Apache SystemML - Declarative Large-Scale Machine Learning

Romeo Kienzler (IBM Waston IoT)
Berthold Reinwald (IBM Almaden Research Center)
Frederick R. Reiss (IBM Almaden Research Center)
Matthias Rieke (IBM Analytics)

Swiss Data Science Conference 16 - ZHAW - Winterthur

“High-level programming”

–Assembler vs. Python?

Why another lib?

- Custom machine learning algorithms
- Declarative ML
- Transparent distribution on data-parallel framework
  - Scale-up
  - Scale-out
- Cost-based optimiser generates low level execution plans

Why on Spark?

- Unification of SQL, Graph, Stream, ML
- Common RDD structure
- General DAG execution engine
  - lazy evaluation
  - distributed in-memory caching

2007-2008: Multiple projects at IBM Research – Almaden involving machine learning on Hadoop.

2009: We form a **dedicated team** for scalable ML

2009-2010: Through engagements with customers, we observe how data scientists create **ML solutions**.

2007 2008 2009 2010

Research

2011 2012 2013 2014

June 2015: IBM Announces open-source SystemML

September 2015: Code available on Github

November 2015: SystemML enters Apache incubation

February 2016: First release (0.9) of Apache SystemML

June 2016: Second Apache release (0.10)

2015

2016

SystemML at IBM Watson Health

Moved from Hadoop MapReduce to Spark

SystemML supports both frameworks

Exact same code

300X faster on 1/40th as many nodes

Data Scientist

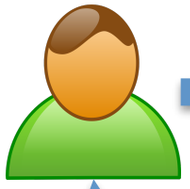

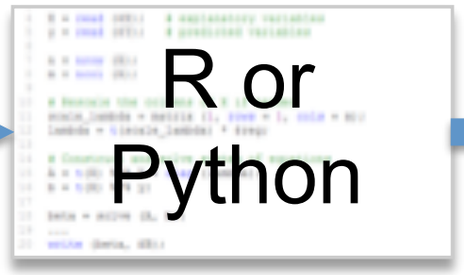

Systems Programmer

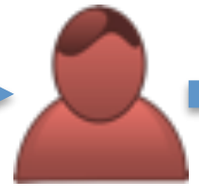

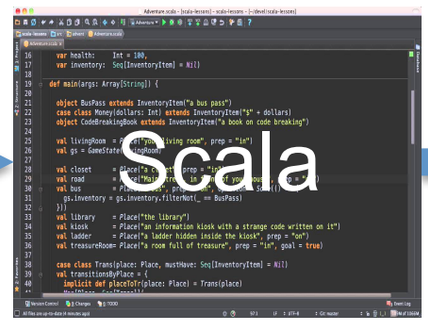

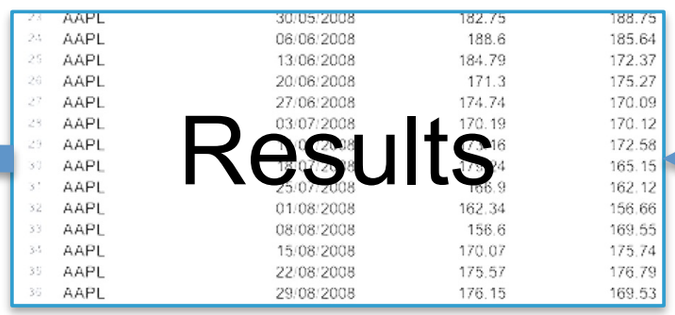

Alternating Least Squares

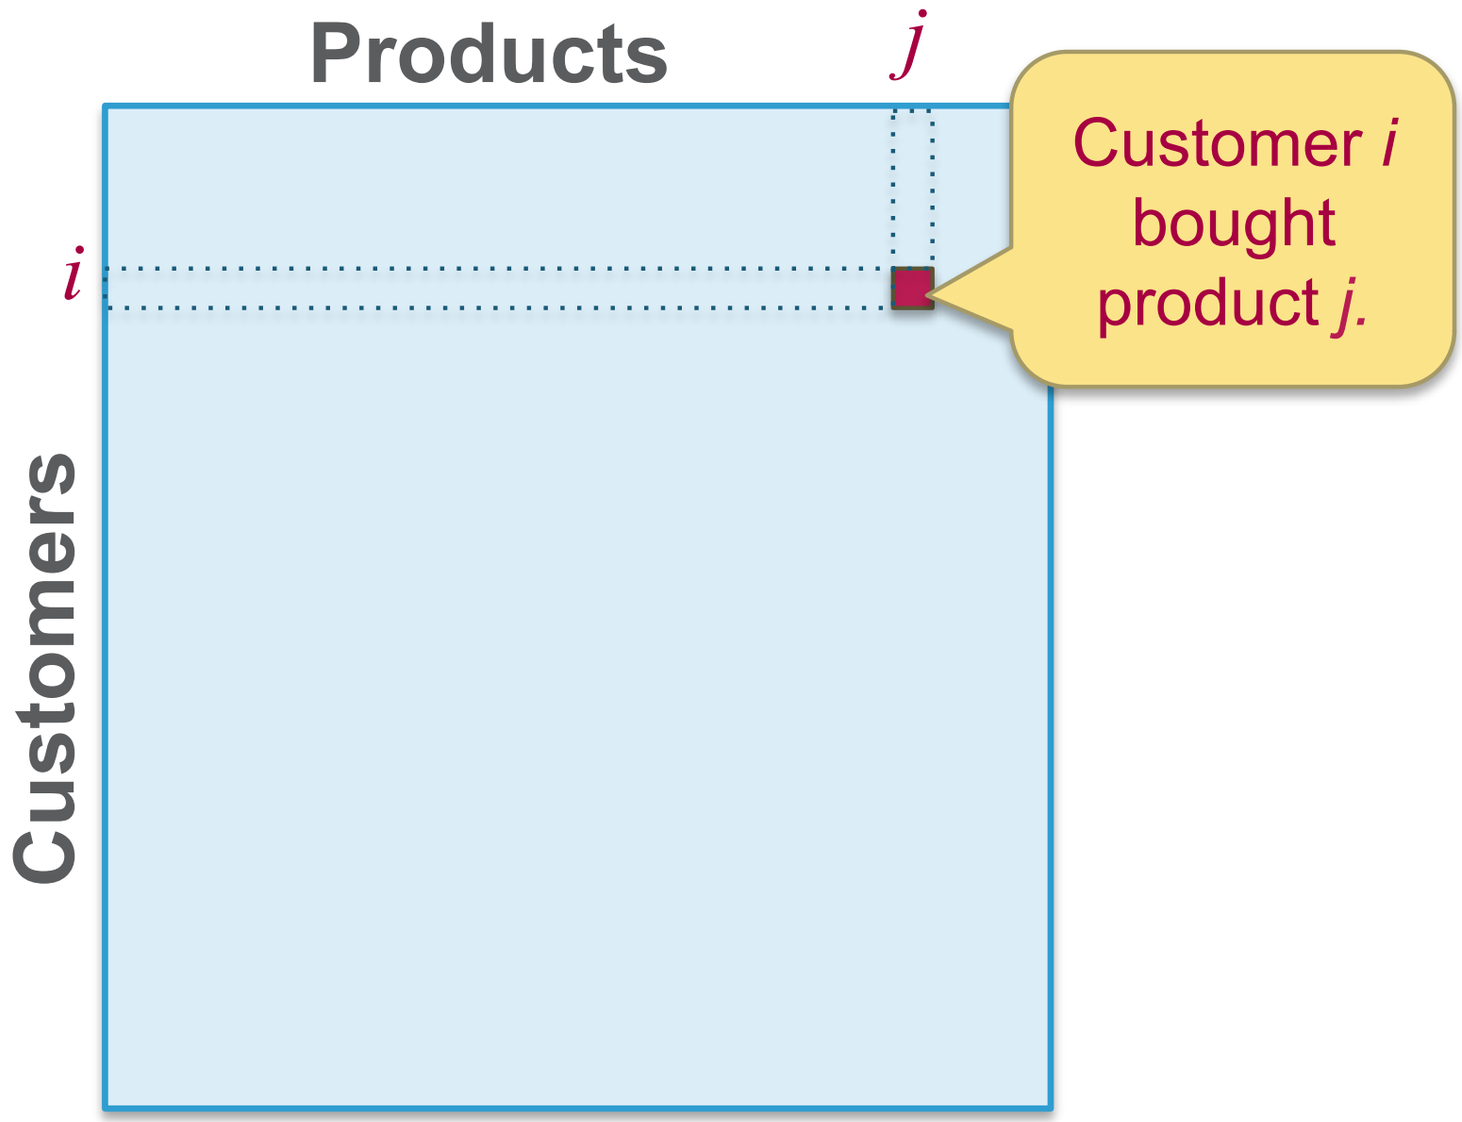

Alternating Least Squares

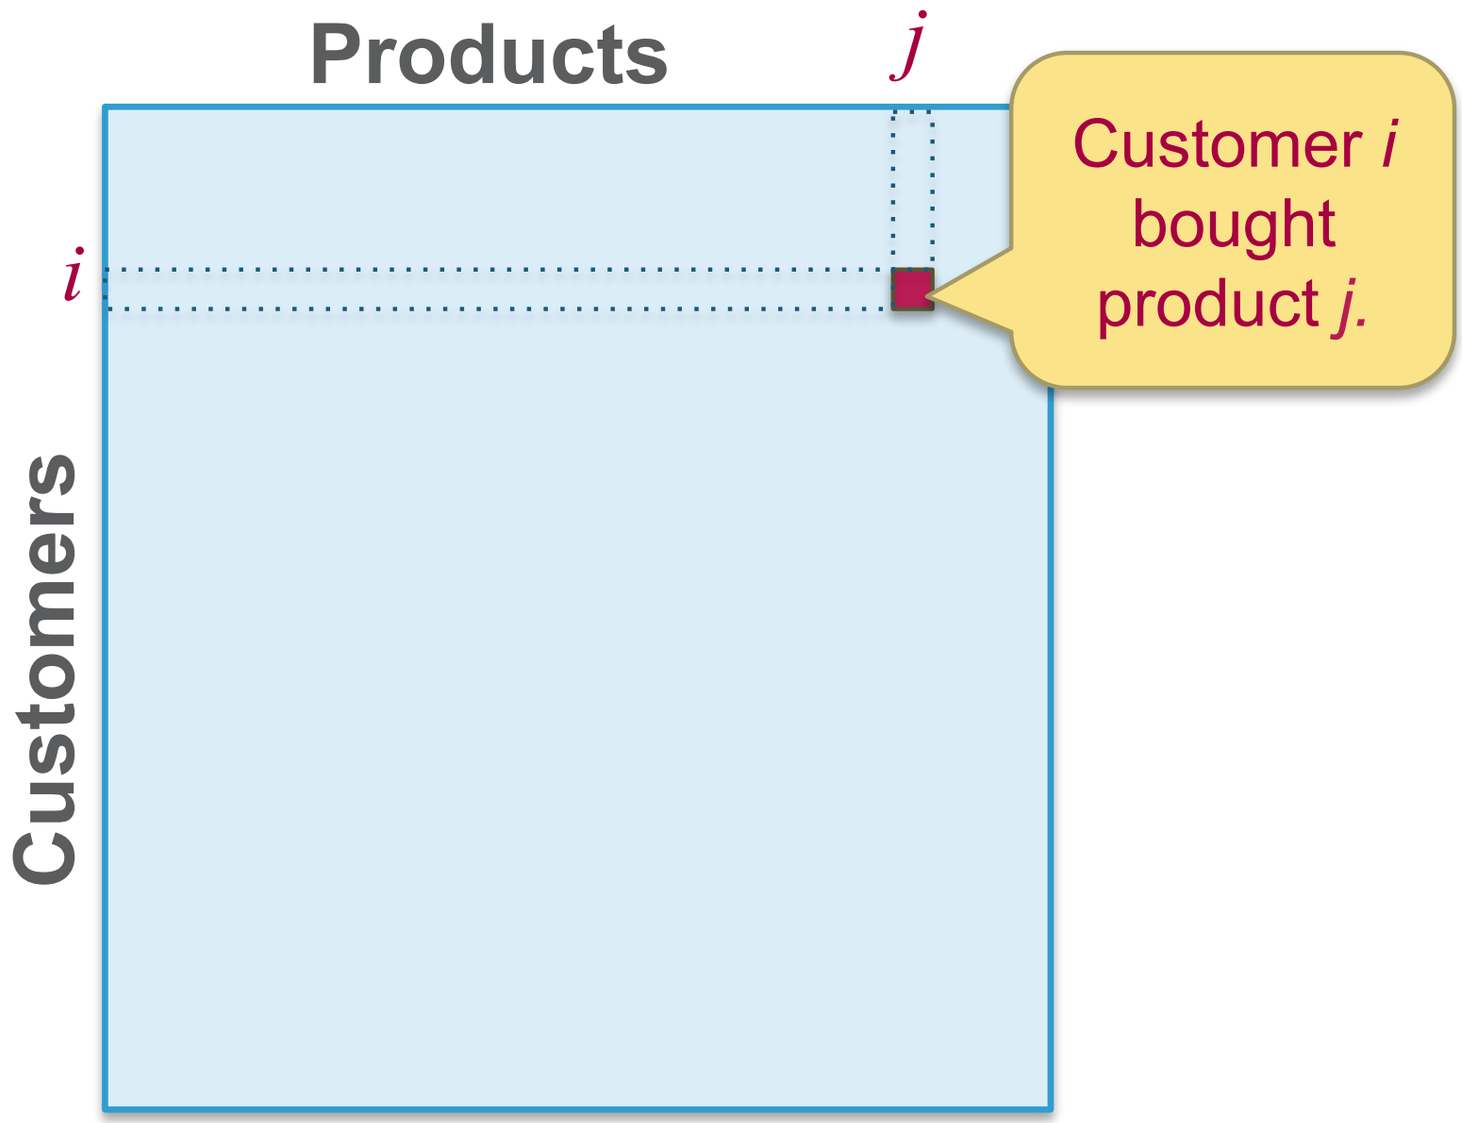

Alternating Least Squares

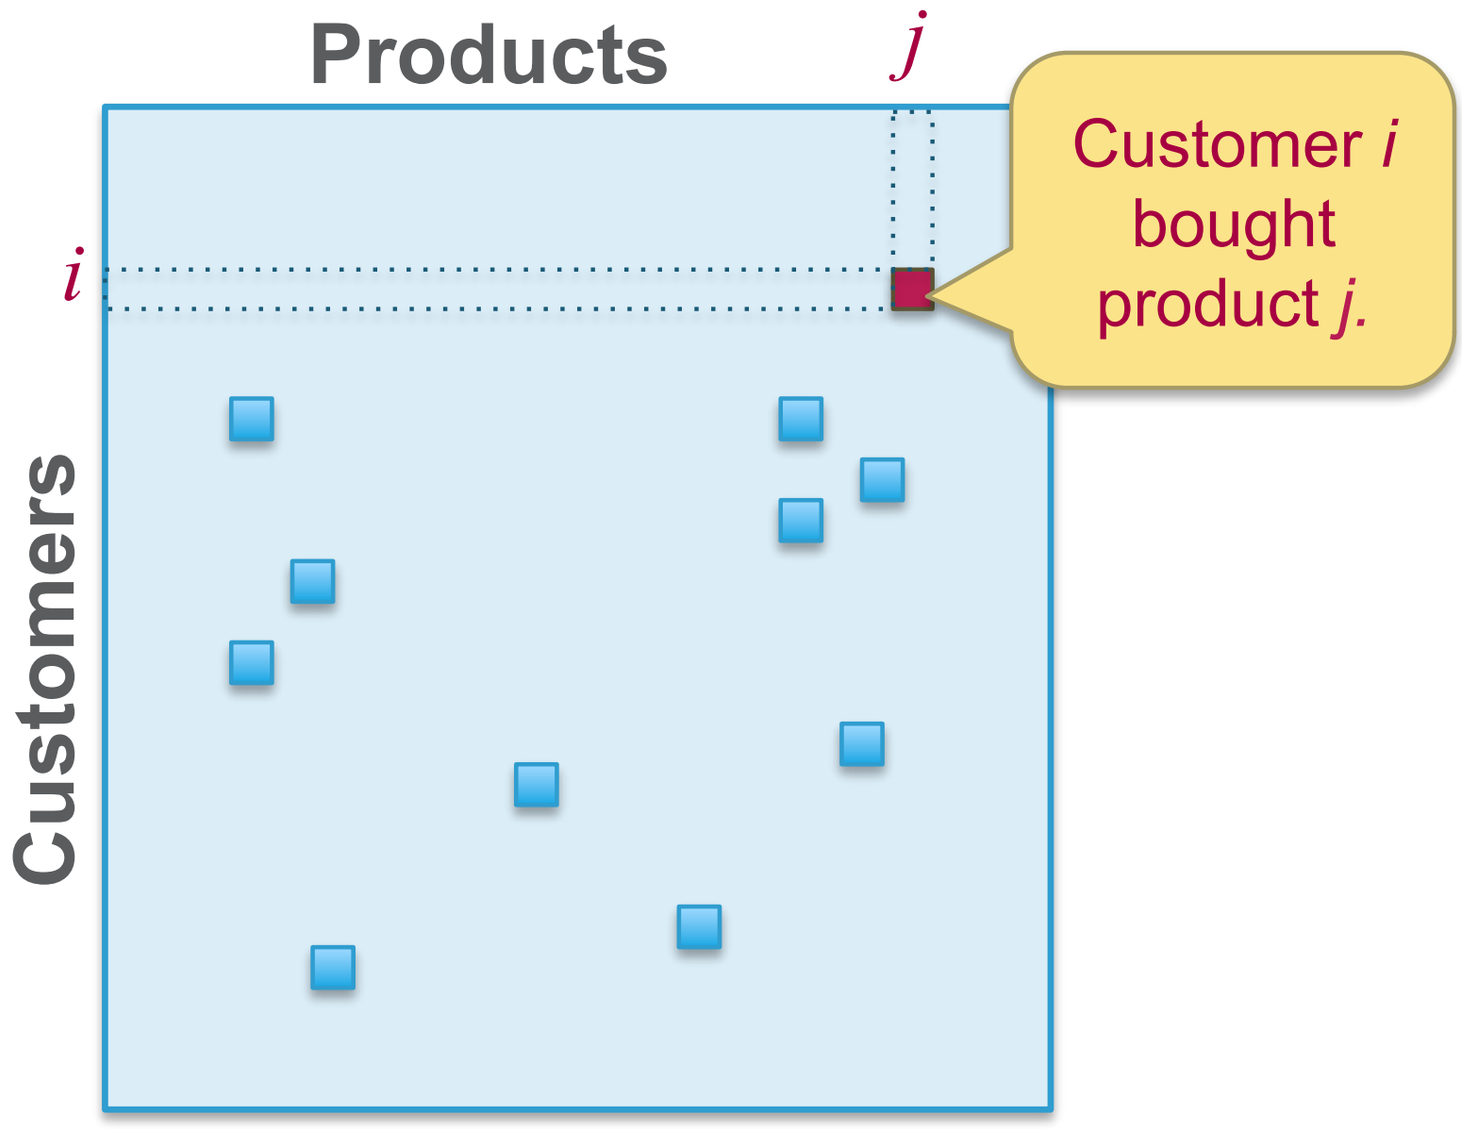

Products Factor

Customers Factor

p_u

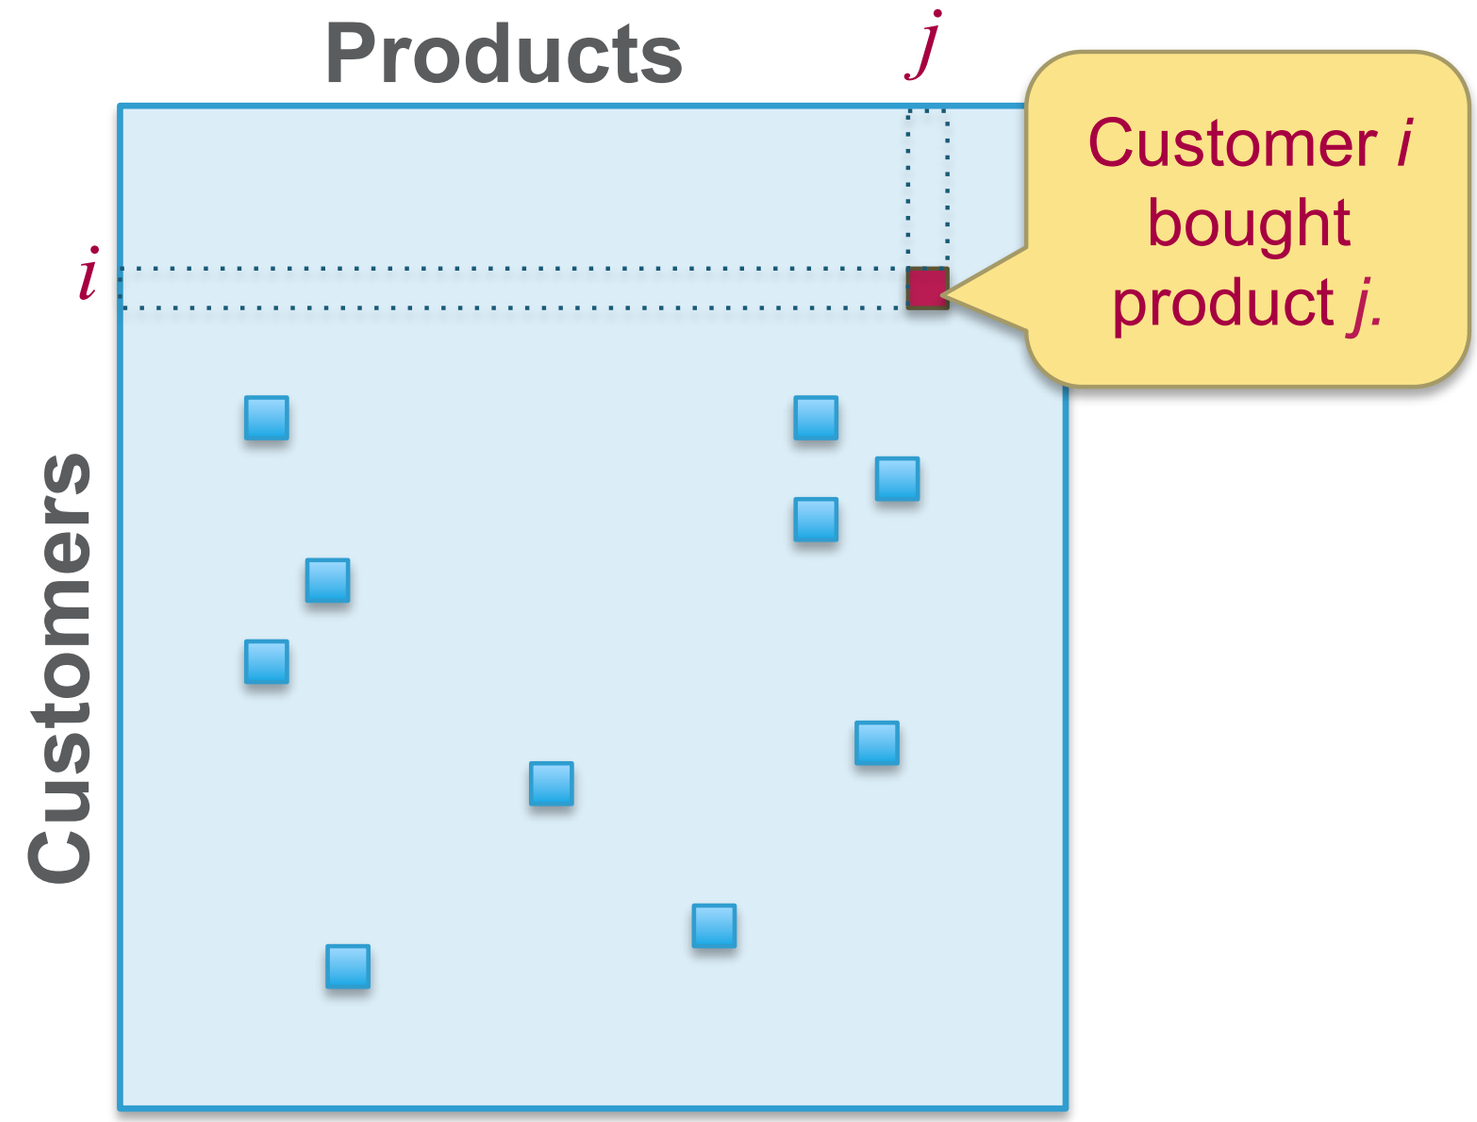

Products Factor

q_i

Customers Factor

p_u

Products

Customers

j

i

Customer i bought product j.

