

EasyRec: An Easy-to-Use, Extendable and Efficient Framework for Building Industrial Recommendation Systems

Mengli Cheng, Yue Gao, Guoqiang Liu, HongSheng Jin

Alibaba Group

{mengli.cml, dawn.gy, yancheng.lgq, hongsheng.jhs}@alibaba-inc.com

Abstract

We present EasyRec, an easy-to-use, extendable and efficient recommendation framework for building industrial recommendation systems. Our EasyRec framework is superior in the following aspects: first, EasyRec adopts a modular and pluggable design pattern to reduce the efforts to build custom models; second, EasyRec implements hyper-parameter optimization and feature selection algorithms to improve model performance automatically; third, EasyRec applies online learning to fast adapt to the ever-changing data distribution. The code is released: <https://github.com/alibaba/EasyRec>.

Introduction

Recommendation systems are widely used in customer-oriented business platforms as they improve user engagement and platform revenues. However, building practical recommendation systems requires large amounts of effort. Due to the large number of items, recommendation system usually adopts "funnel approach" (Covington, Adams, and Sargin 2016), which consists of multiple stages: candidate matching, pre-ranking, ranking, and probably re-ranking. Each of the stages requires building one or several models. What's more, building the models usually requires training, evaluation, export, and online serving, as illustrated in Fig.1(b). Meanwhile, deep models are computation intensive and require large amounts of training data, therefore, distributed training is necessary for fast convergence. The huge amounts of computation also brings challenges for deploying efficient online serving system. Besides, offline and online features often suffer from inconsistencies, as their generation process may be maintained by different people. In this demo, we present EasyRec, an easy-to-use recommendation framework to address the above challenges.

Overview

In this section, we introduce the framework of EasyRec in detail. As illustrated in Fig.1(a), EasyRec adopts a modular design, and the modules work in a pluggable mode. Thanks to the pluggable design, each module could be extended separately and flexibly without modifying any other modules. For example, to build a new Click-Through-Rate(CTR) estimation model, we only need to define the network structure

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

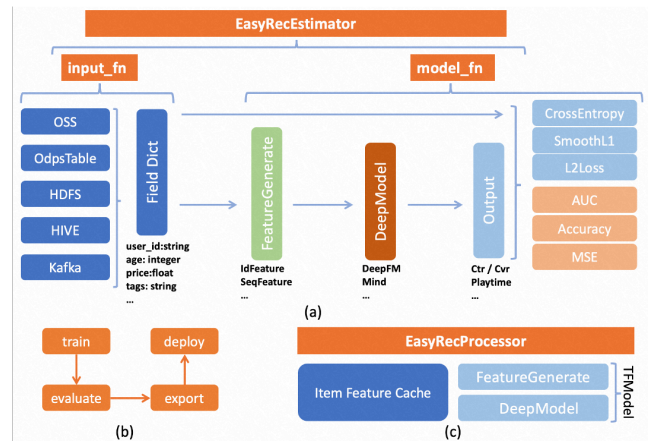


Figure 1: An overview framework of EasyRec: (a) architecture, (b) model development workflow, (c) online serving.

and related parameters. Then the new CTR model can run on all distributed platforms supported by EasyRec, such as the Machine Learning Platform for AI (PAI) of Alibaba Cloud¹. EasyRec consists of the following modules: data loading module, feature generation module, deep learning model module, loss and metric function module. All the modules are configurable with hyper-parameters defined in protobuf text format. In data module, data readers are implemented to support most big data engines, including hive, oss, kafka, etc. EasyRec also supports multiple data formats such as csv, tfrecord, parquet, etc. The core of feature generation module is implemented using C++ for high performance, which is then wrapped with tensorflow operators so that it could be easily integrated into tensorflow graphs. The feature operators are re-used in online serving to ensure consistency between online and offline features. Moreover, EasyRec integrates more than 20 state-of-the-art models for different stages of recommendation systems, such as deep structured semantic models (Huang et al. 2013) for candidate matching and deep interest network (Zhou et al. 2018) for ranking. We further adapt the tensorflow estimator (Cheng et al. 2017) to support distributed training and evaluation on multiple platforms (Kubernetes (Kubeflow 2017), Yarn (Apache 2017), etc), which is named EasyRecEstimator.

