

—— “2007-2017” ECUG 十年高峰论坛 ——

AI算法实现和云平台应用

2017.12.17

About Me

- ✿ HBase/Hadoop contributor
- ✿ OpenStack contributor
- ✿ Seagull(Docker Web UI) author
- ✿ TensorFlow/Kubernetes developer
- ✿ System architect at 4Paradigm



Agenda

- ✿ 人工智能与机器学习介绍
- ✿ 机器学习算法原理与实现
- ✿ 云机器学习平台架构实践

Agenda

- ❖ 人工智能与机器学习介绍
- ❖ 机器学习算法原理与实现
- ❖ 云机器学习平台架构实践

机器学习介绍

- ❖ "A computer program is said to learn from **experience E** with respect to some class of **tasks T** and **performance measure P** if its performance at tasks in T, as measured by P, improves with experience E." - Mitchell
- ❖ 场景 + 指标 + 样本 -> 机器学习模型

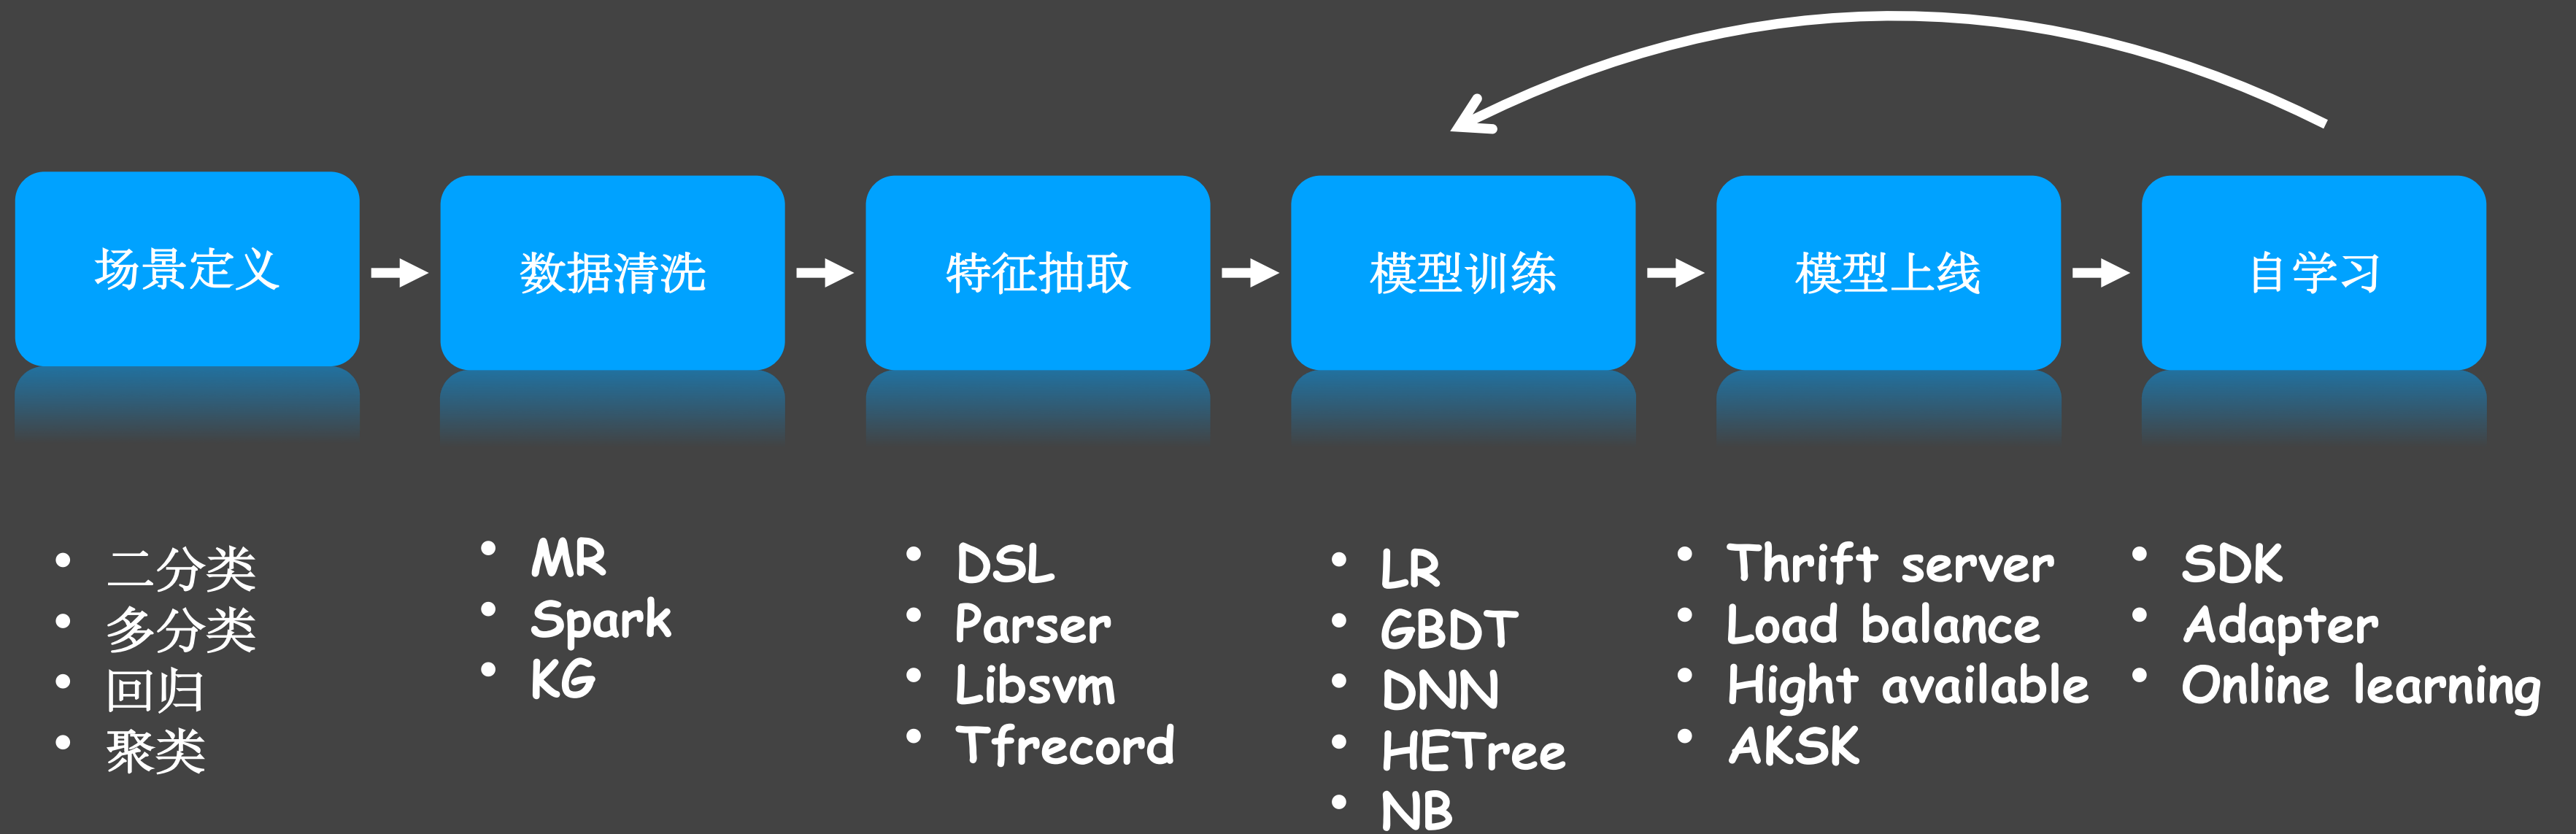
机器学习应用

- ❖ 1000张猫的图片，得到一个识别猫模型
- ❖ 1000万盘围棋棋局，得到一个AlphaGo模型
- ❖ 历史的信用卡操作记录，得到一个反欺诈模型
- ❖

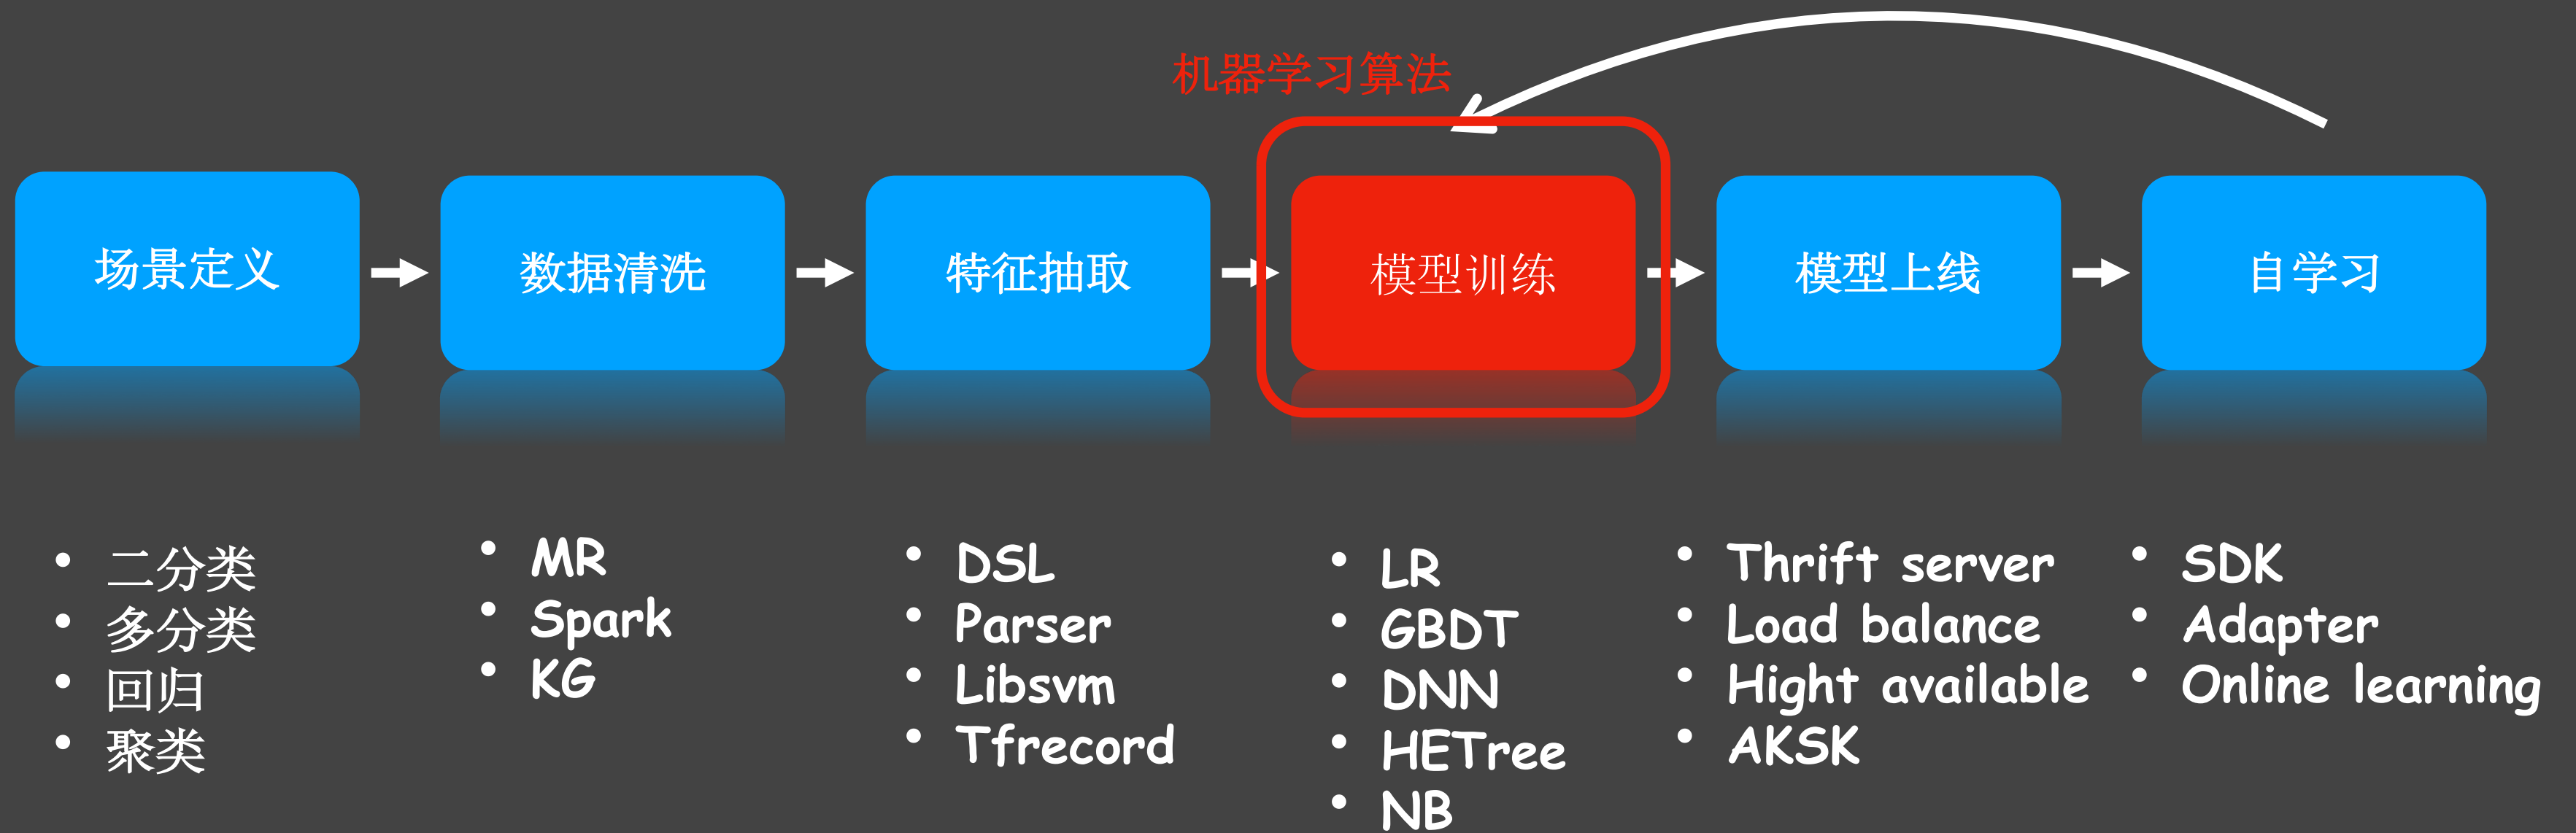
机器学习应用

- ❖ 1000张猫的图片，**得到**一个识别猫模型
- ❖ 1000万盘围棋棋局，**得到**一个AlphaGo模型
- ❖ 所有的信用卡操作记录，**得到**一个反欺诈模型
- ❖

