Amazon: a Playground for Machine Learning & Al

Cedric Archambeau

Principal Applied Scientist

cedrica@amazon.com

Current Trends in Artificial Intelligence Copenhagen, November 2017



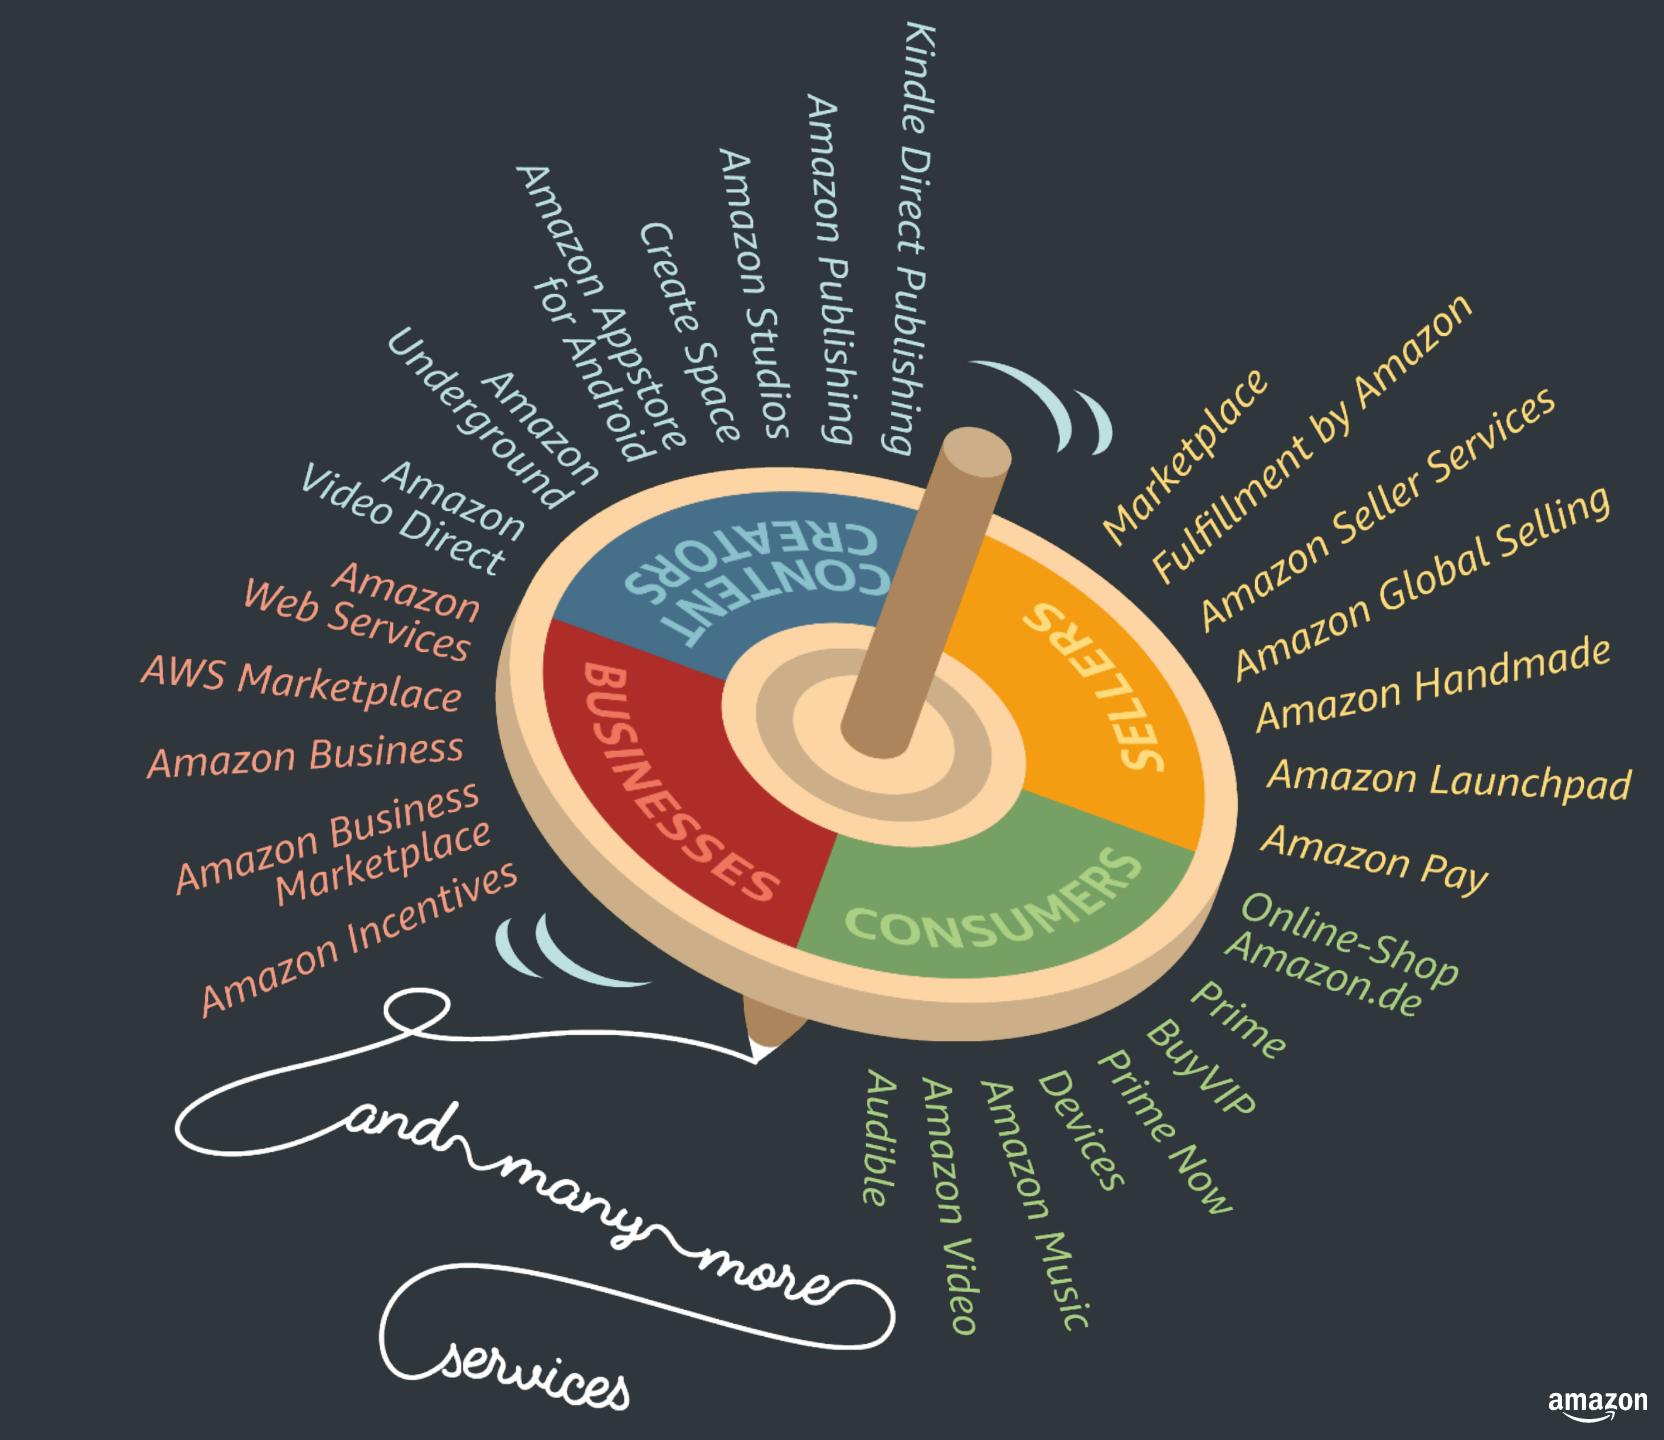
Amazon.com, 2017



Over 300 million active customer accounts worldwide.

More than 2 million active seller accounts.

Over a million active AWS accounts.



Is Amazon a retailer or a tech company? A marketplace or a logistics company? A publisher or a movie studio?



Today, machine learning is creating a paradigm shift.

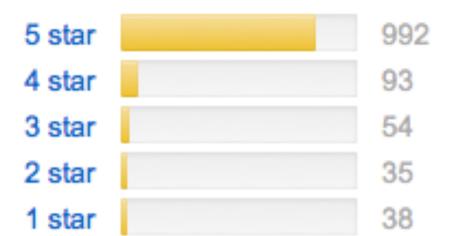
"It is a golden age. Machine learning and AI is a horizontal enabling layer. It will empower and improve every business, every government organization, every philanthropy."

Jeff Bezos in Geekwire (May 6, 2017)

Customer Reviews

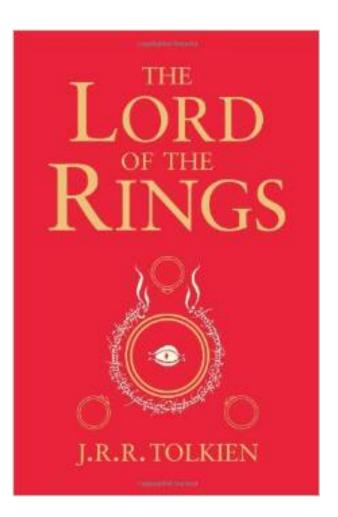


4.6 out of 5 stars



Share your thoughts with other customers

Write a customer review



See all 1,212 customer reviews >

Most Helpful Customer Reviews

38 of 39 people found the following review helpful



By A Customer on 18 Mar 2002

Format: Paperback

JRR Tolkien's book "The Lord of the Rings" needs no introduction. It's a timeless classic that should at least be given a try by absolutley everyone! That said, this edition is just great! I was quite surprised when I first saw this edition with my own eyes. The books are large, pages are thick, the covers and the illustrations are a feast to look at and the text is big and clear. If you are looking for an edition of "The Lord of the Rings" that is a bit more special than the cheaper ones, this is the one to get!

Comment

Was this review helpful to you?





78 of 81 people found the following review helpful



★★★★★ Trully unbelievable!

