



## Due Diligence for Investments in AI



**Kashyap Kompella, CFA**  
CEO, RPA2AI Research  
[kashyap.kompella@rpa2ai.com](mailto:kashyap.kompella@rpa2ai.com)

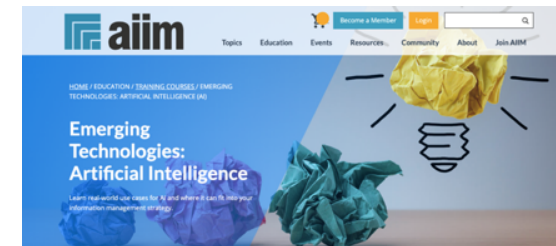
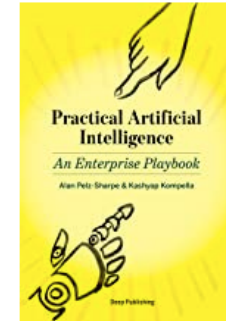
CFA Society, Bangalore 22/Feb/2020

Practical. Unbiased. Trusted.

# Kashyap Kompella: My Background

- CEO and Chief Analyst of Industry Analyst firm RPA2AI Research
- Chief-AI-Officer on-demand to leading organizations / start-ups
- Advisor to Venture Capital / Private Equity Firms
- Technology Industry Analyst
- Co-founder, Computer Vision start-up (2010-2012)
  
- Co-author of “Practical AI”
- Developed courses on Emerging Tech for managers
  
- Visiting Faculty, Institute of Directors
- Member, IEEE Committee on Ethical AI for Business

Best seller

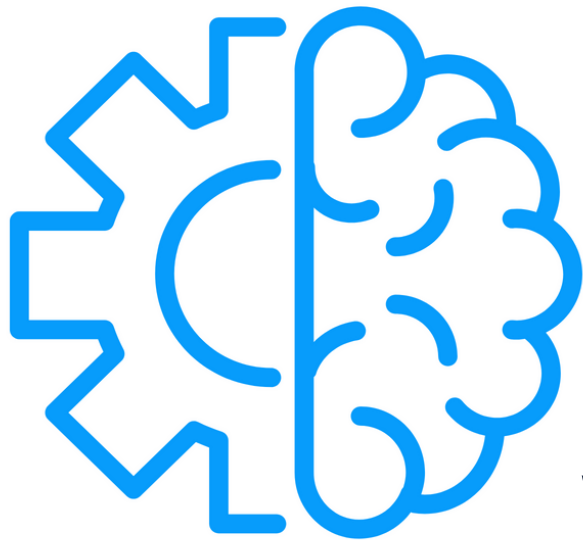


“10 Most Innovative Global Executives in AI” - Analytics Insight  
“ Top 10 Thought Leaders on AI ” – Thinkers360



Context: AI Gold Rush





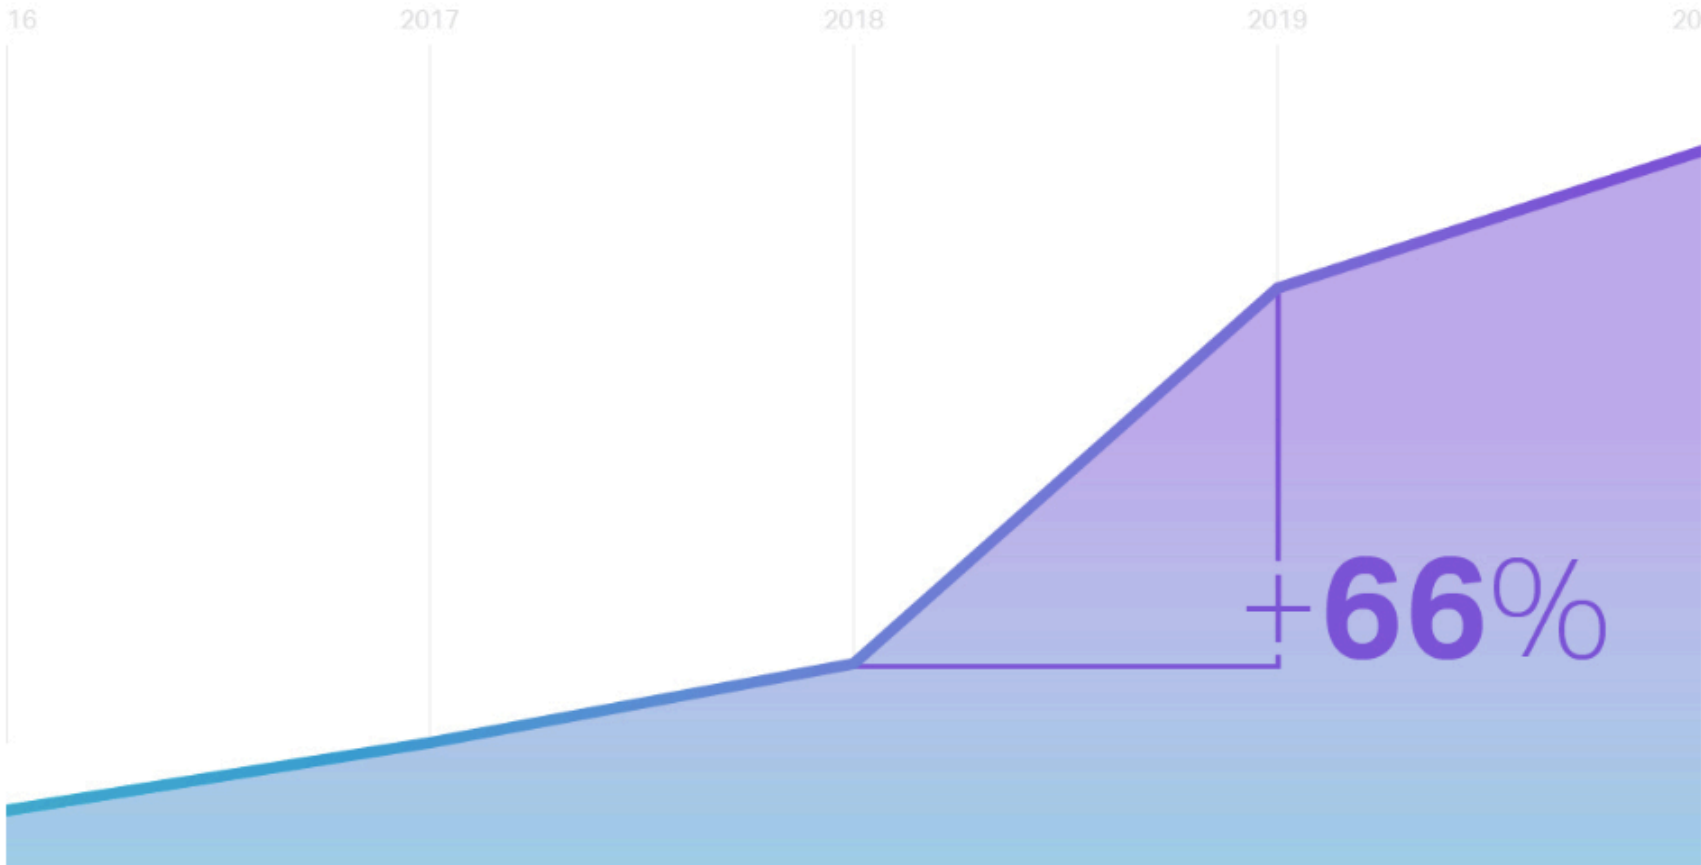
**When you're fundraising, it's AI.**

**When you're hiring, it's ML.**

**When you're implementing, it's regression.**



# Rise in Self-Reported AI Expertise



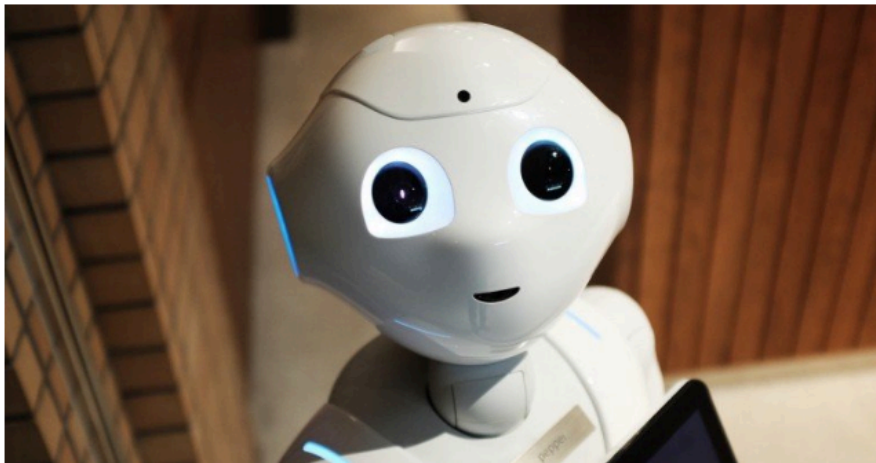
# Context: Everyone wants to jump on the bandwagon

---

## MIT Technology Review

Artificial Intelligence Mar 5

### About 40% of Europe's "AI companies" don't use any AI at all

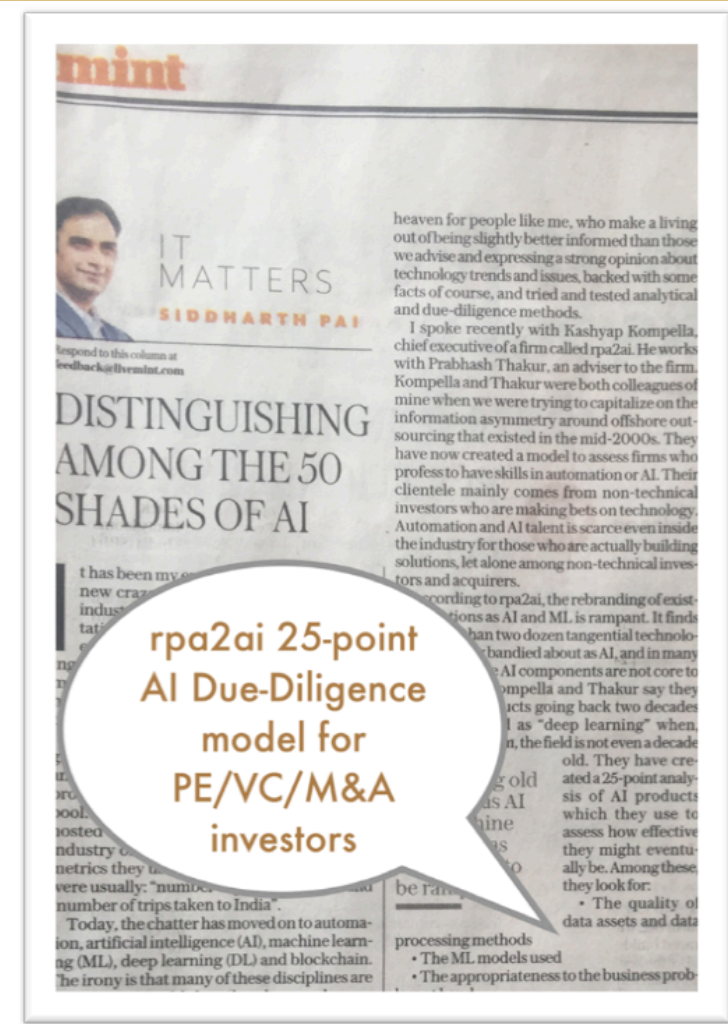


A surprising number of firms are jumping on the artificial-intelligence bandwagon—without actually investing in any AI.

**What?** Yes, I'm afraid so. MMC Ventures, an investment firm in London, looked at 2,830 European companies that claim to make use of AI. It found that a whopping 40% of them are not using any machine learning, a field of AI that has taken off in recent years both academically and commercially.

# Fifty Shades of Artificial Intelligence

- Do you often see old products being rebranded as AI?
- Do you think AI inside is not really core to the product?
- Do you need to filter out useful AI from vaporware AI?



# Key Points

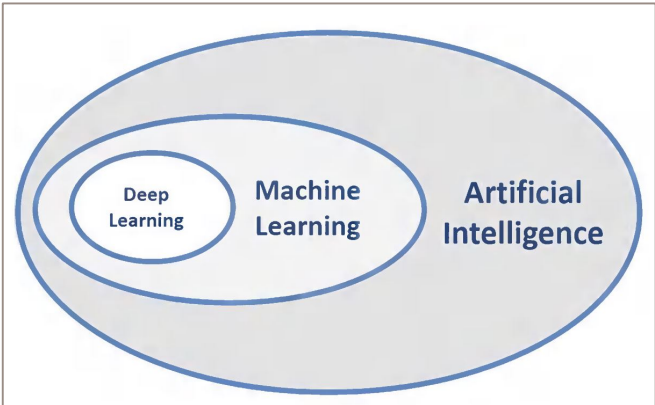
- ① AI Interest and Investments on the rise
- ② Due Diligence drives Go/No-Go & Valuations
- ③ SaaS DD is mature; AI DD is not
- ④ DD on AI is Art + Science
- ⑤ Data, Model, Process, Infra, Team, Risks
- ⑥ Domain Benchmarks for AI vs Traditional
- ⑦ Quality, Performance, Cost drive Viability
- ⑧ Premium for “AI” products and startups
- ⑨ Biz problem influences VC decisions
- ⑩ Simple Tech + Big Opportunity = Unicorn?