得到机器学习应用

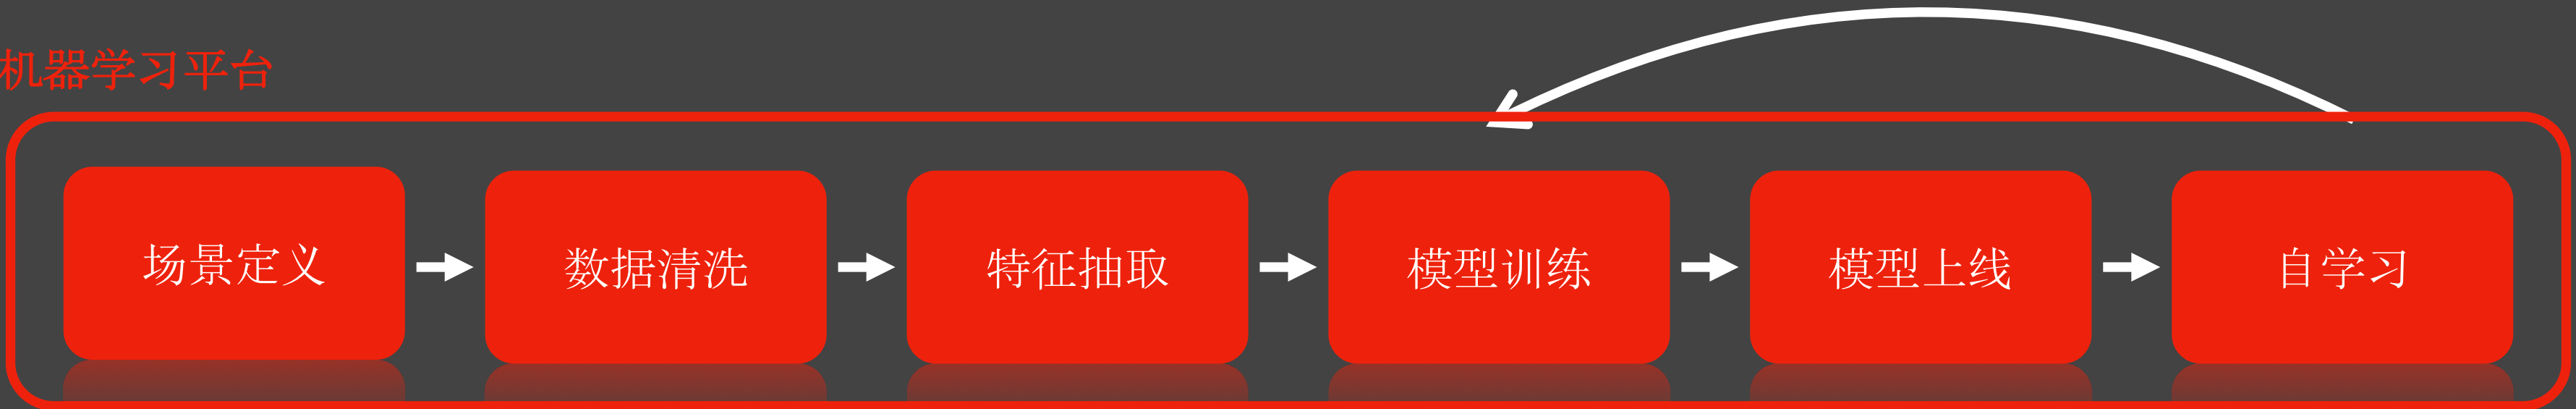


得到机器学习应用



得到机器学习应用

机器学习平台



- 二分类
- 多分类
- 回归
- 聚类

- MR
- Spark
- KG

- DSL
- Parser
- Libsvm
- Tfrecored

- LR
- GBDT
- DNN
- HETree
- NB

- Thrift server
- Load balance
- Hight available
- AKSK

- SDK
- Adapter
- Online learning

Agenda

- ✿ 人工智能与机器学习介绍
- ✿ 机器学习算法原理与实现
- ✿ 云机器学习平台架构实践

机器学习算法 - 逻辑回归

逻辑回归: 简单而强大的机器学习算法, 广泛用于推荐系统、CTR等场景

- ✿ 监督型模型
- ✿ 线性模型
- ✿ 高性能
- ✿ 可解析性强
- ✿ 可拓展性强

机器学习算法 - 逻辑回归

| | 性别:男 | 性别:女 | 年龄:0-30 | 年龄:30-60 | 年龄:60-90 | 年龄:80+ | 收入>50 |
|-------------|------------|------------|------------|------------|------------|-------------|-------|
| 样本1 | 1 | 0 | 0 | 1 | 0 | 0 | 1 |
| 样本2 | 0 | 1 | 0 | 0 | 0 | 1 | 0 |
| ... | ... | ... | ... | ... | ... | ... | ... |
| LR模型 | 0.7 | 0.4 | 0.6 | 0.9 | 0.3 | 0.01 | |

机器学习算法 - 逻辑回归

性别:男 性别:女 年龄:0-30 年龄:30-60 年龄:60-90 年龄:80+ 收入>50

LR模型 0.7 0.4 0.6 0.9 0.3 0.01

新样本 1 0 1 0 0 0

计算得分 0.7 0 0.6 0 0 0

$$\text{activation}(0.7 + 0.6) = 1$$

机器学习算法 - 逻辑回归

```

import tensorflow as tf
import numpy as np
from tensorflow.examples.tutorials.mnist import input_data

# Train data
mnist = input_data.read_data_sets("./mnist/", one_hot=True)

# Hyper parameters
learning_rate = 0.1
batch_size = 10
train_epoch_number = 100
display_interval = 10

# The variables to compute
x = tf.placeholder(tf.float32, [None, 784])
y = tf.placeholder(tf.float32, [None, 10])
W = tf.Variable(tf.zeros([784, 10]), name="weight")
b = tf.Variable(tf.zeros([10]), name="bias")

pred = tf.nn.softmax(tf.matmul(x, W) + b)
loss = tf.reduce_mean(-tf.reduce_sum(y * tf.log(pred), reduction_indices=1))
train_op = tf.train.GradientDescentOptimizer(learning_rate).minimize(loss)

init_op = tf.initialize_all_variables()

run_op = tf.train.SessionRunner(1, 1, tf.Graph().get_operations())

```

机器学习算法 - 逻辑回归

逻辑回归: 简单而强大的机器学习算法, 广泛用于推荐系统、CTR等场景


- ✿ 使用特征工程生成非线性Feature
- ✿ 模型为一维数组 (长度与Feature维度相同)
- ✿ 定义Loss函数作为指标 (如CrossEntropy)
- ✿ 使用梯度下降算法训练 (如Adagrad优化器)
- ✿ 通过Parameter server实现分布式训练