By A Customer on 26 Feb 2000

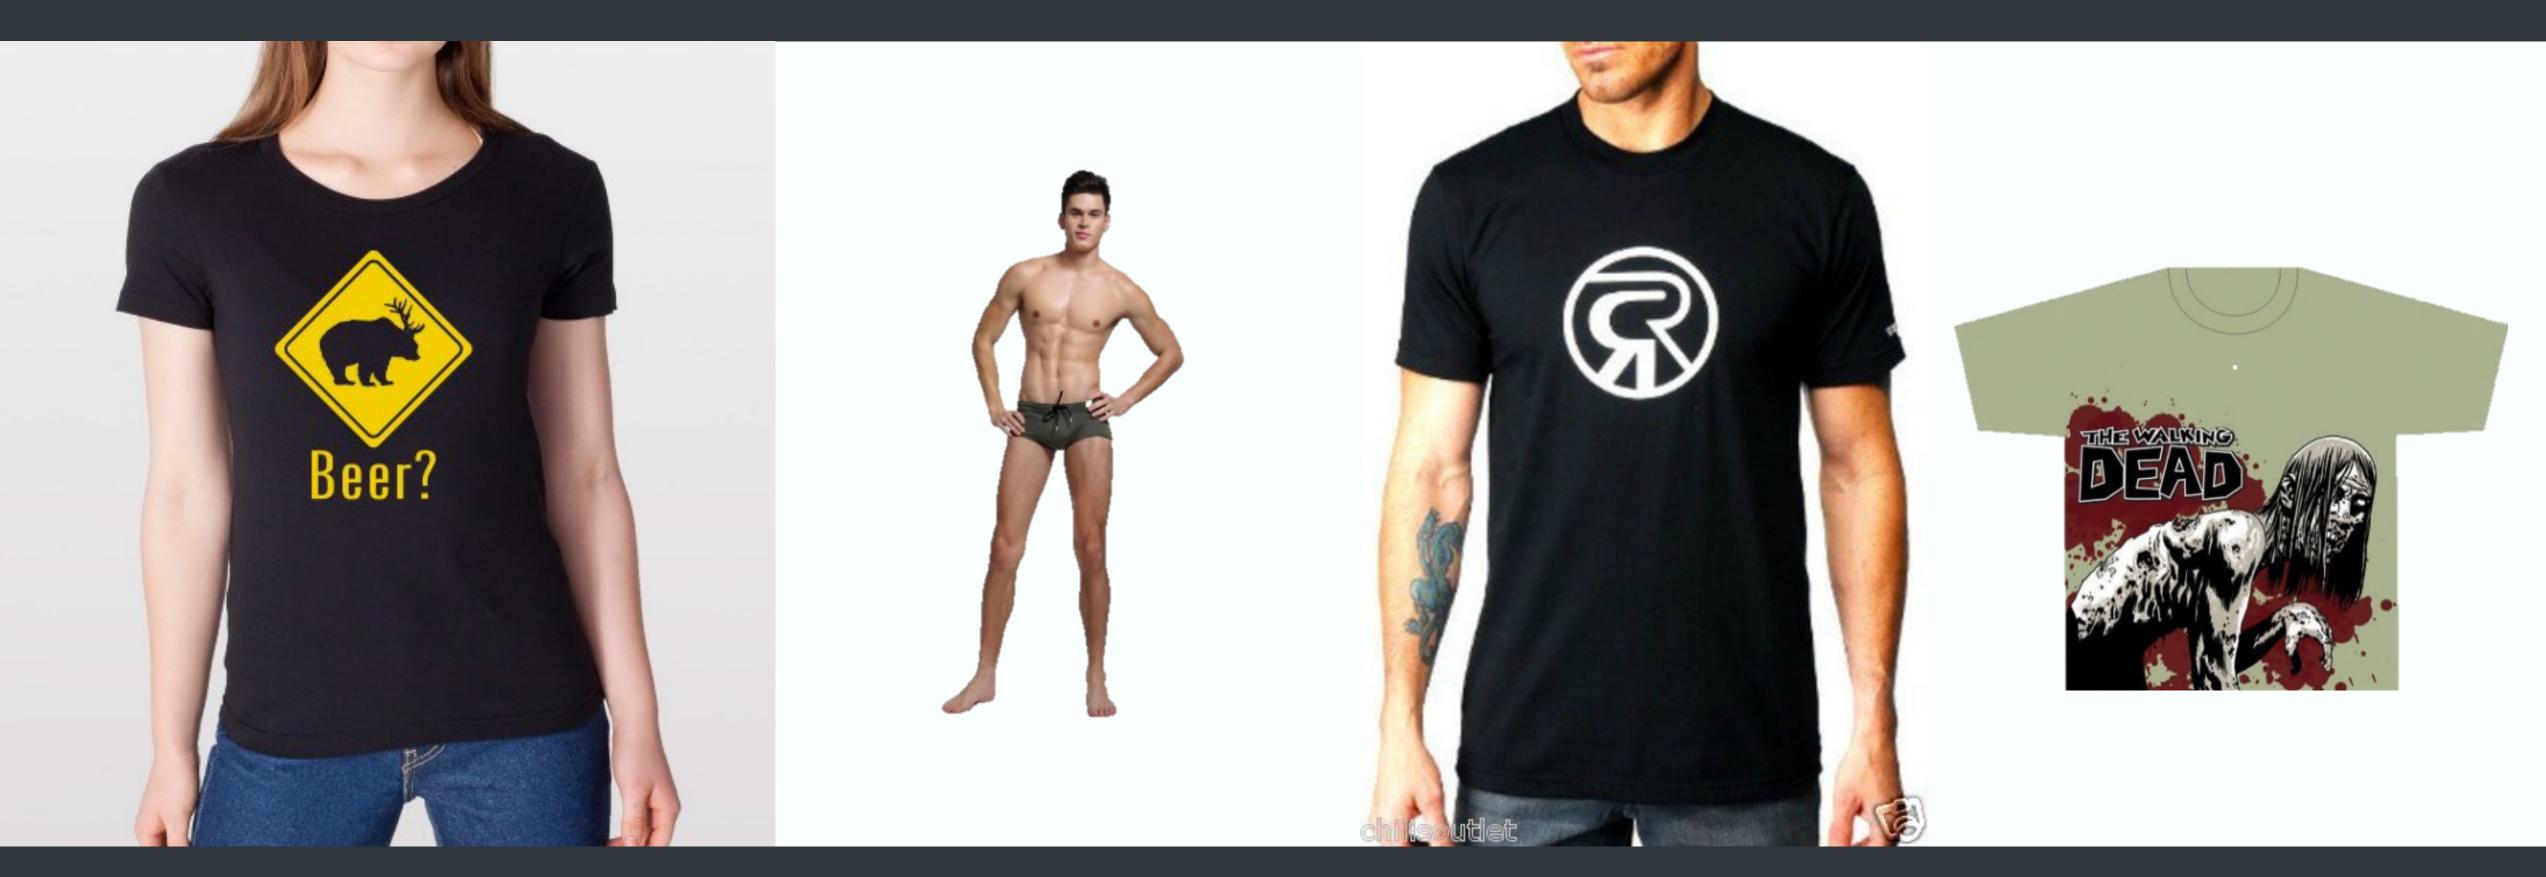
Format: Misc. Supplies

I know, this might look like an expensive buy, but believe me, once you get to the last CD of this brilliant dramatisation of Lord of the Rings, you won't even care about the money you've spent! This is how good this radio version is I would give it 1000 stars if i could I

