



GraphLab
Create™

Speeding up Machine Learning on GPU

Danny Bickson
Co-Founder

GraphLab Project History

GraphLab
(2009)

GraphChi
(2011)

GraphLab
Create
(2014)

Carnegie Mellon



GraphLab Conferences

2012



2013



2014



Saturday, October 11, 2014

News from Recsys - GraphLab Create based solution wins second place

Yesterday I met Robi Palovicz at ACM Recsys. Robi is a PhD from The Hungarian Academy of Sciences, supervised by my friend and Colleague [Andras Benczur](#). Robi was leading the team who won 2nd place at the [Recsys challenge](#) 2014.




His solution was based on Graphlab Create, and more interesting, it takes only 30 seconds to compute the winning solution with GraphLab Create!

Robi has promised to share their winning solution using an Ipython Notebook. Once it is ready



The algorithms are reported in the following

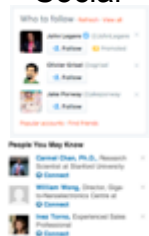
Data-driven, *predictive data apps* are making our world amazing...



Recommenders



Social



Industrial Apps



Pers. Medicine



Forecasters



Fraud & Anomaly Detection

Human Sensing



Sentiment Analysis & other Text Apps



Was using 217 business rules
hoping world doesn't change

Have an inspiring idea to
reinvent their business

Building a predictive app

Key pains:

Noisy Space of Tools



Data scientists use a **variety of tools**, across
different programming languages...
require a lot of **context-switching**...
affects productivity and impedes reproducibility.