¹<https://www.alibabacloud.com/en/product/machine-learning>

Model Train

EasyRec provides easy-to-use interfaces for multiple distributed platforms. For example, we could submit a distributed training job on PAI with just a simple command:

```
pai -name easy_rec_ext -Dcmd=train
-Dconfig=oss://easyrec/test/deepfm.config
-Dtrain_tables=odps://test/tables/train
-Deval_tables=odps://test/tables/test
-Dcluster='{ "ps": { "count": 1, "cpu": 1000 },
  "worker": { "count": 3, "cpu": 800 } }'
-Dmodel_dir=oss://easyrec/models/deepfm/
```

Here, we created a distributed training job of DeepFM (Guo et al. 2017) with 1 parameter server and 3 workers. The configuration file `deepfm.config` defines all the parameters, which consists of `data.config`, `feature.config`, `model.config`, `train.config`, and `eval.config`.

EasyRec also provides graphical interfaces for the Kubernetes platform to start training jobs with just a few clicks². EasyRec supports two kinds of distribution strategies: parameter server strategy and multi-worker mirrored strategy. We further optimize the later with sparse operation kit (Nvidia 2017), which speeds up training process by more than 3 times on criteo dataset³.

Hyper-Parameter Optimization

Recommendation models are usually faced with data-sparsity problem (Li et al. 2019) due to the high-dimensional feature space of categorical features. Therefore, the hyper-parameters must be carefully tuned to avoid over-fitting over training data. In EasyRec, we address this problem by providing hyper-parameter optimization through seamless integration with NNI (Microsoft 2018). We also employ early stopping strategies (Golovin et al. 2017) to improve search efficiency. As described above, EasyRec defines all the hyper-parameters in protobuf text format and optimizes them with simple definitions of parameter names and search spaces. For example, we could optimize the regularization parameter on feature embedding with the following config:

```
"model_config.embedding_regularization":
{"_type": "uniform", "_value": [1e-6, 1e-4]}
```

Feature Selection

Typical industrial recommendation systems usually generate hundreds to thousands of features. Too many features not only cause the data-sparsity problem, but also require large amounts of computation. EasyRec provides systematic methods to evaluate the importance of features, and filter out features of lower importance. The importance of features is evaluated by variational dropout (Ma et al. 2021), which is improved by optimizing the dropout hyper-parameters for both training and test data, preventing over-fitting over training data, and leading to better performance on test data. We apply our improved method to a typical scene with more than 1000 features: about half of the features are deleted,

²easyrec.readthedocs.io/en/latest/quick_start/dlc_tutorial.html

³www.kaggle.com/competitions/criteo-display-ad-challenge

and the inference speed of online serving is improved by 2 times without degrading performance.

Online Serving

Online serving faces two main challenges: i) large numbers of Queries Per Second(QPS); ii) Request Time(RT) is usually limited to less than 0.5 seconds. EasyRec optimizes online serving performance by the following techniques: first, a LRU cache is maintained to ensure fast access to frequent items; second, we integrate feature generation into tensorflow graph by wrapping them as tensorflow operators, leading to interleaved execution among different features; third, we optimize embedding lookup with vector instructions to speed up embedding addition and multiplication operations. With the above optimizations, the RT of online serving is reduced to 1/4 of the original implementations.

Online Learning

Online learning improves performance by fast adaption to the changing distribution of features and targets, which is caused by the introduction of new items and users. However, deploying an online learning service mainly faces two challenges: first, a real-time sample streams must be created; second, deep models are generally too large, making fast updating of whole models difficult. EasyRec solves the first challenge by building a flink-based template to create real-time sample streams, which consists of three modules: events aggregation, label generation, join of label and feature stream. EasyRec overcomes the second challenge by incremental update. Specifically, we record the parameters updated during training, and periodically send changed parameters to message queues. Online service retrieves changes from the message queues, and applies them to EasyRec models. In a typical scene, the number of parameters updated is reduced to less than 5% of all parameters. The incremental updates are applied within seconds, compared to the 10+ minutes delay in updating all parameters.