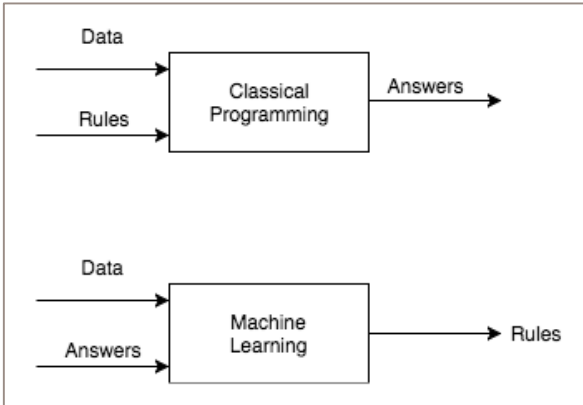


# What do people mean when they say “AI”?

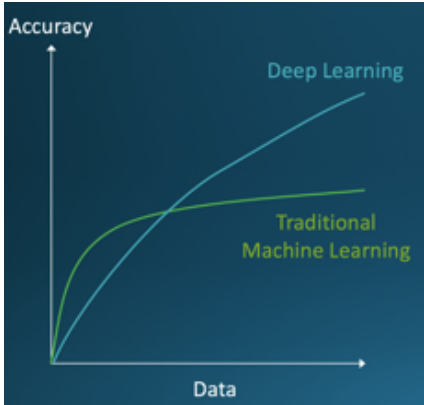
1



2

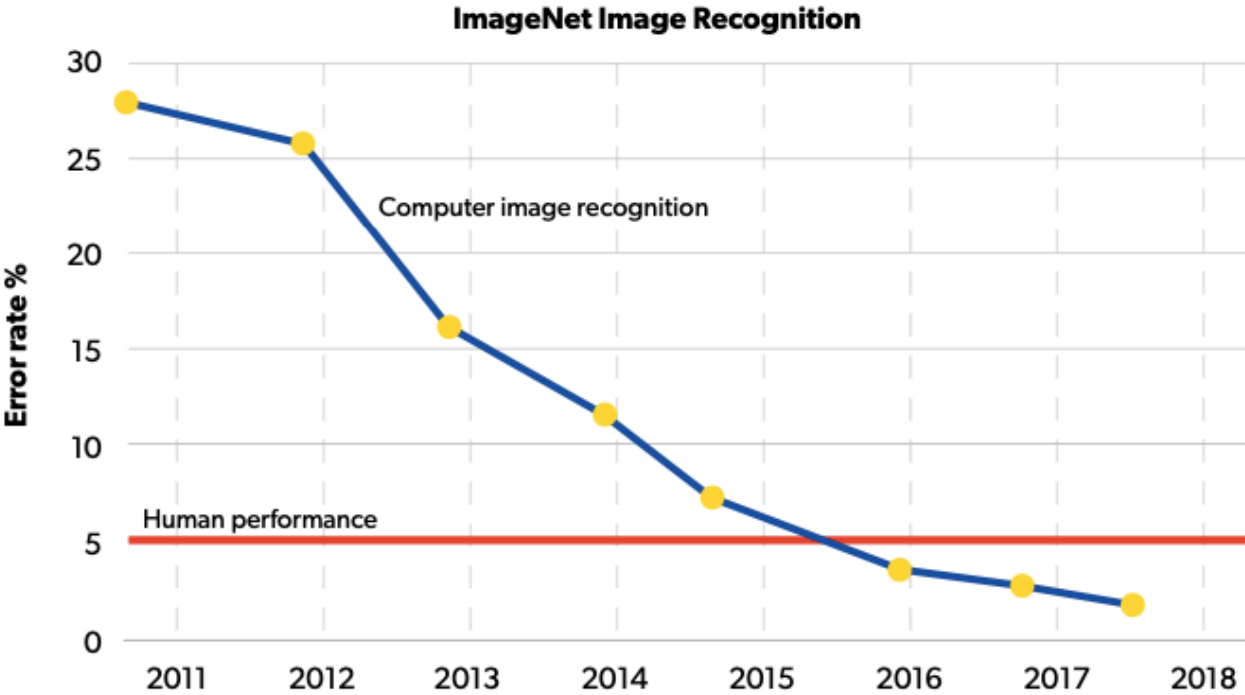


3



# ML Performance Improvements - Illustration

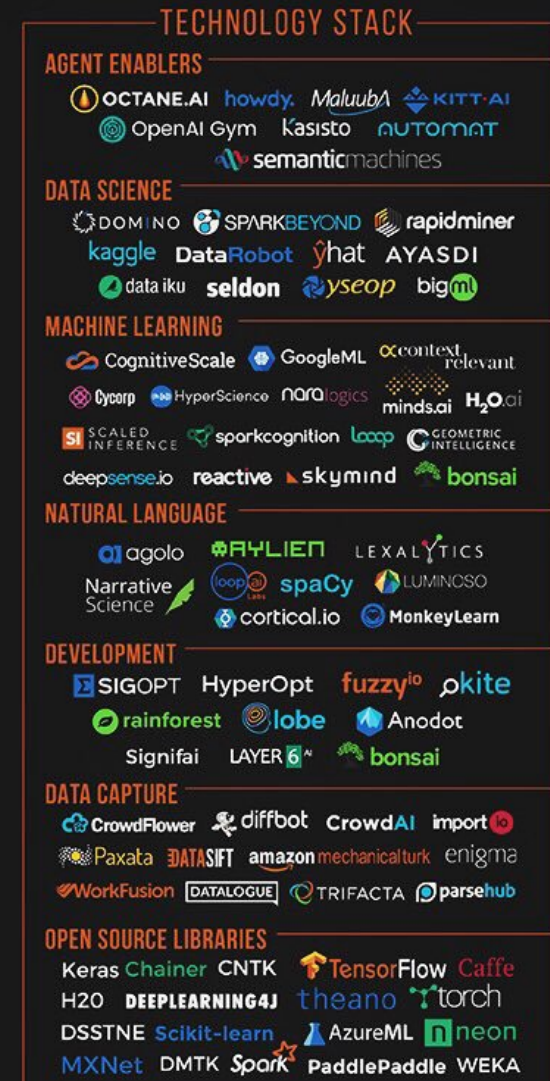
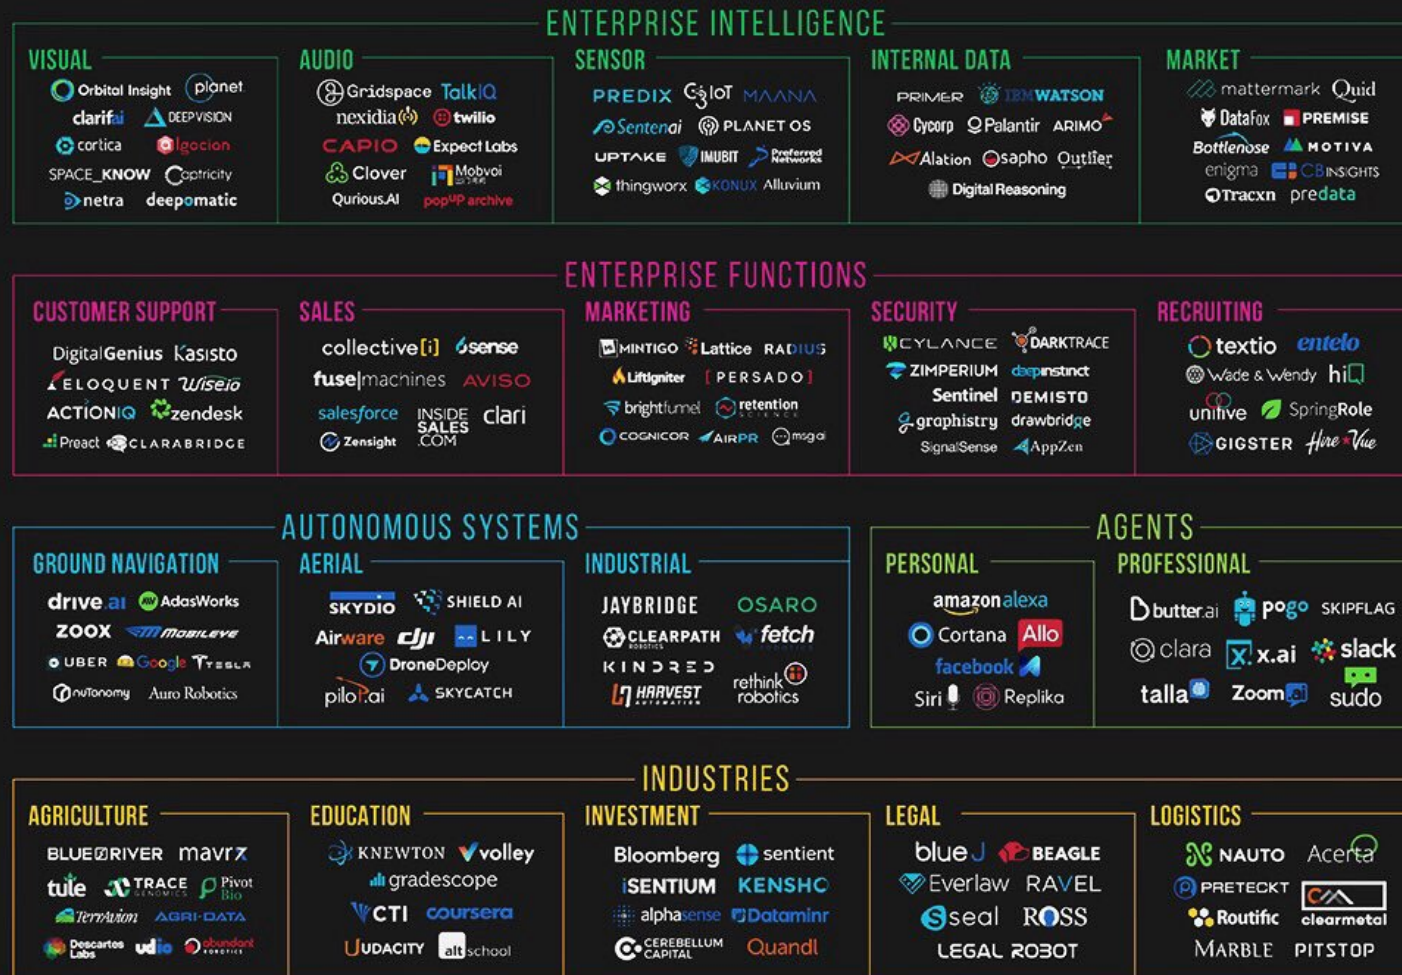
Convolutional neural networks are delivering human-level image recognition



