

Big Data Analytics

Lars Schmidt-Thieme

Information Systems and Machine Learning Lab (ISMLL)
Institute of Computer Science
University of Hildesheim, Germany

C. Distributed Computing Environments / 3. Computational Graphs (TensorFlow)

Syllabus

- | | | |
|------------|------|---|
| Tue. 9.4. | (1) | 0. Introduction |
| | | A. Parallel Computing |
| Tue. 16.4. | (2) | A.1 Threads |
| Tue. 23.4. | (3) | A.2 Message Passing Interface (MPI) |
| Tue. 30.4. | (4) | A.3 Graphical Processing Units (GPUs) |
| | | B. Distributed Storage |
| Tue. 7.5. | (5) | B.1 Distributed File Systems |
| Tue. 14.5. | (6) | B.2 Partitioning of Relational Databases |
| Tue. 21.5. | (7) | B.3 NoSQL Databases |
| | | C. Distributed Computing Environments |
| Tue. 28.5. | (8) | C.1 Map-Reduce |
| Tue. 4.6. | (9) | C.2 Resilient Distributed Datasets (Spark) |
| Tue. 11.6. | — | — <i>Pentecoste Break</i> — |
| Tue. 18.6. | (10) | C.3 Computational Graphs (TensorFlow) |
| | | D. Distributed Machine Learning Algorithms |
| Tue. 25.6. | (11) | D.1 Distributed Stochastic Gradient Descent |
| Tue. 2.7. | (12) | D.2 Distributed Matrix Factorization |
| Tue. 9.7. | (13) | Questions and Answers |

Outline

1. The Computational Graph
2. Variables
3. Example: Linear Regression
4. Automatic Gradients
5. Large Data I: Feeding
6. Large Data II: Reader Nodes
7. Debugging

Outline

1. The Computational Graph
2. Variables
3. Example: Linear Regression
4. Automatic Gradients
5. Large Data I: Feeding
6. Large Data II: Reader Nodes
7. Debugging

TensorFlow

- ▶ Computational framework
- ▶ multi-device, distributed
- ▶ Core in C/C++, standard interface in Python
 - ▶ several further language bindings, e.g., R, Java
- ▶ open source
 - ▶ developed by Google
 - ▶ initially released Nov. 2015
 - ▶ 2nd generation framework
 - ▶ 1st generation framework was called DistBelief
- ▶ alternative: pytorch
 - ▶ developed by facebook, since 2016, open source

Tensors

- ▶ tensor = multidimensional array
 - ▶ rank = number of dimensions
 - ▶ shape = vector of sizes,
one size for each dimension.

rank	common name	shape
0	scalar	()
1	vector	(size)
2	matrix	(numrows, numcols)
≥ 3	tensor of higher order	(numdim ₁ , numdim ₂ , ..., numdim _r)

- ▶ examples:

$$A = \begin{pmatrix} 1.0 & -3.0 & 2.3 & 1.7 \\ 5.6 & 0.0 & -1.3 & 3.4 \\ -7.7 & -3.3 & -2.1 & 5.2 \end{pmatrix}, \quad \text{shape}(A) = (3, 4), \quad \text{rank}(A) = 2$$

Computational Graphs

- ▶ TensorFlow organizes a computation as a directed graph.
- ▶ Nodes represent a tensor.
 - ▶ or a list of tensors.
- ▶ Tensors can be:
 - ▶ stored tensors
 - ▶ immutable, value provided at creation time: **tf.constant**
 - ▶ immutable, value provided when running the graph: **tf.placeholder**
 - ▶ mutable: **tf.Variable**
 - ▶ computed tensors (**operations**):
 - ▶ having one or more input tensors
 - ▶ having one or more output tensors
 - output index: **port**
- ▶ Edges represent dependencies.
 - ▶ Edge $x \rightarrow y$ if y is computed and x one of its inputs.

Sessions

- ▶ A session represents the state of an ongoing computation on a computational graph.
- ▶ create with default constructor `tf.Session`.
- ▶ compute the value of a tensor node with `run`.