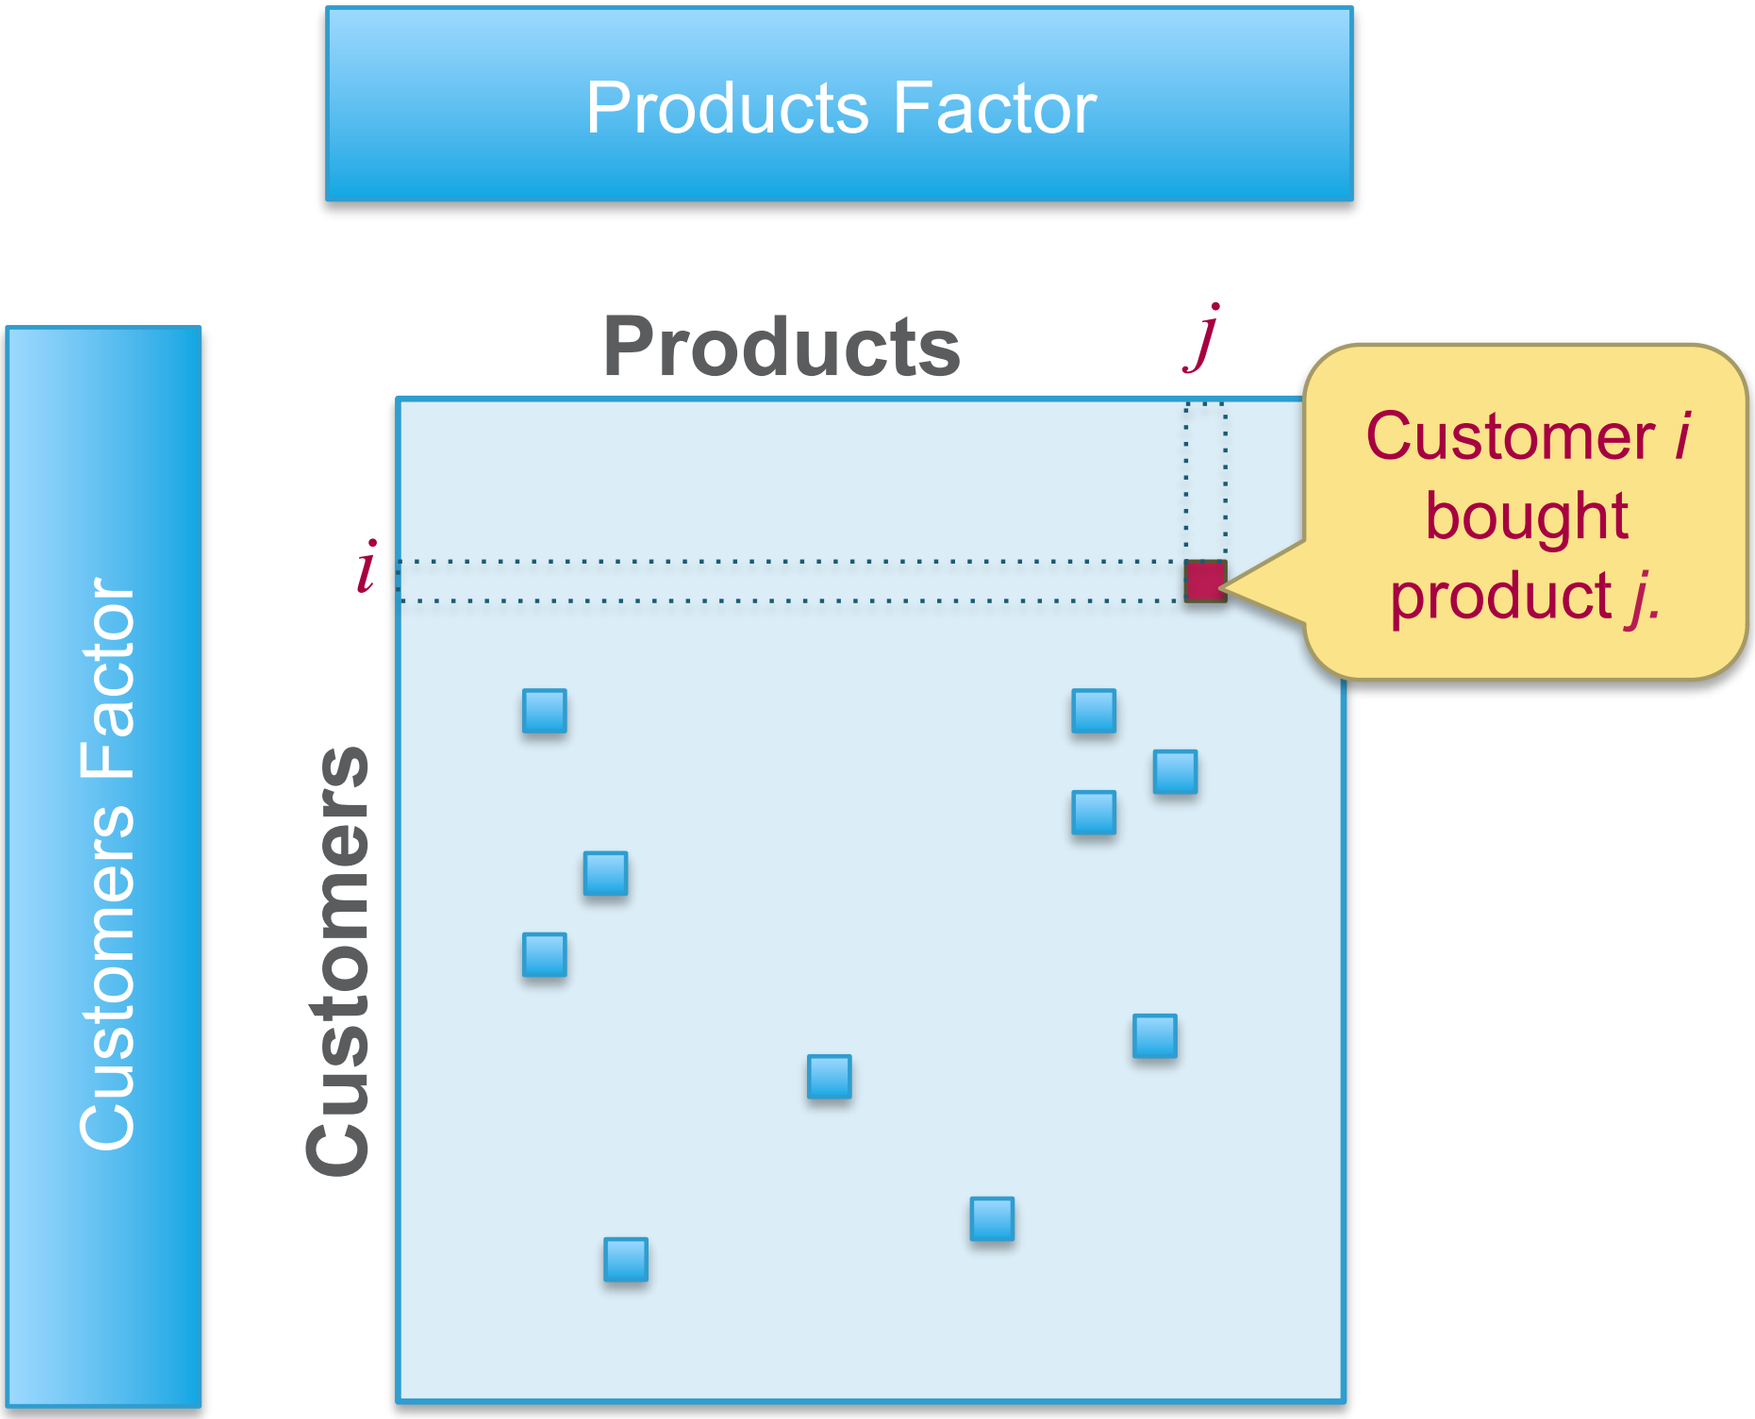

$$min_{q,p} \sum_{u,i} (r_{ui} - p_u^T q_i)^2$$

Multiply these two factors to produce a less-sparse matrix.

×

Products Factor

Customers Factor

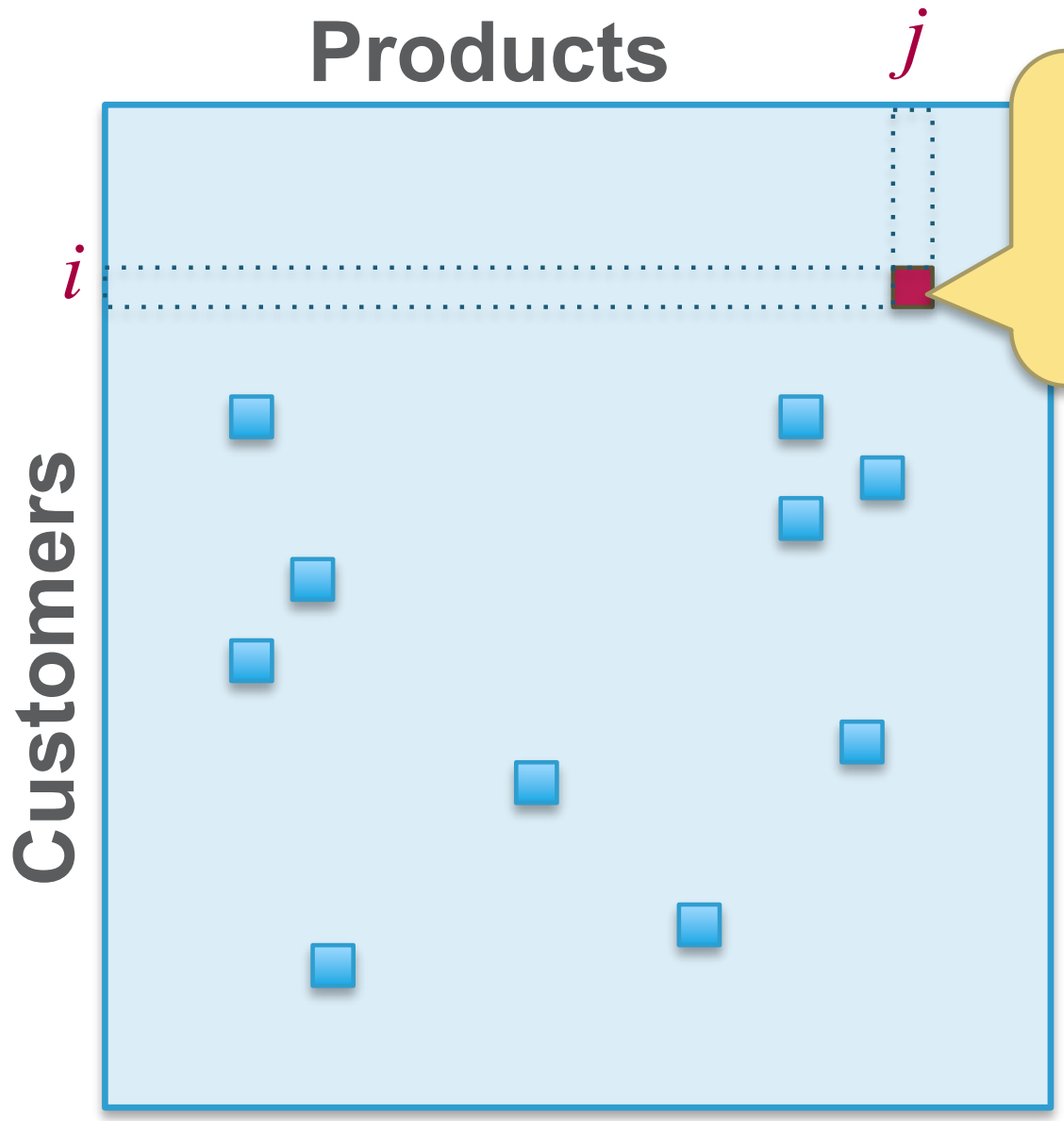

$$r'_{ui} = p_u^T q_i$$

$$min_{q,p} \sum_{u,i} (r_{ui} - p_u^T q_i)^2$$

Multiply these two factors to produce a less-sparse matrix.

×

Products Factor

Customers Factor

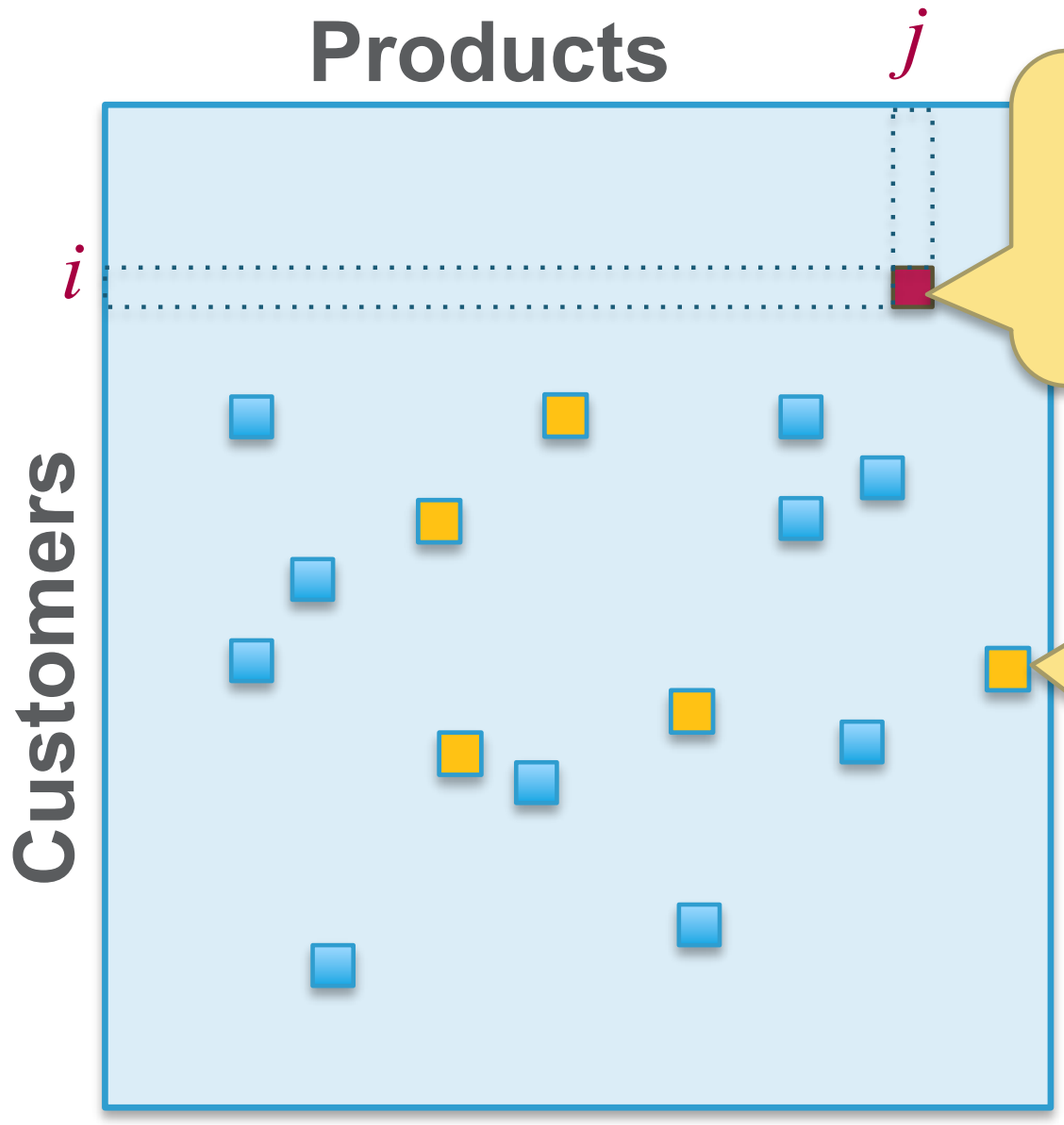

Customer *i* bought product *j*.