机器学习算法 - AlphaZero

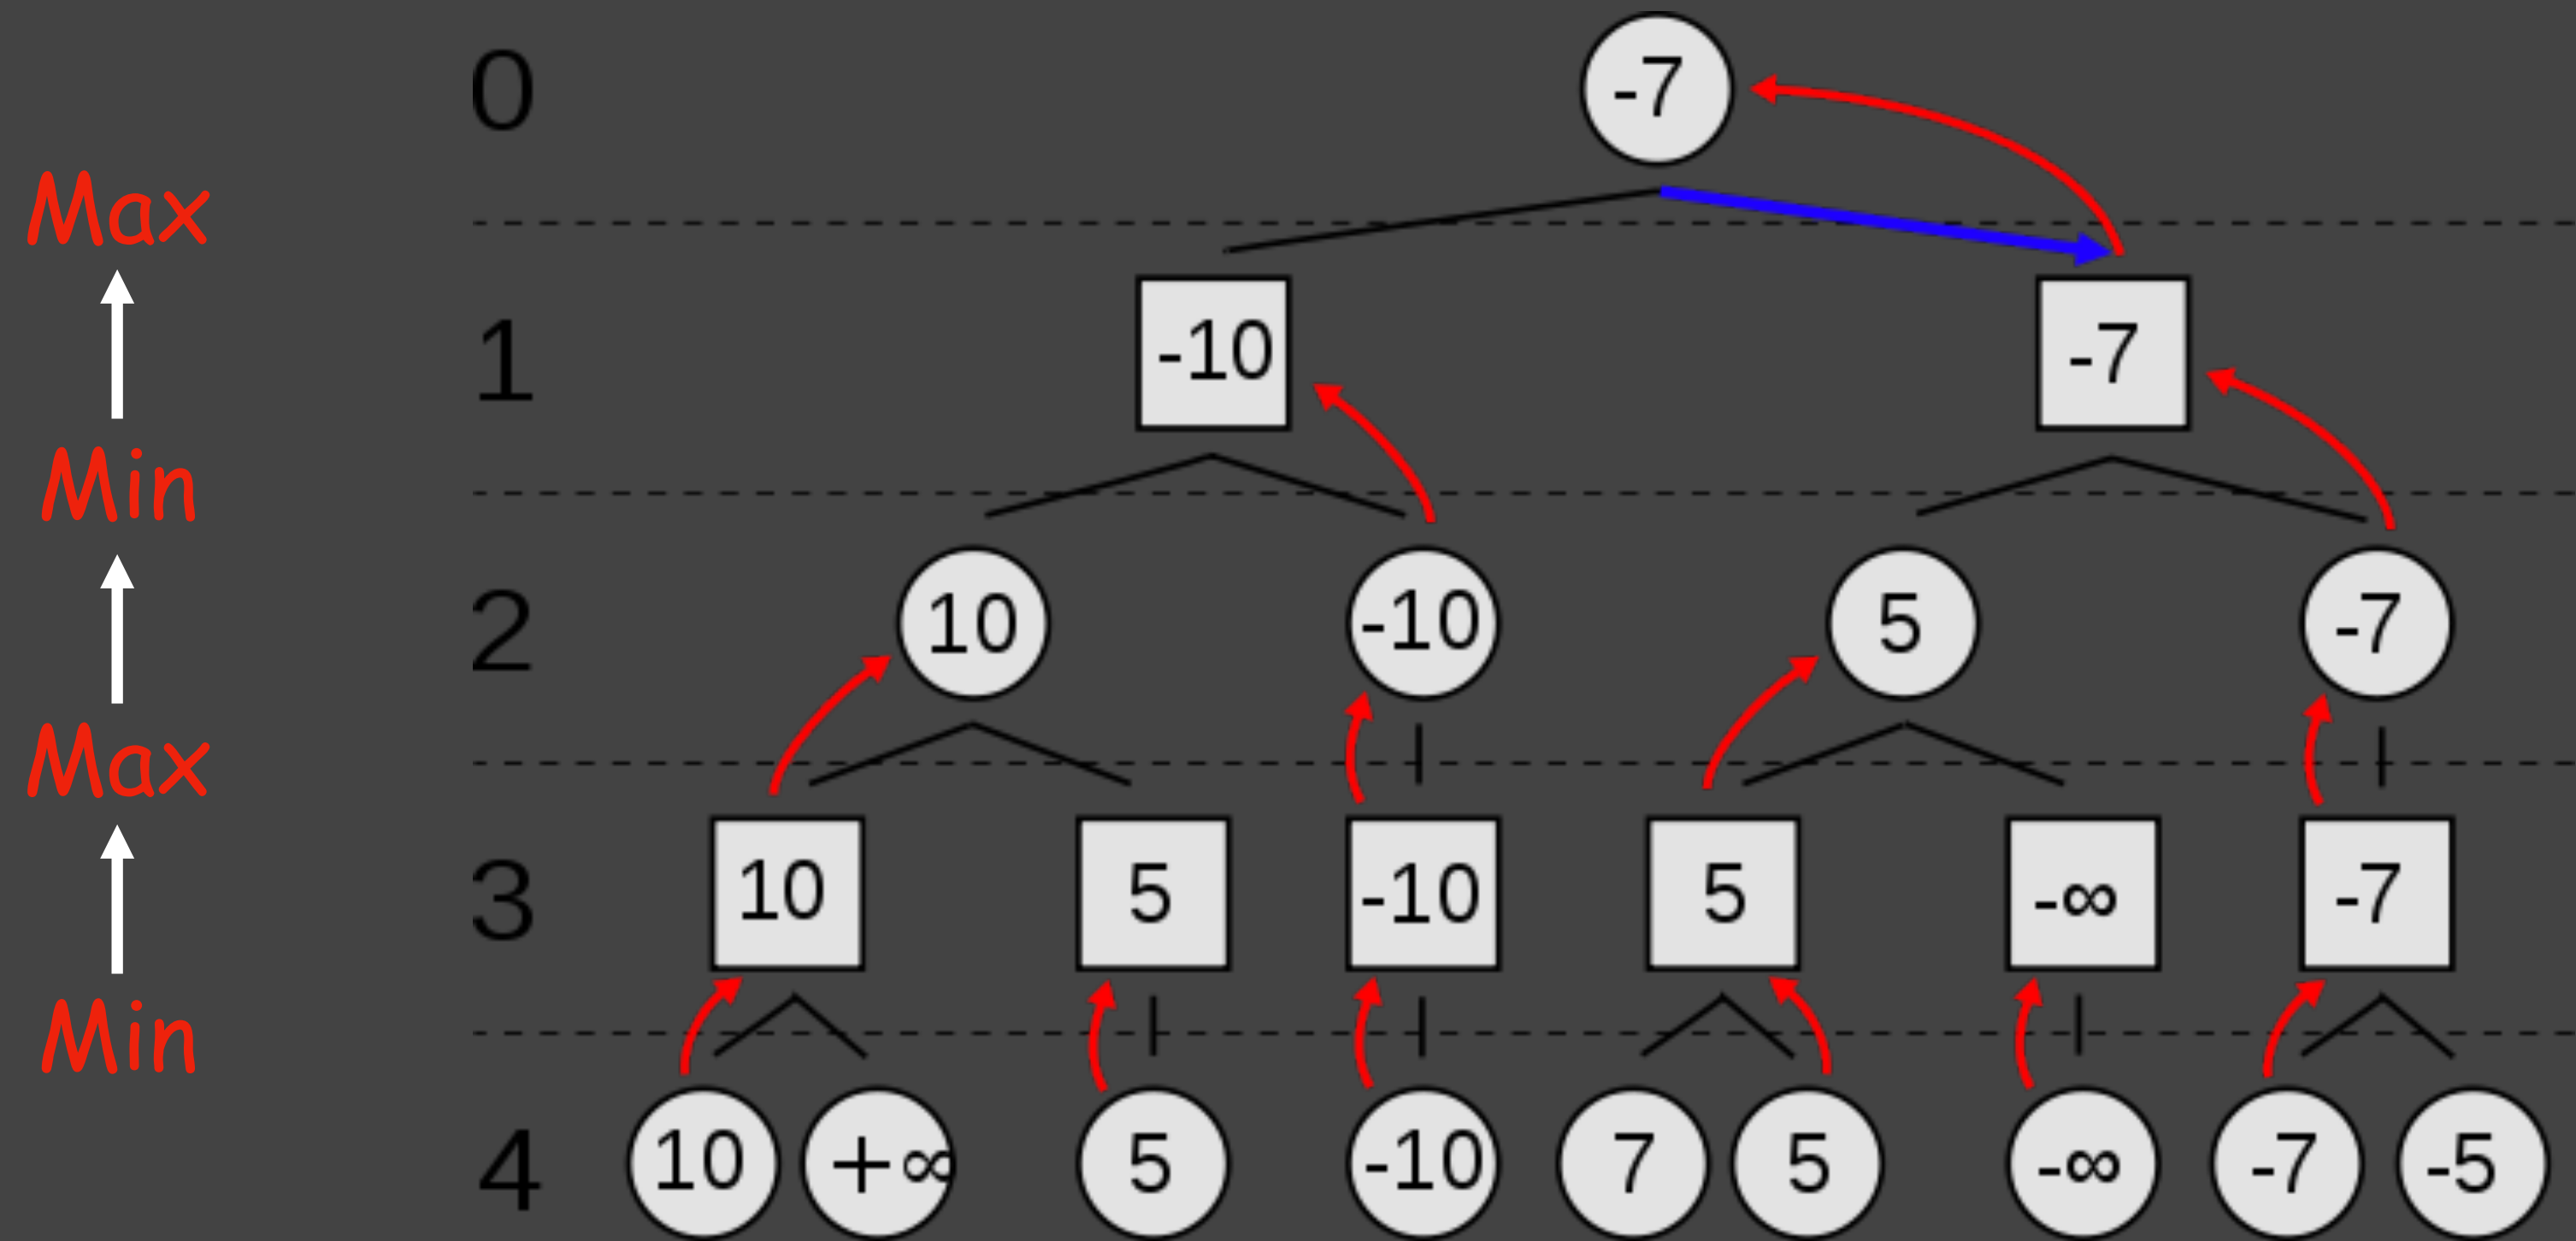
AlphaZero: 基于蒙特卡罗树搜索和神经网络的机器学习模型

- ✿ 基于蒙特卡罗树结构定义围棋规则（非自学习理解）
- ✿ 有监督学习，学习蒙特卡罗树搜索结果
- ✿ 使用Policy gradient（增强学习）生成更优的样本数据
- ✿ 模型为ResNet（输入Feature输出Probability）
- ✿ 适用于Combination game（零和、完全信息、无随机）