Two Phases

1. Construct the Computational Graph
 - ▶ create tensor nodes
 - ▶ possibly referencing other tensor nodes as inputs
2. Compute values of a node of the Computation Graph (**running**)
 - ▶ usually specify target tensor(s)
 - ▶ computes all intermediate tensors required for this tensor
 - ▶ yield the value of the target tensor(s)

Hello TensorFlow: Add two Constants

```
1 import tensorflow as tf
2
3 a = tf.constant(3.0)
4 b = tf.constant(4.0)
5 x = tf.add(a, b)
6
7 print(a)
8 print(b)
9 print(x)
10
11 sess = tf.Session()
12 x_val = sess.run(x)
13 print(x_val)
```

Output:

```
1 Tensor("Const:0", shape=(), dtype=float32)
2 Tensor("Const_1:0", shape=(), dtype=float32)
3 Tensor("Add:0", shape=(), dtype=float32)
4 7
```

Hello TensorFlow: Add two Constants

```
1 import tensorflow as tf
2
3 a = tf.constant([3.0, -2.7, 1.2])
4 b = tf.constant([4.0, 5.1, -1.7])
5 x = tf.add(a, b)
6
7 print(a)
8 print(b)
9 print(x)
10
11 sess = tf.Session()
12 x_val = sess.run(x)
13 print(x_val)
```

Output:

```
1 Tensor("Const_2:0", shape=(3,), dtype=float32)
2 Tensor("Const_3:0", shape=(3,), dtype=float32)
3 Tensor("Add_1:0", shape=(3,), dtype=float32)
4 [ 7.0  2.4 -0.5 ]
```

Tensor Types

- ▶ Different element types are represented by `tf.DType`:

dtype	description
<code>tf.float16</code>	16-bit half-precision floating-point
<code>tf.float32</code>	32-bit single-precision floating-point
<code>tf.float64</code>	64-bit double-precision floating-point
<code>tf.bfloat16</code>	16-bit truncated floating-point
<code>tf.complex64</code>	64-bit single-precision complex
<code>tf.complex128</code>	128-bit double-precision complex
<code>tf.int8</code>	8-bit signed integer
<code>tf.uint8</code>	8-bit unsigned integer
<code>tf.uint16</code>	16-bit unsigned integer
<code>tf.int16</code>	16-bit signed integer
<code>tf.int32</code>	32-bit signed integer
<code>tf.int64</code>	64-bit signed integer
<code>tf.bool</code>	Boolean
<code>tf.string</code>	String
<code>tf.qint8</code>	Quantized 8-bit signed integer
<code>tf.quint8</code>	Quantized 8-bit unsigned integer
<code>tf.qint16</code>	Quantized 16-bit signed integer
<code>tf.quint16</code>	Quantized 16-bit unsigned integer
<code>tf.qint32</code>	Quantized 32-bit signed integer
<code>tf.resource</code>	Handle to a mutable resource

- ▶ if omitted, inferred from values:

```

1 a = tf.constant(4.0)
2 b = tf.constant(4)
3 print(a)
4 print(b)

```

Output:

```

1 Tensor("Const:0", shape=(), dtype=float32)
2 Tensor("Const_1:0", shape=(), dtype=int32)

```

Operations: Overloaded Operators

```
1 import tensorflow as tf
2
3 a = tf.constant(3.0)
4 b = tf.constant(4.0)
5 x = a + b
6
7 print(x)
8
9 sess = tf.Session()
10 x_val = sess.run(x)
11 print(x_val)
```

Output:

```
1 Tensor("add_1:0", shape=(), dtype=float32)
2 7
```

operator	operation node
+	tf.add
-	tf.subtract
*	tf.multiply
/	tf.divide

Outline

1. The Computational Graph
- 2. Variables**
3. Example: Linear Regression
4. Automatic Gradients
5. Large Data I: Feeding
6. Large Data II: Reader Nodes
7. Debugging

```
v = tf.Variable(initial_value = None, ..., name = None, ...,
               dtype = None)
```