Source: <https://www.eff.org/ai>

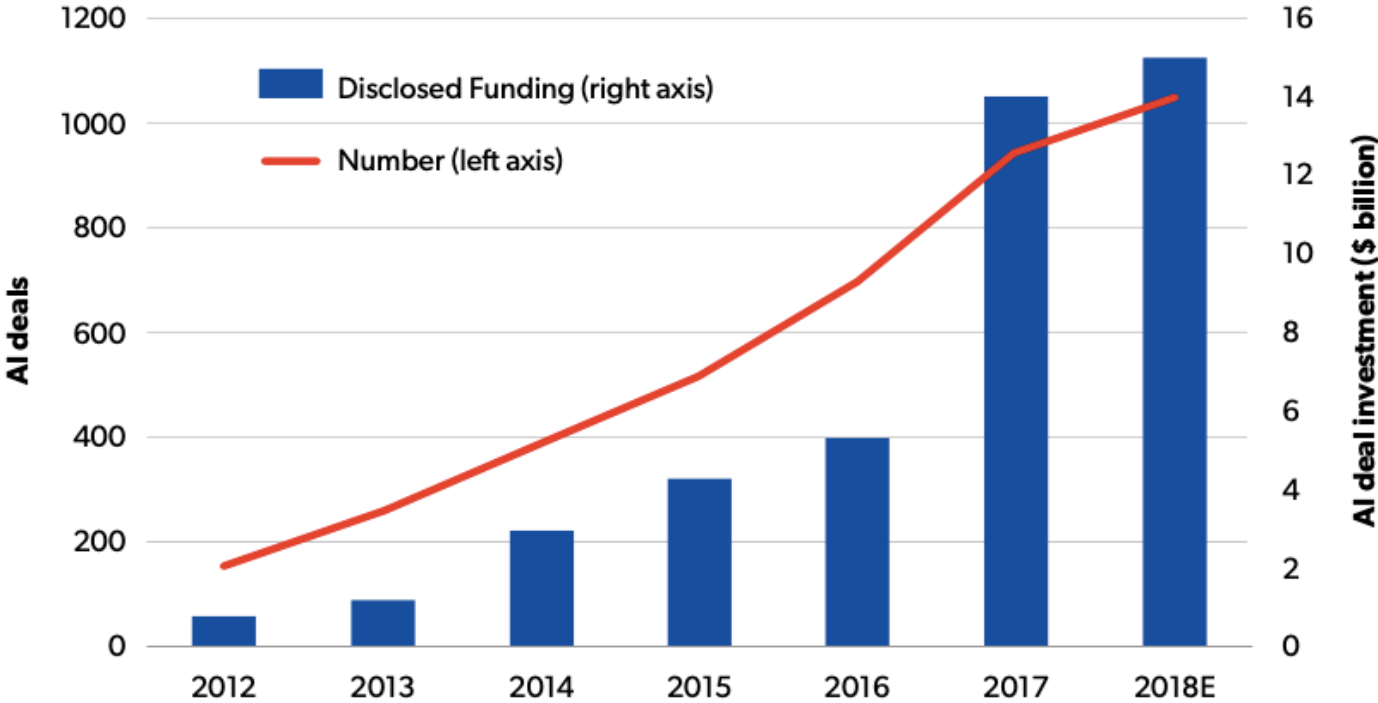
# Hype but genuine progress

## MACHINE INTELLIGENCE 3.0



# VC Investments in AI increasing rapidly

Fig 23. Venture capital investment in AI has increased 15-fold in five years



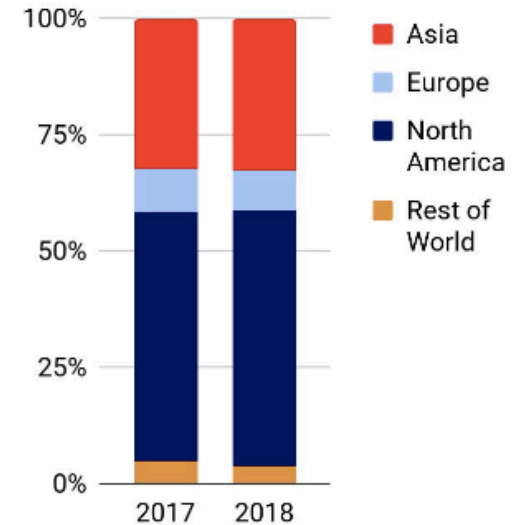
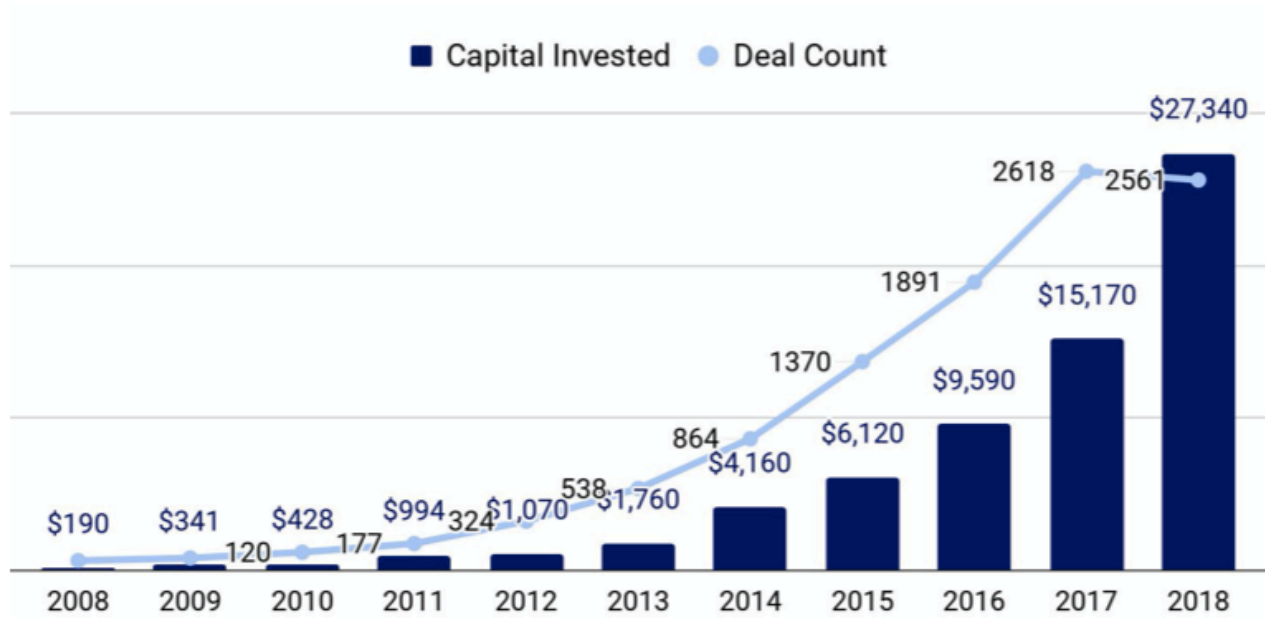
Source: CB Insights, MMC Ventures



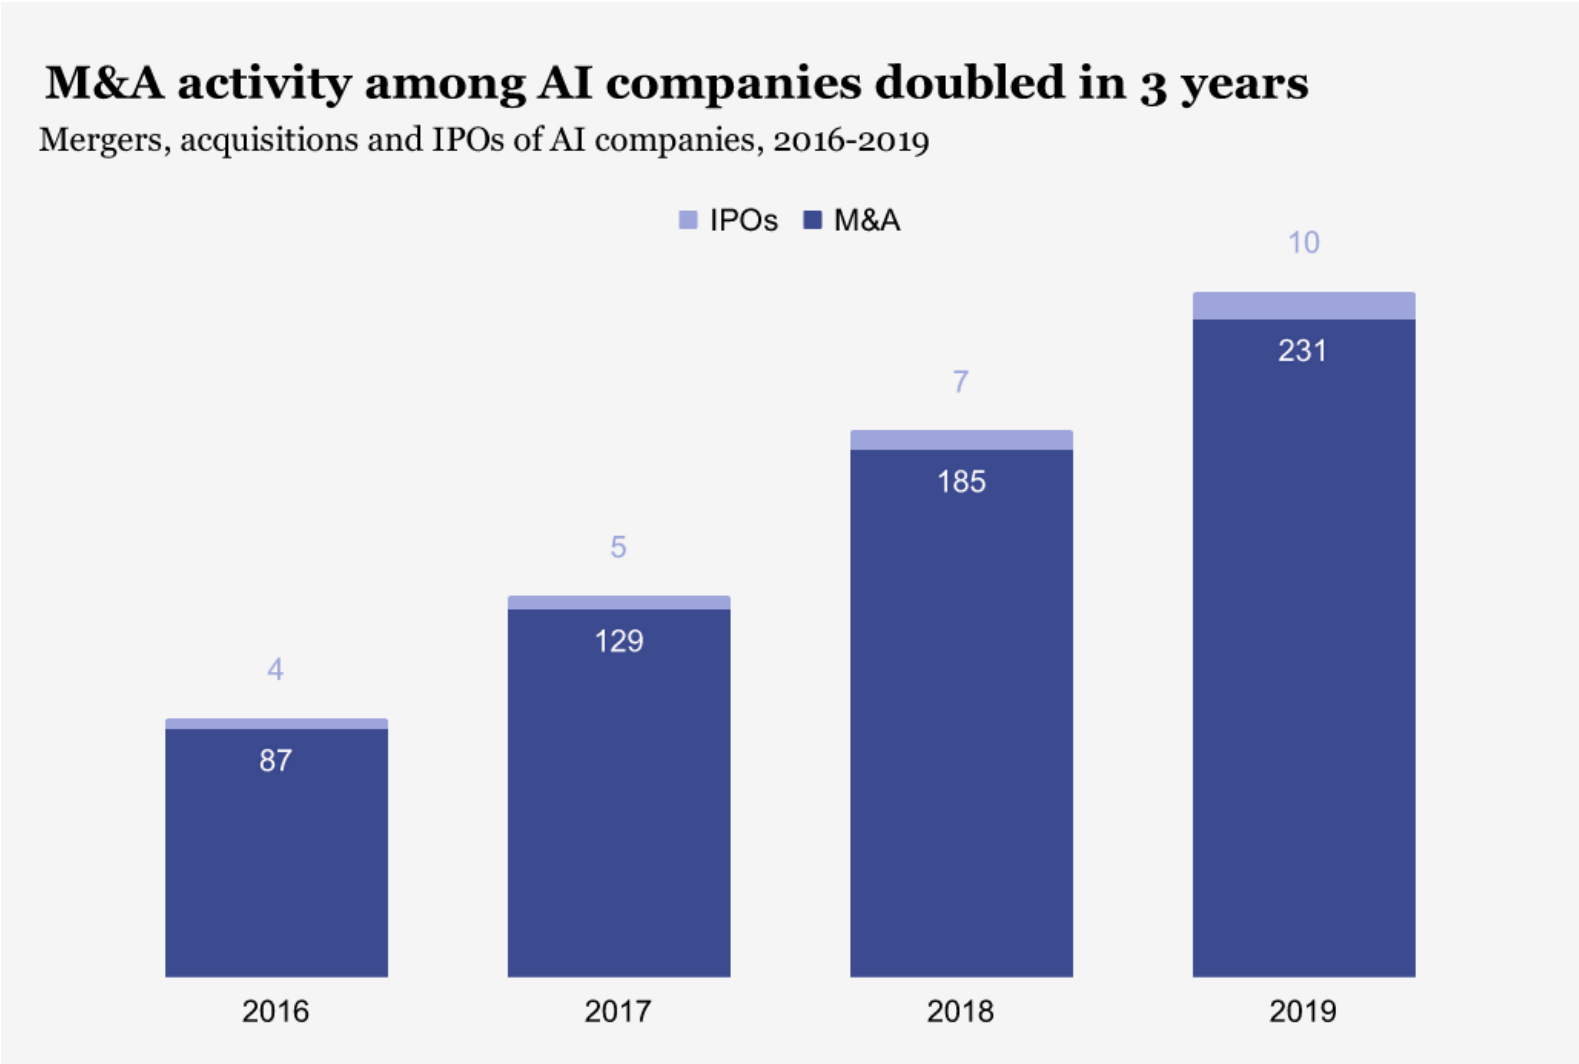
# VC Investments in AI increasing rapidly

## Global venture capital investments in AI themes grow at a clip to reach >\$27B/year

▶ Almost 80% more capital invested in FY18 vs FY17 with North America leading the way at 55% market share.



# VC Investments in AI increasing rapidly



# AI Market Segments / Value Chain



# VC Investments: SaaS vs. AI start-ups

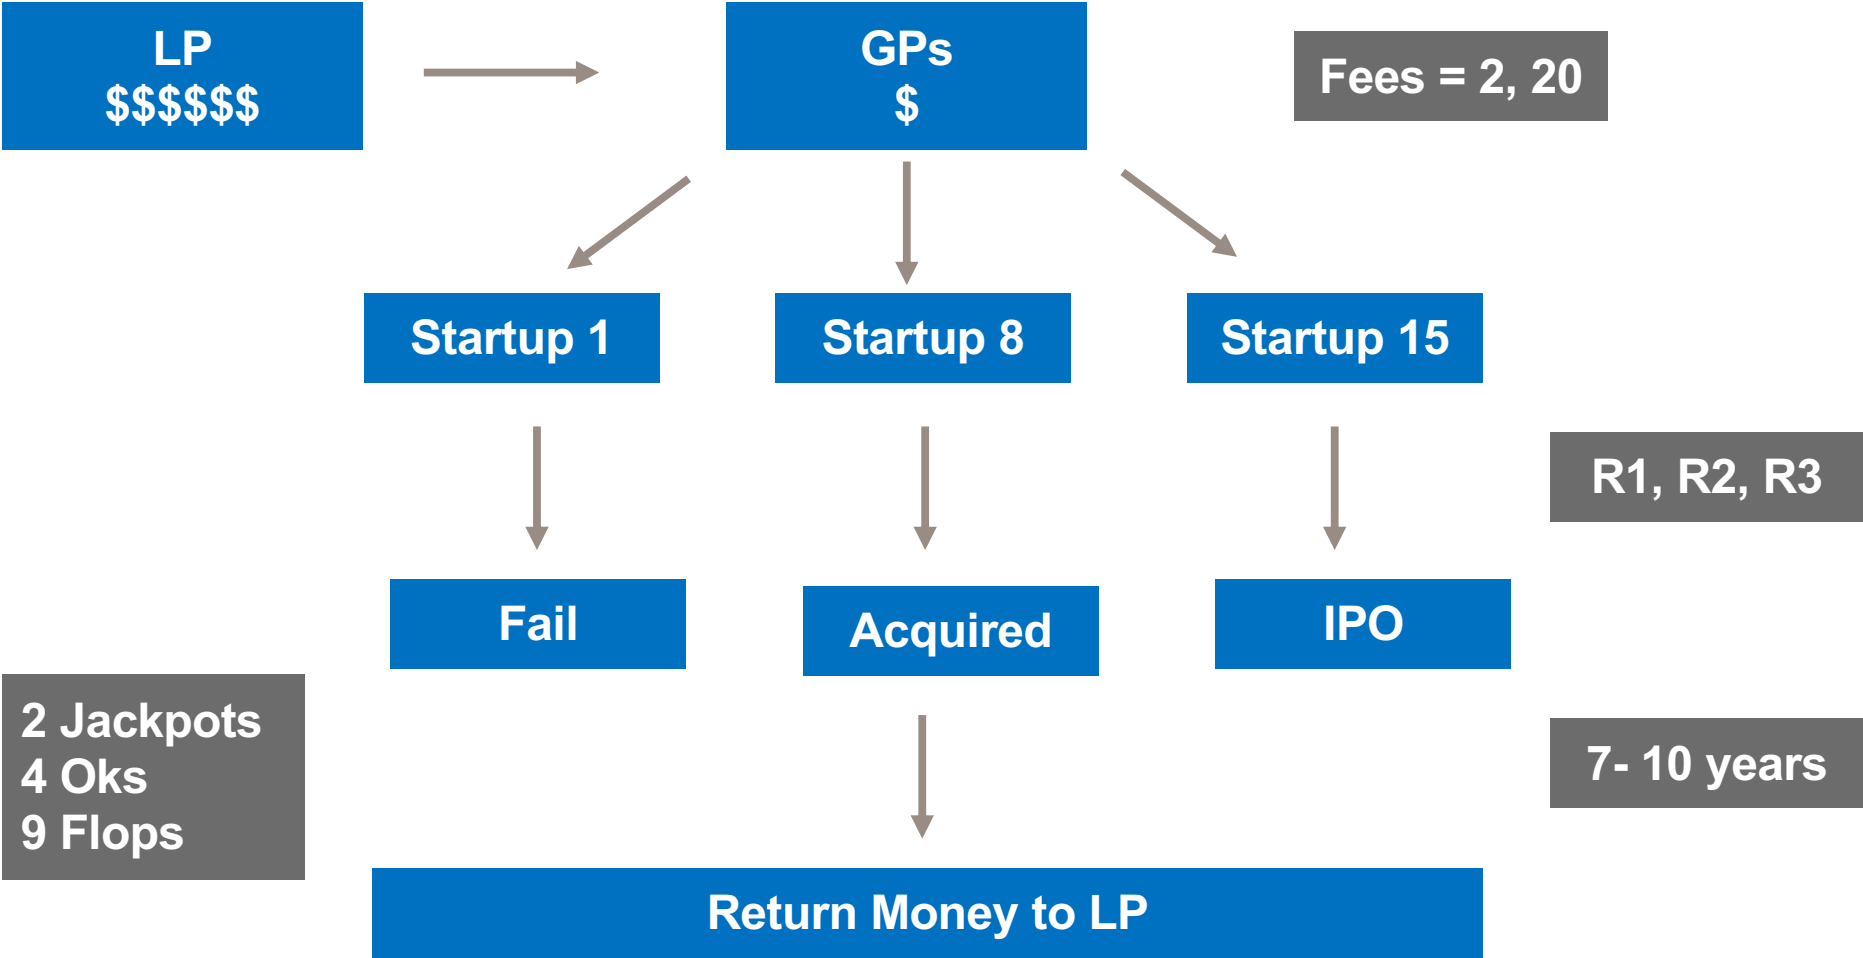
	SaaS	AI / ML
Annual Recurring Revenue (ARR)	\$ 100 K	Paid Pilots
Valuation and Size of Funding Round	\$ 5 M 1 – 2 M	\$ 10 M 2 – 3M
Team	Tech Co-Founder Domain Expertise	Tech Co-Founder Domain Expertise ML Background
Product & Market	Early users feedback Big Market Opportunity	MVP + ML Model Pilot Customers Big Market Opportunity