机器学习算法 - MinMax




机器学习算法 - MinMax



机器学习算法 - MCTS


✿ Exploration

✿ Exploitation

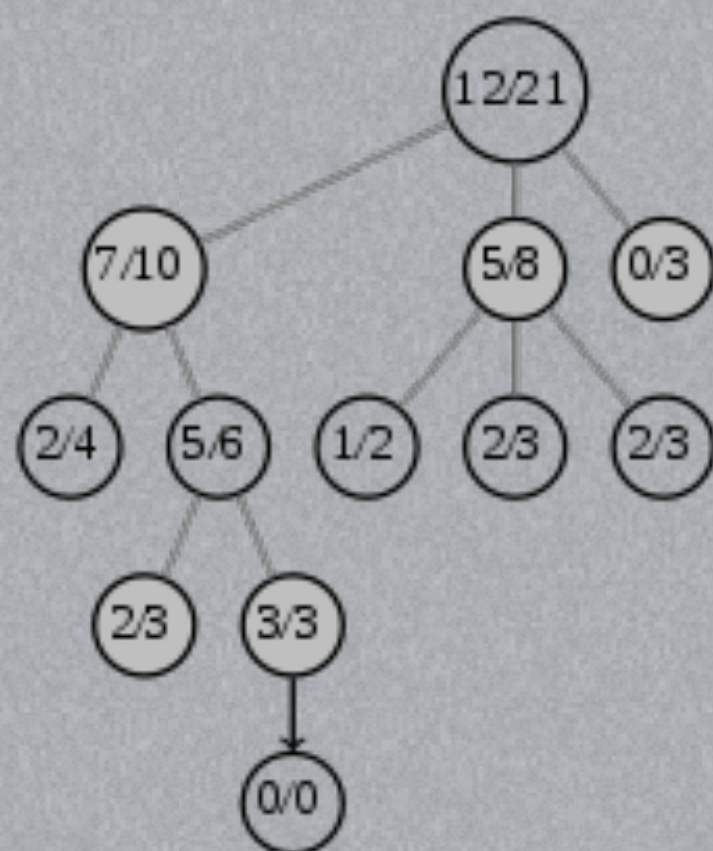


机器学习算法 - MCTS

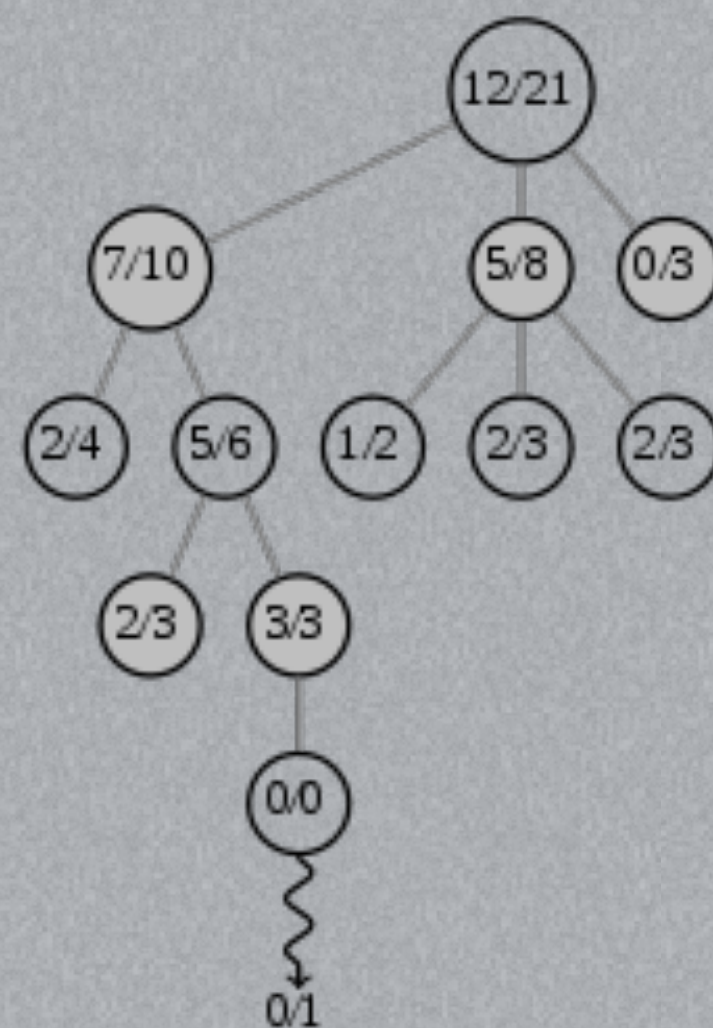
Selection




Expansion



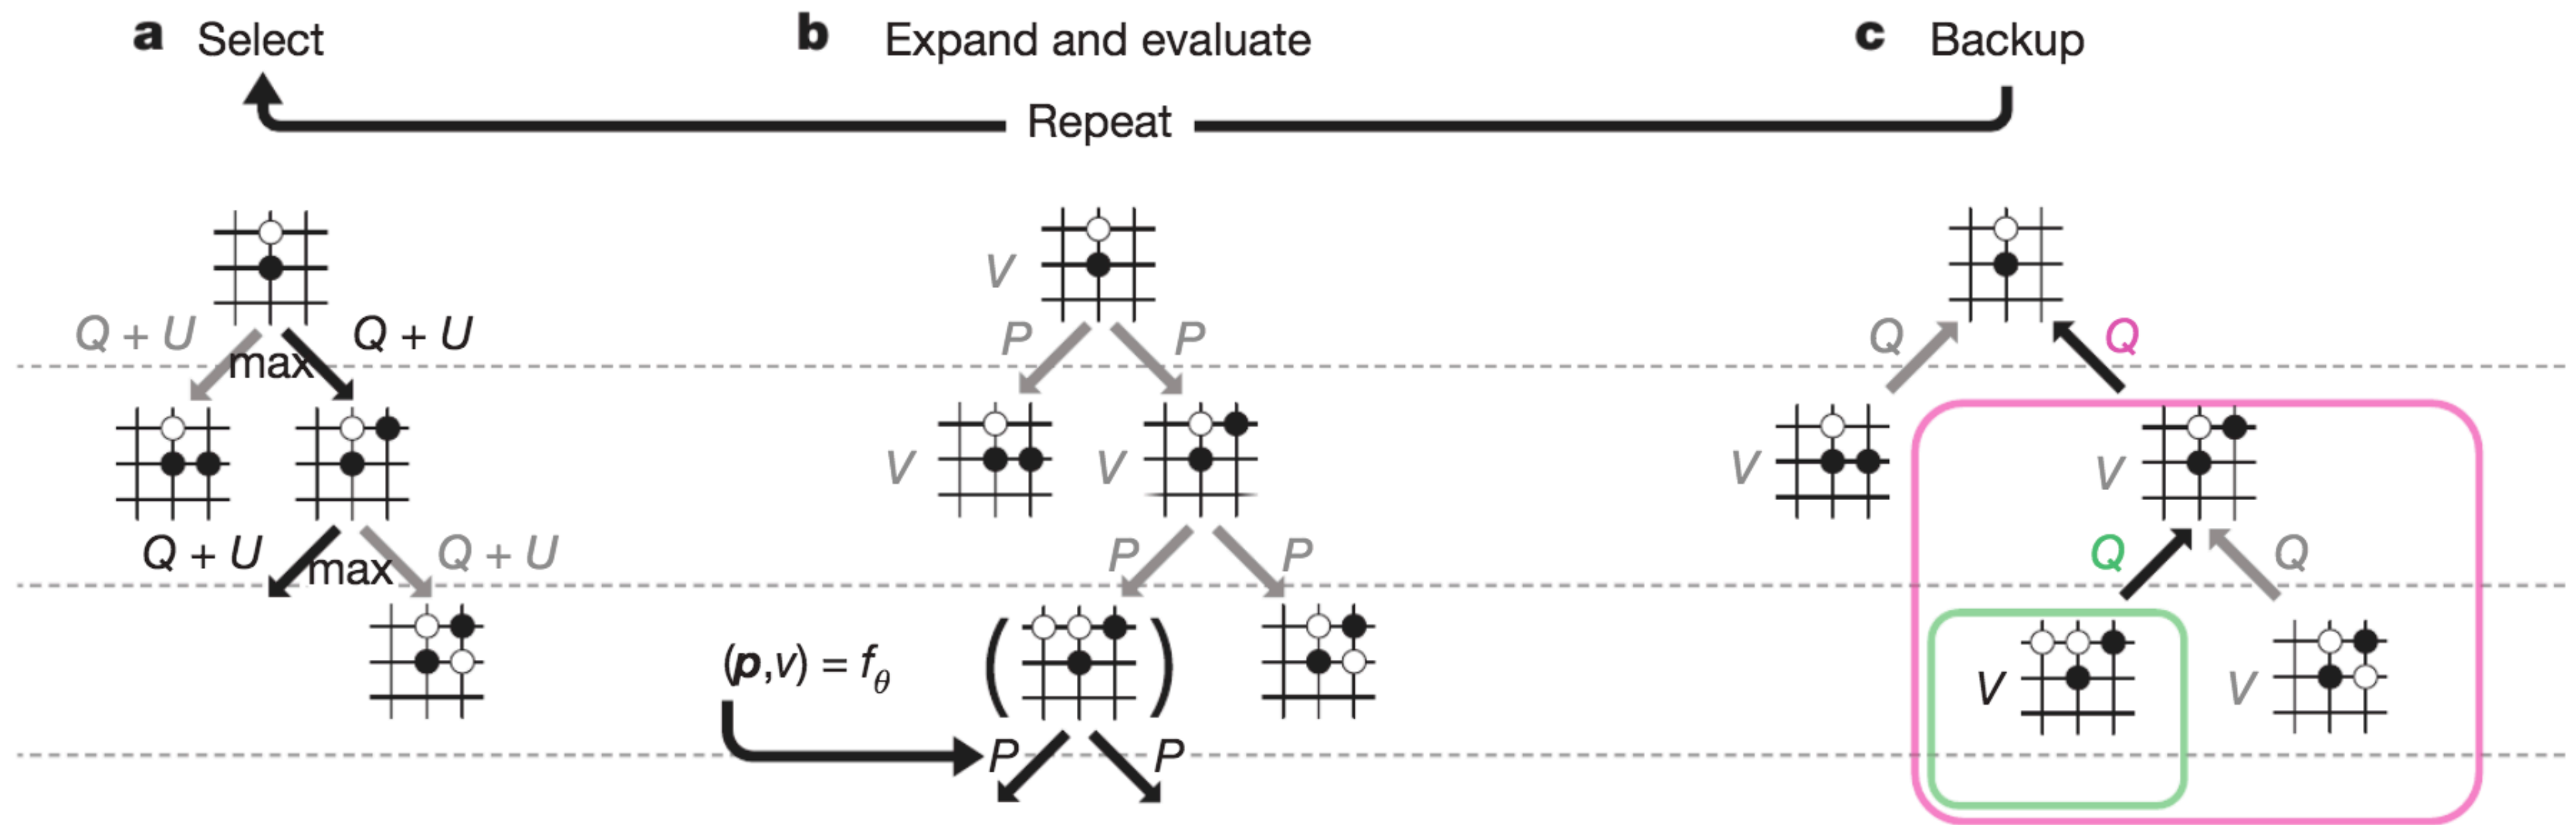
Simulation



Backpropagation



机器学习算法 - AlphaZero



$$(p, v) = f_\theta(s) \text{ and } l = (z - v)^2 - \pi^T \log p + c \|\theta\|^2$$

机器学习算法 - AlphaZero

- ❖ "UCB value": $Q + U$
- ❖ Q : $\text{Sum_of_}V(\text{children}) / N$
- ❖ V : quality value, from neural network
- ❖ N : visit times
- ❖ U : $C * P * \text{"vanila UCB"}$
- ❖ C : Constant parameter to control exploration/exploitation
- ❖ P : probability to play, from neural network

机器学习算法 - AlphaZero

- ❖ "UCB value": $Q + U$
- ❖ Q : $\text{Sum_of_}V(\text{children}) / N$
- ❖ V : quality value, from neural network
- ❖ N : visit times
- ❖ U : $C * P * \text{"vanila UCB"}$
- ❖ C : Constant parameter to control exploration/exploitation
- ❖ P : probability to play, from neural network

机器学习算法 - 自动求导

- ✿ Op求导 (Add/Minus/Multiple/Divide/Square...)
- ✿ 链式法则
- ✿ 求偏导
- ✿ 梯度下降 (Learning rate)
- ✿ Python operator override

机器学习算法 - 自动求导

on GitHub
/tobegit3hub/miniflow

```
class AddOp(Op):
    """
    The addition operation which has only two inputs. The input can be
    primitive, ConstantOp, PlaceholderOp, VariableOp or other ops.
    """

    def __init__(self, input1, input2, name="Add"):
        super(AddOp, self).__init__(name)

        if not isinstance(input1, Op):
            self._op1 = ConstantOp(input1)
        else:
            self._op1 = input1

        if not isinstance(input2, Op):
            self._op2 = ConstantOp(input2)
        else:
            self._op2 = input2

        self._graph = graph.get_default_graph()
        self._graph.add_to_graph(self)

    def forward(self):
        result = self._op1.forward() + self._op2.forward()
        return result

    def grad(self, partial_derivative_opname=None):
        result = self._op1.grad(partial_derivative_opname) + self._op2.grad(
            partial_derivative_opname)
        return result
```

```
class PowerOp(Op):
    def __init__(self, input, power, name="Power"):
        super(PowerOp, self).__init__(name)

        if not isinstance(input, Op):
            self._op = ConstantOp(input)
        else:
            self._op = input

        self._power = power

        self._graph = graph.get_default_graph()
        self._graph.add_to_graph(self)

    def forward(self):
        result = pow(self._op.forward(), self._power)
        return result

    def grad(self, partial_derivative_opname=None):
        if isinstance(self._op, PlaceholderOp) or isinstance(self._op, ConstantOp):
            # op is the constant
            grad = 0
        elif isinstance(self._op, VariableOp):
            # op is the variable
            grad = self._power * pow(self._op.forward(), self._power - 1)
        else:
            # op is other complex operation and use chain rule
            grad = self._power * pow(self._op.forward(), self._power - 1
                ) * self._op.grad(partial_derivative_opname)

        return grad
```

LEFNLU L2Nff
b9Lft9f-qeLTL9fTL6-obu9w6)
L2Nff = 26fL-obJ-d19q(b9Lft9f-qeLTL9fTL6-obu9w6) + 26fL-obZ-d19q(
d6L d19q(26fL-obJ-d19q(b9Lft9f-qeLTL9fTL6-obu9w6)=4016):

LEFNLU d19q
) * 26fL-ob-d19q(b9Lft9f-qeLTL9fTL6-obu9w6)
d19q = 26fL-home1 * hom(26fL-ob-d19q(1), 26fL-home1 - 1

机器学习算法 - 自动求导

on GitHub
/robegit3hub/miniflow

```

basic_operation.py — tensorflow_examples ×
1 #!/usr/bin/env python
2
3 # Copyright 2017 The Authors. All Rights Reserved.
4 #
5 # Licensed under the Apache License, Version 2.0 (the "License")
6 # you may not use this file except in compliance with the License
7 # You may obtain a copy of the License at
8 #
9 #     http://www.apache.org/licenses/LICENSE-2.0
10 #
11 # Unless required by applicable law or agreed to in writing, software
12 # distributed under the License is distributed on an "AS IS" BASIS
13 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either expressed or
14 # implied. See the License for the specific language governing
15 # permissions and limitations under the License.
16
17 import tensorflow as tf
18
19
20 def main():
21     sess = tf.Session()
22
23     # Add
24     a = tf.constant(32.0)
25     b = tf.constant(10.0)
26     c = a + b
27     print(sess.run(c)) # Should be 42.0
28
29     # Minus
30     c = a - b
31     print(sess.run(c)) # Should be 22.0
32
33     # Multiple
34     c = a * b
35     print(sess.run(c)) # Should be 320.0
36
37     # Divide
38     c = a / b
39     print(sess.run(c)) # Should be 3.2
40
41
42 if __name__ == "__main__":
43     main()
44
basic_operation.py — miniflow_examples ×
1 #!/usr/bin/env python
2
3 # Copyright 2017 The Authors. All Rights Reserved.
4 #
5 # Licensed under the Apache License, Version 2.0 (the "License")
6 # you may not use this file except in compliance with the License
7 # You may obtain a copy of the License at
8 #
9 #     http://www.apache.org/licenses/LICENSE-2.0
10 #
11 # Unless required by applicable law or agreed to in writing, software
12 # distributed under the License is distributed on an "AS IS" BASIS
13 # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either expressed or
14 # implied. See the License for the specific language governing
15 # permissions and limitations under the License.
16
17 import miniflow as tf
18
19
20 def main():
21     sess = tf.Session()
22
23     # Add
24     a = tf.constant(32.0)
25     b = tf.constant(10.0)
26     c = a + b
27     print(sess.run(c)) # Should be 42.0
28
29     # Minus
30     c = a - b
31     print(sess.run(c)) # Should be 22.0
32
33     # Multiple
34     c = a * b
35     print(sess.run(c)) # Should be 320.0
36
37     # Divide
38     c = a / b
39     print(sess.run(c)) # Should be 3.2
40
41
42 if __name__ == "__main__":
43     main()
44

```


机器学习算法 - 自动求导

on GitHub
/robegit3hub/miniflow

- ✿ TensorFlow vs MiniFlow(22x speedup)
- ✿ 12.59 > 0.16 | 12.61 > 0.16 | 66.58 > 3.01

```

→ add_operation git:(master) x ./tensorflow_add.py
Benchmark scenario: add operation, epoch: 100000
Run time(s): 12.5948340893
→ add_operation git:(master) x cd ../multiple_operation
miniflow_multiple.py tensorflow_multiple.py
→ multiple_operation git:(master) x ./tensorflow_multiple.py
Benchmark scenario: multiple operation, epoch: 100000
Run time(s): 12.6196570396
→ multiple_operation git:(master) x cd ../linear_regression
miniflow_linear_regression.py tensorflow_linear_regression.py
→ linear_regression git:(master) x ./tensorflow_linear_regression.py
Benchmark scenario: linear regression, epoch: 100000
Run time(s): 66.5865750313