New nonzero values become product suggestions.

$$r'_{ui} = p_u^T q_i$$

```
val model = ALS.train(ratings, rank, numIterations, 0.01)
```

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

Every line has a clear purpose!

https://github.com/apache/spark/blob/master/mllib/src/main/scala/org/apache/spark/mllib/recommendation/ALS.scala

https://github.com/apache/spark/blob/master/mllib/src/main/scala/org/apache/spark/mllib/recommendation/ALS.scala

25 lines' worth of algorithm…

…mixed with 800 lines of performance code

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

```
U = rand(nrow(X), r, min = -1.0, max = 1.0);
V = rand(r, ncol(X), min = -1.0, max = 1.0);
while(i < mi) {
   i = i + 1; ii = 1;
   if (is_U)
      G = (W * (U %*% V - X)) %*% t(V) + lambda * U;
   else
      G = t(U) %*% (W * (U %*% V - X)) + lambda * V;
   norm_G2 = sum(G ^ 2); norm_R2 = norm_G2;
   R = -G; S = R;
   while(norm_R2 > 10E-9 * norm_G2 & ii <= mii) {
     if (is_U) {
       HS = (W * (S %*% V)) %*% t(V) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       U = U + alpha * S;
     } else {
       HS = t(U) %*% (W * (U %*% S)) + lambda * S;
       alpha = norm_R2 / sum (S * HS);
       V = V + alpha * S;
     }
     R = R - alpha * HS;
     old_norm_R2 = norm_R2; norm_R2 = sum(R ^ 2);
     S = R + (norm_R2 / old_norm_R2) * S;
     ii = ii + 1;
   }
   is_U = ! is_U;
}
```

SystemML:
compile and run at scale
no performance code needed!

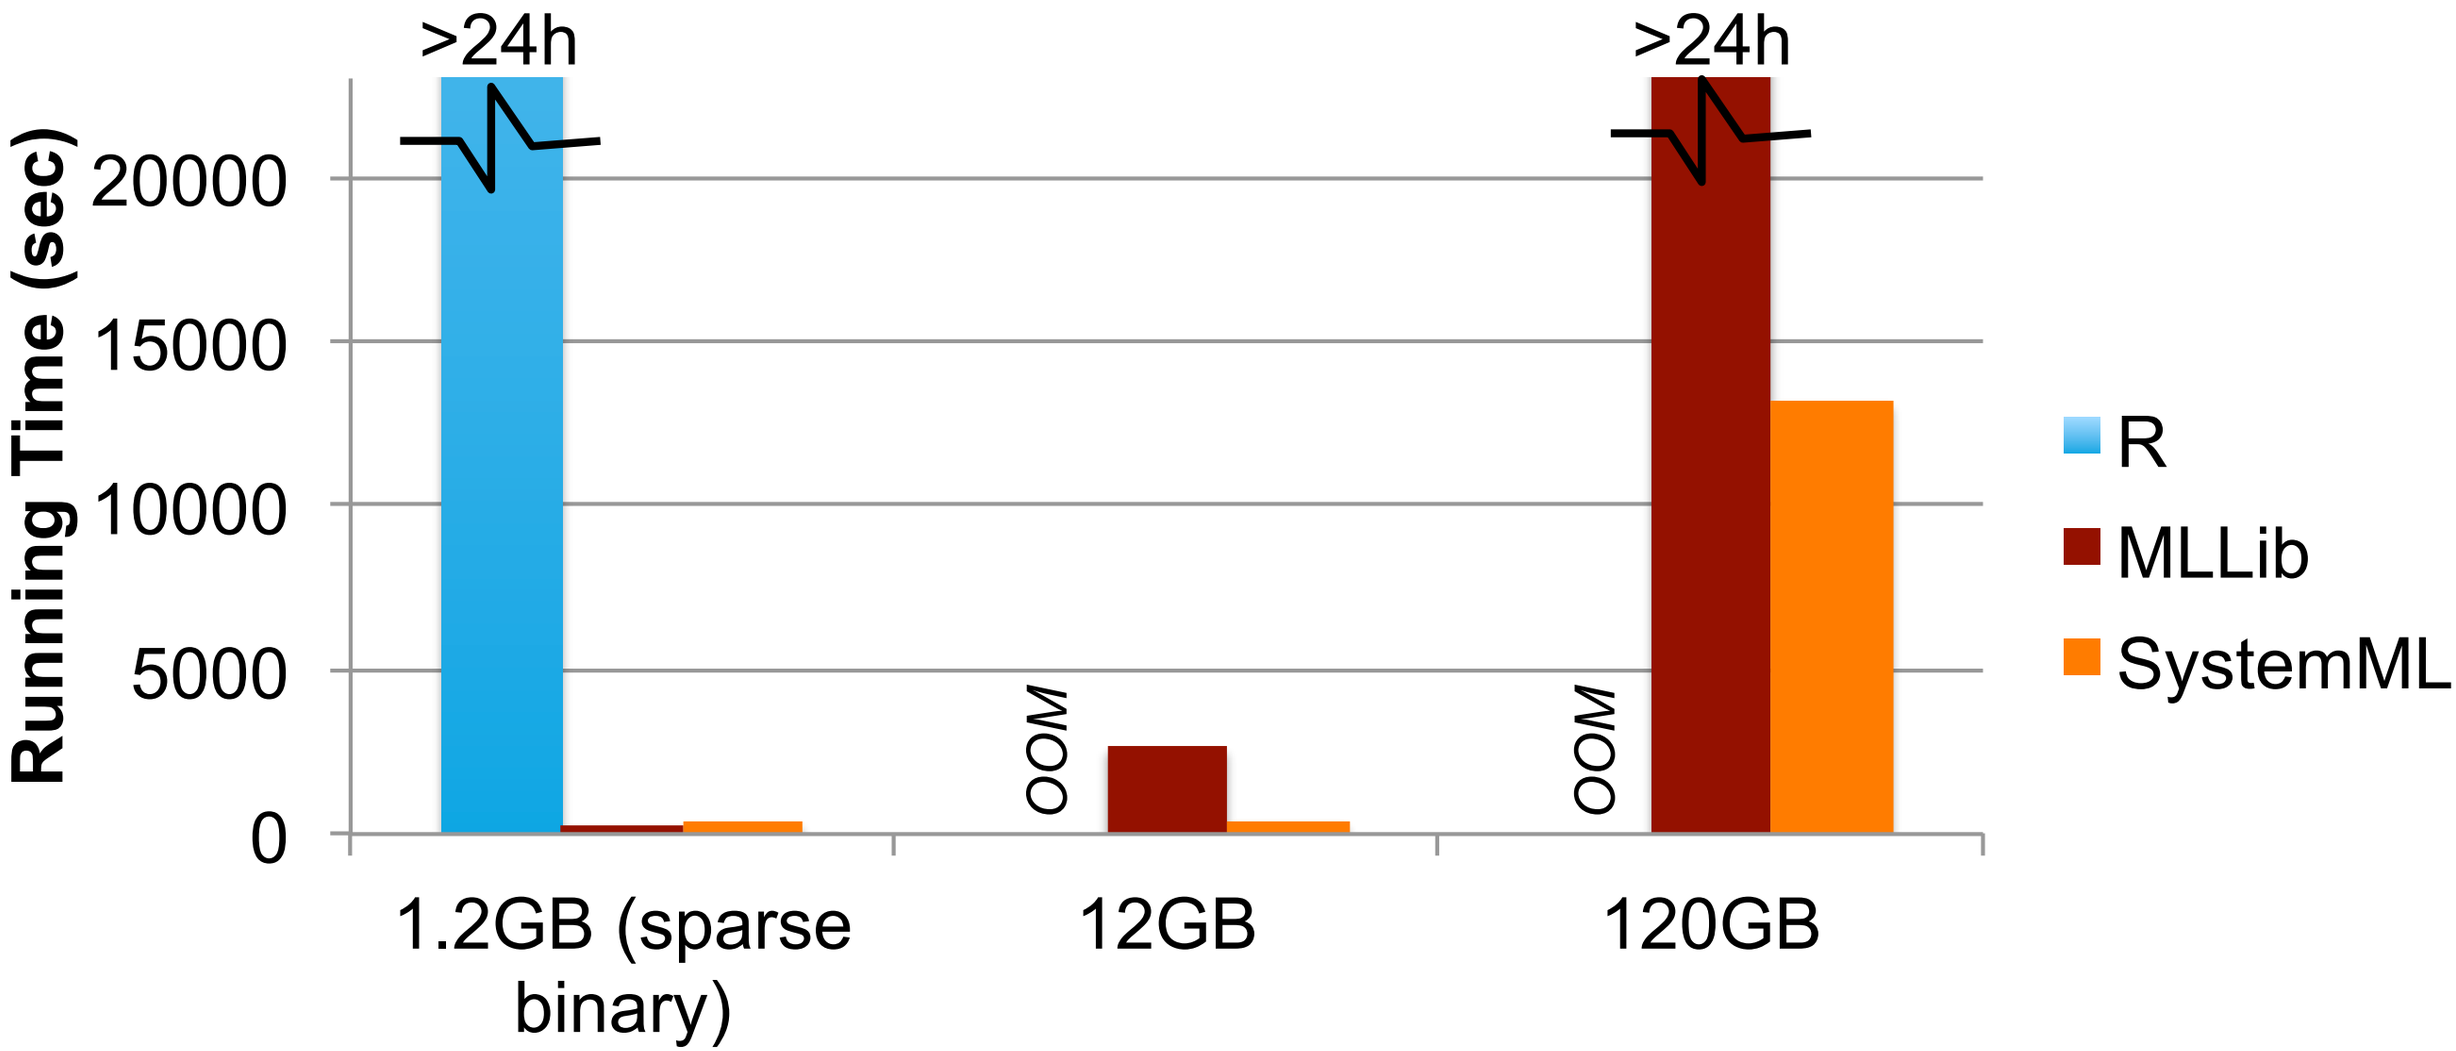
>24h
>24h
Running Time (sec)
20000
15000
10000
5000
0
OOM
OOM
1.2GB (sparse binary)
12GB
120GB
R
MLLib
SystemML