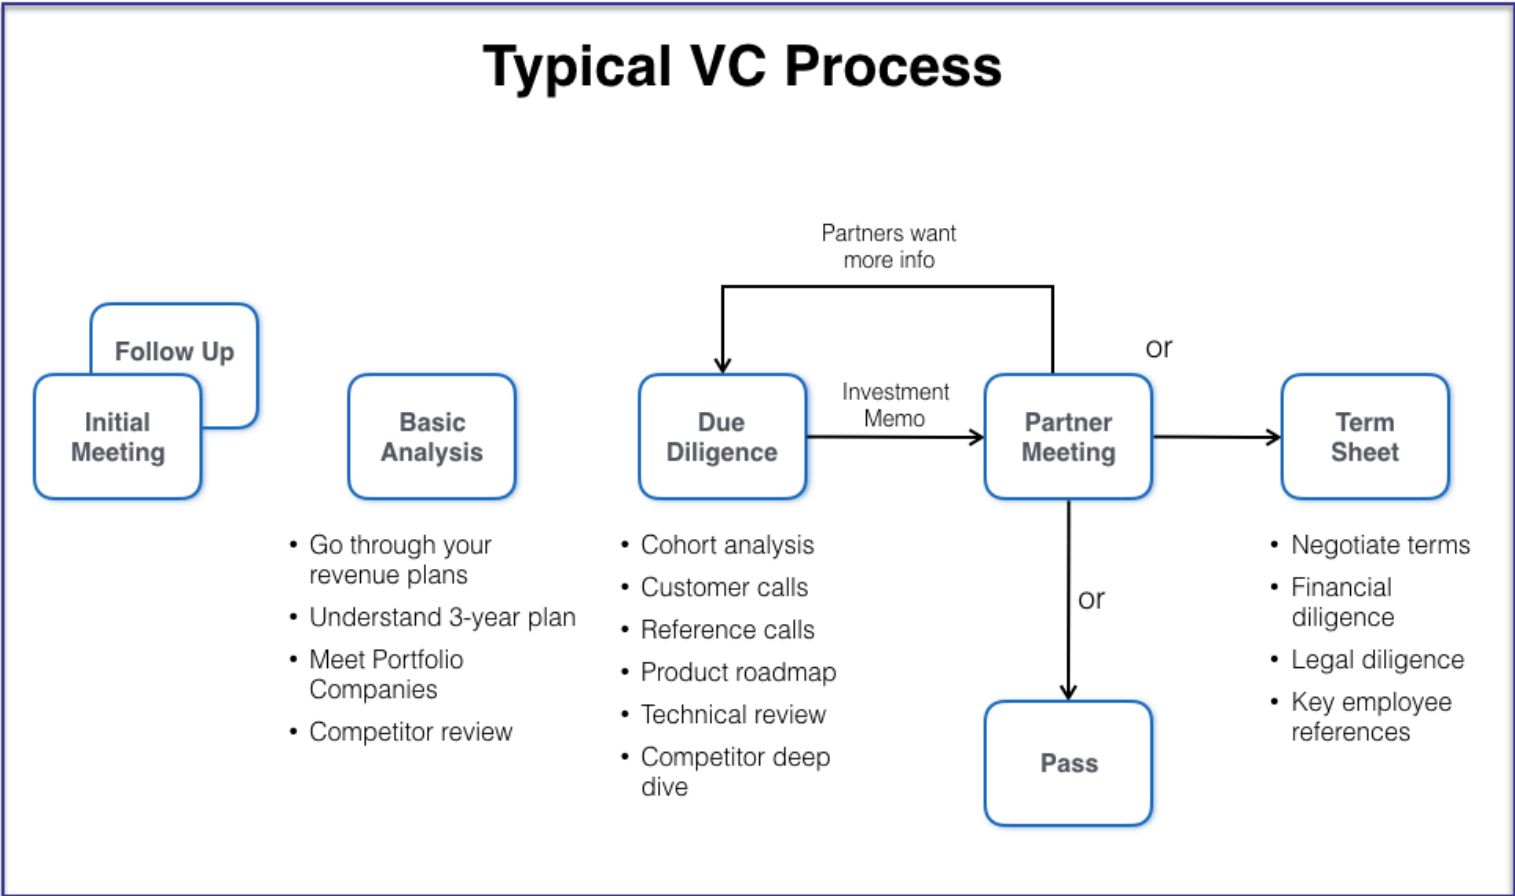
# VC Investments: SaaS vs. AI start-ups (continued)

	SaaS	AI / ML
Sales / Distribution	Go-To-Market Plan	Go-To-Market Plan
Regulation	--	Transparency Compliance
Moat	Plan (too early)	Proprietary Data Data Processing
AI	--	Dataset Model Performance Improvements

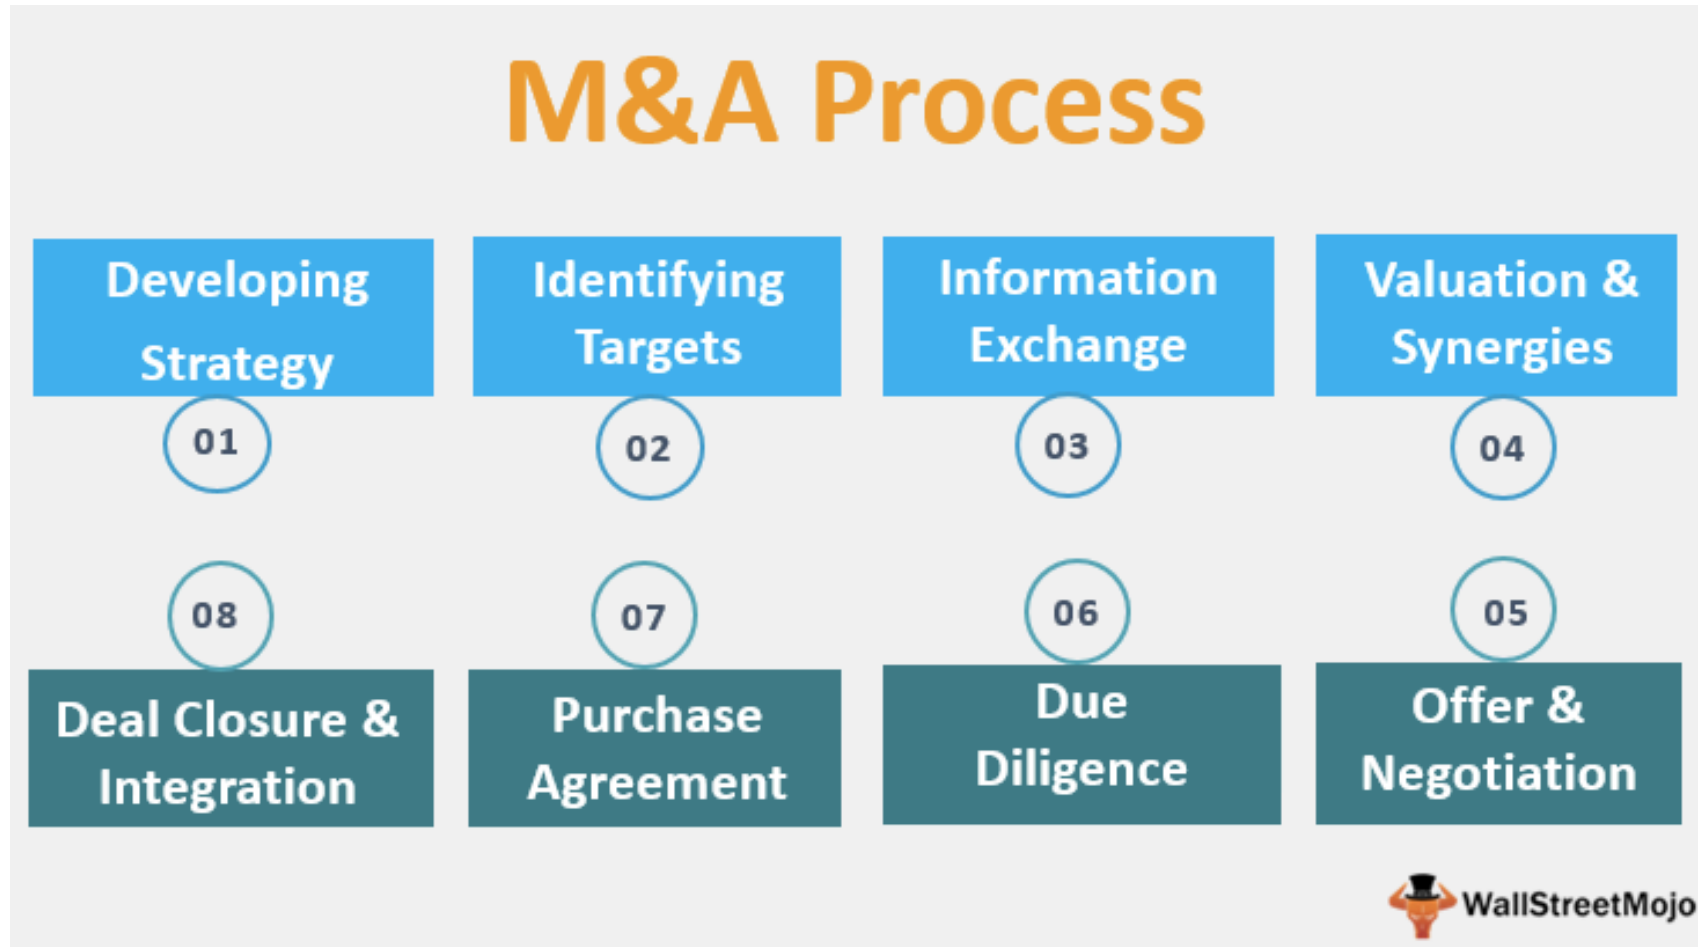
# Intro to Venture Capitalists (Stylized)



# VC Investment Process



# M&A Process





# Evaluating the ML in Technology Products (1)

---

Scope

What features / functionality use machine learning (ML)

What are the benefits vis-a-vis non-ML approaches?

Data

Datasets used to build the ML models

What data is required from the client/user?

Components

List of 3rd party / Open Source ML software used

List of your own components/modules

# Evaluating the ML in Technology Products (2)

---

Architecture

What is architecture of the product (modules/components + flow)

ML Methods

What are main ML methods used?

Justification (research papers that recommend the approach)?

ML Performance  
Evaluation

What are the metrics used to assess the ML model performance?

# Evaluating the ML in Technology Products (3)

---

Continuous Improvement

Does the ML model continue to learn?  
How is model performance monitored?  
When does the model get retrained?