- ▶ has immutable element type
- ▶ has mutable shape (**set_shape**).
- ▶ has mutable element values.
 - ▶ set by **tf.assign**, **tf.assign_add** (operations)
 - ▶ separate values in each session
- ▶ has to be initialized before first use:
 - ▶ run **v.initializer** operation or
 - ▶ run initializers of all variables:

```
1 init = tf.global_variables_initializer()
2 sess.run(init)
```

Variables

```
1 import tensorflow as tf
2
3 x = tf.Variable(3.0)
4 sess = tf.Session()
5 sess.run( x.initializer )
6 x_val = sess.run(x)
7 print(x_val)
8
9 x_plus_one = tf.assign_add(x, 1.0)
10
11 for t in range(5):
12     x_val = sess.run(x_plus_one)
13     print(t, x_val)
```

Output:

```
1 3.0
2
3 0 4.0
4 1 5.0
5 2 6.0
6 3 7.0
7 4 8.0
```


Initializing from Other Variables

```

1 import tensorflow as tf
2
3 x = tf.Variable(3.0)
4 y = tf.Variable( tf.multiply(tf.constant(2.0),
5                             x.initialized_value()) )
6 init = tf.global_variables_initializer()
7 sess = tf.Session()
8 sess.run(init)
9 y_val = sess.run(y)
10 print(y_val)

```

Output:

1 6.0

- ▶ **v.initialized_value** assures that a variable has been initialized before
 - ▶ do not use

```
1 y = tf.Variable( tf.multiply(tf.constant(2.0), x) )
```

as x may be selected to be initialized after y .

Outline

1. The Computational Graph
2. Variables
- 3. Example: Linear Regression**
4. Automatic Gradients
5. Large Data I: Feeding
6. Large Data II: Reader Nodes
7. Debugging

Example: Linear Regression

$$\hat{y} := \beta_0 + X\beta \quad \text{prediction}$$

$$r := y - \hat{y} \quad \text{residuum}$$

$$\ell := \frac{1}{2} \sum_{n=1}^N r_n^2 \quad \text{loss/error}$$

$$-\frac{\partial \ell}{\partial \beta} := X^T r \quad \text{negative gradient w.r.t. } \beta$$

$$-\frac{\partial \ell}{\partial \beta_0} := \sum_{n=1}^N r_n \quad \text{negative gradient w.r.t. } \beta_0$$

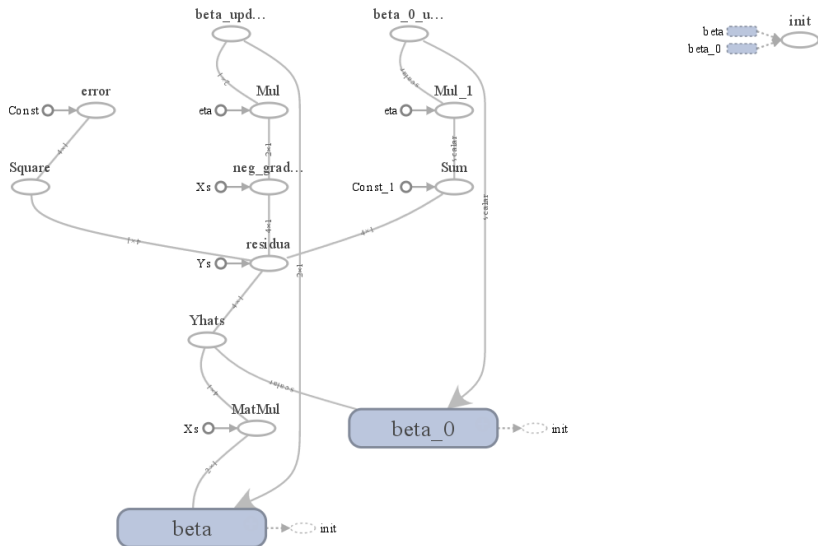
$$\beta^{\text{next}} := \beta - \eta \frac{\partial \ell}{\partial \beta} \quad \text{update of } \beta$$

$$\beta_0^{\text{next}} := \beta_0 - \eta \frac{\partial \ell}{\partial \beta_0} \quad \text{update of } \beta_0$$