Architecture

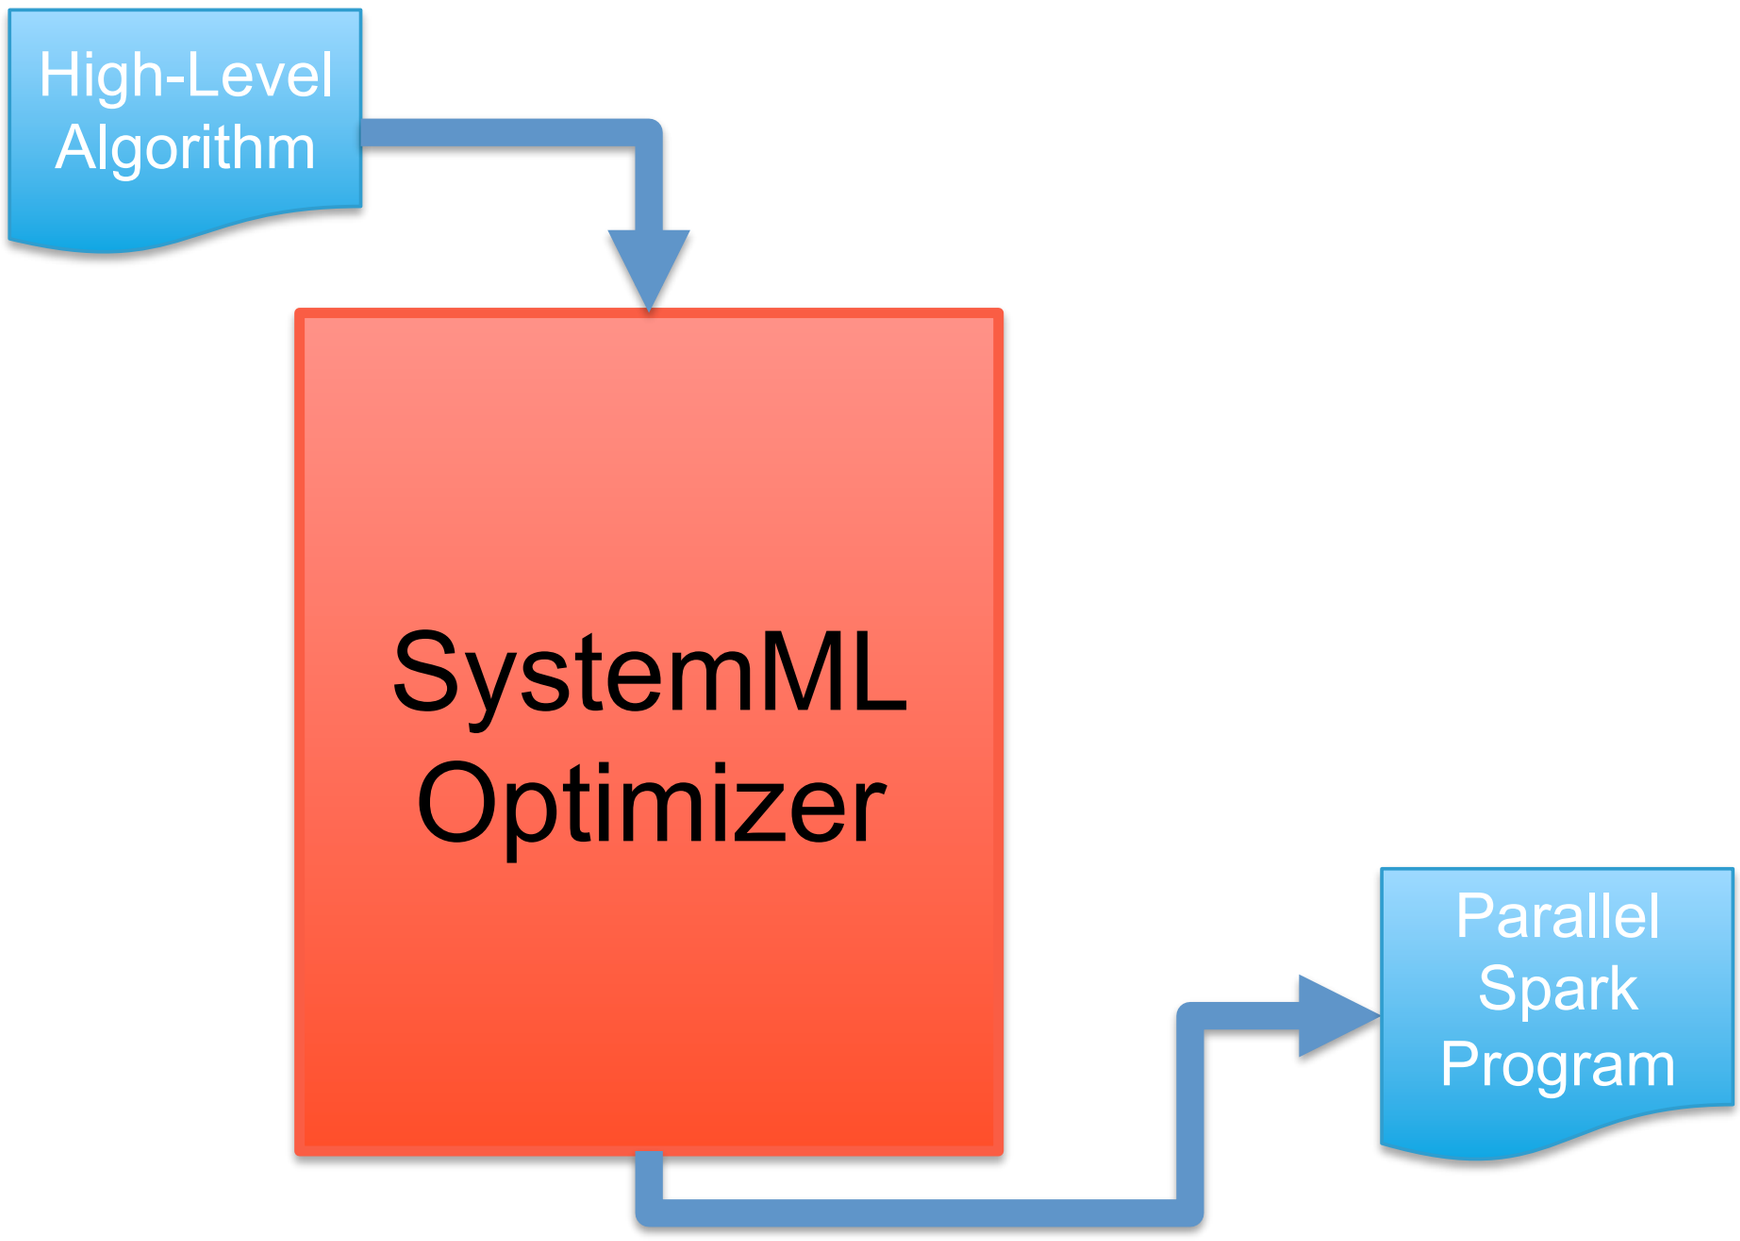

Architecture

High-level language front-ends

Cost Based Optimizer

High-Level Operations (HOPs)

General representation of statements in the data analysis language

Low-Level Operations (LOPs)

General representation of operations in the runtime framework

Multiple execution environments

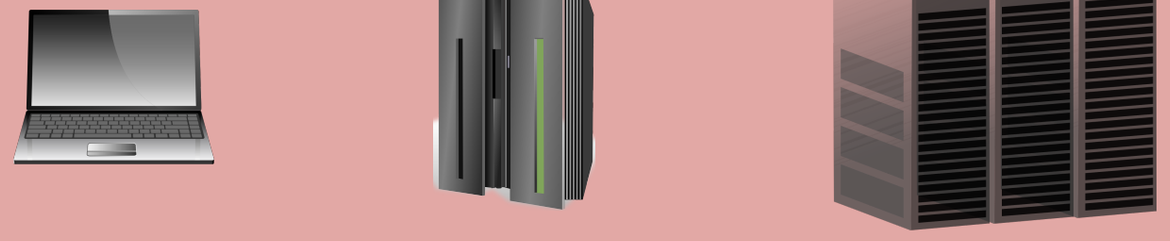

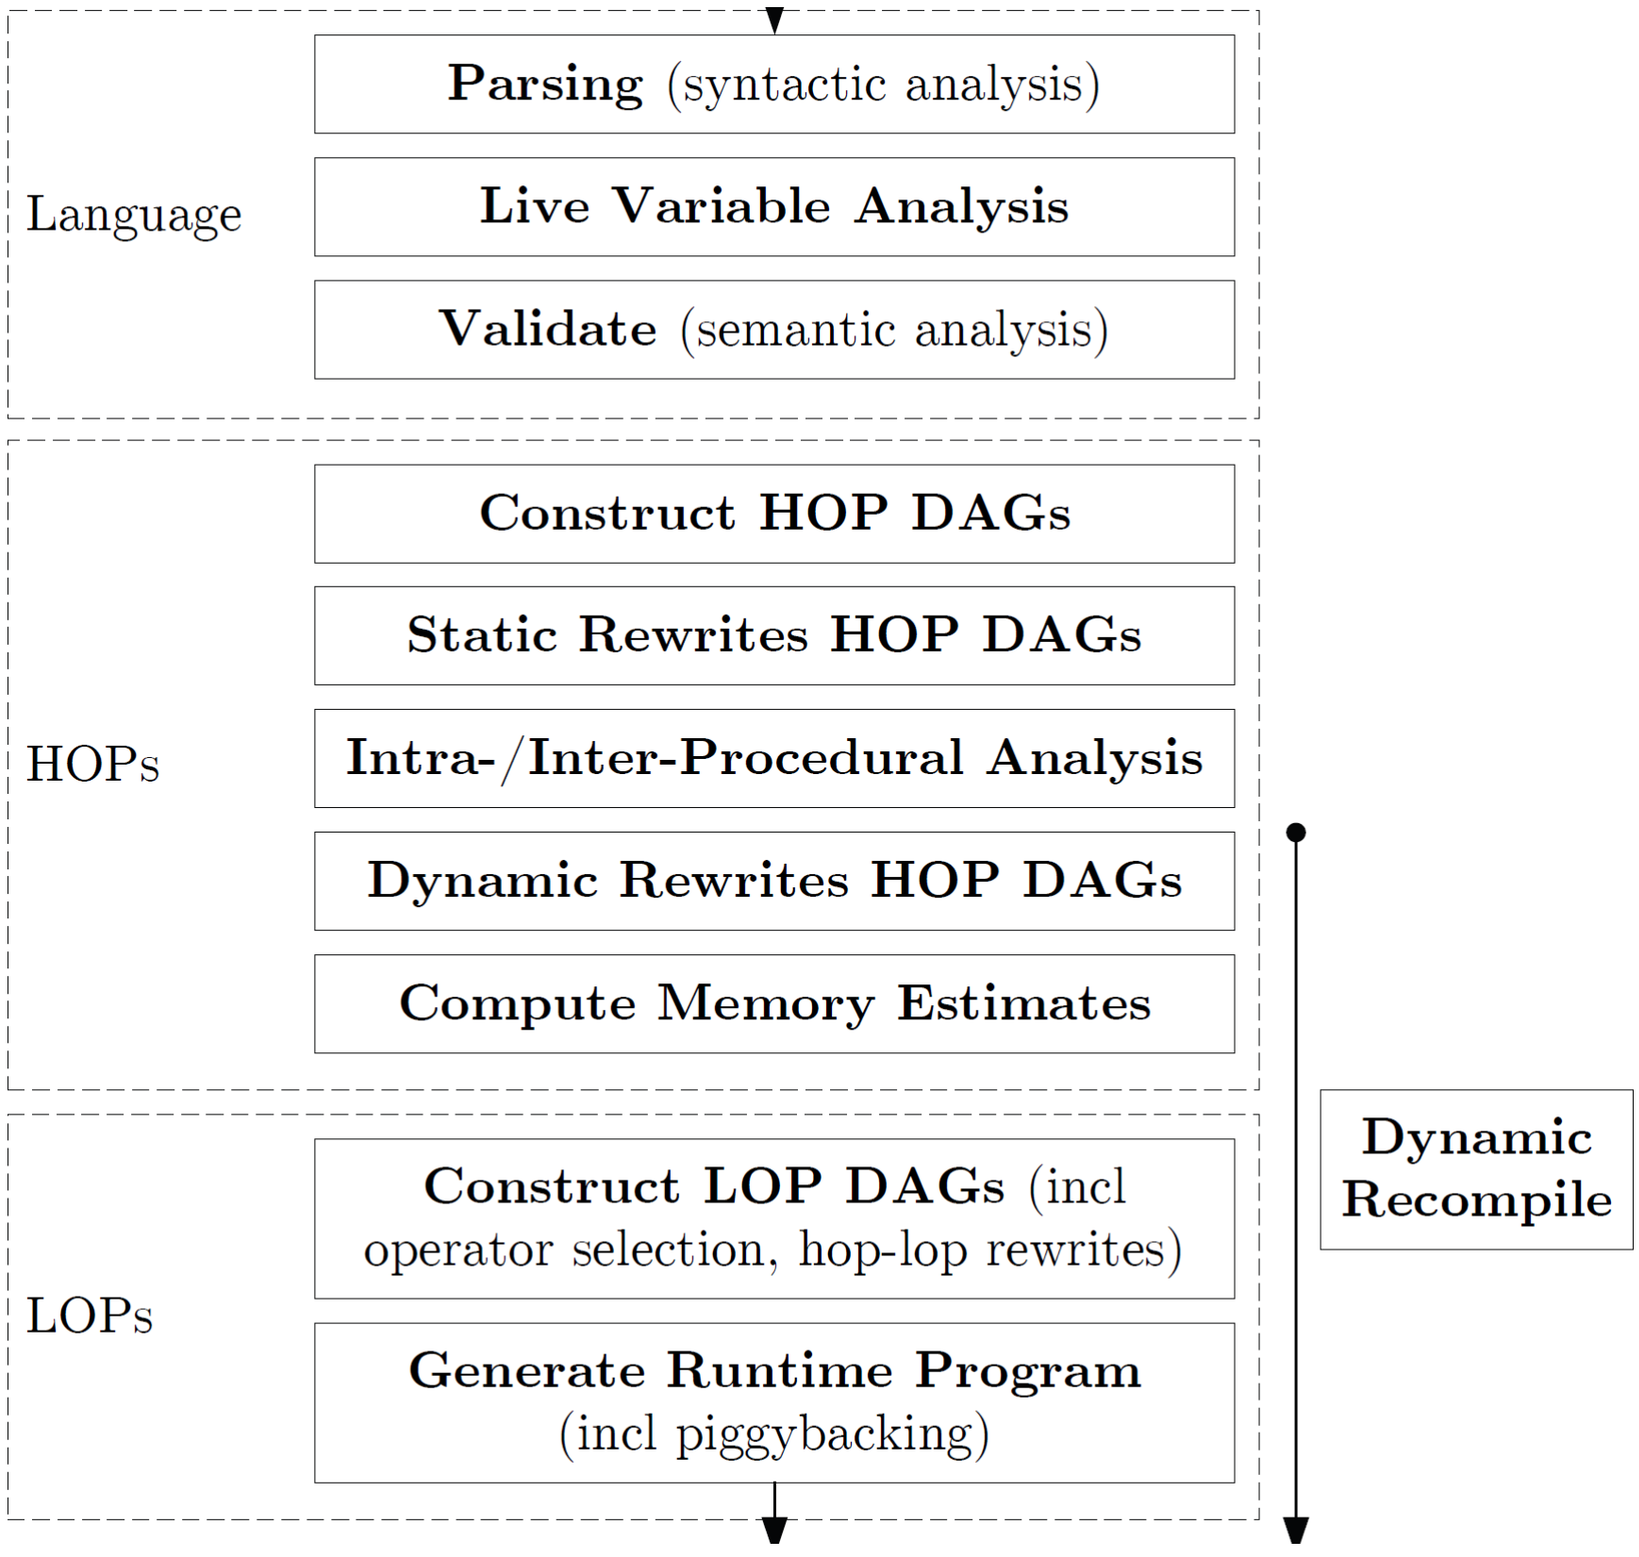
Language
Parsing (syntactic analysis)
Live Variable Analysis
Validate (semantic analysis)
HOPs
Construct HOP DAGs
Static Rewrites HOP DAGs
Intra-/Inter-Procedural Analysis
Dynamic Rewrites HOP DAGs
Compute Memory Estimates
LOPs
Construct LOP DAGs (incl operator selection, hop-lop rewrites)
Generate Runtime Program (incl piggybacking)
Dynamic Recompile