Ben Lorica,
Data Analysis: Just one component of
the Data Science workflow

Hiring Talent

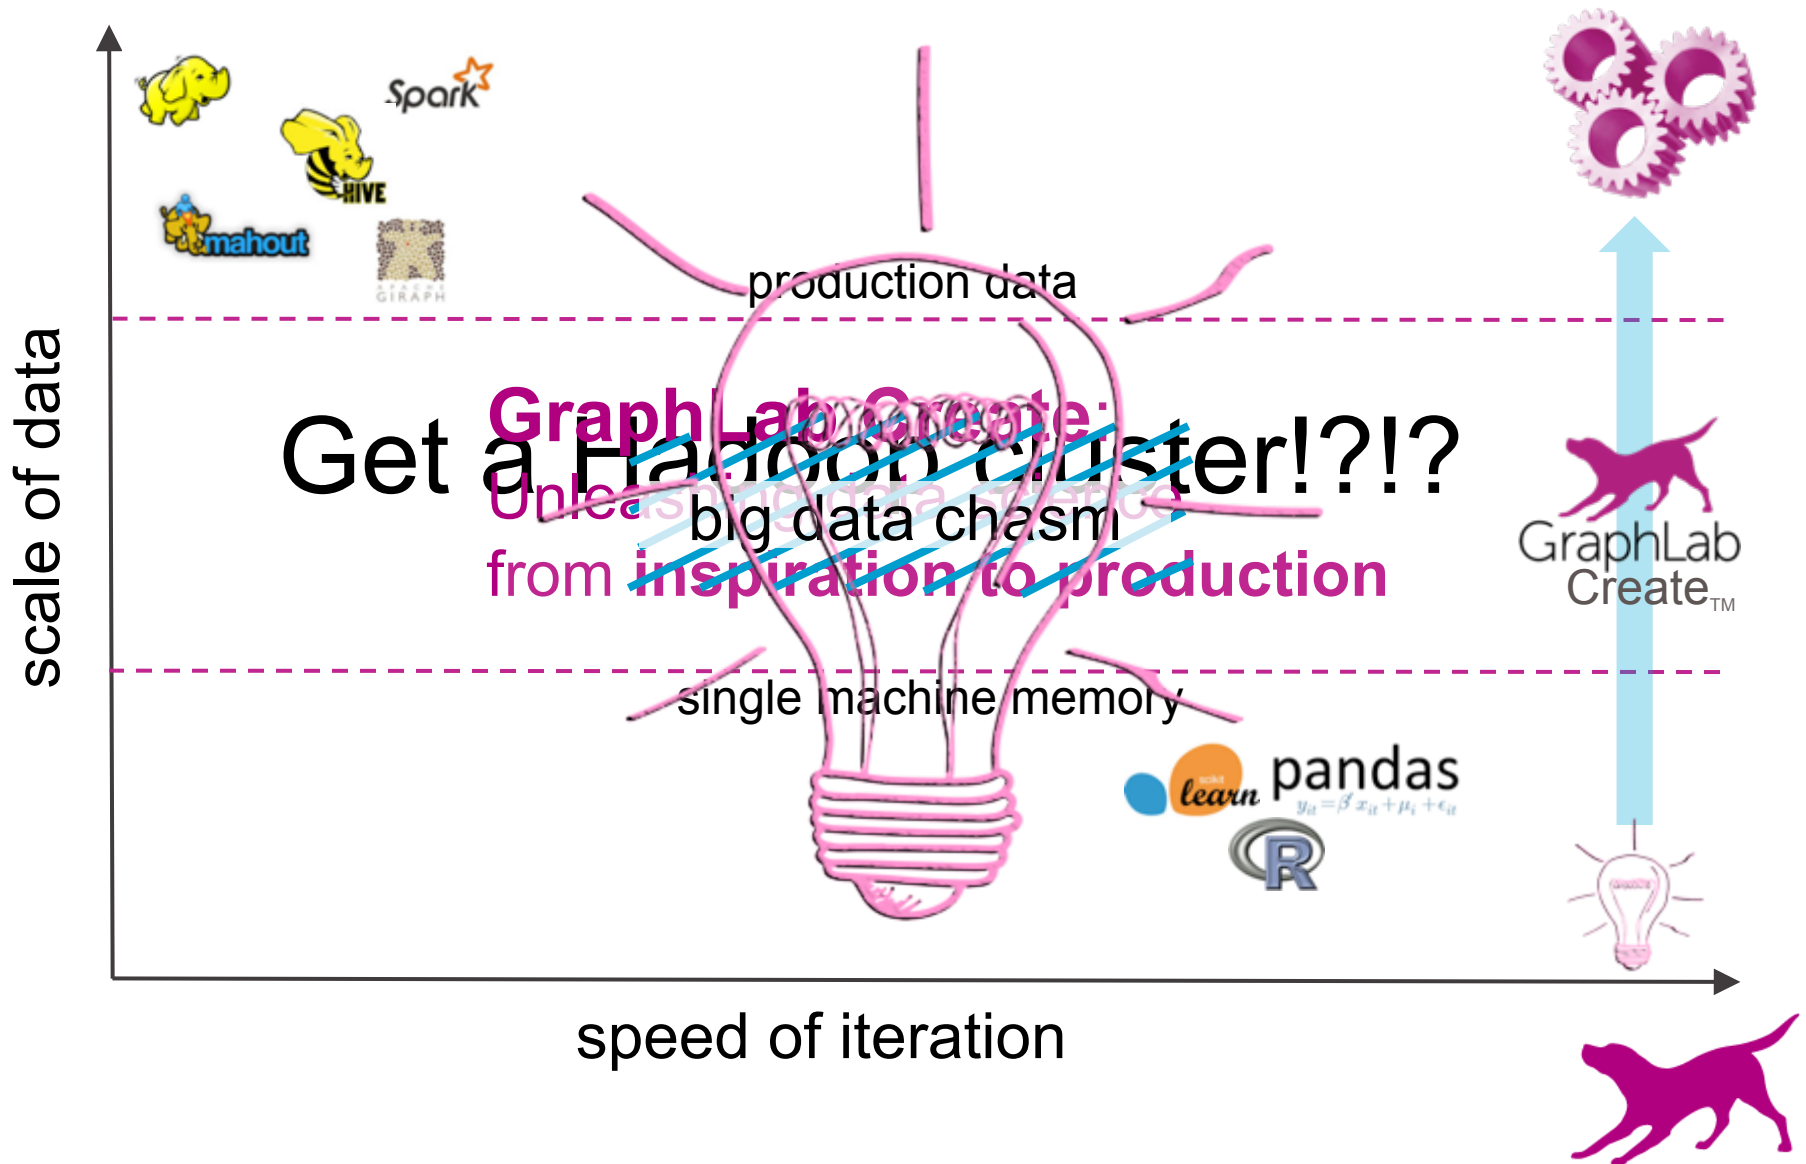


35%

Shortfall in data-savvy workers
needed to make sense out of
big data by 2018 [McKinsey 2011]





Crossing the Big Data Chasm

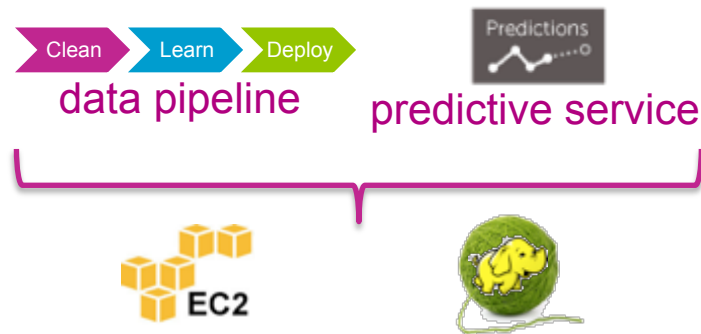


Data scientist: inspiration to production


Prototype

-  Use my laptop
- Variety of data
- Not toy data scales
-  Language I love
- Iterate quickly