Related Works

Previously, many recommendation model libraries are released, such as FuxiCtr (Zhu et al. 2021), TorchRec (Facebook 2021), and DeepCtr (Shen 2017). EasyRec differs from them in that it is not just a model library, it also provides deeply optimized and easy-to-use training and serving services on distributed platforms. It further applies hyper-parameter optimization, feature selection, and online learning to generate high performance models with efficiency.

Conclusion

In this demo, we present EasyRec, an easy-to-use, extendable, and efficient recommendation framework which runs on many distributed platforms. It provides hyper-parameter tuning and feature selection for automatic model optimization. It further implements online learning to fast adapt to the ever-changing data distribution. In the future, we will continue to develop our framework to support more state-of-the-art recommendation models.

References

- Apache. 2017. Introduction to Hadoop YARN. <https://hadoop.apache.org/docs/stable/hadoop-yarn/hadoop-yarn-site/YARN.html>. Accessed: 2022-09-18.
- Cheng, H.-T.; Haque, Z.; Hong, L.; Ispir, M.; Mewald, C.; Polosukhin, I.; Rountos, G.; Sculley, D.; Smith, J.; Soergel, D.; Tang, Y.; Tucker, P.; Wicke, M.; Xia, C.; and Xie, J. 2017. TensorFlow Estimators: Managing Simplicity vs. Flexibility in High-Level Machine Learning Frameworks. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Covington, P.; Adams, J. K.; and Sargin, E. 2016. Deep Neural Networks for YouTube Recommendations. *Proceedings of the 10th ACM Conference on Recommender Systems*.
- Facebook. 2021. Pytorch domain library for recommendation systems. <https://github.com/pytorch/torchrec>. Accessed: 2022-09-20.
- Golovin, D.; Solnik, B.; Moitra, S.; Kochanski, G.; Karro, J. E.; and Sculley, D. 2017. Google Vizier: A Service for Black-Box Optimization. *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Guo, H.; Tang, R.; Ye, Y.; Li, Z.; and He, X. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *IJCAI*.
- Huang, P.-S.; He, X.; Gao, J.; Deng, L.; Acero, A.; and Heck, L. 2013. Learning deep structured semantic models for web search using clickthrough data. *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*.
- Kubeflow. 2017. Kubeflow: The Machine Learning Toolkit for Kubernetes. <https://www.kubeflow.org/>. Accessed: 2022-09-18.
- Li, C.; Liu, Z.; Wu, M.; Xu, Y.; Huang, P.; Zhao, H.; Kang, G.; Chen, Q.; Li, W.; and Lee. 2019. Multi-Interest Network with Dynamic Routing for Recommendation at Tmall. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*.
- Ma, X.; Wang, P.; Zhao, H.; Liu, S.; Zhao, C.; Lin, W.; Chih Lee, K.; Xu, J.; and Zheng, B. 2021. Towards a Better Trade-off between Effectiveness and Efficiency in Pre-Ranking: A Learnable Feature Selection based Approach. *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Microsoft. 2018. An open source AutoML toolkit for automate machine learning lifecycle, including feature engineering, neural architecture search, model compression and hyper-parameter tuning. <https://github.com/microsoft/nni>. Accessed: 2022-09-20.
- Nvidia. 2017. Sparse Operation Kit. https://github.com/NVIDIA-Merlin/HugeCTR/tree/master/sparse_operation_kit. Accessed: 2022-09-18.
- Shen, W. 2017. DeepCTR: Easy-to-use, Modular and Extensible package of deep-learning based CTR models. <https://github.com/shenweichen/deepctr>. Accessed: 2022-09-20.
- Zhou, G.; Song, C.-N.; Zhu, X.; Ma, X.; Yan, Y.; Dai, X.-W.; Zhu, H.; Jin, J.; Li, H.; and Gai, K. 2018. Deep Interest Network for Click-Through Rate Prediction. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- Zhu, J.; Liu, J.; Yang, S.; Zhang, Q.; and He, X. 2021. Open Benchmarking for Click-Through Rate Prediction. *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*.