System Performance and Scalability

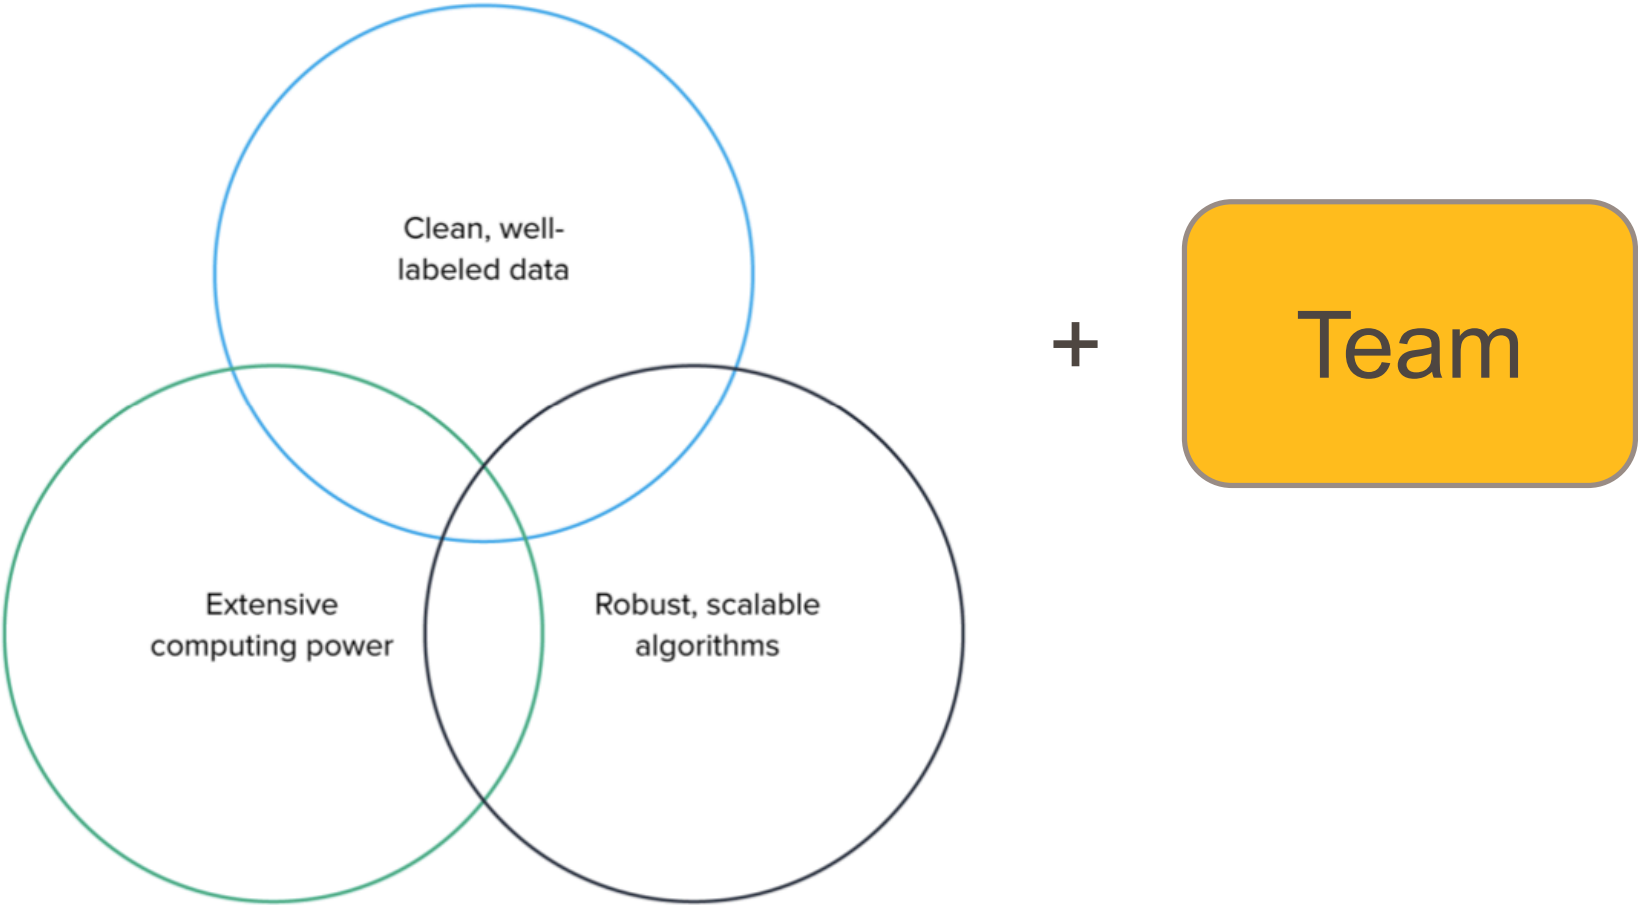
Training time  
Training cost  
Inference time / cost

Data Privacy and Explainability

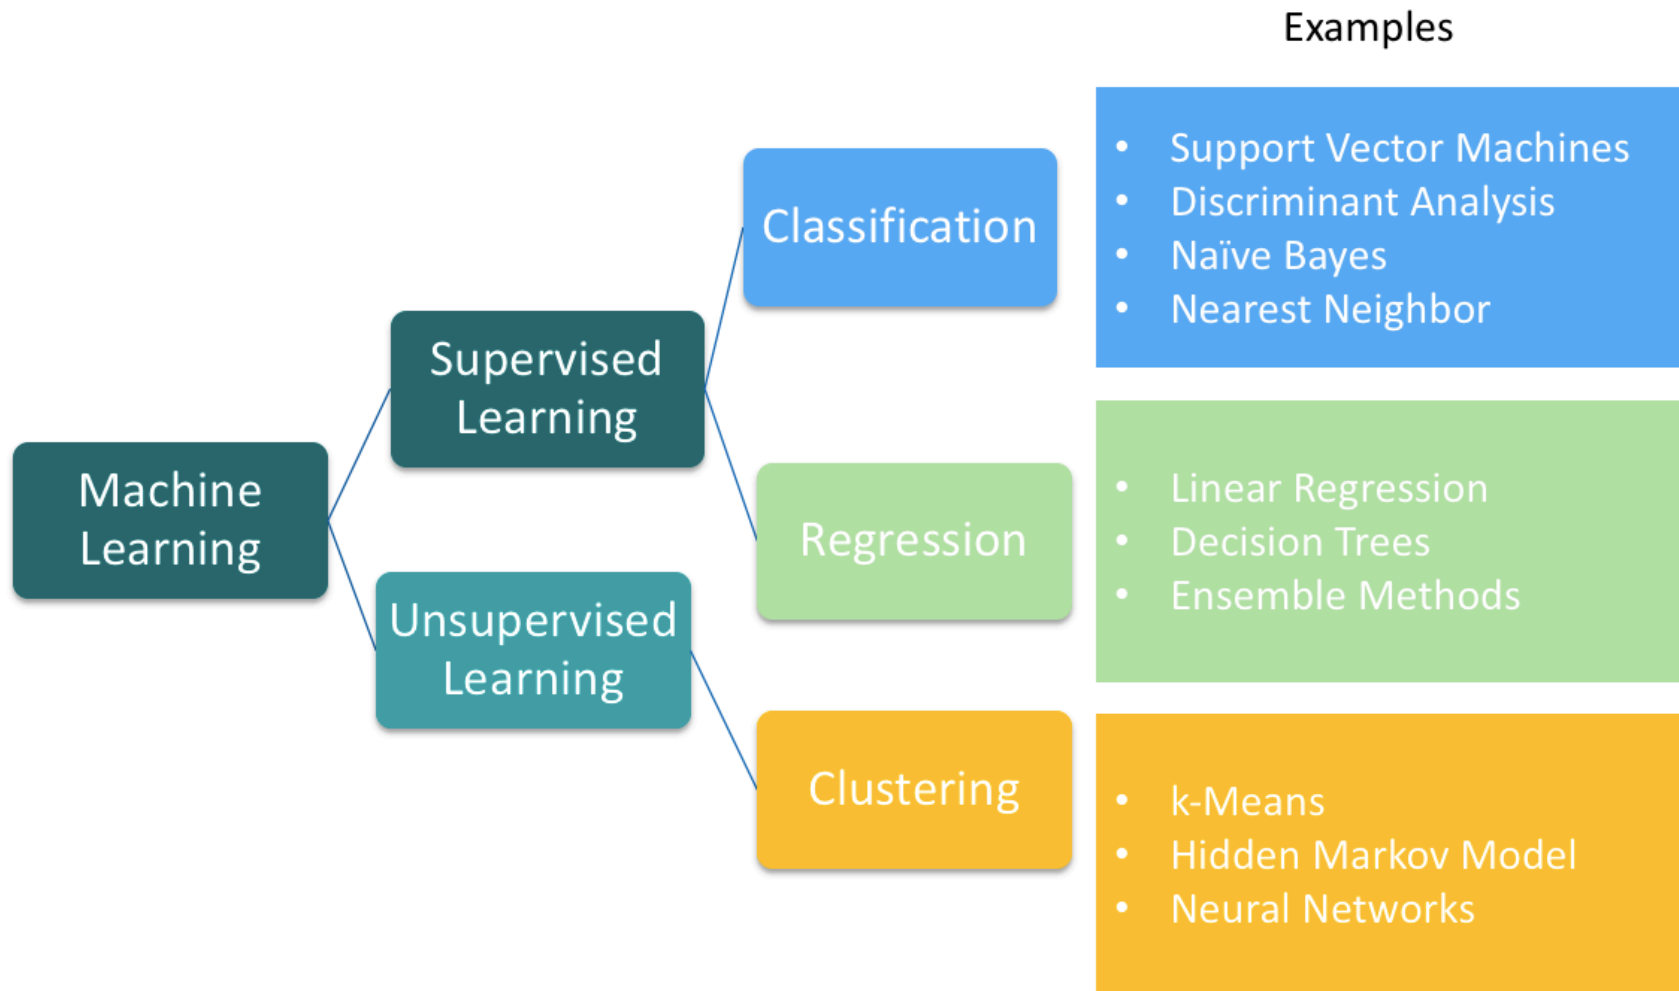
How are data privacy and security handled?  
Does the ML system provide a “why” for the recommendations?

# Technical Feasibility of ML Products

The Three Pillars of Technical Feasibility in AI Startups



# Methods of Machine Learning





# Implementation of Actual Code / Models is becoming easier

---

## *#Decision Trees*

```
from sklearn import tree
model = tree.DecisionTreeClassifier()
model.fit(X_train, y_train)
```

## *#Random Forest*

```
from sklearn import ensemble
model = ensemble.RandomForestClassifier()
model.fit(X_train, y_train)
```

## *#Support Vector Machines*

```
from sklearn import svm
model = svm.SVC()
model.fit(X_train, y_train)
```

## *#Logistic Regression*

```
import sklearn.linear_model as sklm
model = sklm.LogisticRegression()
model.fit(X_train, y_train)
```

## *#Neural Networks*

```
from sklearn.neural_network import MLPClassifier
model = MLPClassifier()
model.fit(X_train, y_train)
```