Automatic review moderation



Automatic Image Quality Assurance

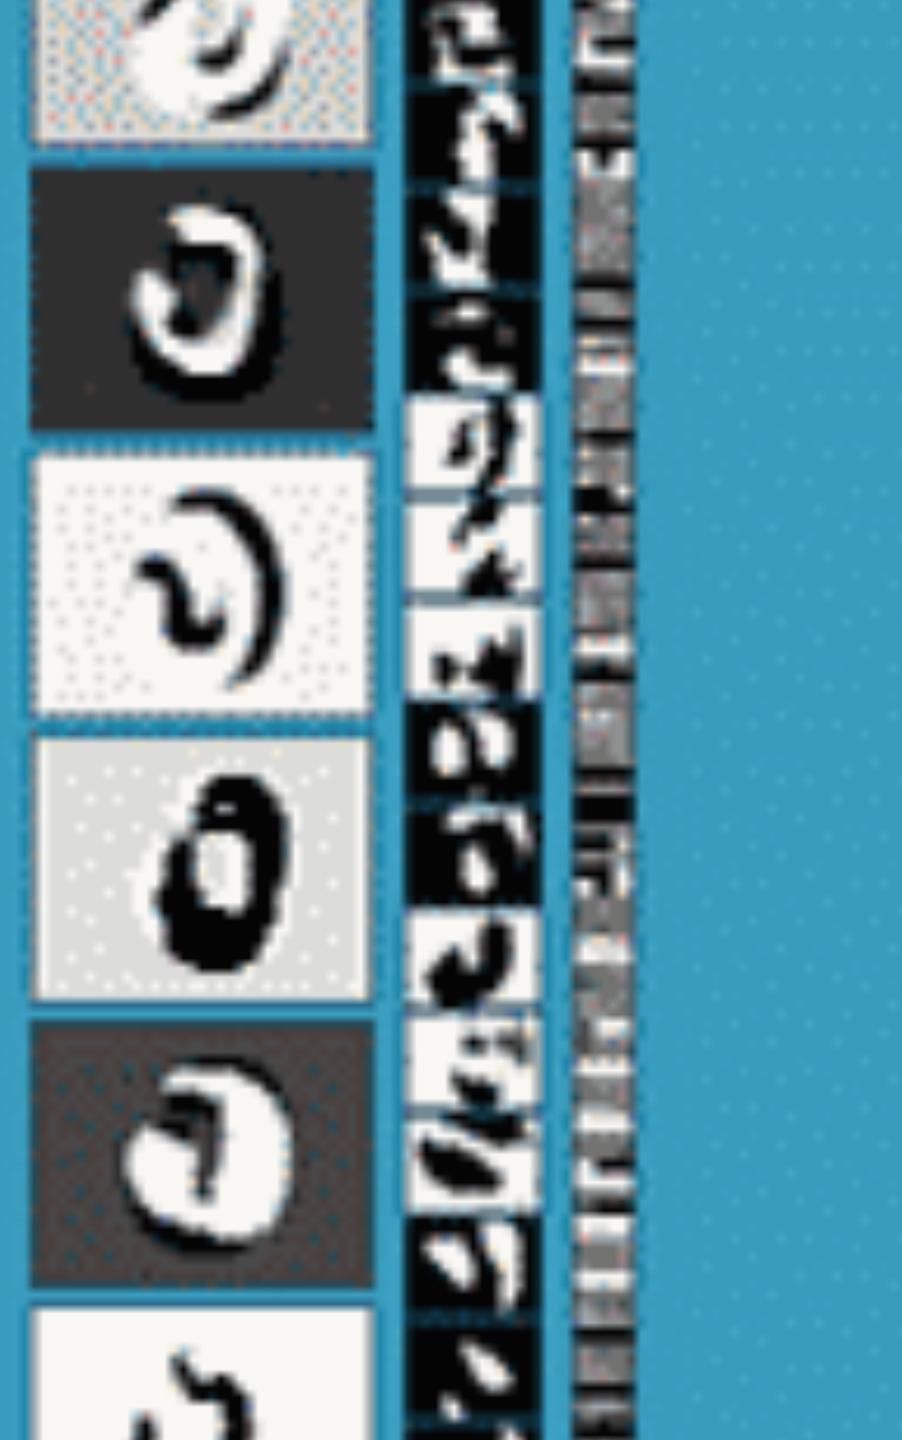


Non-white Background

Product too small

Includes watermark

Not picture, but drawing







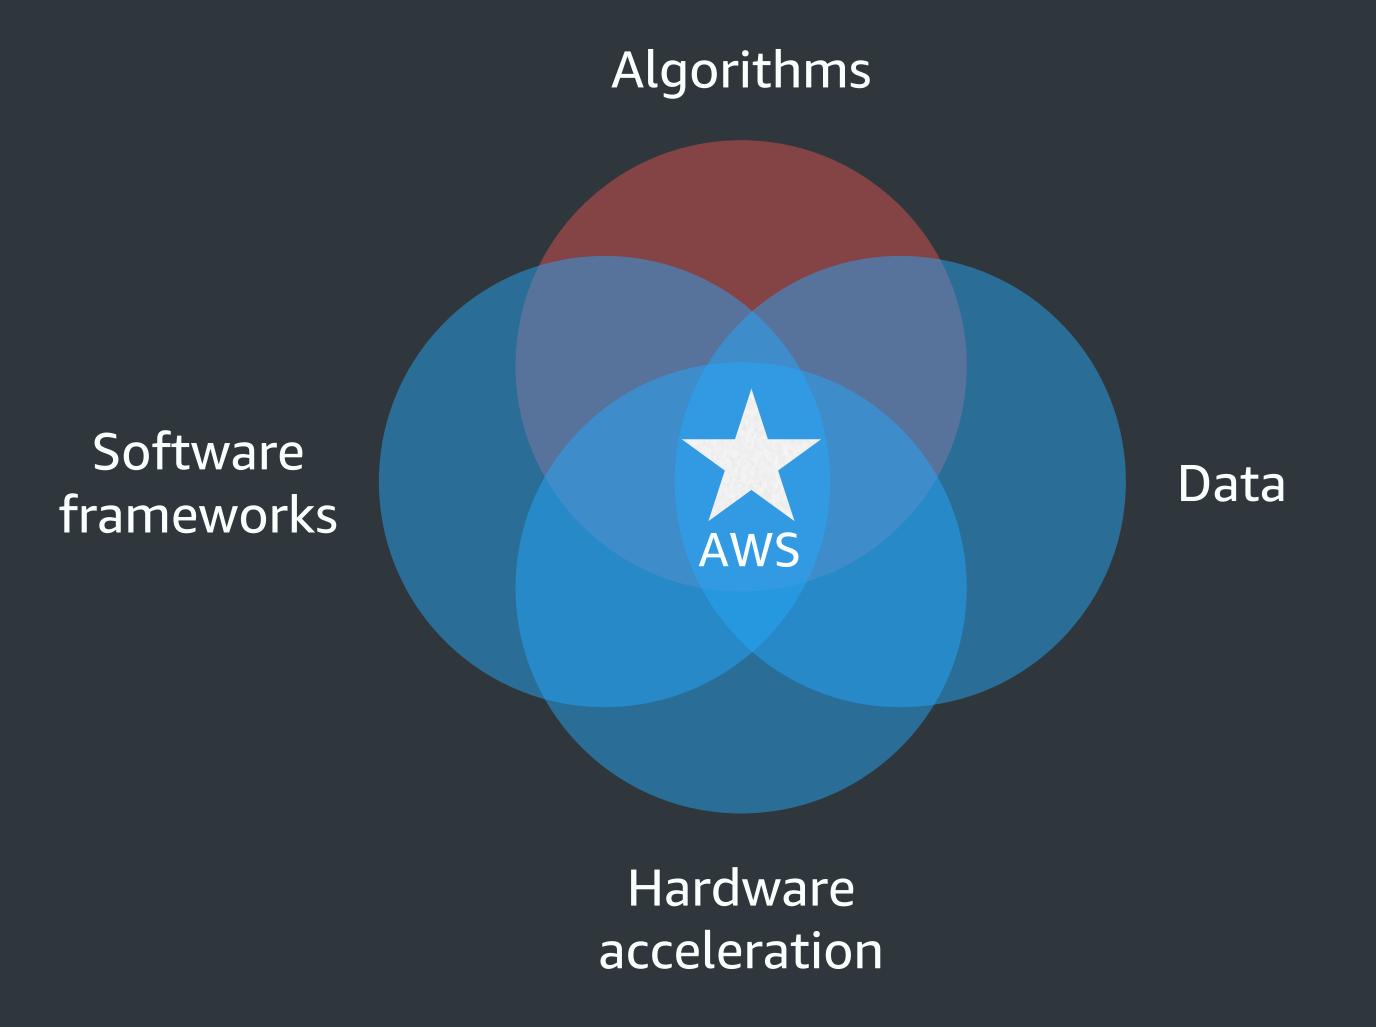


answer:





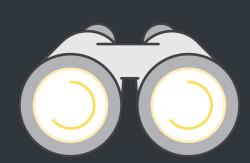
The maturity of deep learning



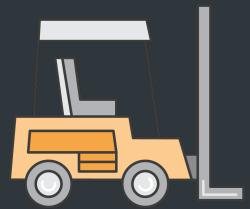


Artificial Intelligence at Amazon

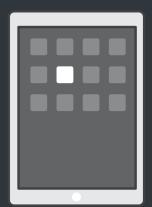
Thousands of employees across the company focused on machine learning & Al