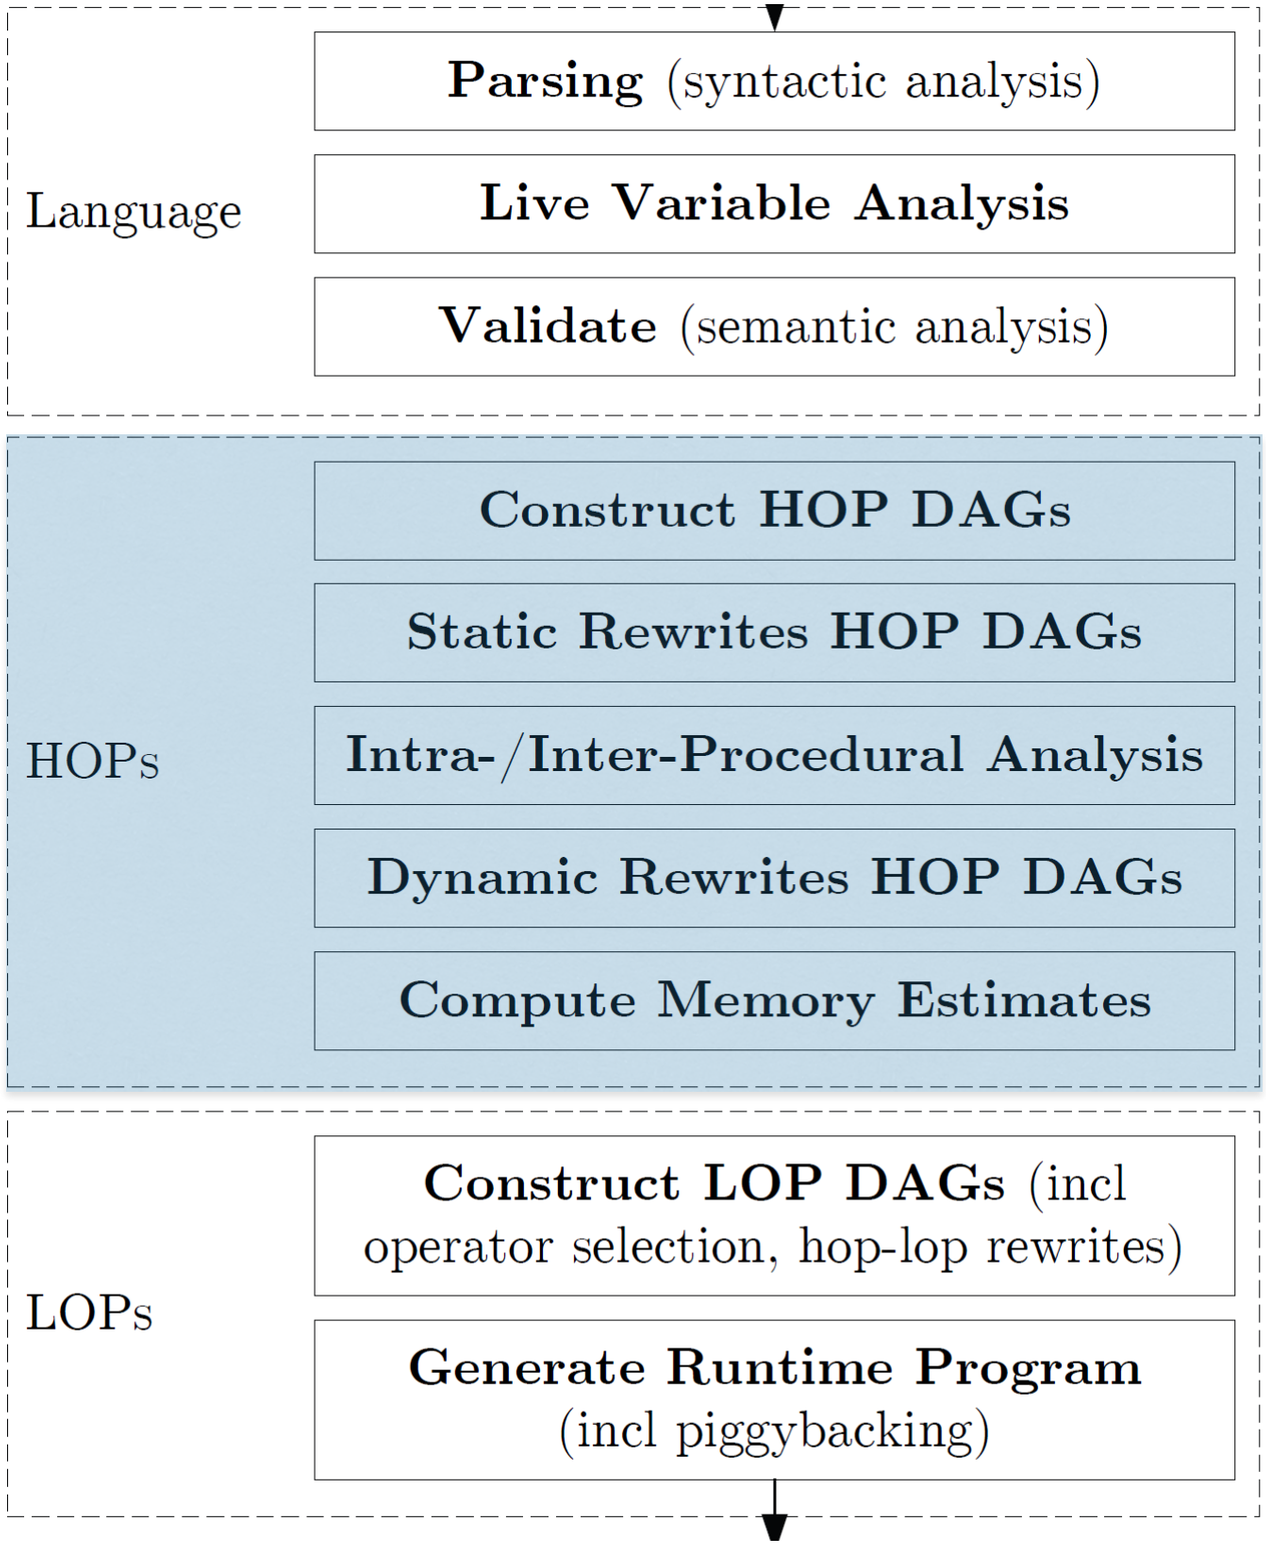

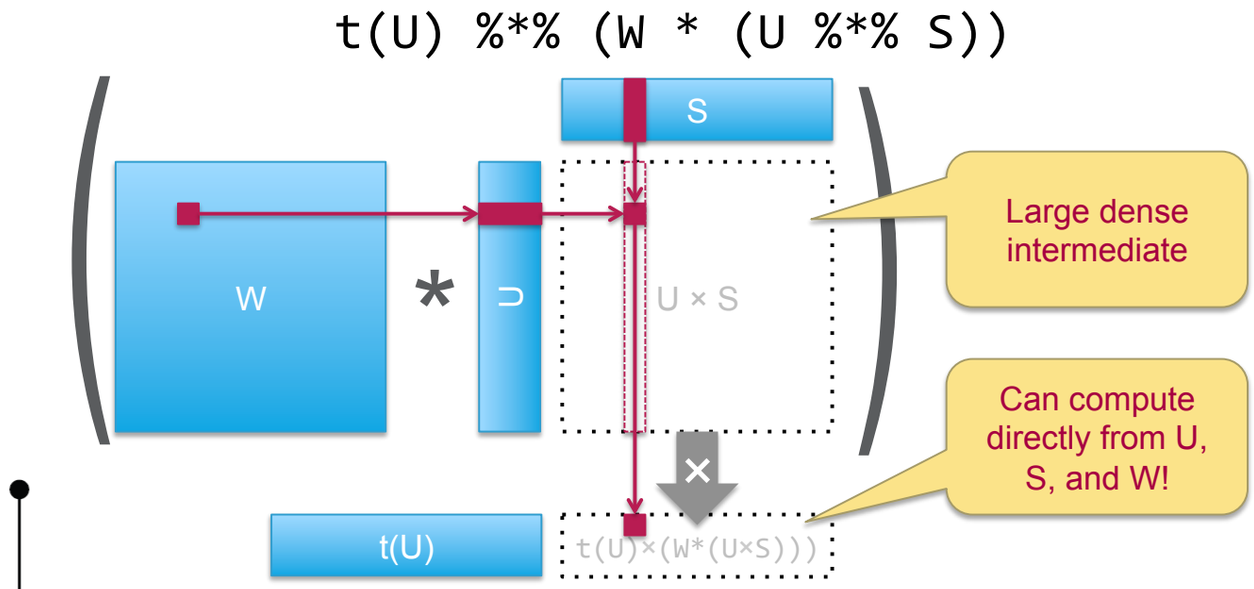

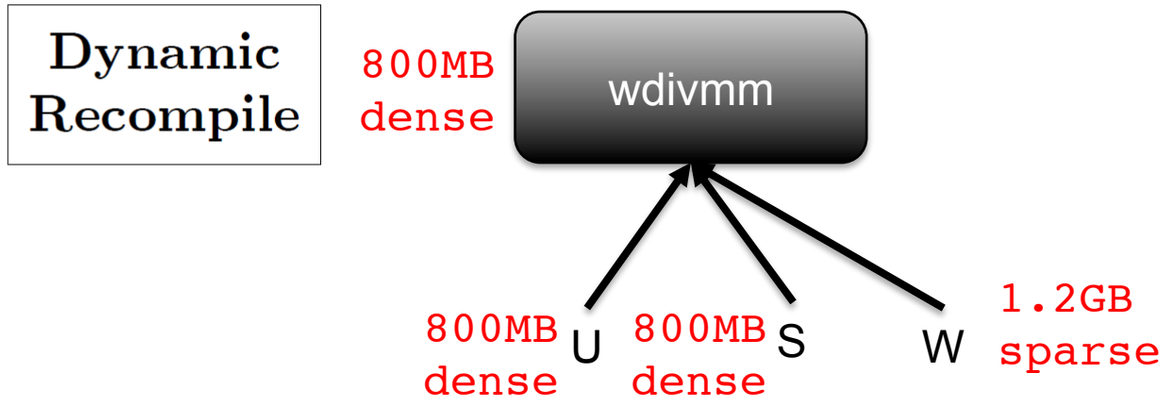

(weighted divide matrix multiplication)

Language

Parsing (syntactic analysis)

Live Variable Analysis

Validate (semantic analysis)

HOPs

Construct HOP DAGs

Static Rewrites HOP DAGs

Intra-/Inter-Procedural Analysis

Dynamic Rewrites HOP DAGs

Compute Memory Estimates

LOPs

Construct LOP DAGs (incl operator selection, hop-lop rewrites)

Generate Runtime Program (incl piggybacking)

Dynamic Recompile

All operands fit into heap → **use one node**

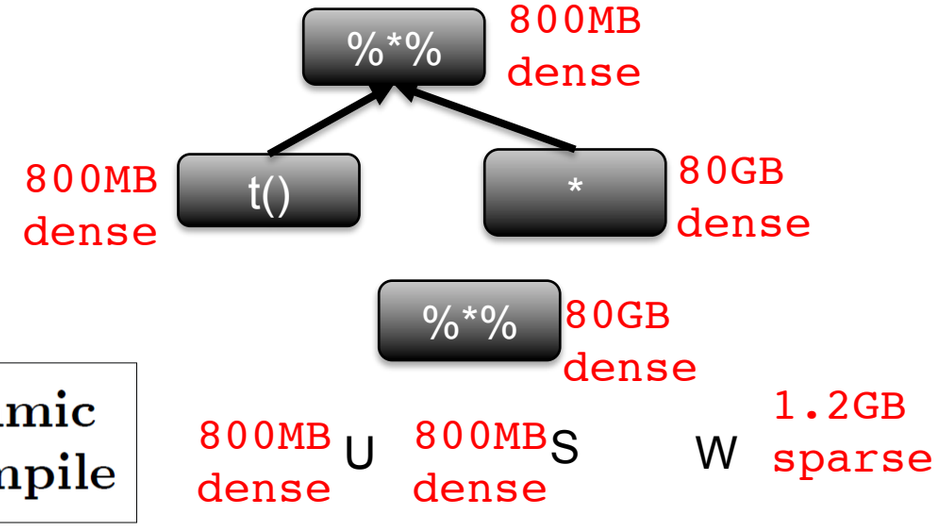

WDivMM

(MapWDivMM)

Apache SystemML

Community GitHub Documentation Download Apache

Apache SystemML

Apache SystemML is a distributed and declarative machine learning platform.