```

```

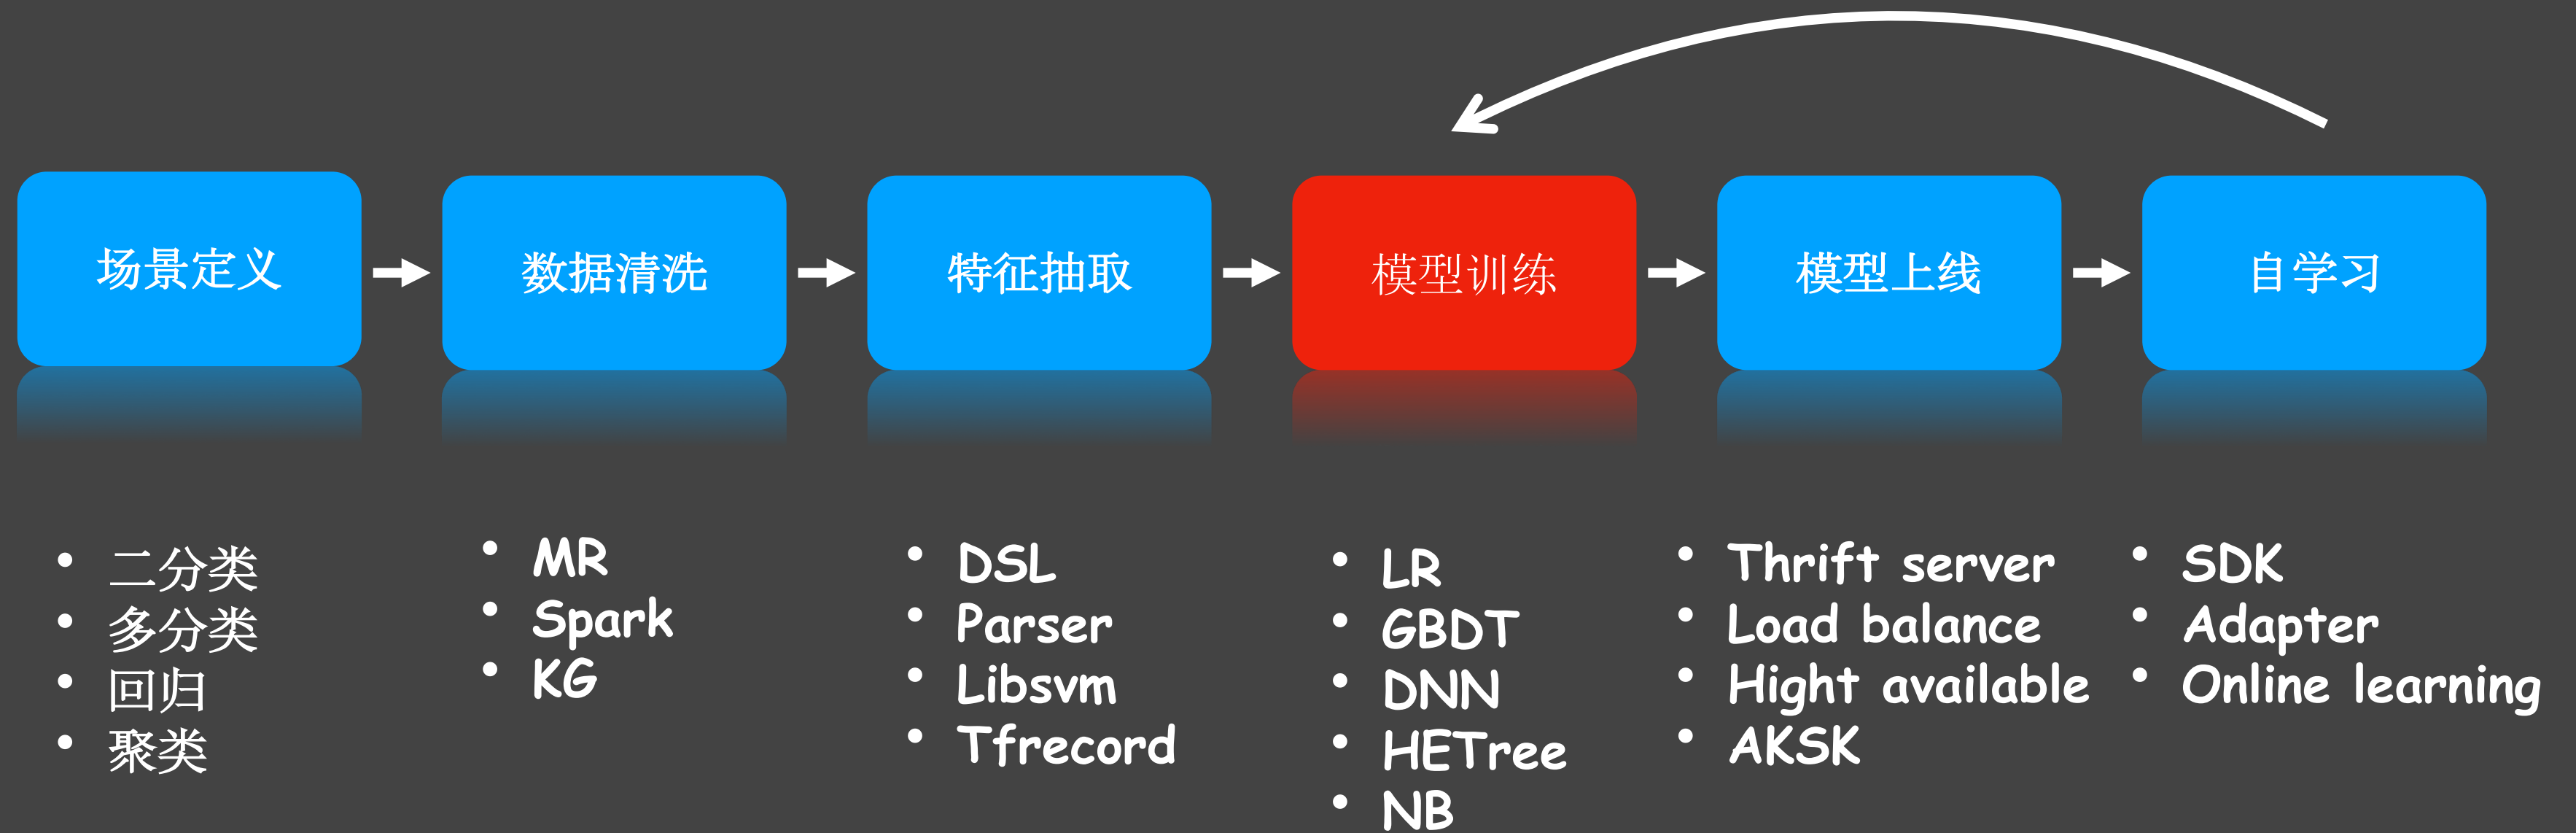
|→ add_operation git:(master) x ./miniflow_add.py
|Benchmark scenario: add operation, epoch: 100000
|Run time(s): 0.163687944412
|→ add_operation git:(master) x cd ../multiple_operation
|miniflow_multiple.py tensorflow_multiple.py
|→ multiple_operation git:(master) x ./miniflow_multiple.py
|Benchmark scenario: multiple operation, epoch: 100000
|Run time(s): 0.165066957474
|→ multiple_operation git:(master) x cd ../linear_regression
|miniflow_linear_regression.py tensorflow_linear_regression.py
|→ linear_regression git:(master) x ./miniflow_linear_regression.py
|Benchmark scenario: linear regression, epoch: 100000
|Run time(s): 3.01005196571

```

Agenda

- ✿ 人工智能与机器学习介绍
- ✿ 机器学习算法原理与实现
- ✿ 云机器学习平台架构实践

云机器学习平台 - 应用流程



云机器学习平台 - 架构设计

- ✿ TensorFlow LR支持10亿维稀疏模型，10万亿维呢？
- ✿ 开源框架与自研框架（机器学习算法）的集成
- ✿ 异构计算集群（CPU、GPU、虚拟机、云平台）的支持
- ✿ 机器学习 workflow 的支持

云机器学习平台 - 架构设计

IaaS

- ✿ AWS VMs
- ✿ AliCloud VMs
- ✿ AliCloud VMs

PaaS

- ✿ Google CloudML
- ✿ Xiaomi CloudML
- ✿ AliCloud PAI

"MLaaS"

- ✿ Azure ML Studio
- ✿ 4Paradigm prophet

图像处理

自然语言处理

语音处理

推荐系统

先知机器学习平台

API / SDK / WEB / Command-line

模型训练

模型服务

数据管理

特征抽取

Distributed
TensorFlow

Distributed MXNet

模型评估

自学习服务

Kubernetes Executor

Yarn / Hadoop Executor

CPU集群

GPU集群

FPGA集群

公有云

私有云

其他服务器

云机器学习平台 - "Google CloudML"

```
FROM ubuntu:16.04

RUN apt-get update && apt-get install -y --no-install-recommends \
    build-essential \
    curl \
    libfreetype6-dev \
    libpng12-dev \
    libzmq3-dev \
    pkg-config \
    python \
    python-dev \
    rsync \
    software-properties-common \
    unzip \
    && \
    apt-get clean && \
    rm -rf /var/lib/apt/lists/*

RUN pip --no-cache-dir install \
    http://storage.googleapis.com/tensorflow/linux/cpu/tensorflow-1.12.0-cp27-none-linux_x86_64.whl

COPY jupyter_notebook_config.py /root/.jupyter/
COPY notebooks /notebooks
COPY run_jupyter.sh /

EXPOSE 6006
EXPOSE 8888

WORKDIR "/notebooks"

CMD ["/run_jupyter.sh", "--allow-root"]
```

Step 1: Build docker image

```
@csrf_exempt
def train(request):
    if request.method == 'POST':

        # Get request data
        body = json.loads(request.body)
        job_name = body.get('job_name')
        job_module_name = body.get('module_name')
        job_trainer_uri = body.get('trainer_uri')
        job_args = body.get('job_args')
        job_train_data_paths = body.get('trainer_data_paths')
        job_eval_data_paths = body.get('job_eval_data_paths')
        job_output_path = body.get('output_path')
        job_master_spec = body.get('master_spec')

        # Generate Kubernetes resources
        time_suffix = datetime.datetime.now().strftime("%m%d%H%M")
        kubernetes_resource_name = job_name + "-" + time_suffix
        train_container_run_commands = ["/run_user_module.py",
                                         job_trainer_uri, job_module_name]

        client = KubernetesClient().get_client()
```