# Machine Learning Applications: Comparing Performance

Applications



Reference Datasets



Reference Model



Performance Metrics

Image Classification



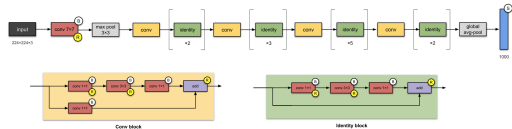
Imagenet



Resnet-50

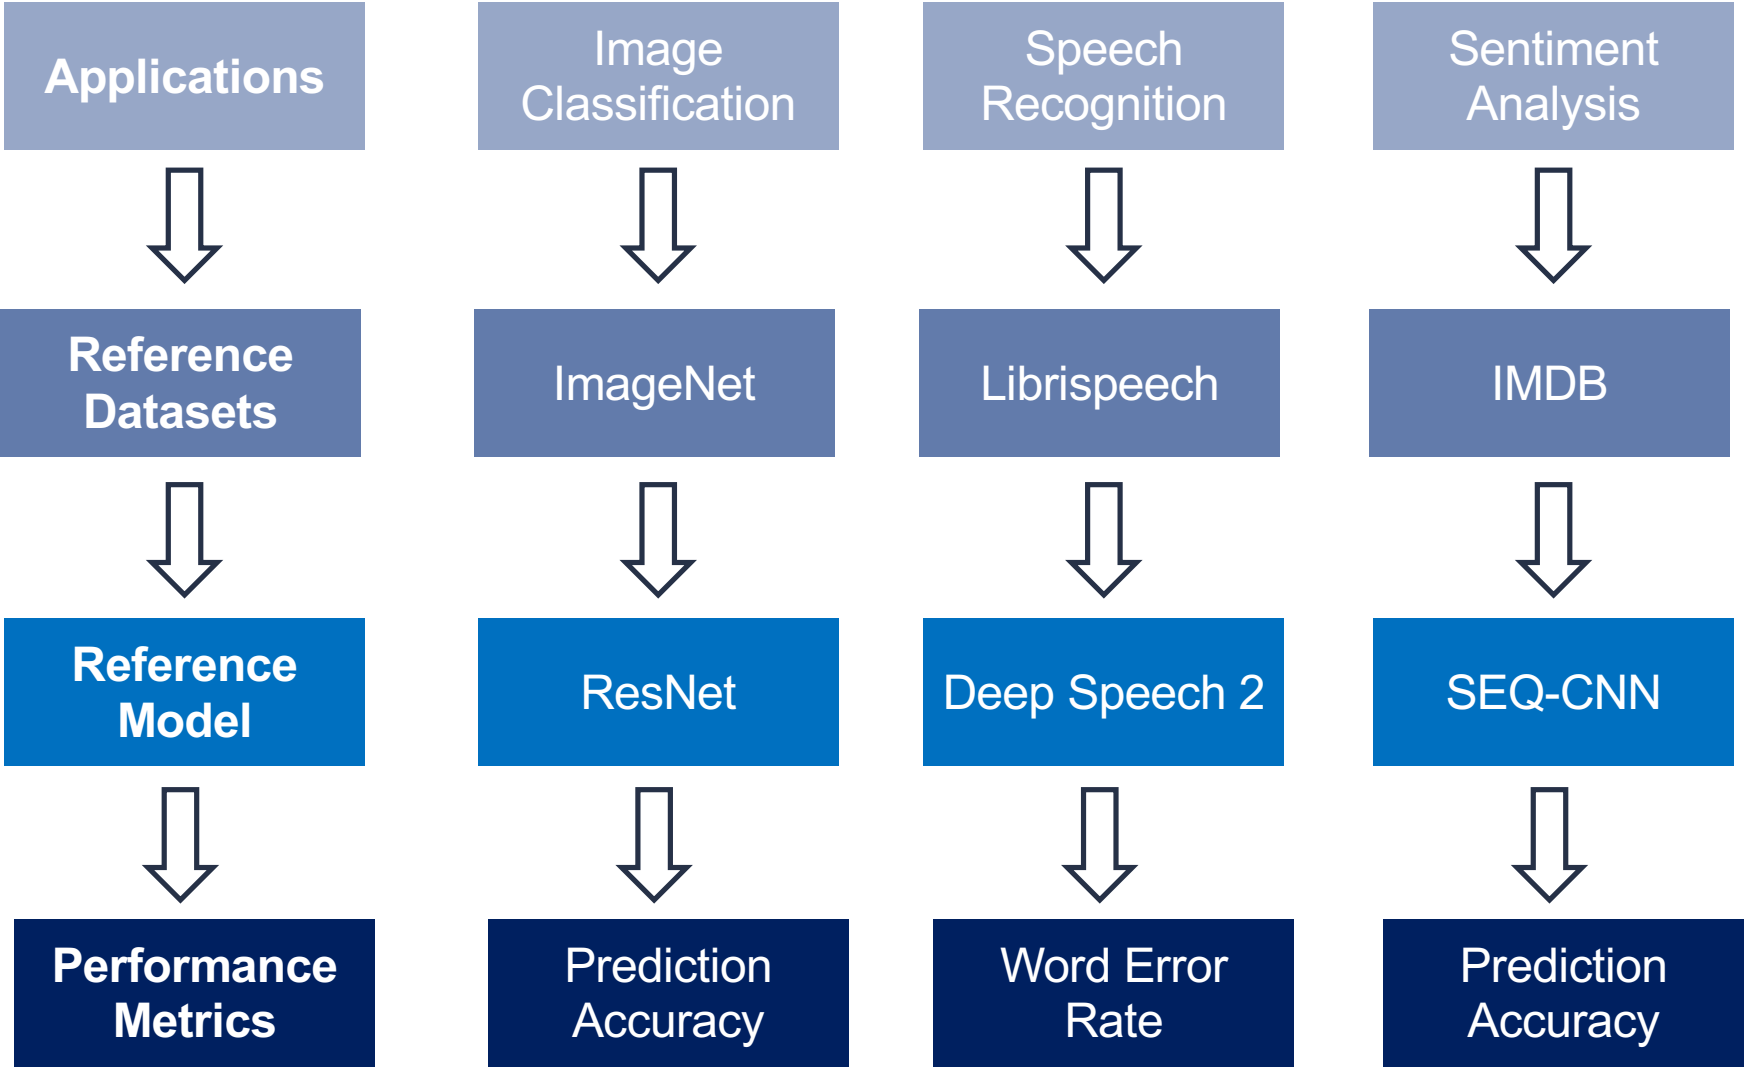


Prediction Accuracy



75%

# Machine Learning Applications: Baseline Performance



# Illustrative List of Datasets & Machine Learning Models

Computer Vision	Audio / Speech	Language / Text	Commerce /Online
<ul style="list-style-type: none"><li>Image Classification</li><li>Object Detection</li><li>Video Analysis</li><li>Face Identification</li><li>Medical Imaging</li><li>Autonomous Vehicles</li></ul>	<ul style="list-style-type: none"><li>Speech Recognition</li><li>Text-to-Speech</li><li>Question Answering</li><li>Keyword Detection</li><li>Speaker ID</li><li>Chatbots</li></ul>	<ul style="list-style-type: none"><li>Translation</li><li>Word Embedding</li><li>Language Model</li></ul>	<ul style="list-style-type: none"><li>Ratings</li><li>Recommendations</li><li>Sentiment Analysis</li><li>Fraud Detection</li><li>Anomaly Detection</li><li>Time Series Pred.</li></ul>
COCO ImageNet	SQuAD LibriSpeech LM-Benchmarrk	WMT (for a given language pair e.g. Eng-Ger)	IMDB MovieLens-20M Amazon
ResNet-50 TF Object Detection Detectron	SQuAD Explorer Deep Speech 2	OpenNMT Transformer	Neural Collab. Filtering Conv. Neural Network
COCO mAP Prediction Accuracy	Word Error Rate Perplexity	Bilingual Evaluation Understudy - BLEU	Prediction Accuracy

COCO = Common Objects in Context ; mAP = mean Average Precision

SQuAD = Stanford Question Answering Dataset; WMT = Workshop on Statistical Machine Translation

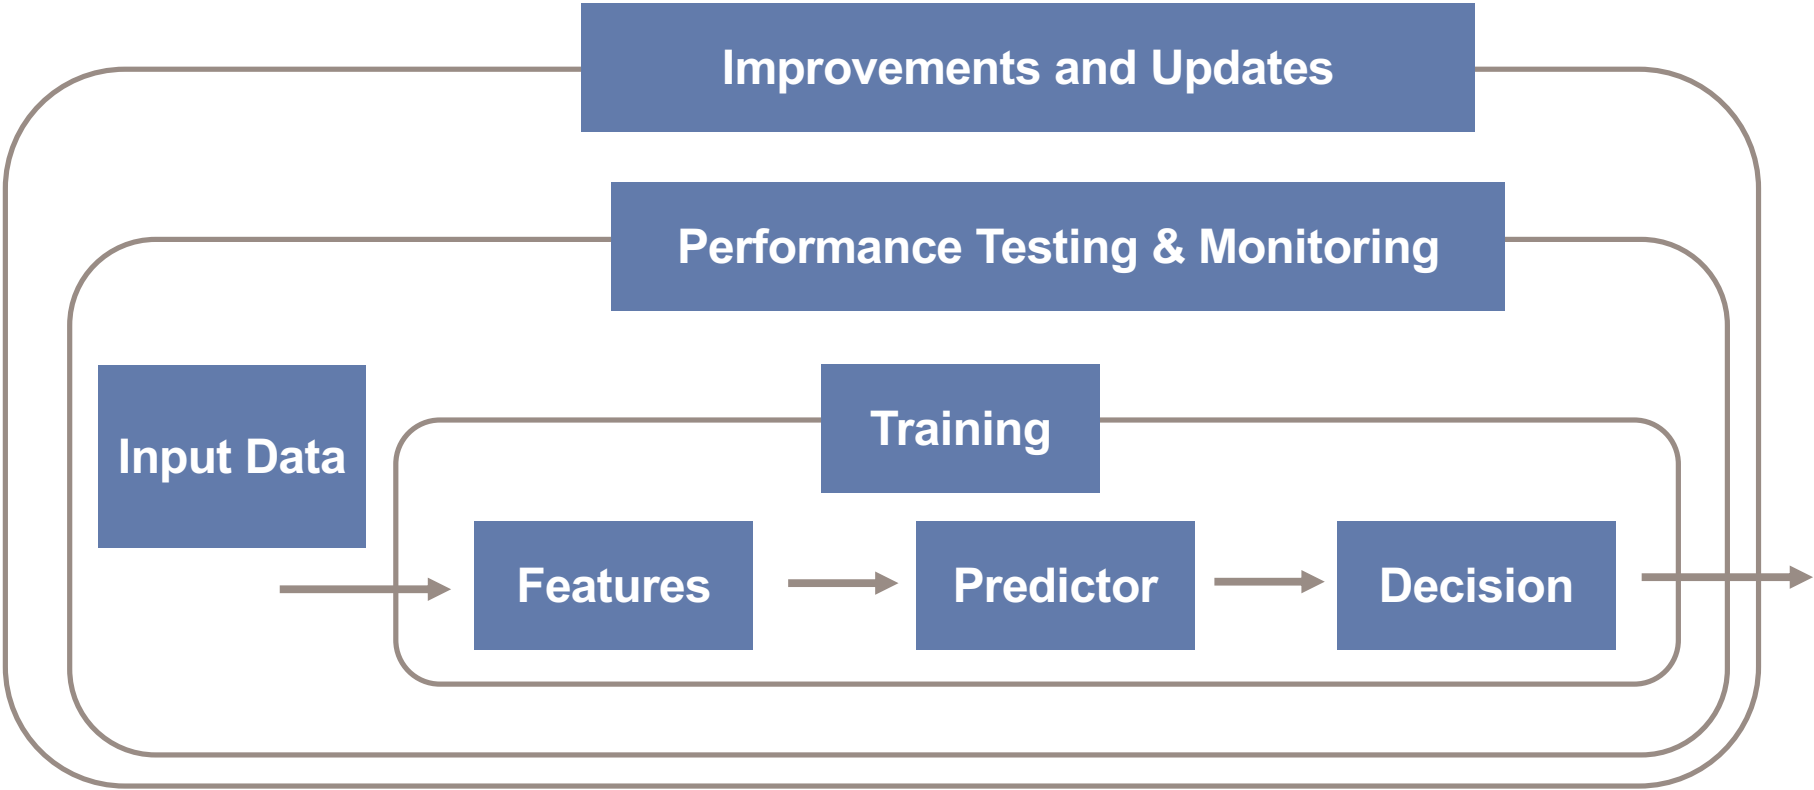
# ML Performance Benchmarks

---



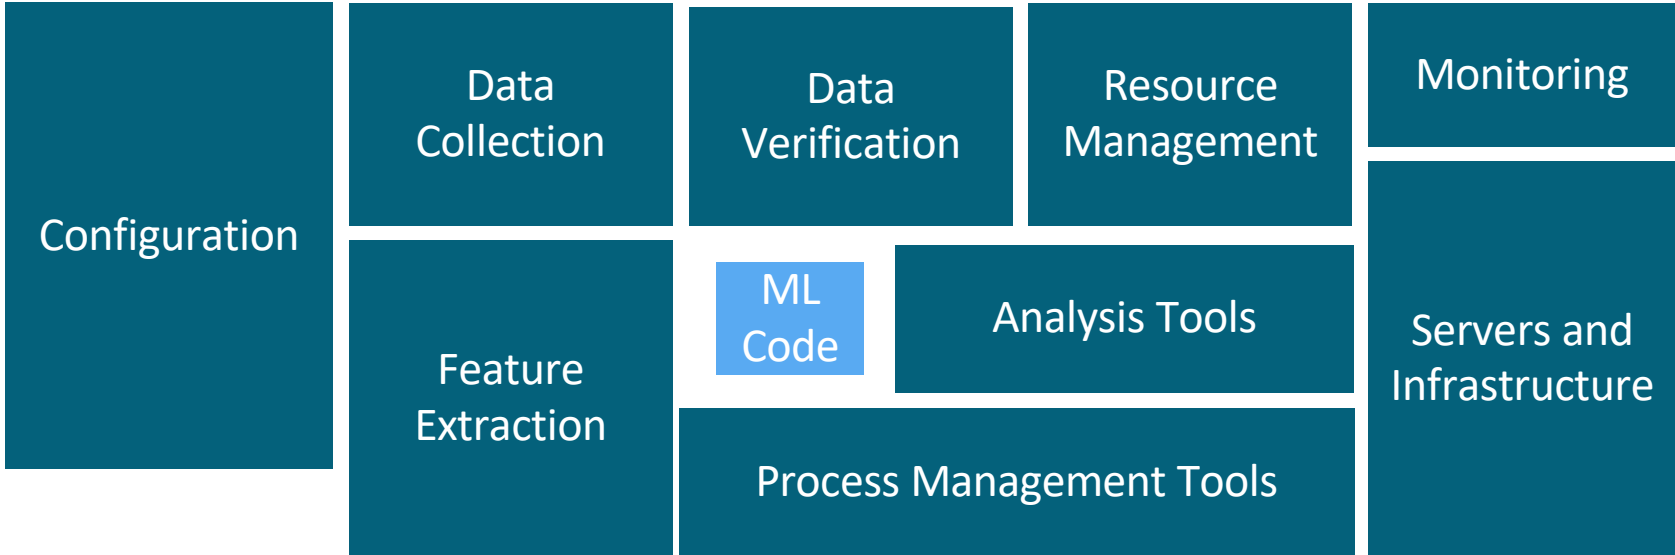
Fair and useful benchmarks for measuring training and inference performance of ML hardware, software, and services.

# ML System Architecture





# ML code in Enterprise Applications



# Machine Learning Application: Viability

---

Quality



Accuracy of Prediction

# ML Performance - Illustration

## Image Classification: ImageNet

Source: AI Index survey and PapersWithCode, 2019.