Discovery & Search



Fulfilment & Logistics



Enhance & Create Products



Bring Machine Learning to All





Artificial Intelligence at Amazon

Thousands of employees across the company focused on machine learning & Al



Discovery & Search



Fulfilment & Logistics



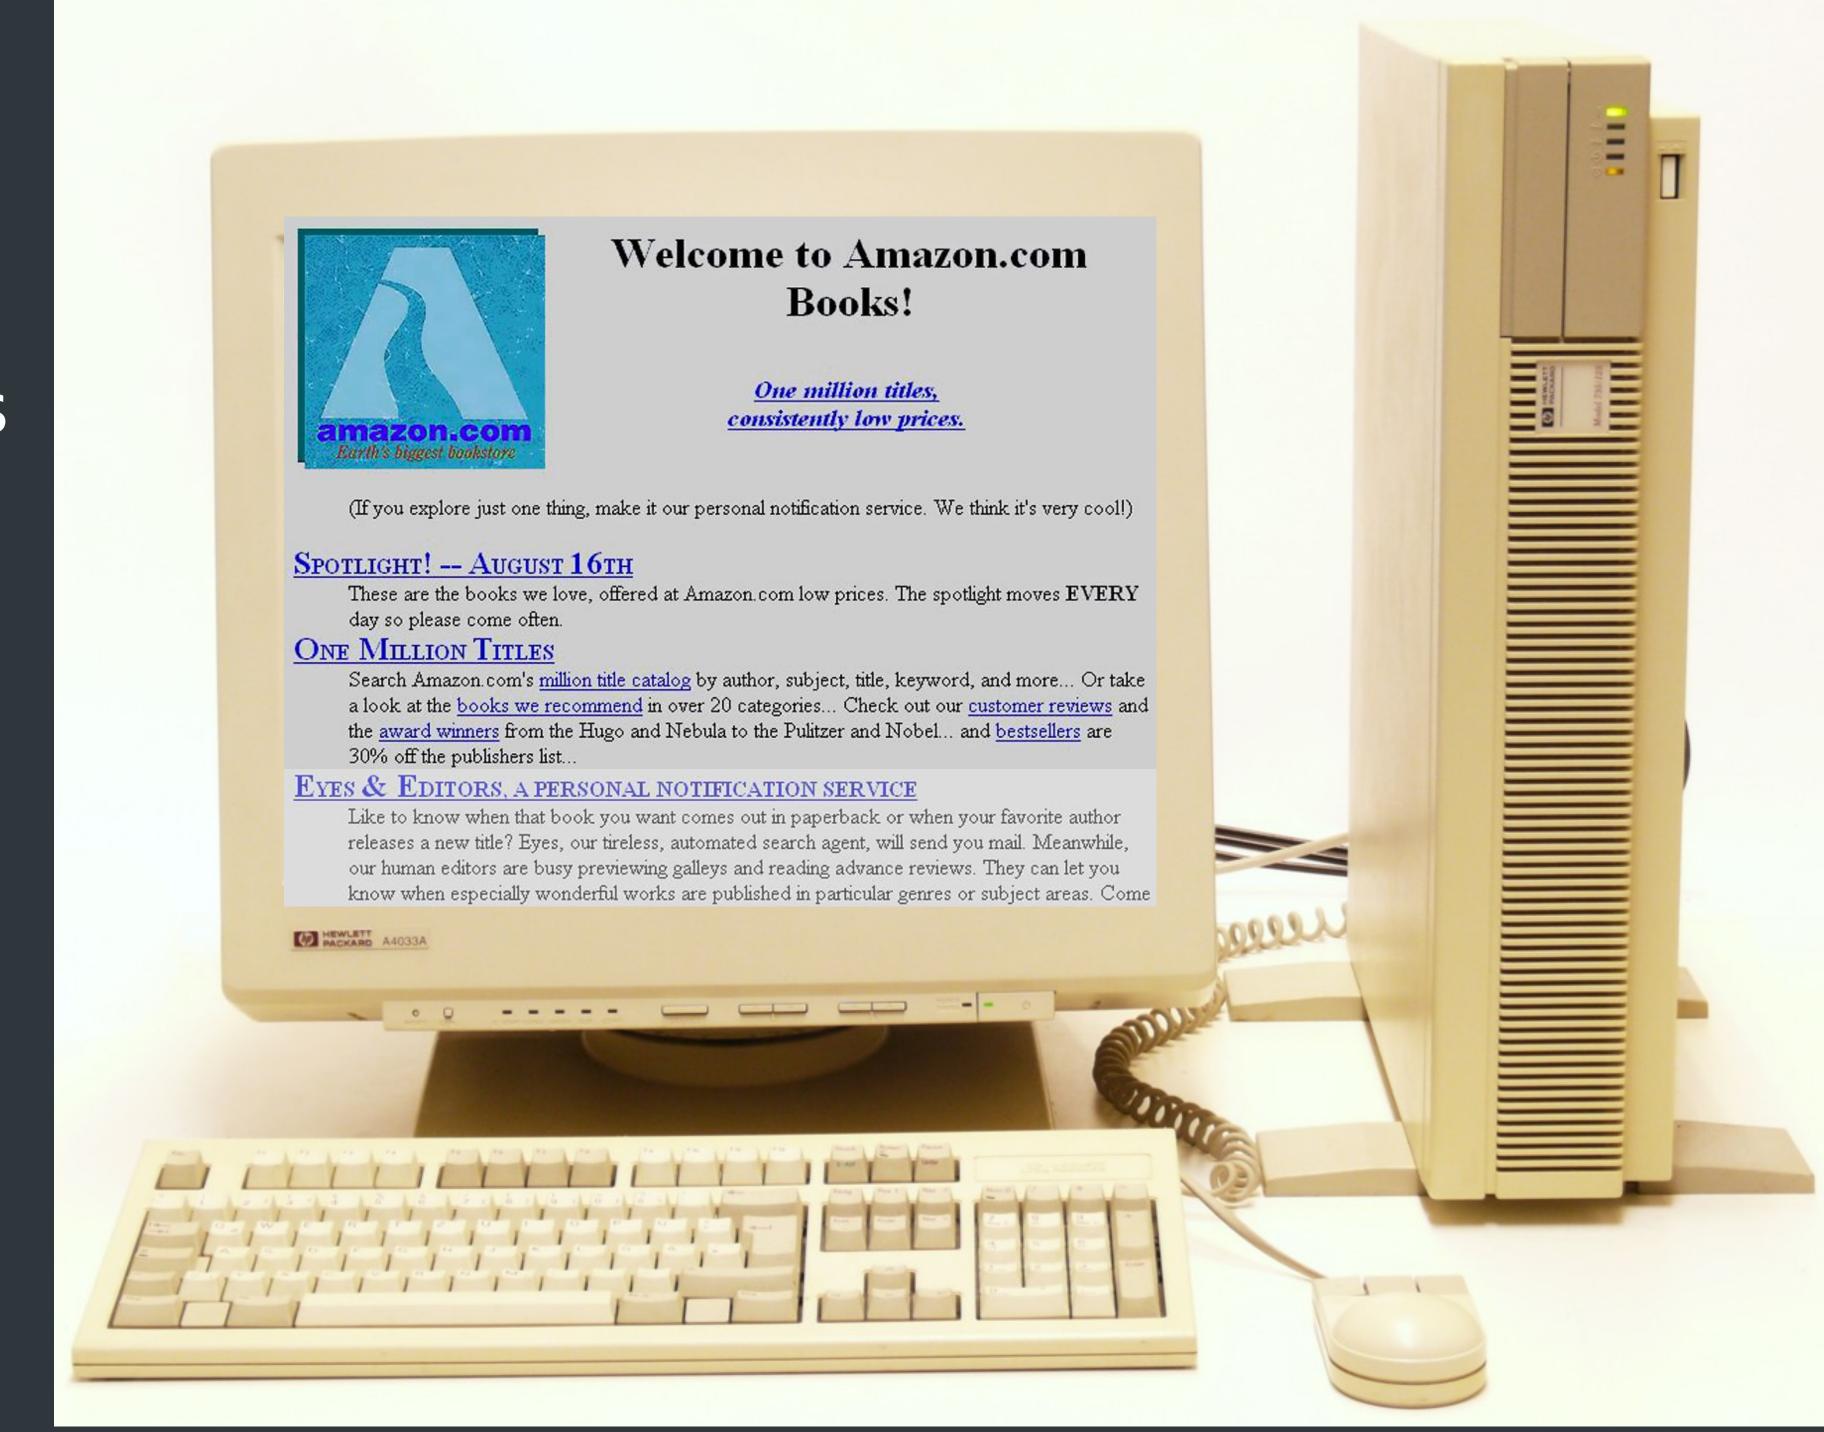
Enhance & Create Products



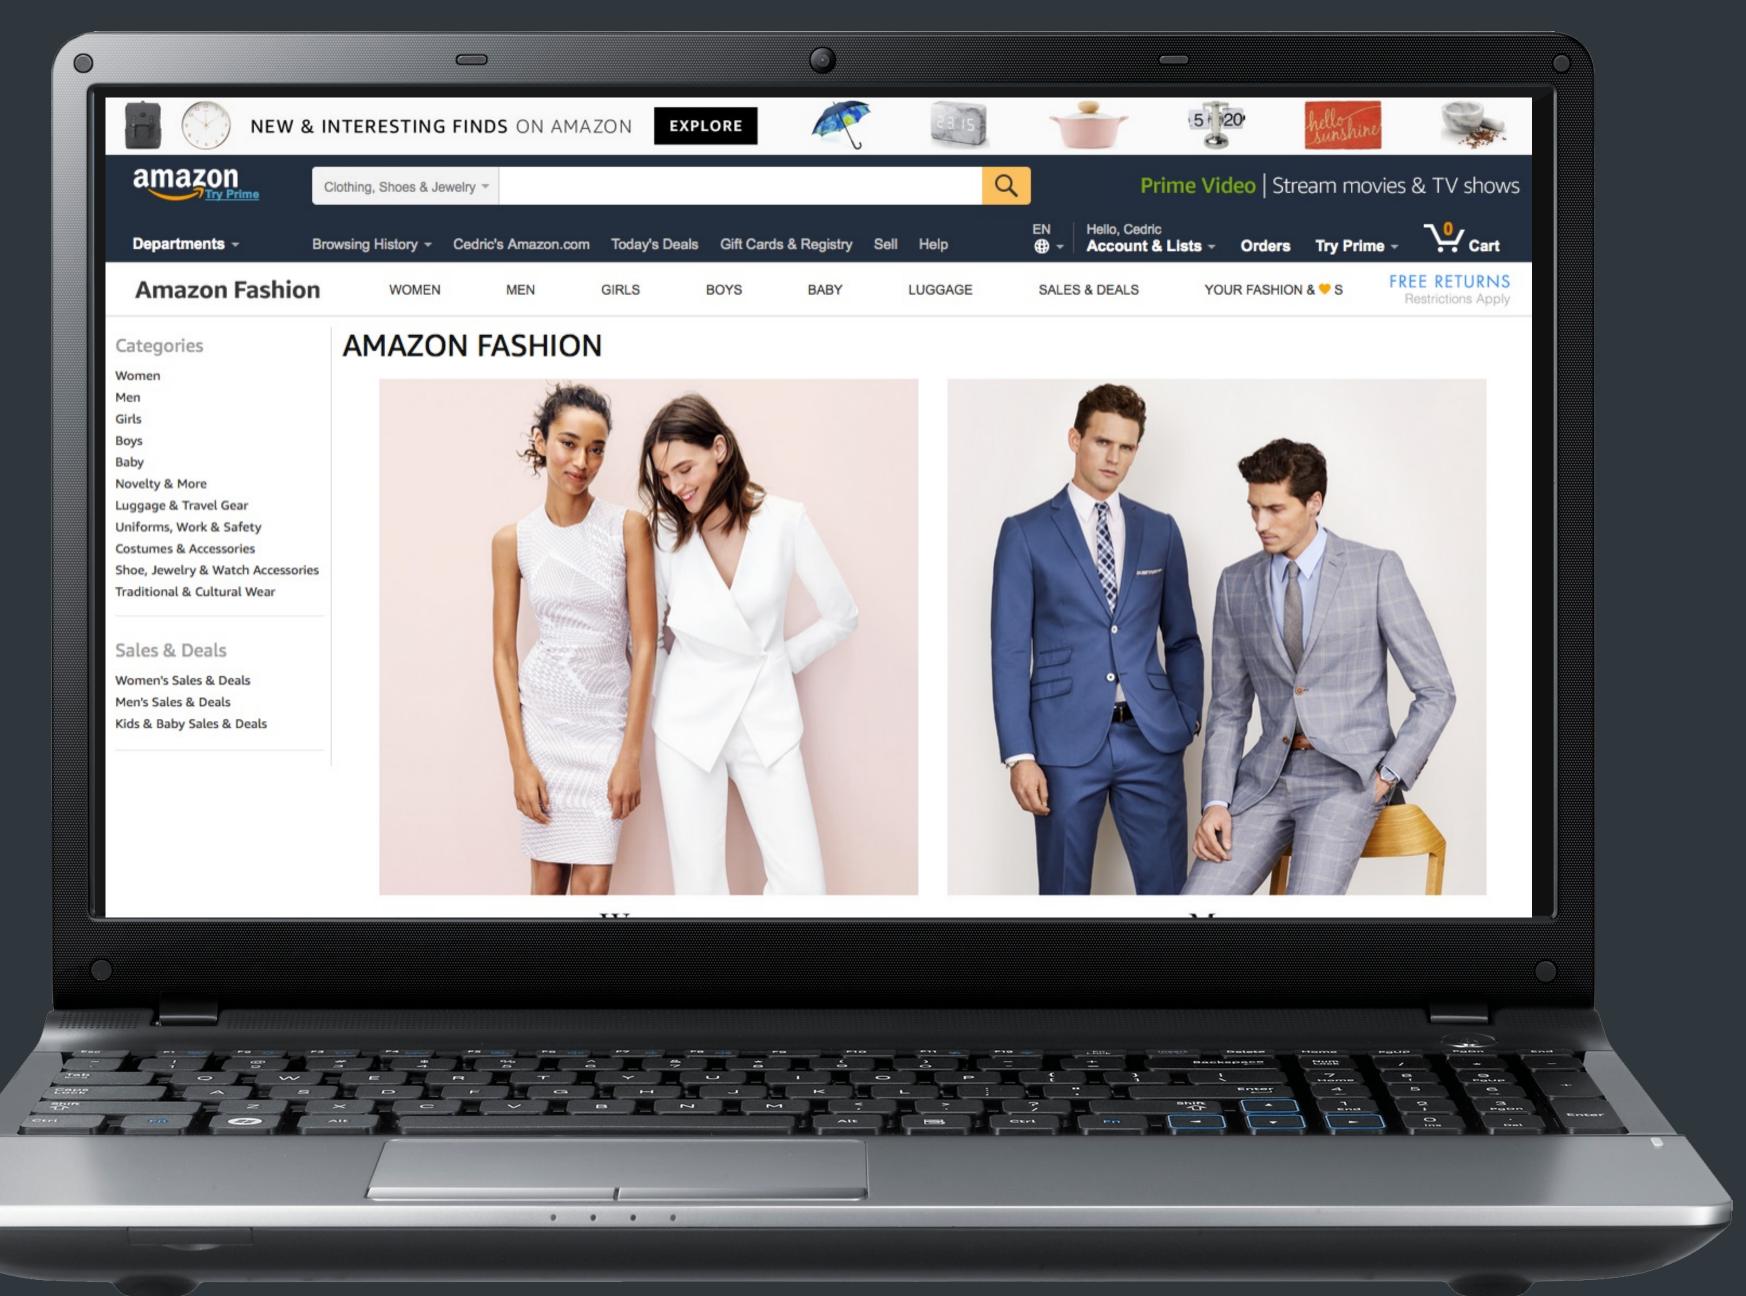
Bring Machine Learning to All