Step 2: Implement API service

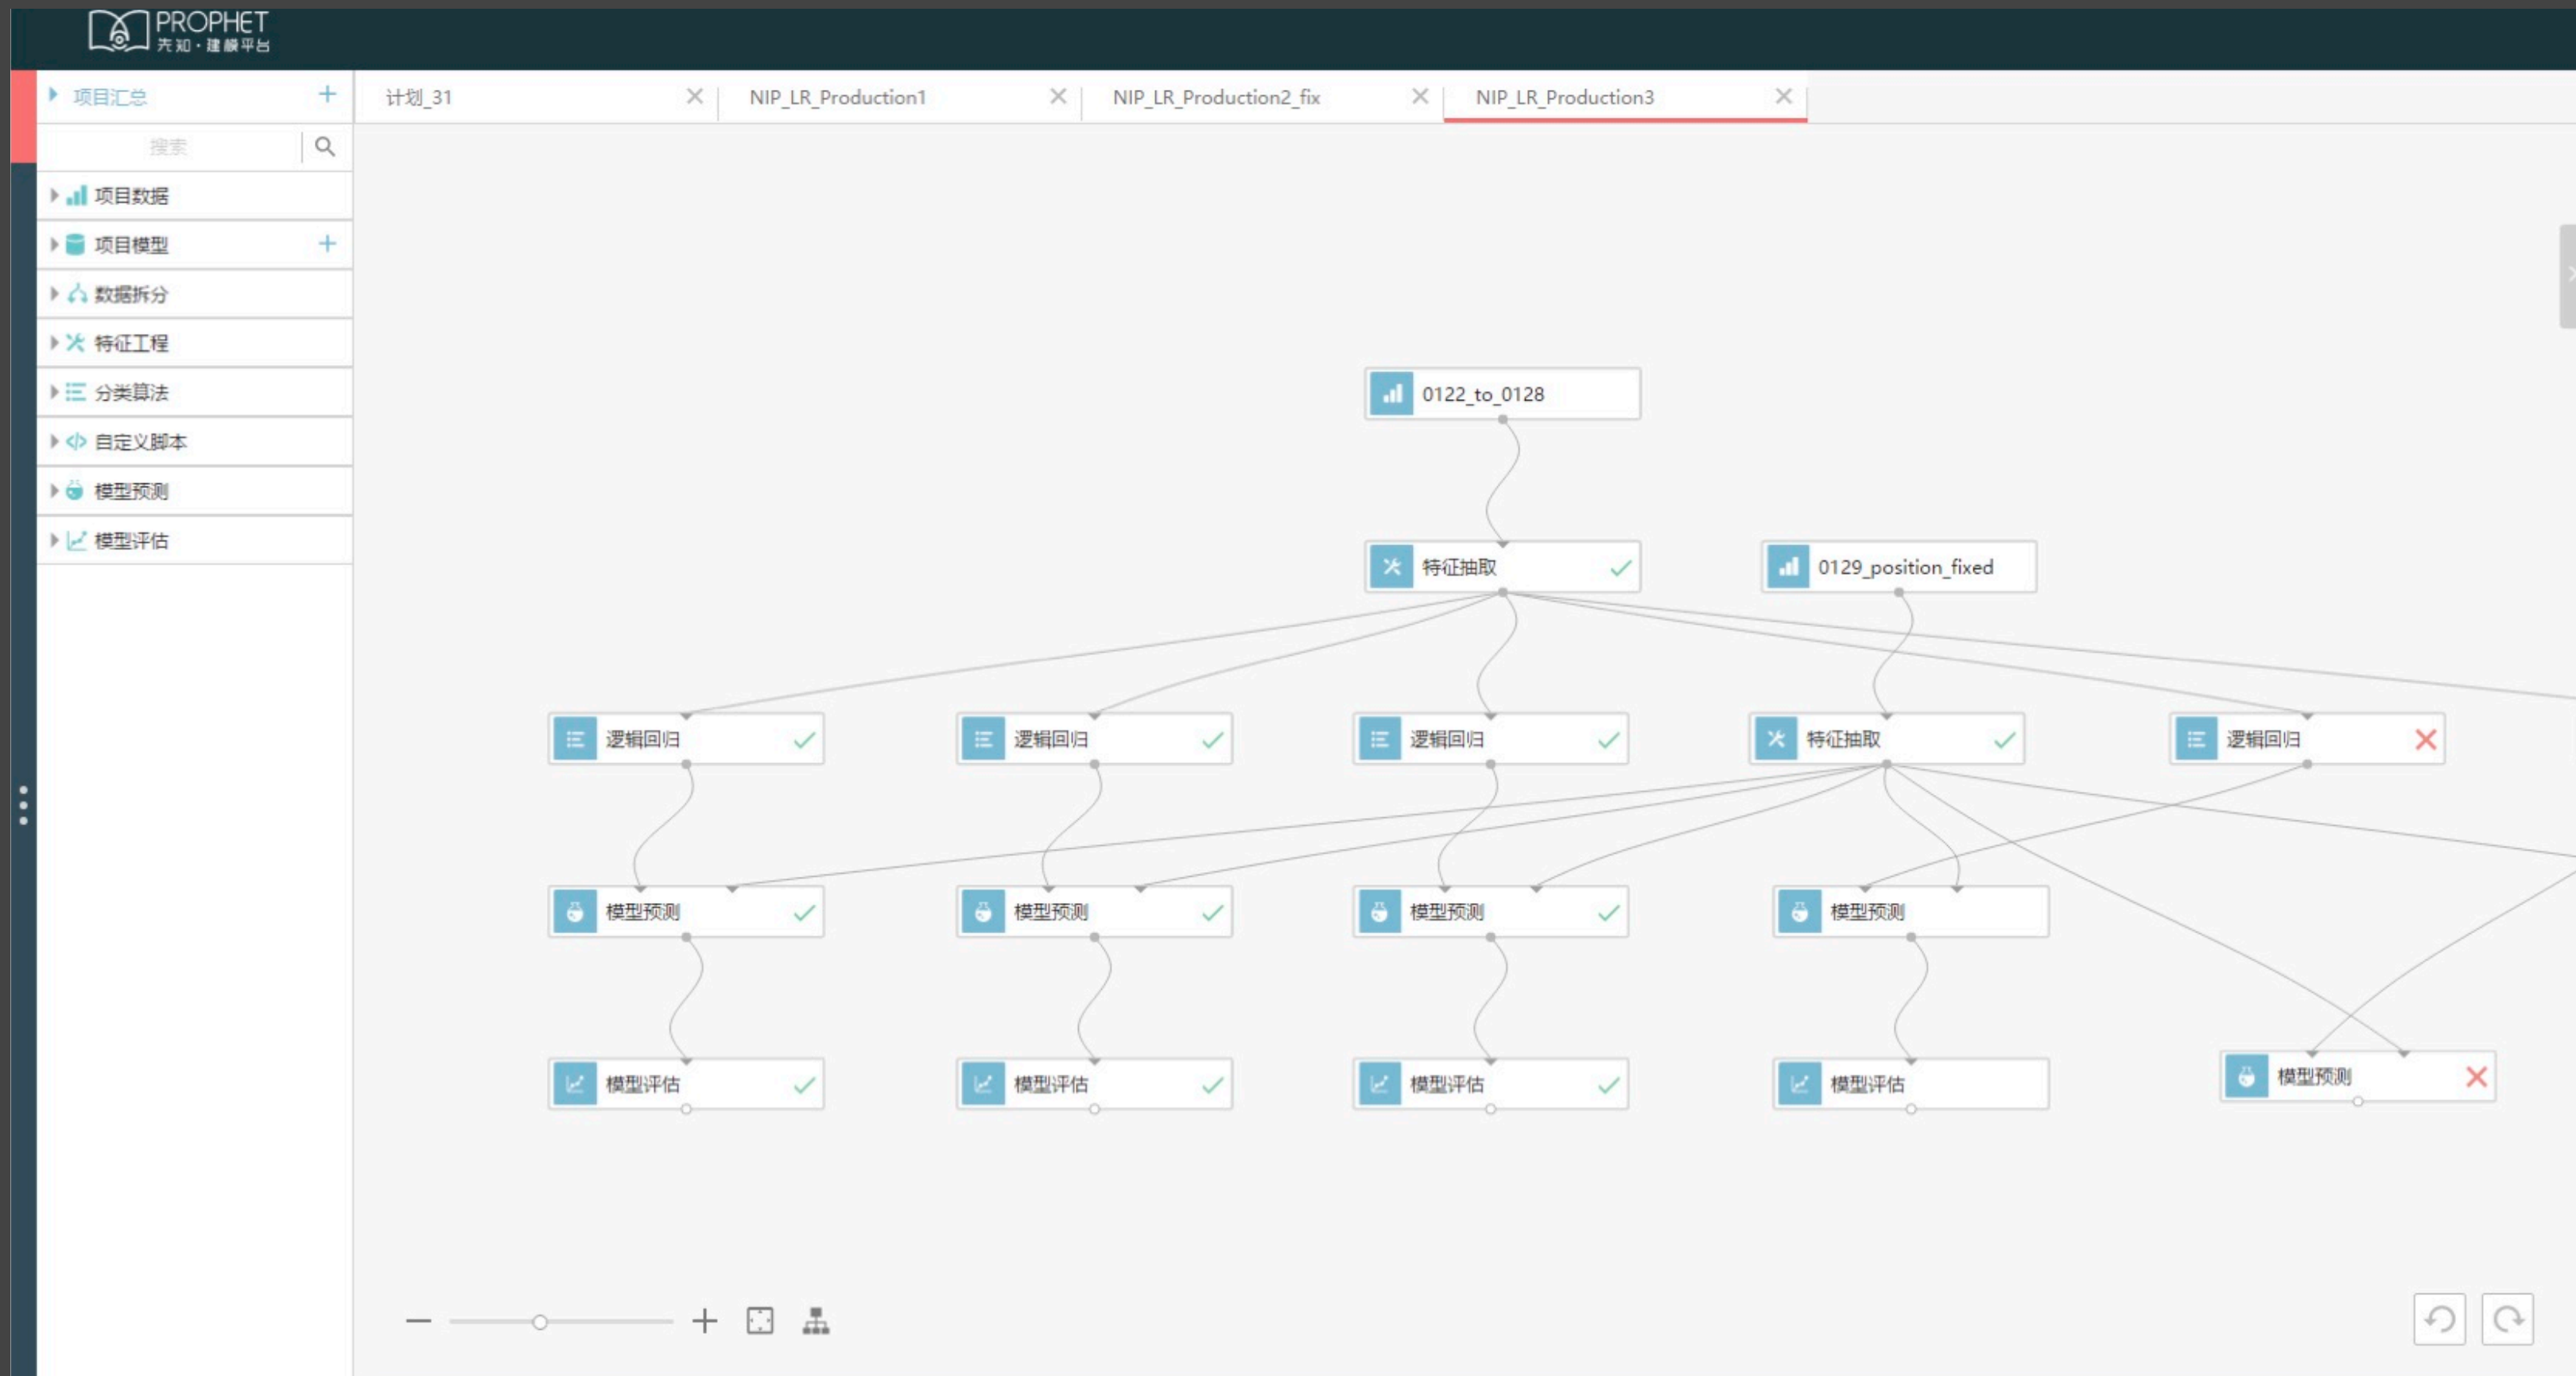
```
deployment_object = {
    "apiVersion": "extensions/v1beta1",
    "kind": "Deployment",
    "spec": {
        "replicas": 1,
        "template": {
            "spec": {
                "containers": [
                    {
                        "name": kubernetes_resource_name,
                        "image": docker_image,
                        "command": container_run_commands,
                        "resources": {
                            "requests": {
                                "cpu": limit_cpu,
                                "memory": limit_memory,
                            }
                        }
                    }
                ]
            }
        }
    }
}
```

Step 3: Submit to Kubernetes

云机器学习平台 - workflow

- ✿ 定义数据引入 -> 特征抽取 -> 模型训练 -> 模型评估 workflow
- ✿ 拖拽式 (Drag-and-drop) 操作方式
- ✿ 自定义功能算子

