchtext.functional as F
10 from torchtext.datasets import SST2
```



```
15
16
17 # URLs.
18 XLM_ROBERTA_VOCAB_URL =
    r"https://download.pytorch.org/models/text/xlmr.vocab.pt"
19 XLM_ROBERTA_MODEL_URL =
    r"https://download.pytorch.org/models/text/xlmr.sentencepiece.bpe.model"
20
21 # Constants.
22 BATCH_SIZE = 16
23 NUM_CLASSES = 2
24 MAX_SEQ_LEN = 256
25 MODEL_INPUT_DIM = 768
26
27 TEXT_FEATURE_NAME = "text"
28 TARGET_FEATURE_NAME = "label"
29
30 DEVICE = torch.device("cpu")
31
32
33 class Tokens(int, Enum):
34     END = 2
35     BEGIN = 0
36     PADDING = 1
37
38
39 # Load data-pipes and transform them into
    data-loaders.
40 train_datapipe = SST2(split="train")
41 valid_datapipe = SST2(split="dev")
42
43 vocab =
    load_state_dict_from_url(XLM_ROBERTA_VOCAB_URL)
44 text_transform = T.Sequential(
45     T.SentencePieceTokenizer(XLM_ROBERTA_MODEL_URL),
46     T.VocabTransform(vocab),
```




```
1 # Define model.
2 classifier_head =
  RobertaClassificationHead(num_classes=NUM_CLASSES, input_dim=MODEL_INPUT_DIM)
3 model =
  XLMR_BASE_ENCODER.get_model(head=classifier_head).to(DEVICE)
4 softmax = nn.Softmax(dim=1)
5
6 # Perform model fine-tuning.
7 criteria = nn.CrossEntropyLoss()
8 optim = AdamW(model.parameters(), lr=1e-5)
9
10
11 def train_iter(input_features, ground_truth):
12     optim.zero_grad()
13     prediction = model(input_features)
14     train_loss = criteria(prediction,
15     ground_truth)
16     train_loss.backward()
17     optim.step()
18
19     print(train_loss.item())
20
21 NUM_EPOCHS = 5
22 NUM_ITERS = 100
23
24
25 def train():
26     for epoch in range(NUM_EPOCHS):
27         model.train()
28         for idx, batch in
29         enumerate(train_dataloader):
30
31             print(f"Batch {idx}")
32
33             # Perform training step.
34             inputs =
35             F.to_tensor(batch["token_ids"],
36             padding_value=Tokens.PADDING).to(DEVICE)
37             targets =
38             torch.tensor(batch["target"]).to(DEVICE)
```




```
40         print(f"Epoch = {epoch}, loss =
    [{loss}], accuracy = {accuracy}")
41
42
43 def valid_iter(input_features, ground_truth):
44     output = model(input_features)
45     valid_loss = float(criteria(output,
    ground_truth).item())
46     return valid_loss, (output.argmax(1) ==
    ground_truth).type(torch.float).sum().item()
47
48
49 def evaluate():
50     model.eval()
51
52     total_loss = 0
53     correct_predictions = 0
54     total_predictions = 0
55     counter = 0
56
57     with torch.no_grad():
58         for valid_batch in valid_dataloader:
59             input_features =
    F.to_tensor(valid_batch["token_ids"],
    padding_value=Tokens.PADDING).to(DEVICE)
60             ground_truth =
    torch.tensor(valid_batch["target"]).to(DEVICE
    )
61             valid_loss, predictions =
    valid_iter(input_features, ground_truth)
62             total_loss += valid_loss
63             correct_predictions +=
    predictions
64             total_predictions +=
    len(ground_truth)
65             counter += 1
66
67     return total_loss / counter,
    correct_predictions / total_predictions
68
69
70 train()
```



```
1 def prediction_function(df):
2     dataloader =
3     get_loader_from_pipe(IterableWrapper(df[TEXT_
4     FEATURE_NAME]), shuffle=False)
5     predictions = list()
6     model.eval()
7     with torch.no_grad():
8         for batch in dataloader:
9             input_features =
10            F.to_tensor(batch["token_ids"],
11            padding_value=Tokens.PADDING).to(DEVICE)
12            logits = model(input_features)
13            probs =
14            softmax(logits).detach().cpu().numpy()
15            predictions.extend(probs)
16
17    return np.array(predictions)
```

giskard_sst2_04_define_prediction_function.py [view raw](#)
hosted with ❤️ by GitHub



5. Wrap model with Giskard

```
1 from giskard import Model
2
3
4 wrapped_model = Model(prediction_function,
5
6     model_type="classification",
7     name="xlm-roberta",
8     feature_names=
9     [TEXT_FEATURE_NAME],
10    classification_labels=
11    [0, 1])
12
13 # Validate wrapped model.
14 wrapped_model.predict(wrapped_data)
```

giskard_sst2_05_model_wrapping.py hosted with ❤️ [view raw](#)
by GitHub

performing the model scan and test-suite generation, we are ready to upload the fine-tuned model and the dataset to the Giskard Server.

```
1 from giskard import GiskardClient
2
3
4 # Uploading the test suite will automatically
  save the model, dataset, tests, slicing &
  transformation functions inside the Giskard
  UI server
5 # Create a Giskard client after having
  install the Giskard server (see
  documentation)
6 token = "" # Find it in Settings in the
  Giskard server
7
8 client = GiskardClient(
9     url="http://localhost:19000", # URL of
  your Giskard instance
10     token=token
11 )
12
13 my_project =
  client.create_project("project_id",
  "PROJECT_NAME", "Project description.")
14
15 # Upload to the current project ☒
16 test_suite.upload(client, "project_id")
```

[giskard_twitter_sentiment_07_artifacts_upload.py](#) [view raw](#)
hosted with ❤️ by GitHub




Conclusion

demonstrated a full pipeline from creating and scanning model to debugging and testing it on the UI. We hope that you find this article helpful. Happy testing!

You will also like






How to deploy a robust HuggingFace model for sentiment analysis into production?




Princy Pappachan



How to test the fairness of ML models? The 80% rule to measure the disparate impact



Rabah Abdul Khalek



How to test ML models? #1 🙌
Introduction



Jean-Marie John-Mathews, Ph.D.



Stay updated with the Giskard Newsletter

Enter your email address

Subscribe

Check out LLMon 🍌 - our new LLM Monitor |

Products

Knowledge

Pricing

Community

Company

★ 1,262

ensure the quality of all AI models.



© GISKARD AI SAS - Built in Europe 🇫🇷
with Trust, Privacy & Safety

[Cookies settings](#)

AI Scan

AI Test

Quickstart

Benefits

News

Tutorials

Community

Documentation

Team

Investors

Jobs

Contact us

