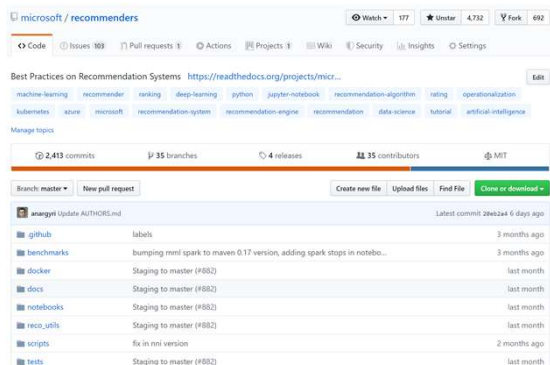


Microsoft Recommenders

Tools to Accelerate Developing Recommender Systems

Scott Graham, Jun Ki Min, Tao Wu
{Scott.Graham, Jun.Min, Tao.Wu} @ microsoft.com

Microsoft Recommenders repository is an open source collection of python utilities and Jupyter notebooks to help accelerate the process of designing, evaluating, and deploying recommender systems. The repository is actively maintained and constantly being improved and extended. It is publicly available on GitHub and licensed through a permissive MIT License to promote widespread use of the tools.



Core components



reco_utils

- Dataset helpers and splitters
- Ranking & rating metrics
- Hyperparameter tuning helpers
- Operationalization utilities



notebooks

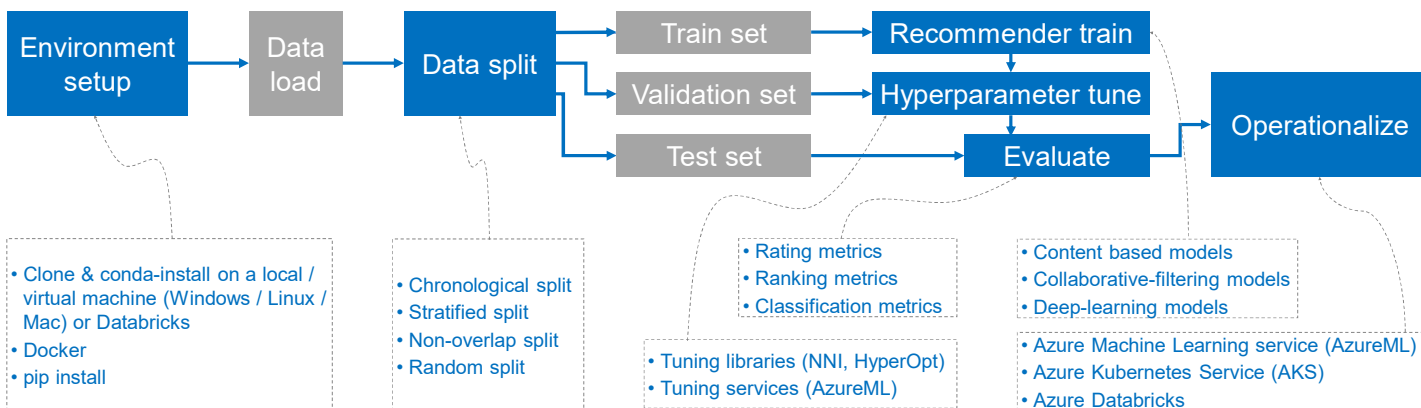
- Spark ALS
- Inhouse: SAR, Vowpal Wabbit, LightGBM, RLRMC, xDeepFM
- Deep learning: NCF, Wide-Deep
- Open source frameworks: Surprise, FastAI



tests

- Unit tests
- Smoke & Integration tests


Workflow



Example: Movie Recommendation Mobile App



- Backend: Flask web service logic / deployment files
- Data: Scripts for setting up the SQL database with MovieLens dataset
- App: Cross platform application (Xamarin.Forms)