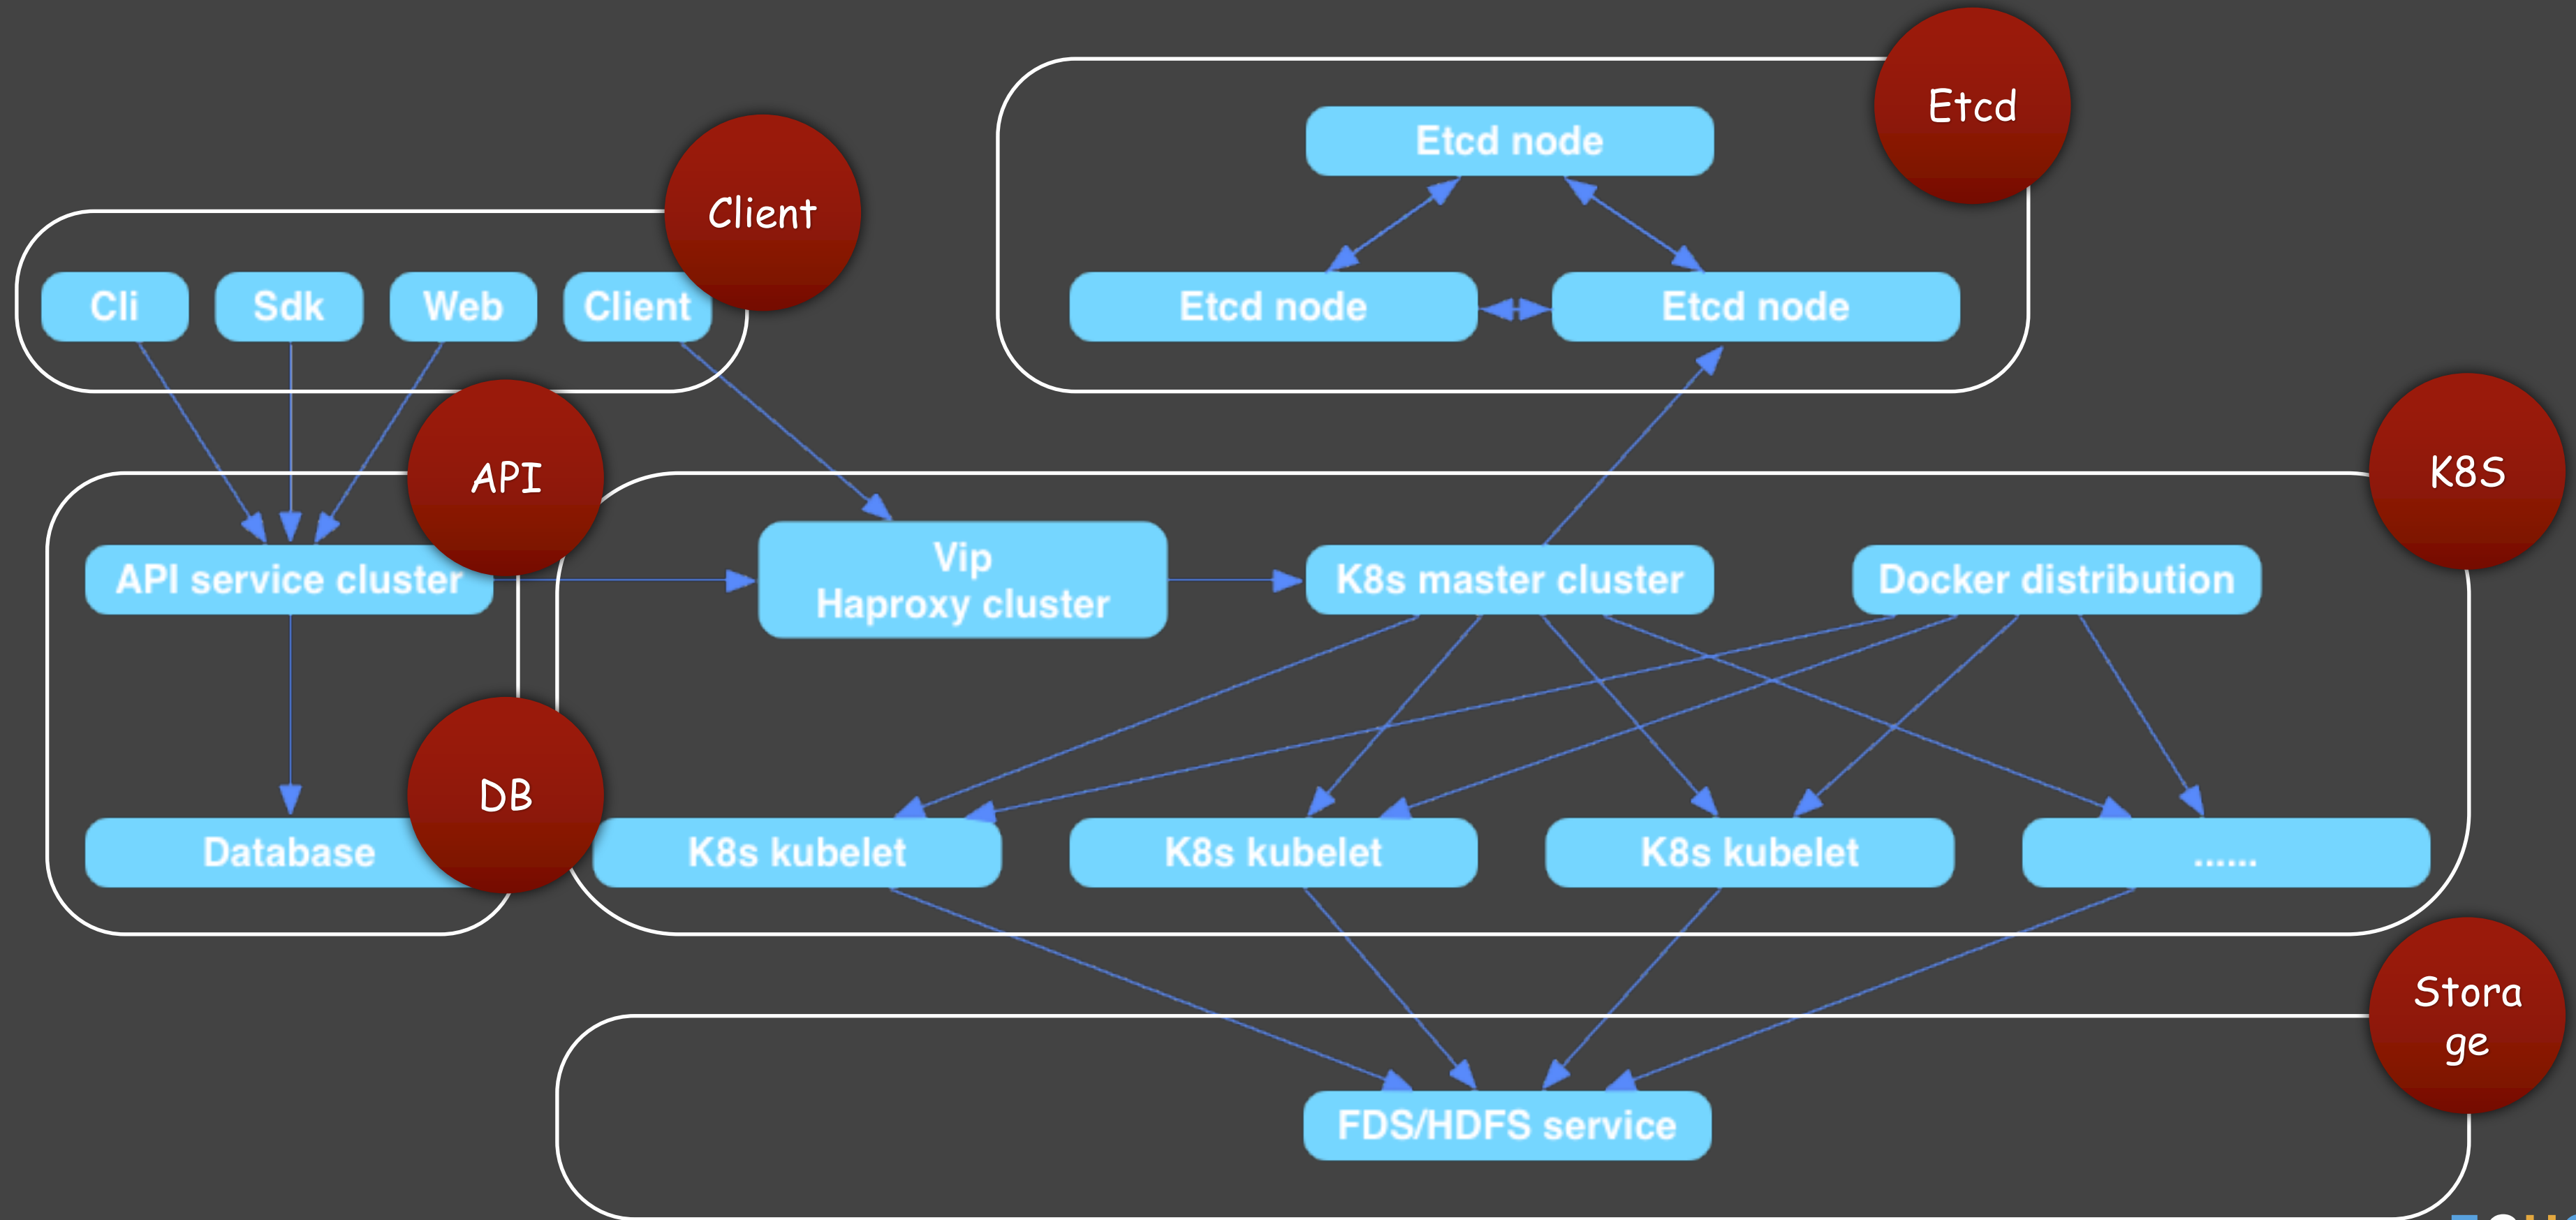
云机器学习平台 - workflow



云机器学习平台 - 算子

- ❖ 最简抽象接口：`abstract void execute()`
- ❖ ~10 数据处理算子，10+ 机器学习训练算子，3 AutoML算子
- ❖ Hadoop/Spark/Kubernetes多集群支持
- ❖ More on <https://www.4paradigm.com/support/help>