Example: Linear Regression

```
1 import tensorflow as tf
2
3 Xs_data = [[2,1], [1,2], [4,3], [3,4]]
4 Ys_data = [[+1], [+1], [-1], [-1]]
5 eta_data = 0.01
6
7 Xs = tf.constant(Xs_data, dtype=tf.float32)
8 Ys = tf.constant(Ys_data, dtype=tf.float32)
9 eta = tf.constant(eta_data)
10
11 beta = tf.Variable([[0], [0]], dtype=tf.float32)
12 beta_0 = tf.Variable(0, dtype=tf.float32)
13
14 Yhats = tf.add(beta_0, tf.matmul(Xs, beta))
15 residua = tf.subtract(Ys, Yhats)
16 error = tf.reduce_sum(tf.square(residua))
17
18 neg_grad_beta = tf.matmul(Xs, residua, adjoint_a=True)
19 beta_update = tf.assign_add(beta, tf.multiply(eta, neg_grad_beta))
20 beta_0_update = tf.assign_add(beta_0, tf.multiply(eta, tf.reduce_sum(residua)))
21
22 init = tf.global_variables_initializer()
23 sess = tf.Session()
24 sess.run( init )
25
26 for t in range(100):
27     error_val, beta_val, beta_0_val = sess.run([error,beta_update,beta_0_update])
28     print(t, error_val, beta_0_val, beta_val[0,0], beta_val[1,0])
```

Example: Linear Regression / Computational Graph



Outline

1. The Computational Graph
2. Variables
3. Example: Linear Regression
- 4. Automatic Gradients**
5. Large Data I: Feeding
6. Large Data II: Reader Nodes
7. Debugging

`tf.gradients(ys, xs, ...)`

- ▶ create operations whose final node computes all the gradients

$$\left(\frac{\partial y_n}{\partial x_m} \right)_{n=1, \dots, N, m=1, \dots, M} \quad y_S = (y_1, \dots, y_N), x_S = (x_1, \dots, x_M)$$

Example: Linear Regression w. Automatic Gradients

```
1 import tensorflow as tf
2
3 Xs_data = [[2,1], [1,2], [4,3], [3,4]]
4 Ys_data = [[+1], [+1], [-1], [-1]]
5 eta_data = 0.01
6
7 Xs = tf.constant(Xs_data, dtype=tf.float32)
8 Ys = tf.constant(Ys_data, dtype=tf.float32)
9 eta = tf.constant(eta_data)
10
11 beta = tf.Variable([[0], [0]], dtype=tf.float32)
12 beta_0 = tf.Variable(0, dtype=tf.float32)
13
14 Yhats = tf.add(beta_0, tf.matmul(Xs, beta))
15 error = tf.reduce_sum(tf.square(tf.subtract(Ys, Yhats)))
16
17 grads = tf.gradients(error, [beta, beta_0])
18 beta_update = tf.assign_sub(beta, tf.multiply(eta, grads[0]))
19 beta_0_update = tf.assign_sub(beta_0, tf.multiply(eta, grads[1]))
20
21 init = tf.global_variables_initializer()
22 sess = tf.Session()
23 sess.run( init )
24
25 for t in range(100):
26     error_val, beta_val, beta_0_val = sess.run([error,beta_update,beta_0_update])
27     print(t, error_val, beta_0_val, beta_val[0,0], beta_val[1,0])
```


How do Automatic Gradients Work?