Real-time movie recommender system with less than 100 lines of code

SAR model training script

```

import logging
from time import time
from sklearn.metrics import roc_auc_score
from sklearn.metrics import precision_score
from sklearn.metrics import recall_score
from sklearn.metrics import f1_score
from sklearn.metrics import accuracy_score
from sklearn.metrics import confusion_matrix
from sklearn.metrics import classification_report
from sklearn.metrics import log_loss
from sklearn.metrics import mean_squared_error
from sklearn.metrics import mean_absolute_error
from sklearn.metrics import mean_absolute_percentage_error
from sklearn.metrics import mean_squared_log_error
from sklearn.metrics import mean_squared_error

# get hold of the current run
run = Run.get_context()
run.set_log_level(logging.INFO)

# create the SAR model
model = SARModel(
    num_users=run.config['num_users'],
    num_items=run.config['num_items'],
    num_factors=run.config['num_factors'],
    num_neighbors=run.config['num_neighbors'],
    num_candidates=run.config['num_candidates'],
    num_threads=run.config['num_threads'],
    num_epochs=run.config['num_epochs'],
    num_validation_epochs=run.config['num_validation_epochs'],
    num_test_epochs=run.config['num_test_epochs'],
    num_test_batches=run.config['num_test_batches'],
    num_test_batch_size=run.config['num_test_batch_size'],
    num_test_batch_delay=run.config['num_test_batch_delay'],
    num_test_batch_timeout=run.config['num_test_batch_timeout'],
    num_test_batch_max_retries=run.config['num_test_batch_max_retries'],
    num_test_batch_retry_delay=run.config['num_test_batch_retry_delay'],
    num_test_batch_retry_timeout=run.config['num_test_batch_retry_timeout'],
    num_test_batch_retry_max_retries=run.config['num_test_batch_retry_max_retries'],
    num_test_batch_retry_delay_timeout=run.config['num_test_batch_retry_delay_timeout'],
    num_test_batch_retry_timeout_max_retries=run.config['num_test_batch_retry_timeout_max_retries'],
    num_test_batch_retry_timeout_retry_delay=run.config['num_test_batch_retry_timeout_retry_delay'],
    num_test_batch_retry_timeout_retry_timeout=run.config['num_test_batch_retry_timeout_retry_timeout'],
    num_test_batch_retry_timeout_retry_max_retries=run.config['num_test_batch_retry_timeout_retry_max_retries'],
    num_test_batch_retry_timeout_retry_delay_timeout_max_retries=run.config['num_test_batch_retry_timeout_retry_delay_timeout_max_retries'],
    num_test_batch_retry_timeout_retry_max_retries_timeout_max_retries=run.config['num_test_batch_retry_timeout_retry_max_retries_timeout_max_retries'],
    num_test_batch_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['num_test_batch_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# train the SAR model
model.train(run.config['train_data'], run.config['validation_data'], run.config['test_data'], run.config['num_epochs'], run.config['num_validation_epochs'], run.config['num_test_epochs'], run.config['num_test_batches'], run.config['num_test_batch_size'], run.config['num_test_batch_delay'], run.config['num_test_batch_timeout'], run.config['num_test_batch_max_retries'], run.config['num_test_batch_retry_delay'], run.config['num_test_batch_retry_timeout'], run.config['num_test_batch_retry_max_retries'], run.config['num_test_batch_retry_delay_timeout'], run.config['num_test_batch_retry_timeout_max_retries'], run.config['num_test_batch_retry_timeout_retry_delay'], run.config['num_test_batch_retry_timeout_retry_timeout'], run.config['num_test_batch_retry_timeout_retry_max_retries'], run.config['num_test_batch_retry_timeout_retry_delay_timeout_max_retries'], run.config['num_test_batch_retry_timeout_retry_max_retries_timeout_max_retries'], run.config['num_test_batch_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'])

# save the model
model.save(run.config['model_path'])

# load the model
model.load(run.config['model_path'])

# create the client with the specified user and configuration
client = Client(run.config['client_host'], run.config['client_port'], run.config['client_timeout'], run.config['client_max_retries'], run.config['client_retry_delay'], run.config['client_retry_timeout'], run.config['client_retry_max_retries'], run.config['client_retry_delay_timeout'], run.config['client_retry_timeout_max_retries'], run.config['client_retry_timeout_retry_delay'], run.config['client_retry_timeout_retry_timeout'], run.config['client_retry_timeout_retry_max_retries'], run.config['client_retry_timeout_retry_delay_timeout_max_retries'], run.config['client_retry_timeout_retry_max_retries_timeout_max_retries'], run.config['client_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'])

# create the recommender
recommender = Recommender(model, client, run.config['recommender_host'], run.config['recommender_port'], run.config['recommender_timeout'], run.config['recommender_max_retries'], run.config['recommender_retry_delay'], run.config['recommender_retry_timeout'], run.config['recommender_retry_max_retries'], run.config['recommender_retry_delay_timeout'], run.config['recommender_retry_timeout_max_retries'], run.config['recommender_retry_timeout_retry_delay'], run.config['recommender_retry_timeout_retry_timeout'], run.config['recommender_retry_timeout_retry_max_retries'], run.config['recommender_retry_timeout_retry_delay_timeout_max_retries'], run.config['recommender_retry_timeout_retry_max_retries_timeout_max_retries'], run.config['recommender_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'])

# create the deployment
app = Flask('recommender', template_folder=run.config['template_folder'])
app.config['RECOMMENDER_HOST'] = run.config['recommender_host']
app.config['RECOMMENDER_PORT'] = run.config['recommender_port']
app.config['RECOMMENDER_TIMEOUT'] = run.config['recommender_timeout']
app.config['RECOMMENDER_MAX_RETRIES'] = run.config['recommender_max_retries']
app.config['RECOMMENDER_RETRY_DELAY'] = run.config['recommender_retry_delay']
app.config['RECOMMENDER_RETRY_TIMEOUT'] = run.config['recommender_retry_timeout']
app.config['RECOMMENDER_RETRY_MAX_RETRIES'] = run.config['recommender_retry_max_retries']
app.config['RECOMMENDER_RETRY_DELAY_TIMEOUT'] = run.config['recommender_retry_delay_timeout']
app.config['RECOMMENDER_RETRY_TIMEOUT_MAX_RETRIES'] = run.config['recommender_retry_timeout_max_retries']
app.config['RECOMMENDER_RETRY_TIMEOUT_RETRY_DELAY'] = run.config['recommender_retry_timeout_retry_delay']
app.config['RECOMMENDER_RETRY_TIMEOUT_RETRY_TIMEOUT'] = run.config['recommender_retry_timeout_retry_timeout']
app.config['RECOMMENDER_RETRY_TIMEOUT_RETRY_MAX_RETRIES'] = run.config['recommender_retry_timeout_retry_max_retries']
app.config['RECOMMENDER_RETRY_TIMEOUT_RETRY_DELAY_TIMEOUT_MAX_RETRIES'] = run.config['recommender_retry_timeout_retry_delay_timeout_max_retries']
app.config['RECOMMENDER_RETRY_TIMEOUT_RETRY_MAX_RETRIES_TIMEOUT_MAX_RETRIES'] = run.config['recommender_retry_timeout_retry_max_retries_timeout_max_retries']
app.config['RECOMMENDER_RETRY_TIMEOUT_RETRY_MAX_RETRIES_TIMEOUT_MAX_RETRIES_TIMEOUT_MAX_RETRIES'] = run.config['recommender_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries']