云机器学习平台 - 高可用



云机器学习平台 - 多租户

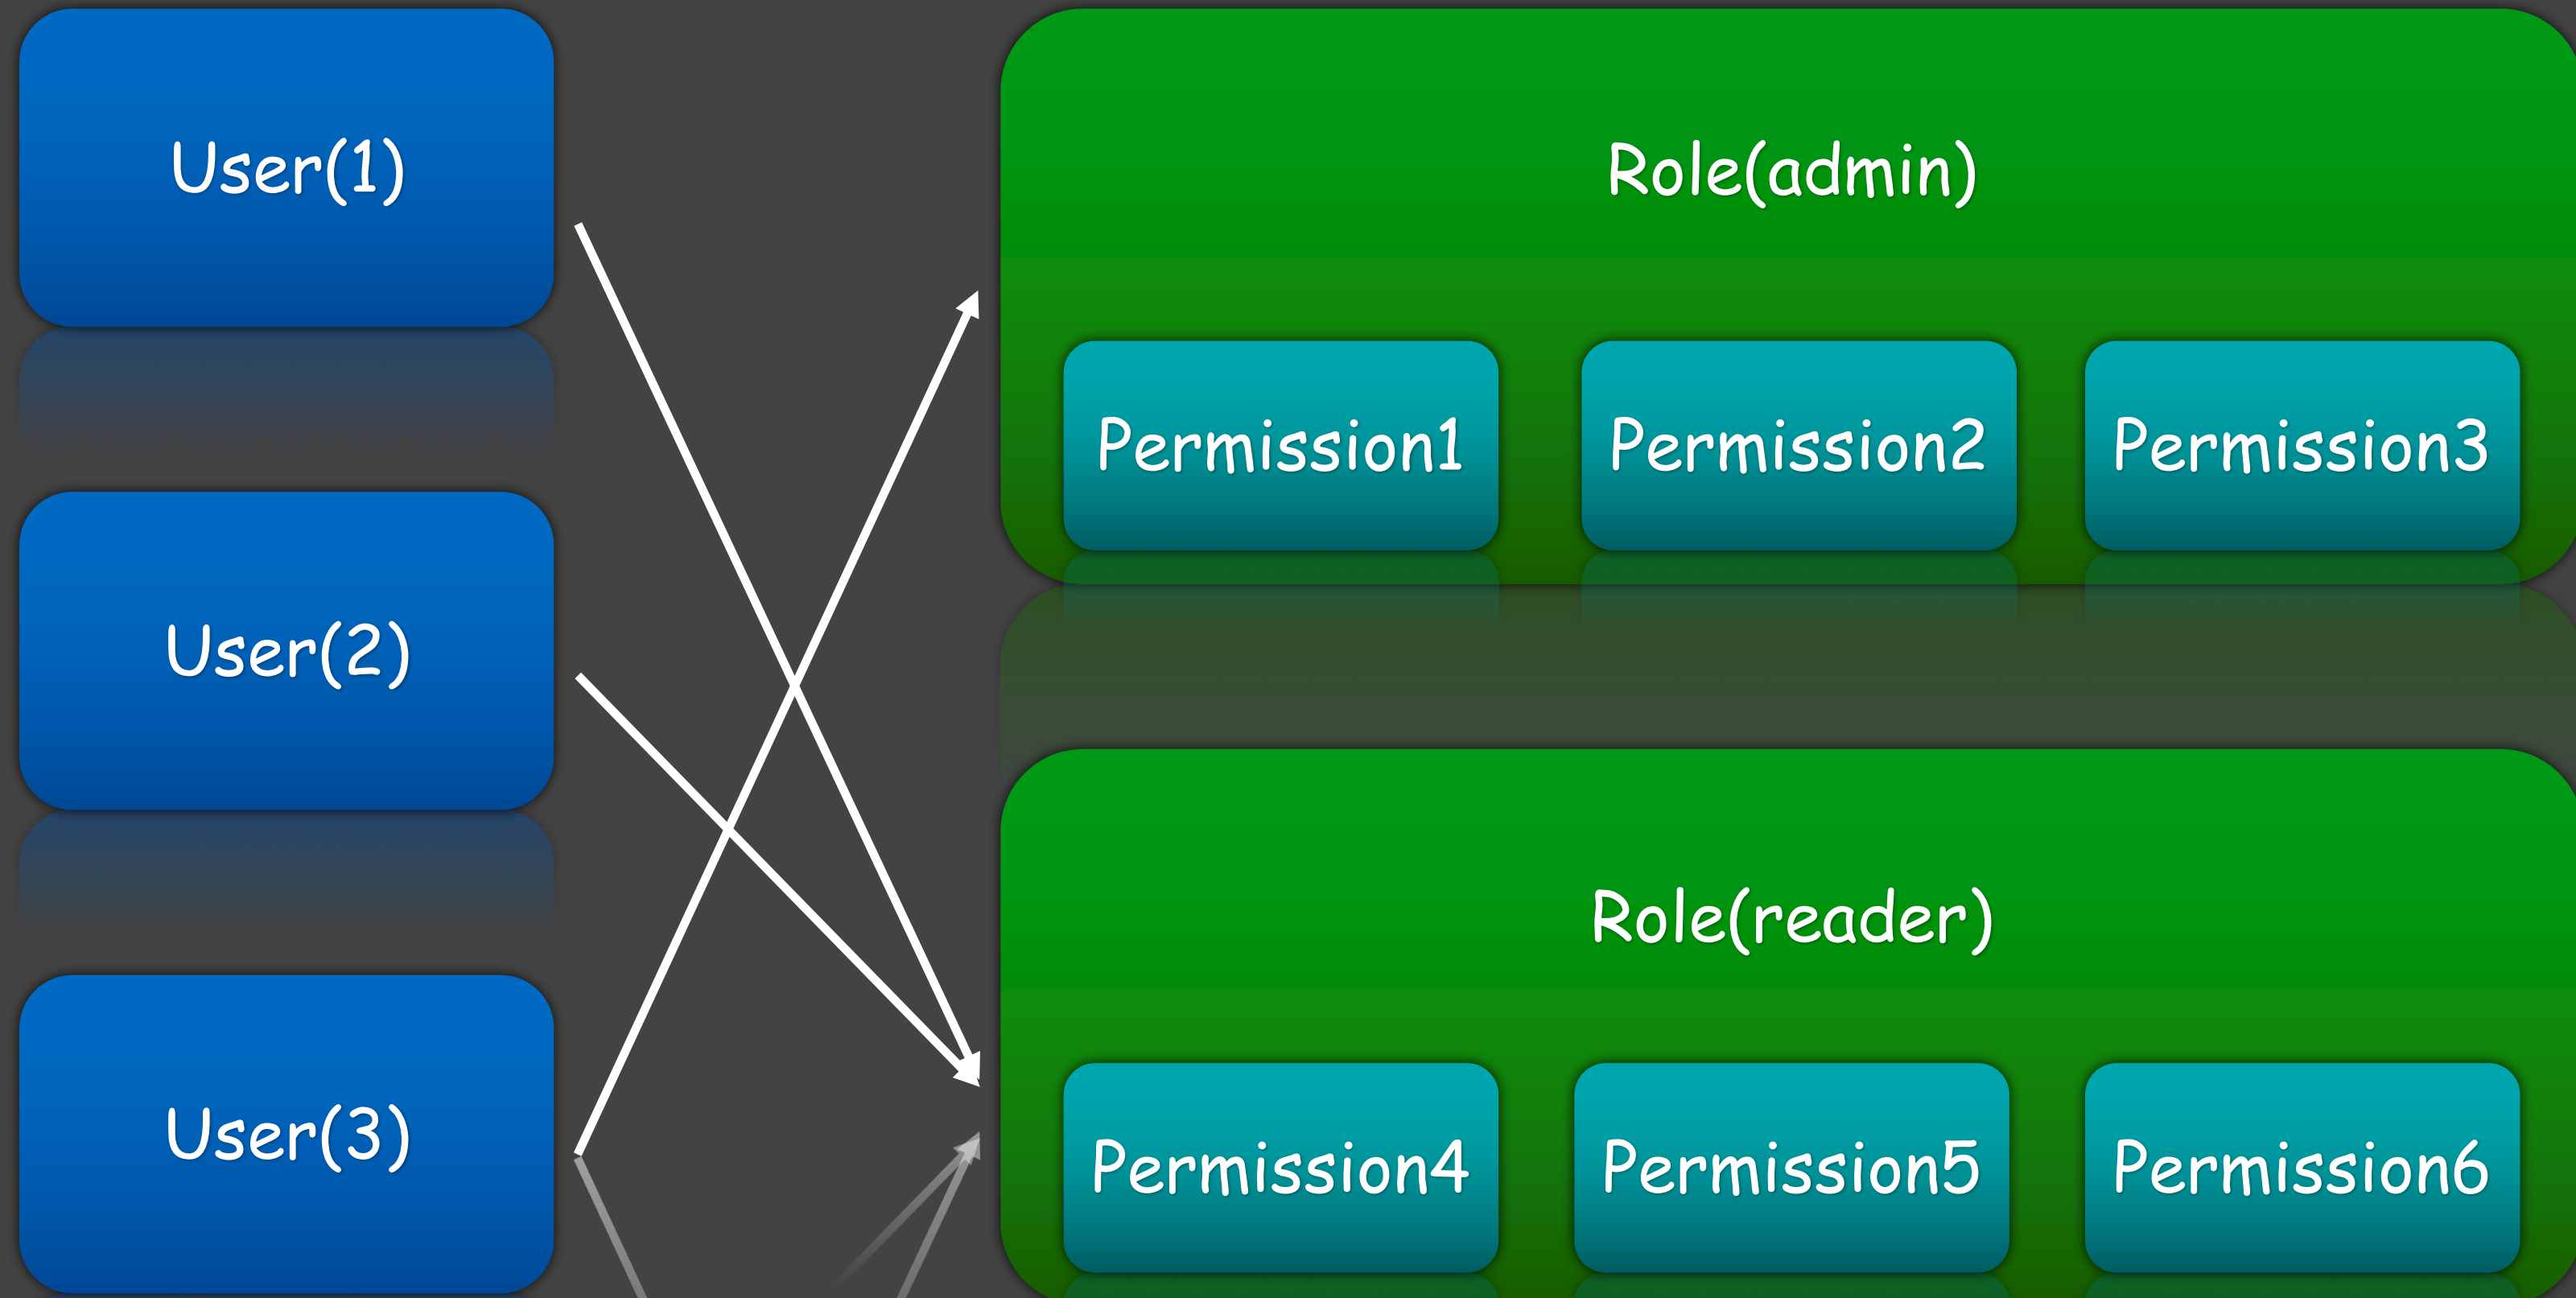
✿ Authentication

- Username / Password
- Access key / Secret key
- LDAP

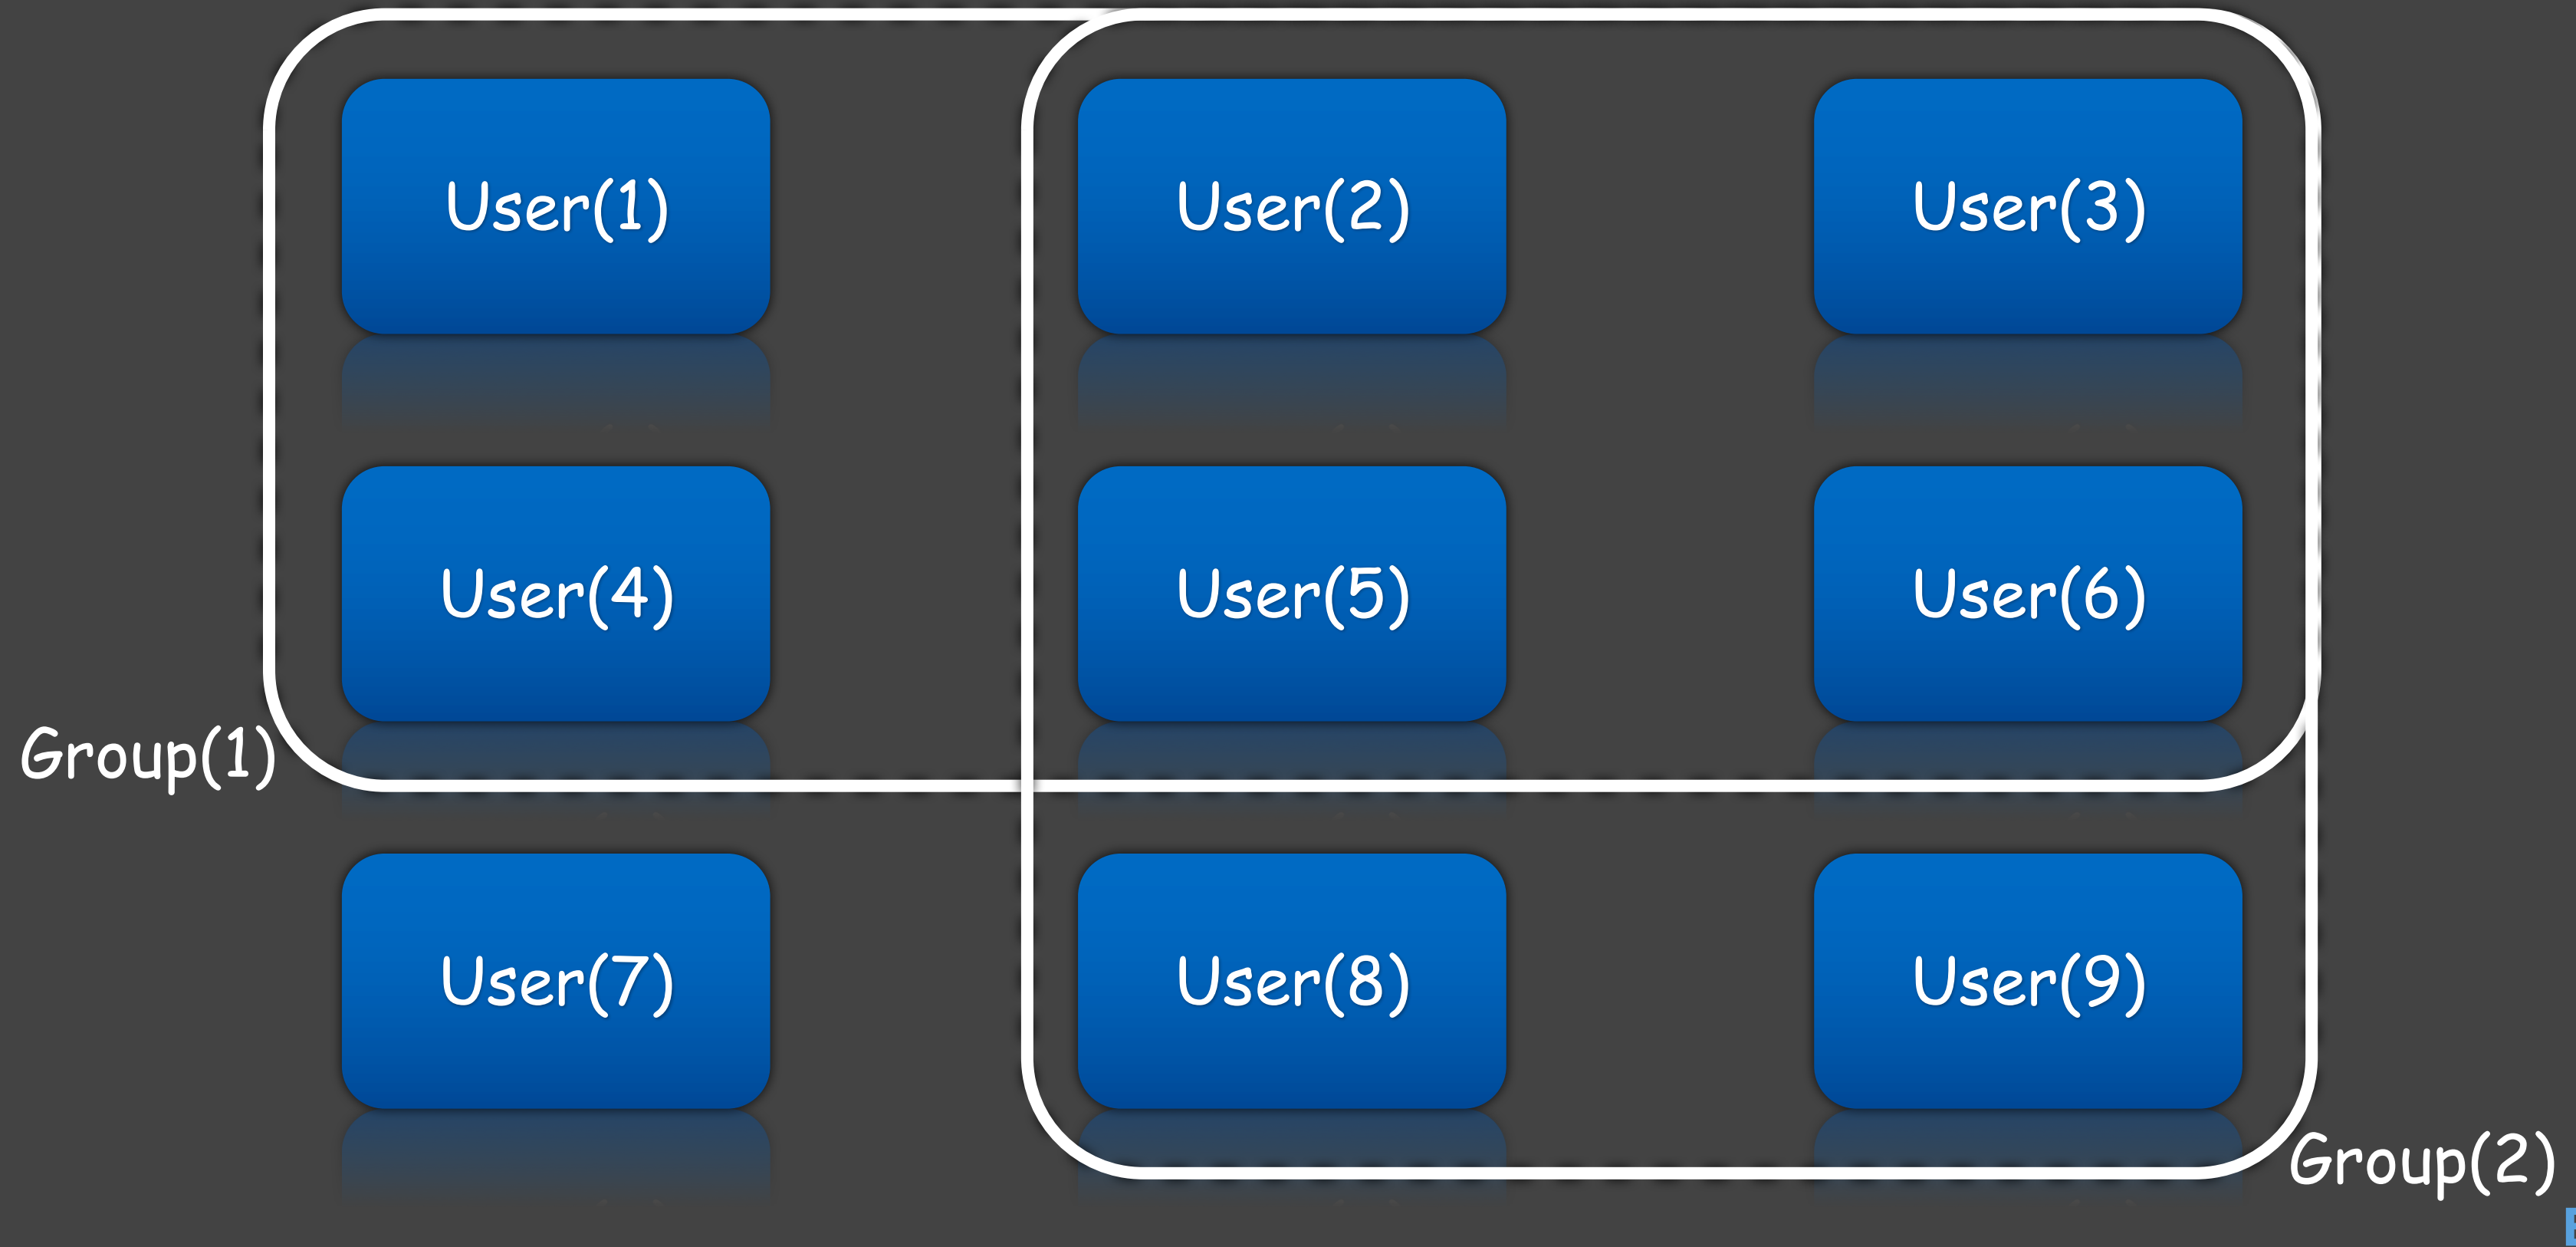
✿ Authorization

- Role-based access control

云机器学习平台 - 多租户



云机器学习平台 - 多租户



Conclusion

- ❖ 人工智能与机器学习介绍
- ❖ 机器学习算法原理与实现
- ❖ 云机器学习平台架构实践

Thanks