Get SystemML

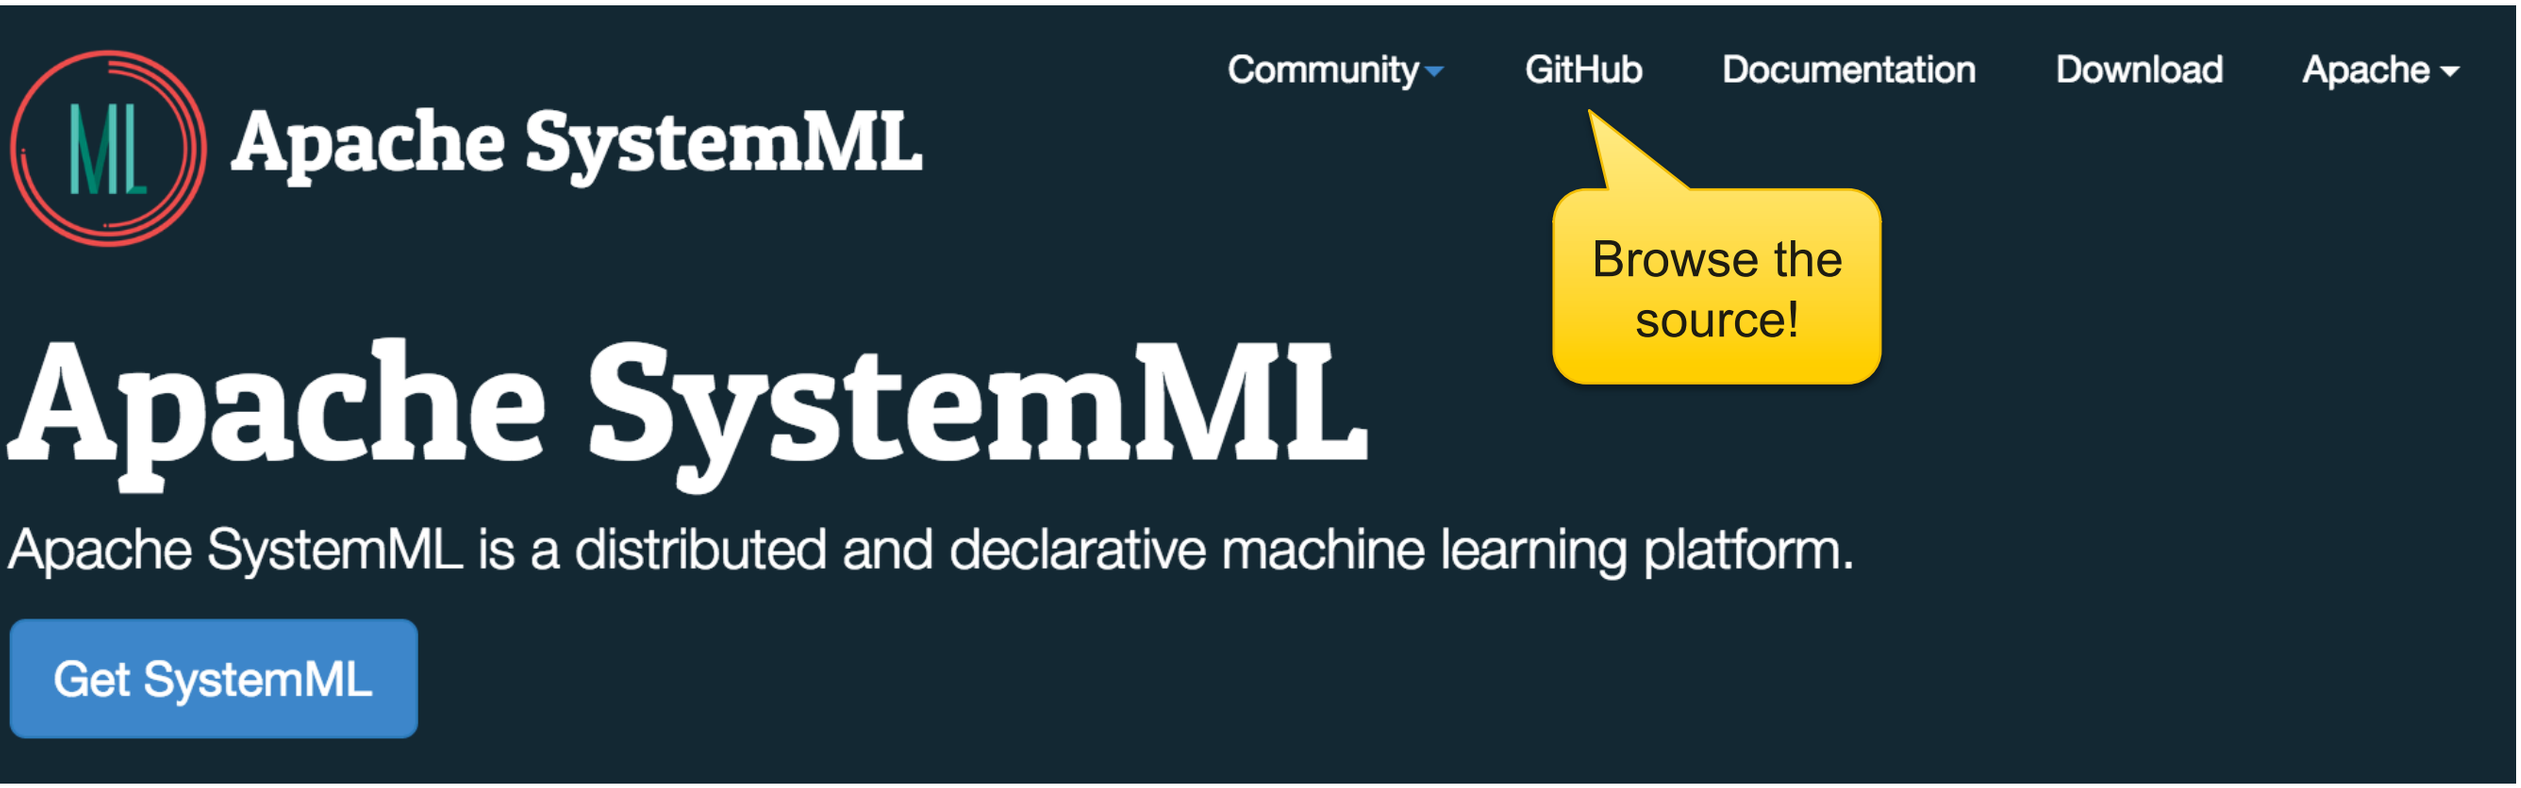
Apache SystemML
Community
GitHub
Documentation
Download
Apache
Browse the source!
Apache SystemML
Apache SystemML is a distributed and declarative machine learning platform.
Get SystemML

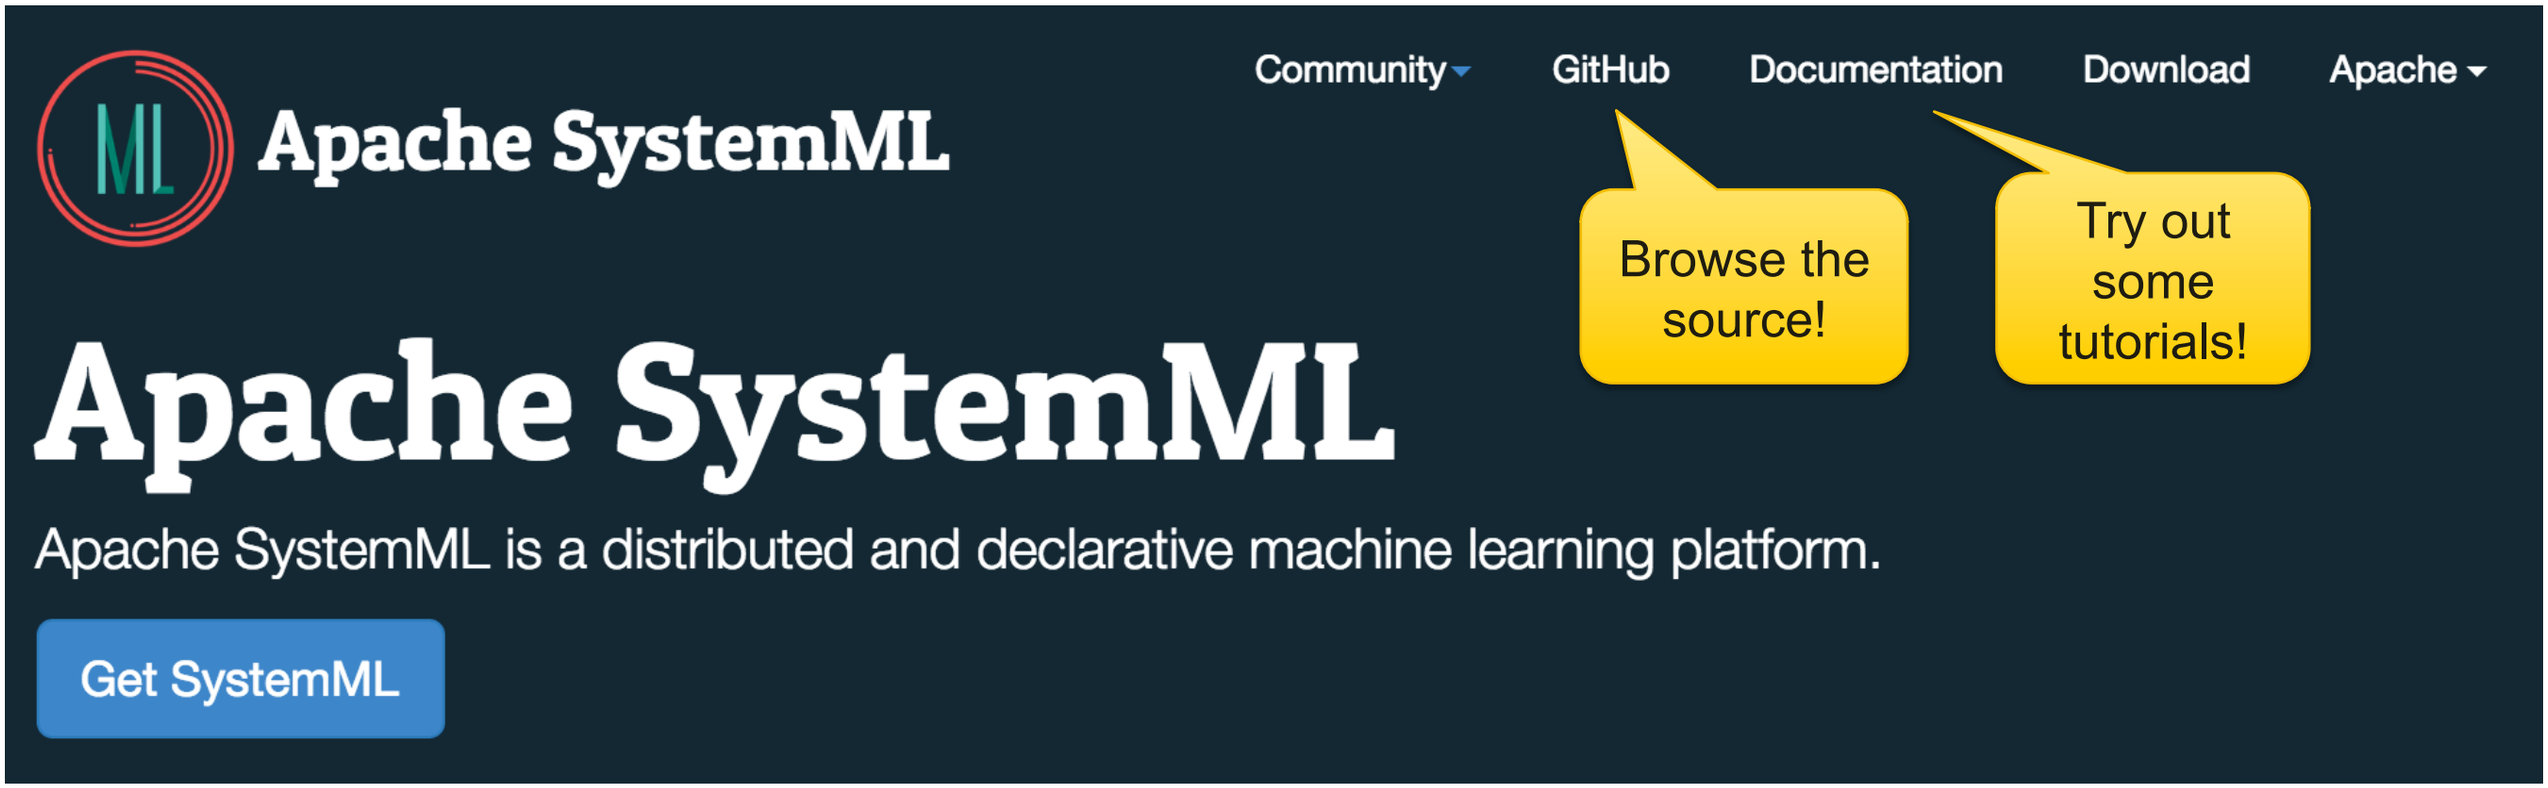
Apache SystemML
Community
GitHub
Documentation
Download
Apache
Browse the source!
Try out some tutorials!
Apache SystemML
Apache SystemML is a distributed and declarative machine learning platform.
Get SystemML