```

Model training experiment

Model deployment to Kubernetes

```

# Register the model
model = run.register_model(name=run.config['model_name'], version=run.config['model_version'], uri=run.config['model_uri'], uri_auth_token=run.config['uri_auth_token'], uri_auth_header=run.config['uri_auth_header'], uri_auth_method=run.config['uri_auth_method'], uri_auth_params=run.config['uri_auth_params'], uri_auth_timeout=run.config['uri_auth_timeout'], uri_auth_max_retries=run.config['uri_auth_max_retries'], uri_auth_retry_delay=run.config['uri_auth_retry_delay'], uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'], uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'], uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'], uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'], uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'], uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'], uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'], uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'], uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'], uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'])

# Create the container
image_config = ContainerImageConfig(
    name=run.config['image_name'],
    uri=run.config['image_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# Create the container
container_target = ContainerTarget(
    name=run.config['container_name'],
    uri=run.config['container_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# Deploy the app
app = run.deploy_app(
    name=run.config['app_name'],
    uri=run.config['app_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# Create the deployment
deployment_target = DeploymentTarget(
    name=run.config['deployment_name'],
    uri=run.config['deployment_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# Create the deployment
deployment = Deployment(
    name=run.config['deployment_name'],
    uri=run.config['deployment_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# Create the deployment
deployment_target = DeploymentTarget(
    name=run.config['deployment_name'],
    uri=run.config['deployment_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

# Create the deployment
deployment = Deployment(
    name=run.config['deployment_name'],
    uri=run.config['deployment_uri'],
    uri_auth_token=run.config['uri_auth_token'],
    uri_auth_header=run.config['uri_auth_header'],
    uri_auth_method=run.config['uri_auth_method'],
    uri_auth_params=run.config['uri_auth_params'],
    uri_auth_timeout=run.config['uri_auth_timeout'],
    uri_auth_max_retries=run.config['uri_auth_max_retries'],
    uri_auth_retry_delay=run.config['uri_auth_retry_delay'],
    uri_auth_retry_timeout=run.config['uri_auth_retry_timeout'],
    uri_auth_retry_max_retries=run.config['uri_auth_retry_max_retries'],
    uri_auth_retry_delay_timeout=run.config['uri_auth_retry_delay_timeout'],
    uri_auth_retry_timeout_max_retries=run.config['uri_auth_retry_timeout_max_retries'],
    uri_auth_retry_timeout_retry_delay=run.config['uri_auth_retry_timeout_retry_delay'],
    uri_auth_retry_timeout_retry_timeout=run.config['uri_auth_retry_timeout_retry_timeout'],
    uri_auth_retry_timeout_retry_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries'],
    uri_auth_retry_timeout_retry_delay_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_delay_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries'],
    uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries=run.config['uri_auth_retry_timeout_retry_max_retries_timeout_max_retries_timeout_max_retries'],
)

```

Acknowledgements

Aaron He, Andreas Argyriou, Bamdev Mishra, Beth Zeranski, Dan Ciborowski, Daniel Schneider, Gianluca Campanella, Heather Spetalnick (Shapiro), Jeremy Reynolds, Jianxun Lian, Le Zhang, Markus Cosowicz, Max Kaznady, Miguel Gonzalez-Fierro, Micro Milletari, Nicolas Hug, Nikhil Joglekar, Pratik Jawanpuria, Robert Alexander, Yassine Khelifi, Zhenhui Xu, ... and 2019 Microsoft NERD Garage intern team!

Call for collaborations - The repository maintainers encourage data scientists and developers to contribute their own algorithm implementations, tools, and best practices to continuously improve the content and quality of tools available to the recommender community. <https://github.com/Microsoft/Recommenders>