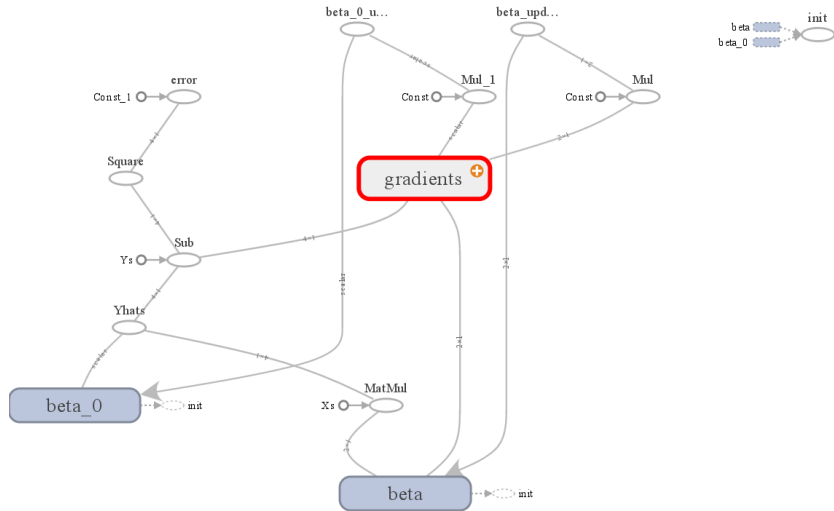
to compute $\frac{\partial y}{\partial x}$:

- ▶ find all paths $p^1, \dots, p^K \in G^*$ in the graph G from x to y
- ▶ use chain rule:


$$\frac{\partial y}{\partial x} = \sum_{k=1}^K \prod_{l=|p^k|}^2 \frac{\partial p_l^k}{\partial p_{l-1}^k}$$

- ▶ each operation $p_l^k =: o$ has to provide its gradient $\frac{\partial o}{\partial i}$ for each of its inputs i .
 - ▶ then $\frac{\partial p_l^k}{\partial p_{l-1}^k} = \frac{\partial o}{\partial i}$ for $i = p_{l-1}^k$.

Example: LinReg w. Auto Grads / Computational Graph



Example: LinReg w. Auto Grads / Computational Graph



Outline

1. The Computational Graph
2. Variables
3. Example: Linear Regression
4. Automatic Gradients
- 5. Large Data I: Feeding**
6. Large Data II: Reader Nodes
7. Debugging

- ▶ representing large data as a whole as a constant is not so useful
 - ▶ e.g., if its size exceeds GPU memory, it cannot be deployed to GPU at all.
- ▶ better break data into smaller pieces
 - ▶ e.g., single instances or minibatches
 - ▶ batch GD \rightarrow SGD
- ▶ build a graph for a single instance / minibatch
- ▶ create placeholder nodes for the instance / minibatch
- ▶ placeholders are filled with the `feed_dict` parameter of `run`.

Placeholder Nodes and Feeding

```
1 import tensorflow as tf
2
3 a = tf.placeholder(shape=(), dtype=tf.float32)
4 b = tf.placeholder(shape=(), dtype=tf.float32)
5 x = a + b
6
7 sess = tf.Session()
8 print( sess.run(x, {a: 3, b: 7}) )
9 print( sess.run(x, {a: -2, b: 4}) )
```

Output:

```
1 10.0
2 2.0
```

Example: Feeding SGD

```

1 import tensorflow as tf
2
3 Xs_data = [[2,1], [1,2], [4,3], [3,4]]
4 Ys_data = [+1, +1, -1, -1]
5 eta_data = 0.01
6
7 X = tf.placeholder(shape=(2), dtype=tf.float32)
8 Y = tf.placeholder(shape=(), dtype=tf.float32)
9 eta = tf.constant(eta_data)
10
11 beta = tf.Variable([0, 0], dtype=tf.float32)
12 beta_0 = tf.Variable(0, dtype=tf.float32)
13
14 Yhat = tf.add(beta_0, tf.reduce_sum(tf.multiply(X, beta)))
15 error = tf.reduce_sum(tf.square(tf.subtract(Y, Yhat)))
16
17 grads = tf.gradients(error, [beta, beta_0])
18 beta_update = tf.assign_sub(beta, tf.multiply(eta, grads[0]))
19 beta_0_update = tf.assign_sub(beta_0, tf.multiply(eta, grads[1]))
20
21 init = tf.global_variables_initializer()
22 sess = tf.Session()
23 sess.run( init )
24
25 for t in range(100):
26     error_epoch = 0
27     for X_data, Y_data in zip(Xs_data, Ys_data):
28         error_val, beta_val, beta_0_val = sess.run([error, beta_update, beta_0_update],
29             { X: X_data, Y: Y_data })
30         error_epoch += error_val
31     print(t, error_epoch, beta_0_val, beta_val[0], beta_val[1])