Production



Monitor



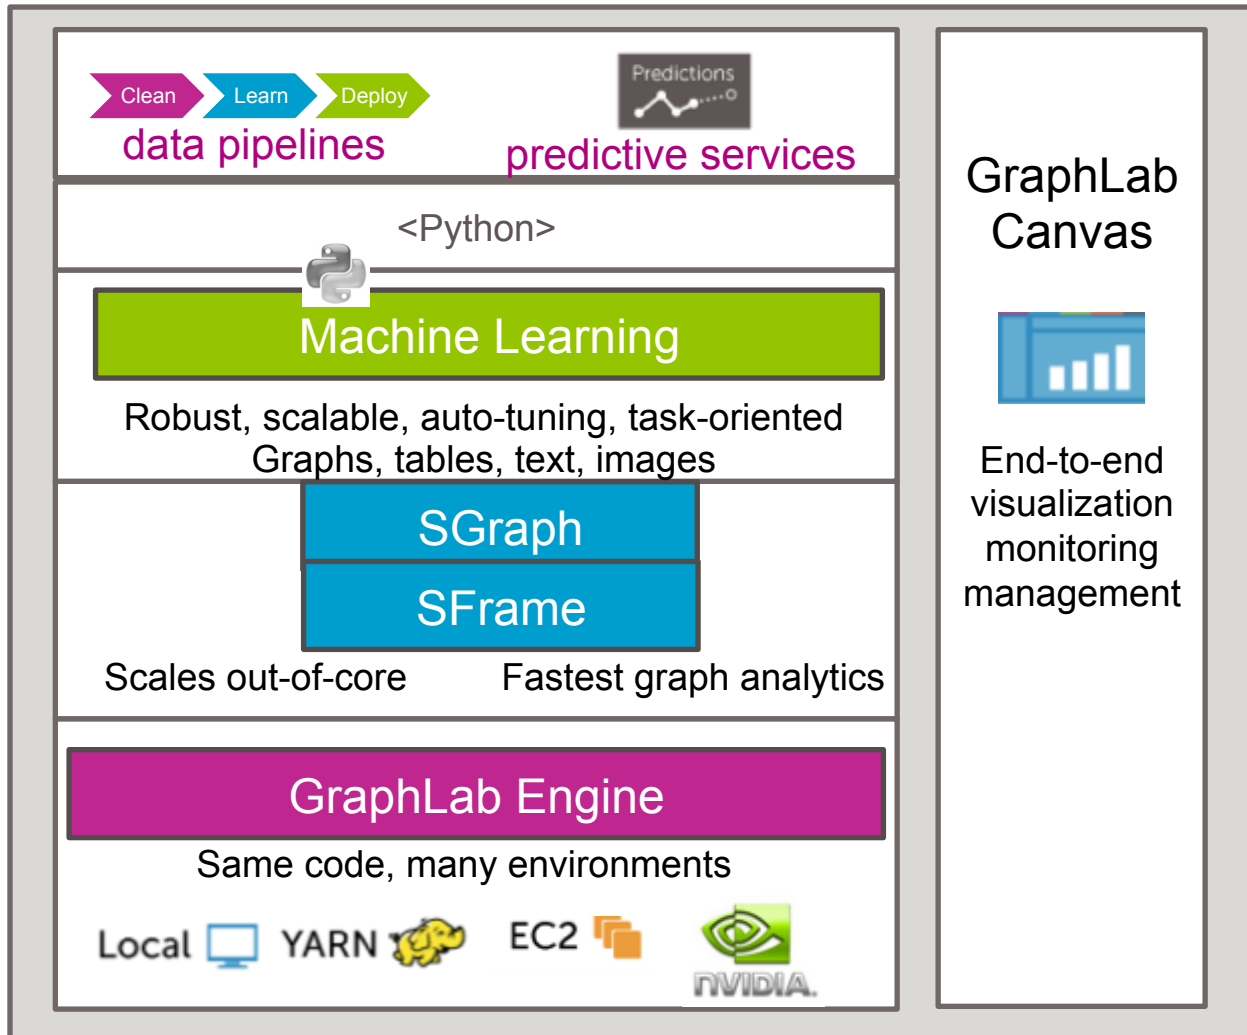

GraphLab Create



Analyze **big data** on **one machine**
graphs, tables, text, images
in **Python**
doesn't have to fit in memory

Distribute in production
with same code
on EC2, Yarn,...

GraphLab Canvas: Monitor & visualize
from prototype to production







```
pip install graphlab-create
```

@graphlabteam





TECHNOLOGY INNOVATION ROUNDTABLE ISRAEL

NOV 5, 2014

Government of Israel



Economic Mission
to the West Coast

Foreign Trade
Administration Israel



HP
(Platinum Sponsor)



Lenovo
(Gold Sponsor)



DELL
(Gold sponsor)