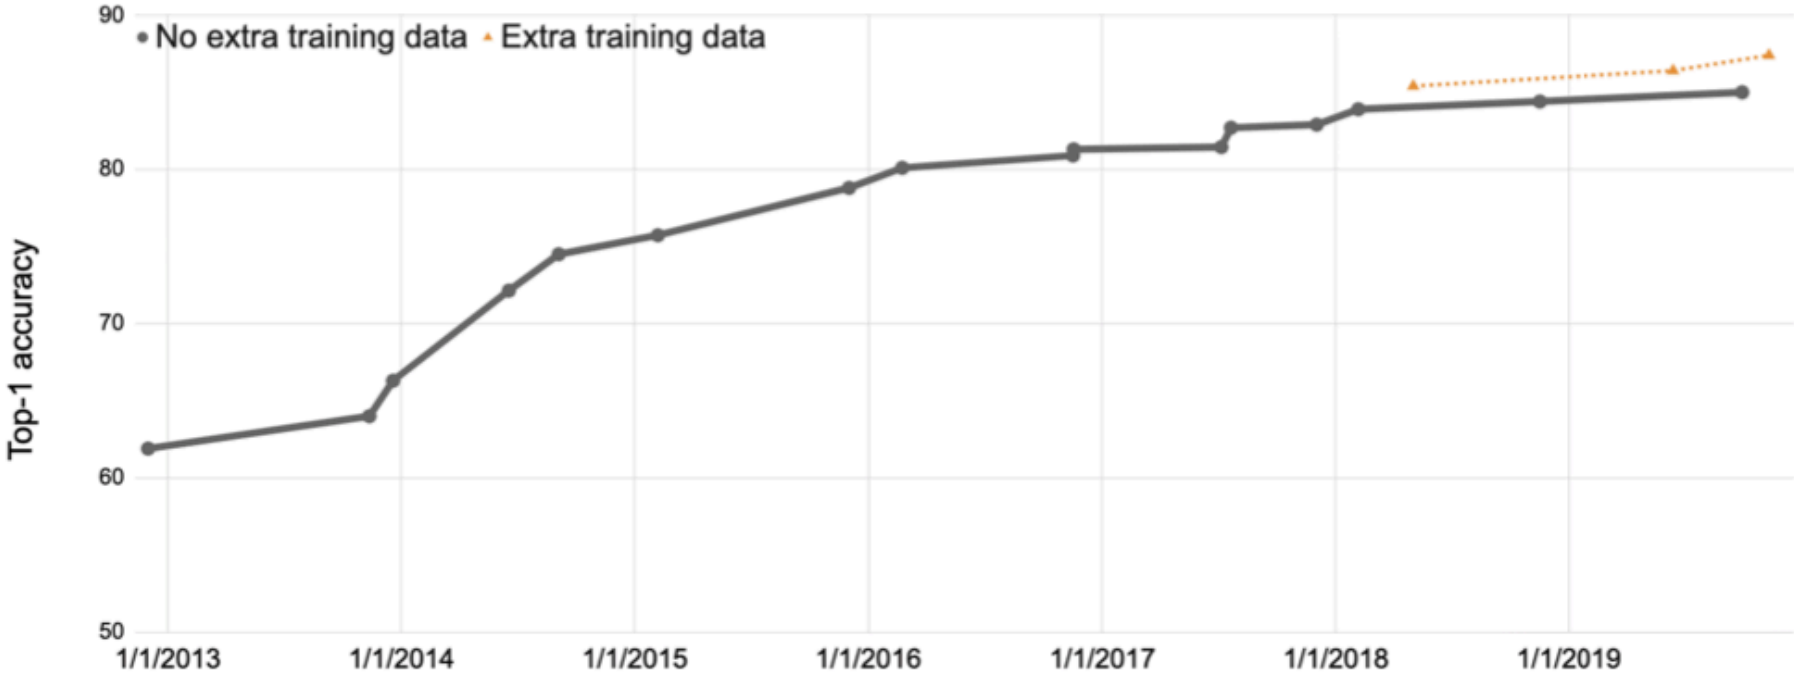
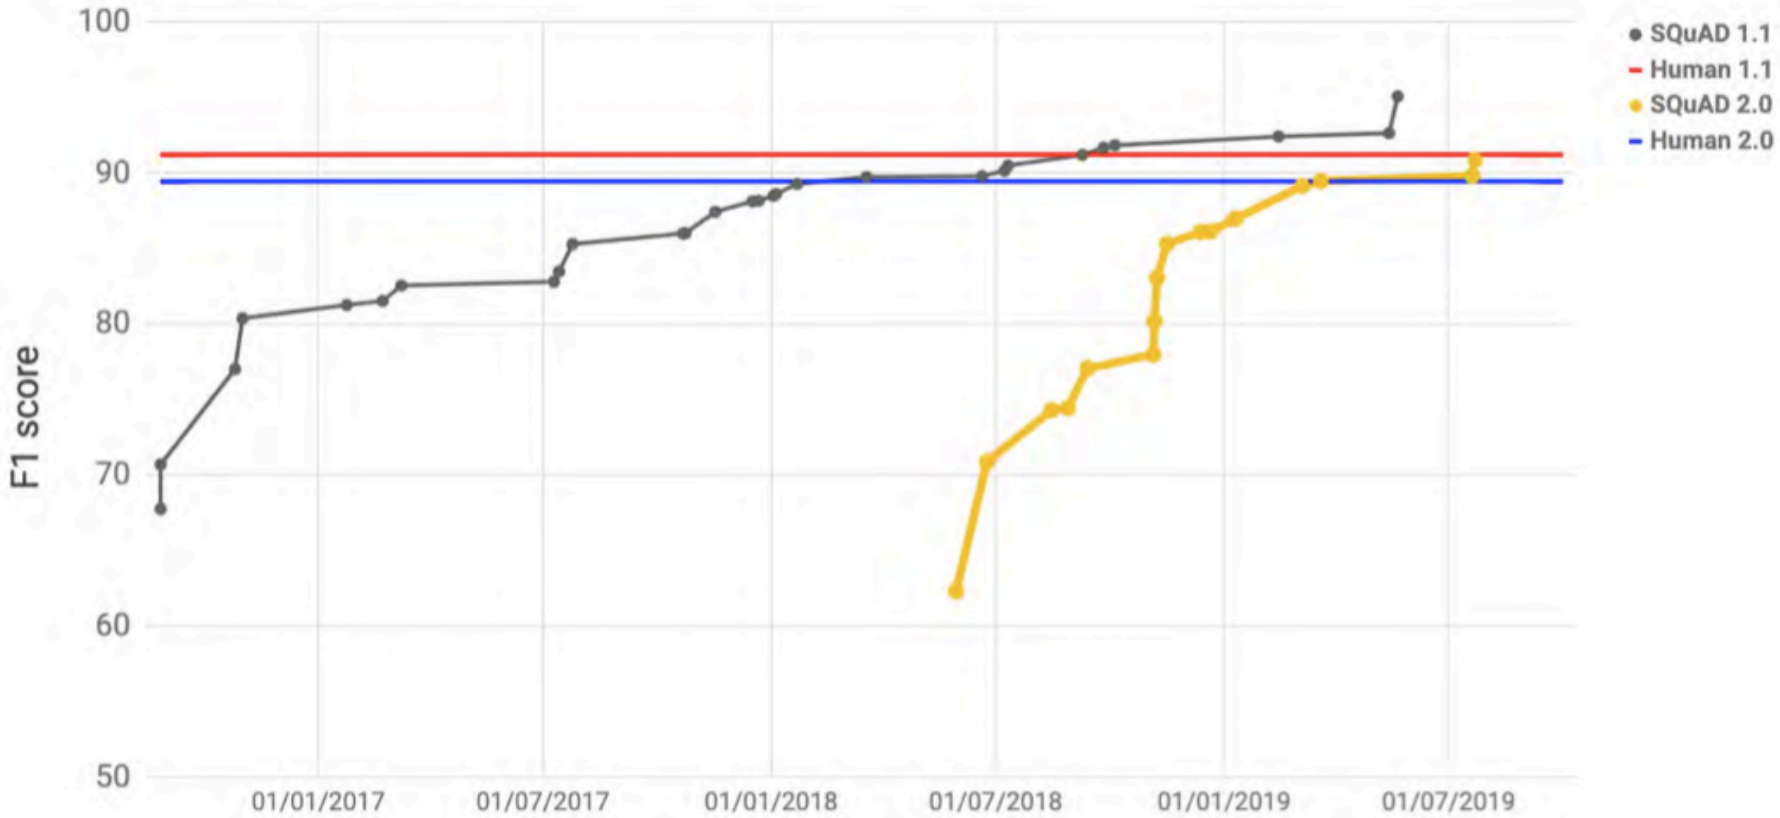


Fig. 3.1.

# ML Performance – Q&A Use Case

SQuAD 1.1 and SQuAD 2.0 - F1 score

Source: CodaLab Worksheets, 2019.



# AI Research Papers on Arxiv

Number of AI papers on arXiv, 2010-2019

Source: arXiv, 2019.

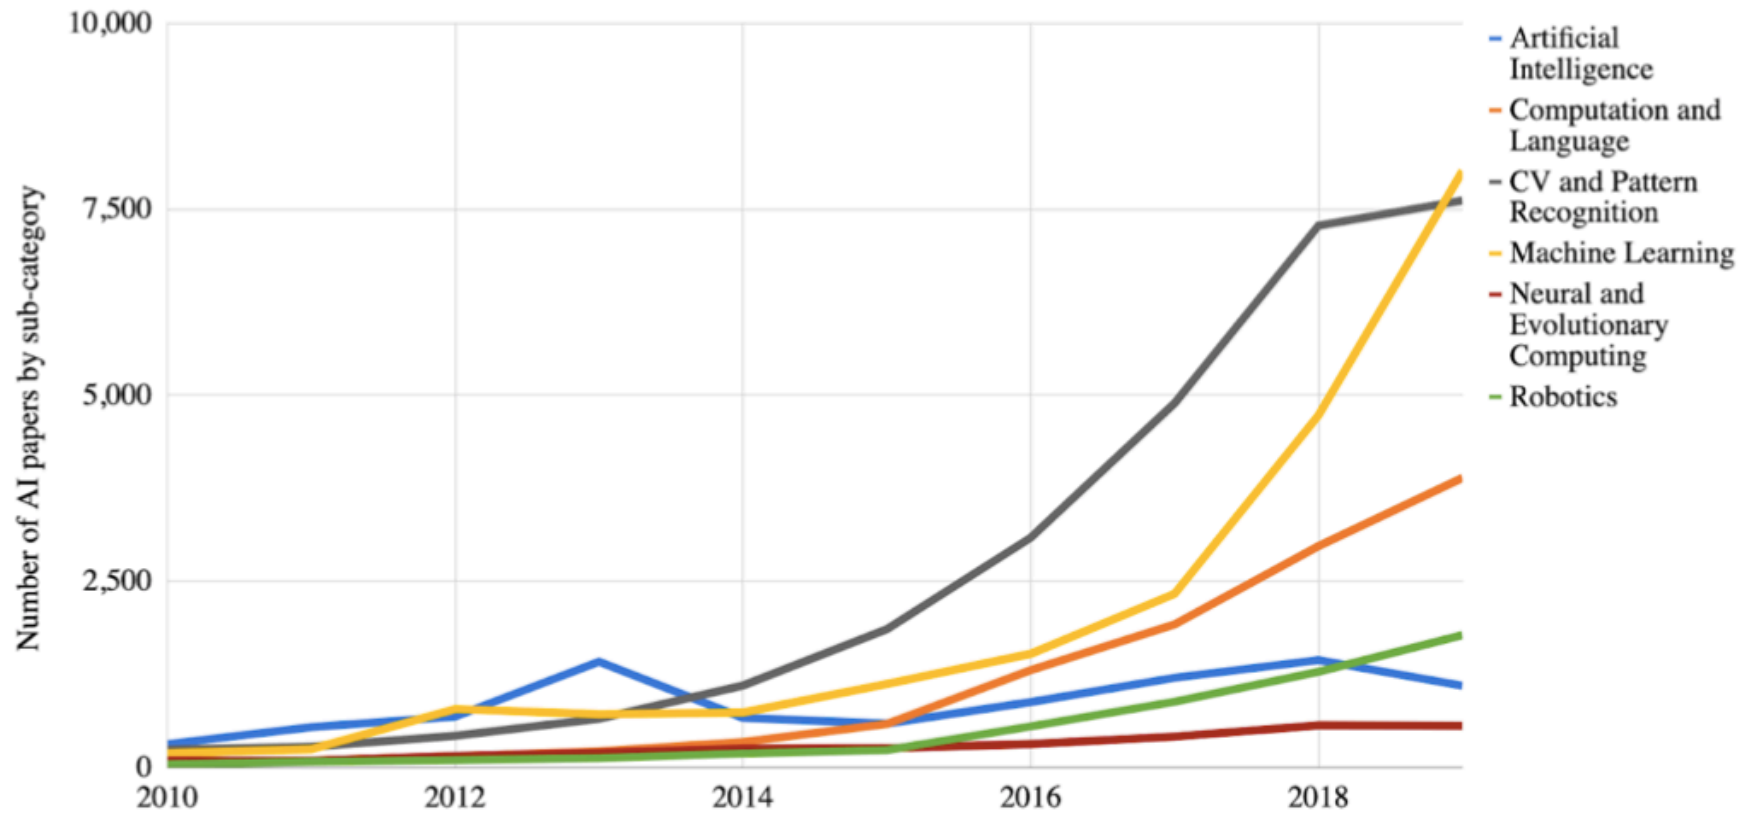


Fig. 1.6.

# Machine Learning Application: Viability

---

**Quality**

**Performance**



Accuracy of Prediction

Speed of Training  
Speed of Inference



# ML Performance – Training Time

ImageNet training time (October 2017 – November 2019)

Source: Stanford DAWN Project, 2019.

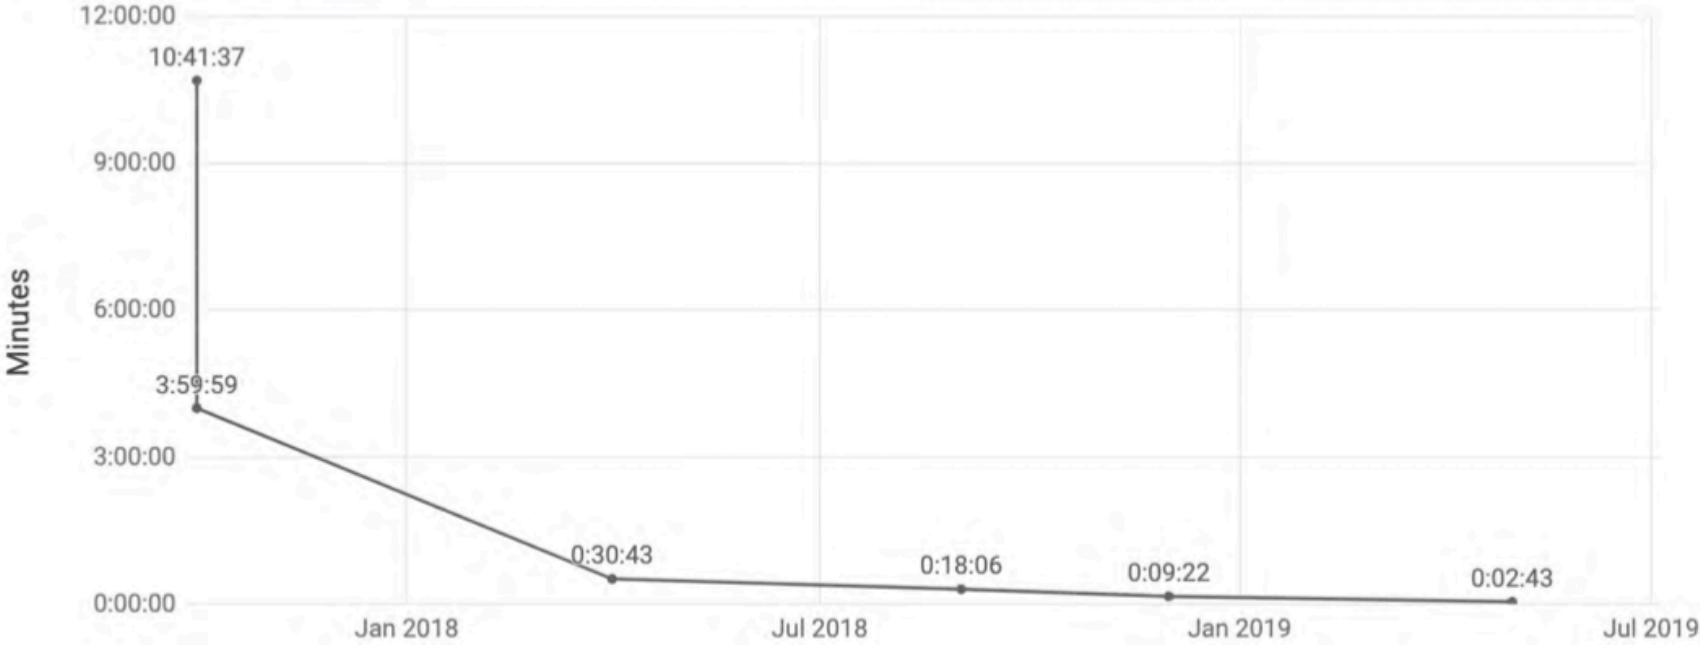


Fig. 3.2a.