```

Outline

1. The Computational Graph
2. Variables
3. Example: Linear Regression
4. Automatic Gradients
5. Large Data I: Feeding
- 6. Large Data II: Reader Nodes**
7. Debugging

Reader Node

```

1 import tensorflow as tf
2
3 data_files = ['lr-data.csv']; eta_data = 0.01
4
5 filename_queue = tf.train.string_input_producer(data_files)
6 reader = tf.TextLineReader(skip_header_lines=1)
7 _, line = reader.read(filename_queue)
8
9 sess = tf.Session()
10 coord = tf.train.Coordinator()
11 threads = tf.train.start_queue_runners(coord=coord, sess=sess)
12
13 for t in range(6):
14     line_val = sess.run(line)
15     print(t, line_val)
16
17 coord.request_stop()
18 coord.join(threads)

```

file lr-data.csv:

```

1 X1, X2, Y
2 2, 1, +1
3 1, 2, +1
4 4, 3, -1
5 3, 4, -1

```

Output:

```

1 0 b'2,\u001,\u001+'
2 1 b'1,\u002,\u001+'
3 2 b'4,\u003,\u00-1+'
4 3 b'3,\u004,\u00-1+'
5 4 b'2,\u001,\u001+'
6 5 b'1,\u002,\u001+'

```

Example: SGD Reading On The Fly

```

1 import tensorflow as tf
2 data_files = ['lr-data.csv']; eta_data = 0.01
3
4 filename_queue = tf.train.string_input_producer(data_files)
5 reader = tf.TextLineReader(skip_header_lines=1)
6 key, value = reader.read(filename_queue)
7 X1, X2, Y = tf.decode_csv(value, record_defaults=[[0.0],[0.0],[0.0]])
8 X = tf.stack([X1, X2])
9 eta = tf.constant(eta_data)
10
11 beta = tf.Variable([0, 0], dtype=tf.float32)
12 beta_0 = tf.Variable(0, dtype=tf.float32)
13 Yhat = tf.add(beta_0, tf.reduce_sum(tf.multiply(X, beta)))
14 error = tf.reduce_sum(tf.square(tf.subtract(Y, Yhat)))
15 grads = tf.gradients(error, [beta, beta_0])
16 beta_update = tf.assign_sub(beta, tf.multiply(eta, grads[0]))
17 beta_0_update = tf.assign_sub(beta_0, tf.multiply(eta, grads[1]))
18 init = tf.global_variables_initializer()
19 sess = tf.Session() ; sess.run( init )
20 coord = tf.train.Coordinator()
21 threads = tf.train.start_queue_runners(coord=coord, sess=sess)
22
23 error_epoch = 0
24 for t in range(400):
25     error_val, beta_val, beta_0_val = sess.run([error,beta_update,beta_0_update])
26     error_epoch += error_val
27     if t % 10 == 0:
28         print(t, error_epoch, beta_0_val, beta_val[0], beta_val[1])
29         error_epoch = 0
30 coord.request_stop()
31 coord.join(threads)

```

Outline

1. The Computational Graph
2. Variables
3. Example: Linear Regression
4. Automatic Gradients
5. Large Data I: Feeding
6. Large Data II: Reader Nodes
- 7. Debugging**