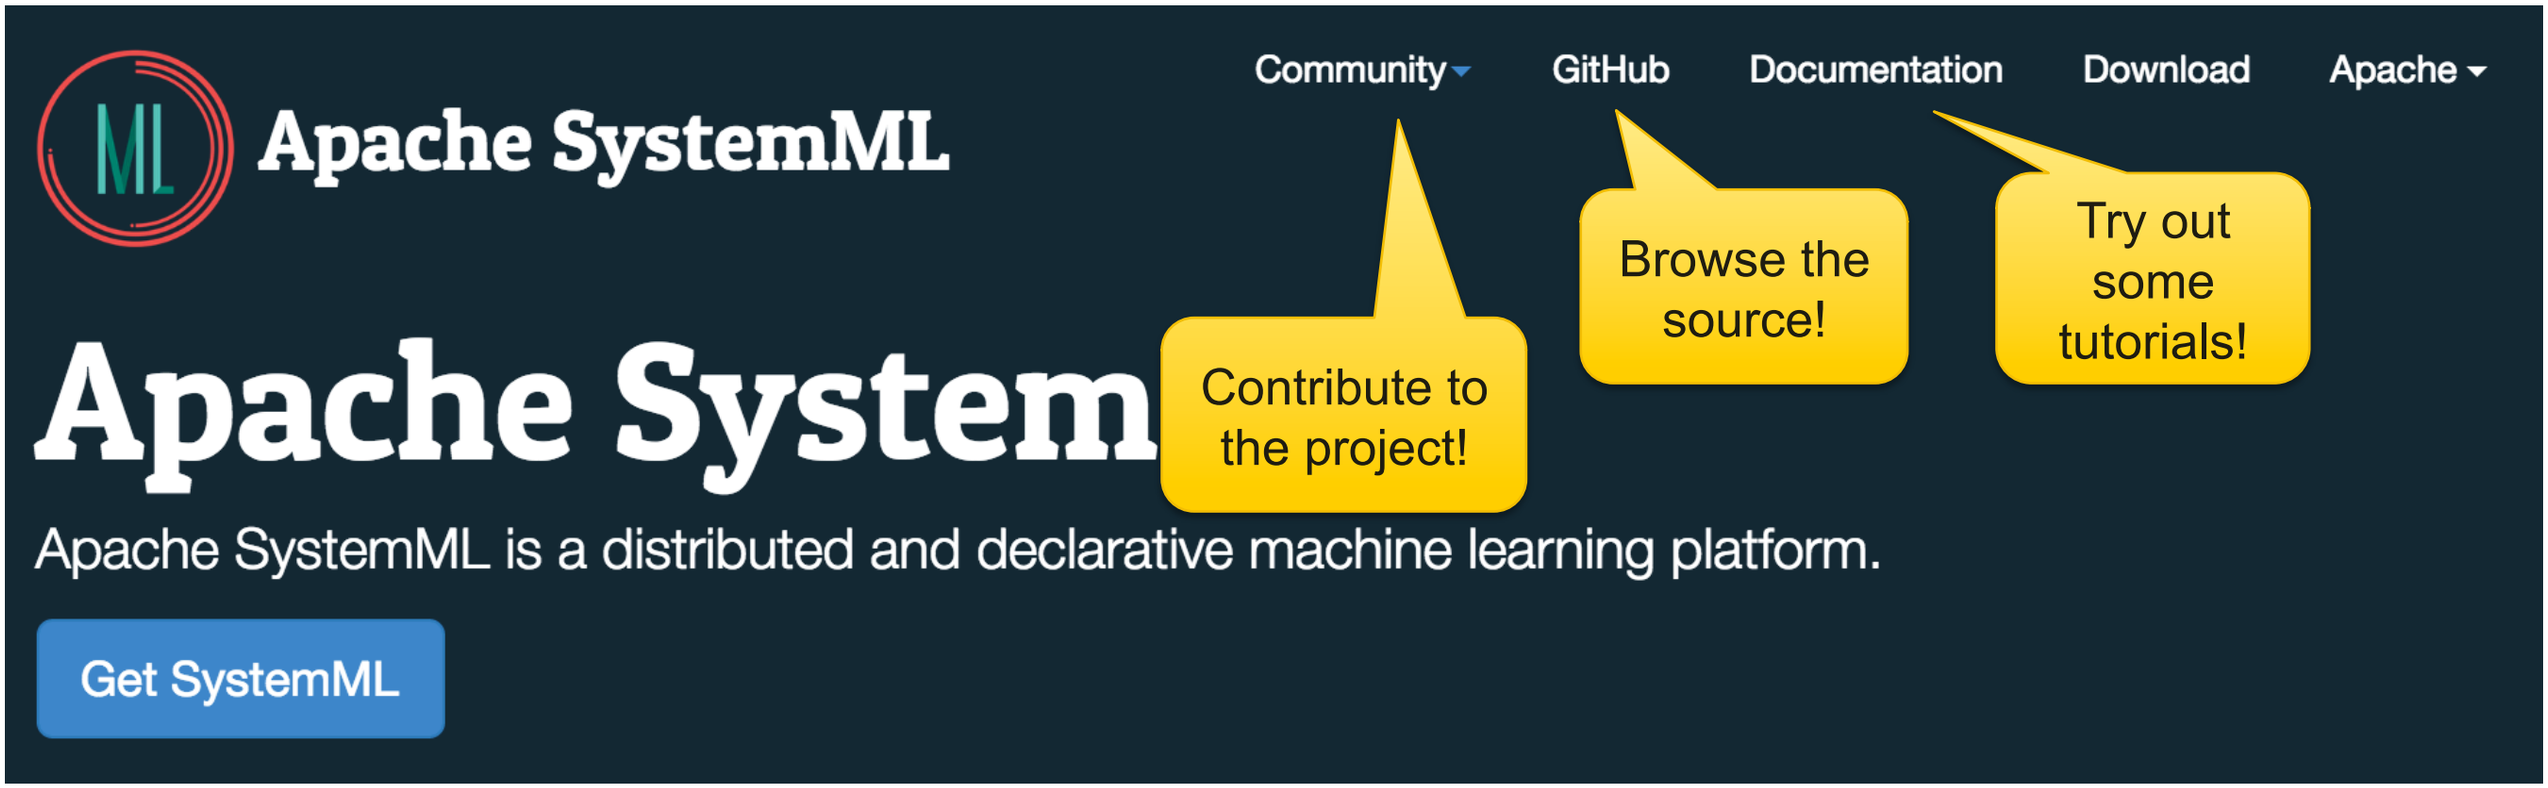
Apache SystemML
Community
GitHub
Documentation
Download
Apache
Apache System
Contribute to the project!
Browse the source!
Try out some tutorials!
Apache SystemML is a distributed and declarative machine learning platform.
Get SystemML

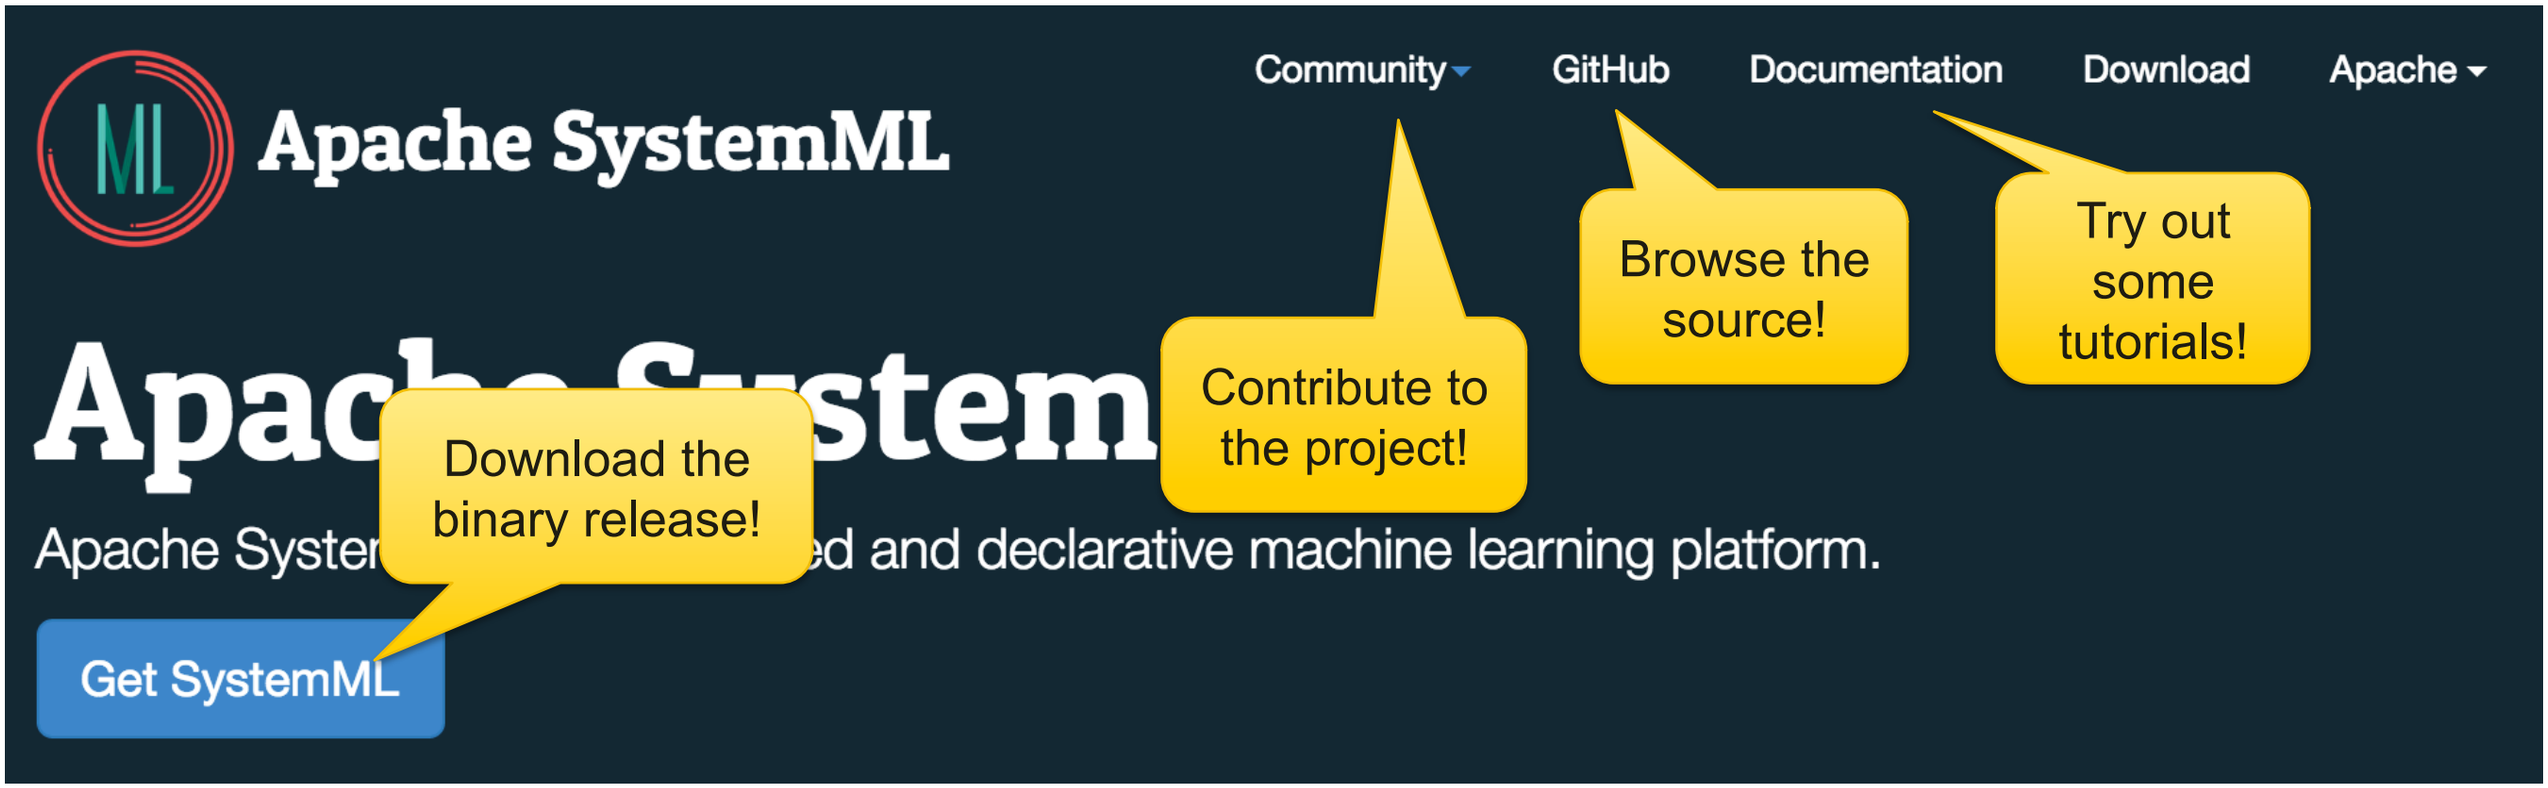
Apache SystemML
Community
GitHub
Documentation
Download
Apache
Apac stem
Apache Syste ed and declarative machine learning platform.
Get SystemML
Download the binary release!
Contribute to the project!
Browse the source!
Try out some tutorials!

Demo