Note: [DAWNBench](#) will migrate to [MLperf](#). The latest point estimate (not shown) from ML Perf is from July, 2019 at 1 minute and 28 seconds uses Top-1 accuracy versus Top-5 accuracy benchmark shown in the graph above.

# Machine Learning Application: Viability

---

**Quality**

**Performance**

**Cost**



Accuracy of Prediction

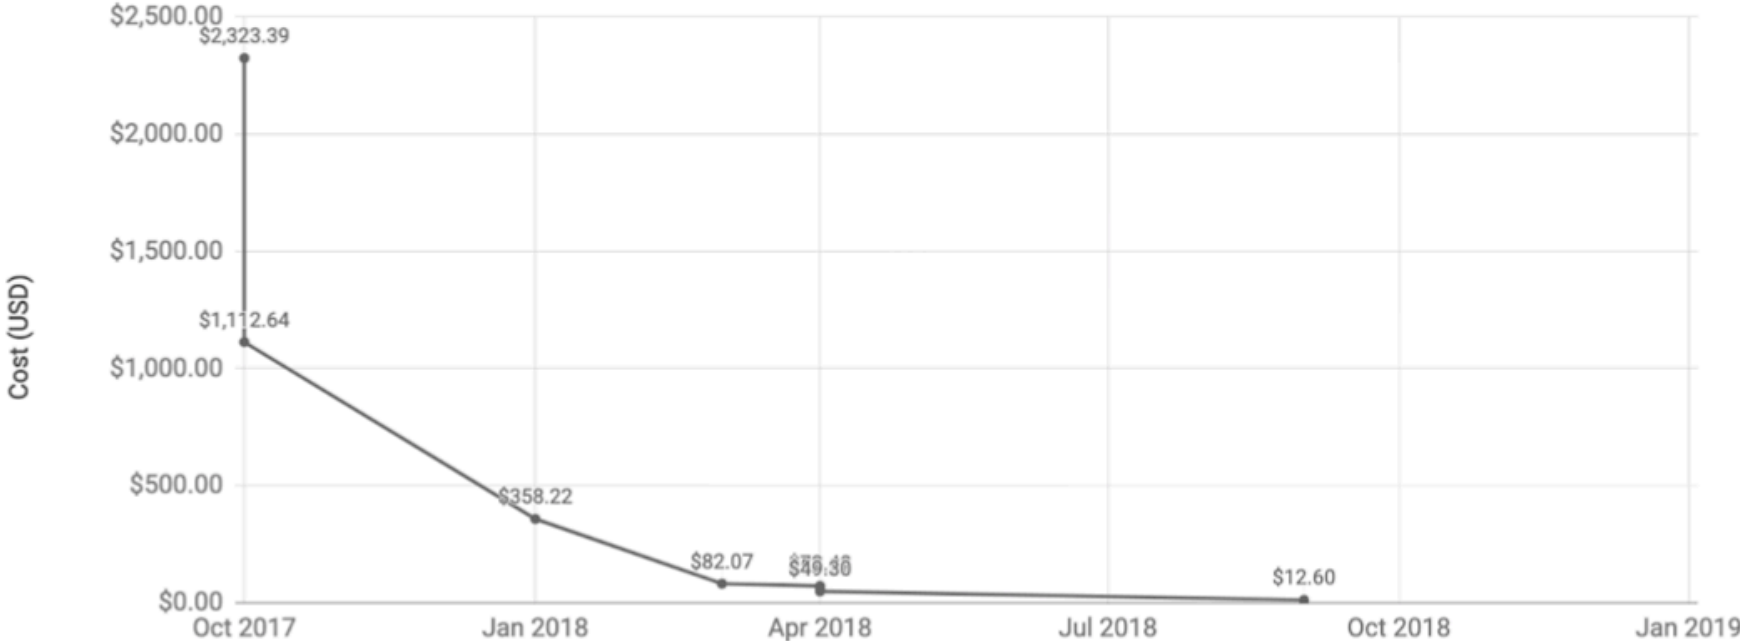
Speed of Training  
Speed of Inference

Public Cloud Cost  
Power consumption

# ML Performance – Training Costs

### ImageNet Training Cost

Source: Stanford DAWN Bench, 2019.



# ML Performance – Training Costs

## Meena Conversation 10

**Human:** *Hi!*  
**Meena:** *Hi, how's it going?*  
**Human:** *All good! It's cold*  
**Meena:** *It's pretty cold here too.*  
**Human:** *where do you stay?*  
**Meena:** *Wisconsin. How about you?*  
**Human:** *San Jose, Cali*  
**Meena:** *Oh nice! I've always wanted to go to Cali.*  
**Human:** *Have you been here before?*  
**Meena:** *No, I haven't. Why?*  
**Human:** *Just curious*  
**Meena:** *Oh, okay. What's it like?*  
**Human:** *It is full of natural beauty*  
**Meena:** *Oh, nice! I've always wanted to go there!*  
**Human:** *you should come visit*  
**Meena:** *I would love to! :D*  
**Human:** *plan a trip soon*  
**Meena:** *I would love to! Where should I go? :D*  
**Human:** *Bay area*  
**Meena:** *Oh, I've always wanted to go! :D*



Elliot Turner  
@eturner303

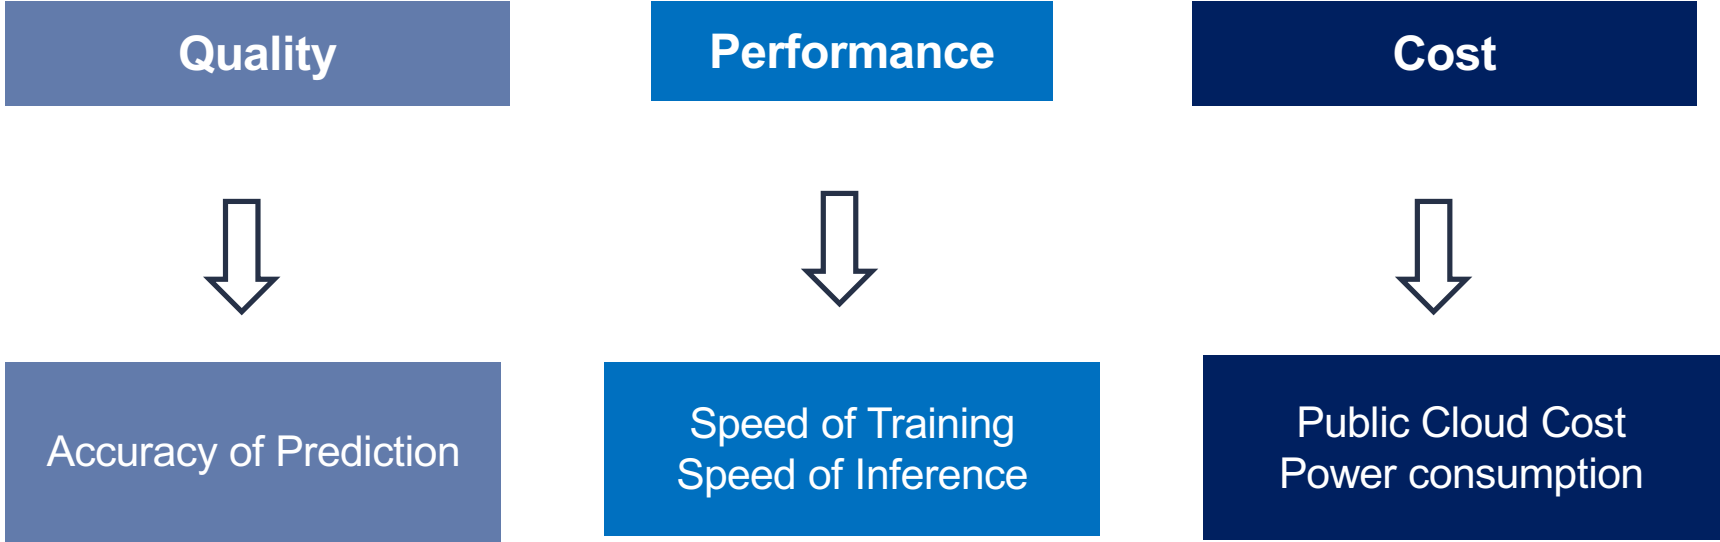
Wow: Google's "Meena" chatbot was trained on a full TPUv3 pod (2048 TPU cores) for **\*\*30 full days\*\*** - That's more than \$1,400,000 of compute time to train this chatbot model. (! 100+ petaflops of sustained compute !)

## Meena Conversation 10

**Human:** *Hi!*  
**Meena:** *Hi, how's it going?*  
**Human:** *All good! It's cold*  
**Meena:** *It's pretty cold here too.*  
**Human:** *where do you stay?*  
**Meena:** *Wisconsin. How about you?*  
**Human:** *San Jose, Cali*  
**Meena:** *Oh nice! I've always wanted to go to Cali.*  
**Human:** *Have you been here before?*  
**Meena:** *No, I haven't. Why?*  
**Human:** *Just curious*  
**Meena:** *Oh, okay. What's it like?*  
**Human:** *It is full of natural beauty*  
**Meena:** *Oh, nice! I've always wanted to go there!*  
**Human:** *you should come visit*  
**Meena:** *I would love to! :D*  
**Human:** *plan a trip soon*  
**Meena:** *I would love to! Where should I go? :D*  
**Human:** *Bay area*

# Machine Learning Application: Viability

---



**Quality, Performance, and Cost metrics important for commercial viability**

# Further Reading: Success Factors for ML Companies (MMC Ventures)





# Further Reading: AI & M&A (pwc)

---



# Further Reading: 10 Best Practices (Law 360)

---

## 10 Best Practices For Due Diligence In AI Transactions

By Lee Tiedrich and Daniel Gurman

(September 21, 2018, 11:39 AM EDT)

Artificial intelligence is projected to contribute over \$15 trillion to the global economy by 2030,[1] which in turn suggests that the number of AI-focused mergers and acquisitions, investments and other types of transactions will continue to grow. Here are 10 best practices for approaching AI-focused transactions, taking into consideration the nature of the technology today, the anticipated technological developments and the evolving legal landscape pertaining to AI.[2]

### 1. Understanding the Transaction

As in any transaction, it is important to identify and understand at the outset of an AI transaction the core assets that support the valuation of the target's business. For example, at its core, there are three central components to AI — the hardware, the software and the data (including training data and AI outputs). Knowing upfront which of these components is material to the target's business will help tailor the buyer's diligence, and it may influence the terms of the purchase agreement and inform the buyer of actions it may want to undertake post-closing. The following best practices provide guidance on how to tailor the due diligence to focus on the core AI assets and can inform the buyer of steps it may want to take during the negotiations or before or after the closing.



Lee Tiedrich



# Further Reading: AI vs. Traditional Software Business (a16z)

---

## The New Business of AI (and How It's Different From Traditional Software)

by Martin Casado and Matt Bornstein

AI, machine & deep learning •  
enterprise & SaaS •  
Company Building 101 •  
economics



SUBSCRIBE TO OUR  
NEWSLETTER!

Email Address\*



At a technical level, artificial intelligence seems to be the future of software. AI is showing remarkable progress on a range of difficult computer science problems, and the job of software developers – who now work with data as much as source code – is changing fundamentally in the process.

Many AI companies (and investors) are betting that this relationship will extend beyond just technology – that AI *businesses* will resemble traditional software companies as well. Based on our experience working with AI companies, we're not so sure.

We are huge believers in the power of AI to transform business: We've put our money behind that thesis, and we will continue to invest heavily in both applied AI companies and AI infrastructure. However, we have noticed in many cases that AI companies simply don't have the same economic construction as software businesses. At times, they can even look more like traditional services companies. In particular, many AI companies have:

1. **Lower gross margins** due to heavy cloud infrastructure usage and ongoing human support;
2. **Scaling challenges** due to the thorny problem of edge cases;
3. **Weaker defensive moats** due to the commoditization of AI models and challenges with data network effects.

# Key Points

- ① AI Interest and Investments on the rise
- ② Due Diligence drives Go/No-Go & Valuations
- ③ SaaS DD is mature; AI DD is not
- ④ DD on AI is Art + Science
- ⑤ Data, Model, Process, Infra, Team, Risks
- ⑥ Domain Benchmarks for AI vs Traditional
- ⑦ Quality, Performance, Cost drive Viability
- ⑧ Premium for “AI” products and startups
- ⑨ Biz problem influences VC decisions
- ⑩ Simple Tech + Big Opportunity = Unicorn?



# Due Diligence Framework