Debugging: Visualize Computational Graph

1. Create a `summary.FileWriter` for the session and graph before running the session:

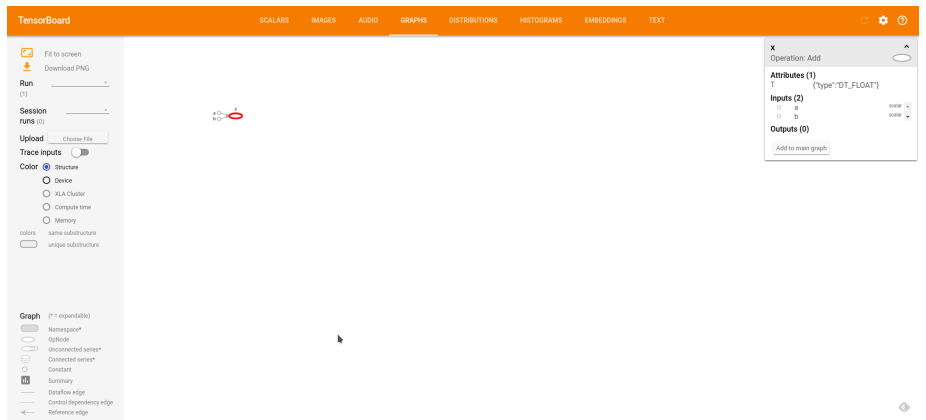
```
1 import tensorflow as tf
2
3 a = tf.constant(3.0, name='a')
4 b = tf.constant(4.0, name='b')
5 x = tf.add(a, b, name='x')
6
7 print(a)
8
9 sess = tf.Session()
10 log = tf.summary.FileWriter('logs/add-two-constants.log', sess.graph)
11 x_val = sess.run(x)
12 log.close()
13 print(x_val)
```

2. run tensorboard on the logdir:

```
1 > tensorboard --logdir logs/add-two-constants.log
```

3. open localhost:6006 in your browser

Debugging: Visualize Computational Graph



The screenshot displays the TensorBoard interface for visualizing a computational graph. The top navigation bar includes tabs for SCALARS, IMAGES, AUDIO, GRAPHS (selected), DISTRIBUTIONS, HISTOGRAMS, EMBEDDINGS, and TEXT. On the left sidebar, there are options for 'Fit to screen', 'Download PNG', 'Run (1)', 'Session runs (0)', 'Upload', 'Trace inputs', 'Color' (set to Structure), and a 'Graph' legend. The main area shows a small computational graph with a red circle highlighting a specific node. On the right, a detailed view of a node is shown, including its operation ('Add'), attributes ('T', 'DT_FLOAT'), and inputs ('a', 'b').

Summary (1/3)

- ▶ TensorFlow represents computations as **graphs**.
 - ▶ nodes representing (a list of) **tensors**.
 - ▶ stored:
 - immutable: **constant**, **placeholder**
 - mutable: **variable**
 - ▶ computed: **operation**
 - ▶ edges representing dependencies
 - ▶ $x \rightarrow y$: y is computed and x is one of its inputs
- ▶ Two phases:
 - ▶ graph construction
 - ▶ executing (parts of) the graph (**running**)

Summary (2/3)

- ▶ Nodes can be **distributed over different devices**.
 - ▶ cores of a CPU, GPUs, different compute nodes
 - ▶ **automatic placement** based on cost heuristics
 - ▶ eligible: sufficient memory available
 - ▶ expected runtime
 - based on cost heuristics
 - possibly also based on past runs
 - ▶ expected time for data movement between devices
- ▶ Operations can be assembled from dozens of **elementary operations**.
 - ▶ elementary math: add, subtract, multiply, divide
 - ▶ elementwise functions: log, exp, etc.
 - ▶ matrix operations: matrix product, inversion, etc.
 - ▶ structural tensor operations: slicing, stacking etc.

Summary (3/3)

- ▶ **Gradients** can be computed automatically.
 - ▶ simply using the chain rule
 - ▶ and explicit gradients for all elementary operations.
 - ▶ gradients add nodes to the graph.
- ▶ Medium-sized data should be broken into parts and **fed into a placeholder** for parts
 - ▶ e.g., SGD: single instances or minibatches
 - ▶ medium-sized data:
 - ▶ too large for the GPU
 - ▶ still can be read on a single data node
- ▶ Large data must be read by **reader nodes** as part of the graph execution.
 - ▶ large data: must be read on different data nodes in a distributed fashion

Further Readings

- ▶ TensorFlow white paper:
 - ▶ Abadi et al. [2016]
 - ▶ not yet fully complete: evaluation section is missing

References I

Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.