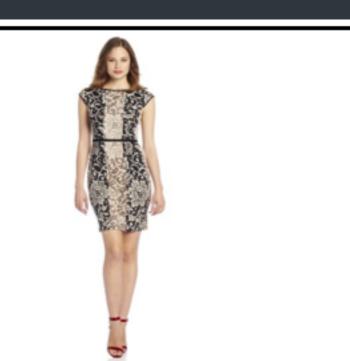
Recommendations & Search



Understanding Fashion & Style





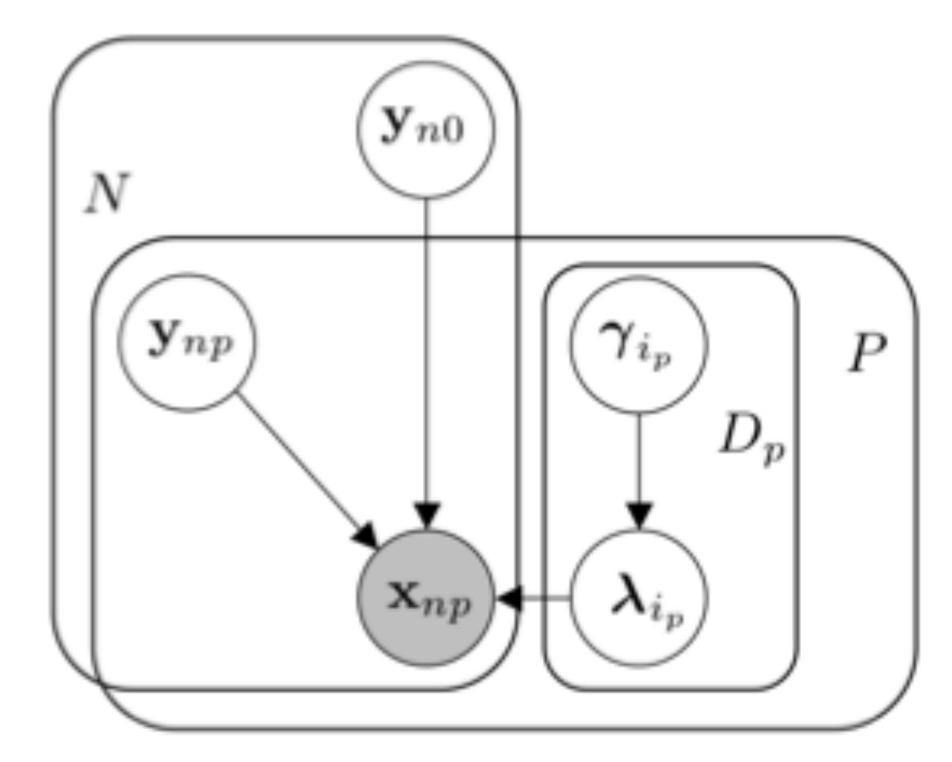


B00DGQ5H3A

Maggy London

Maggy London Women's Petite Printed Ponte Shadow Lace Dress



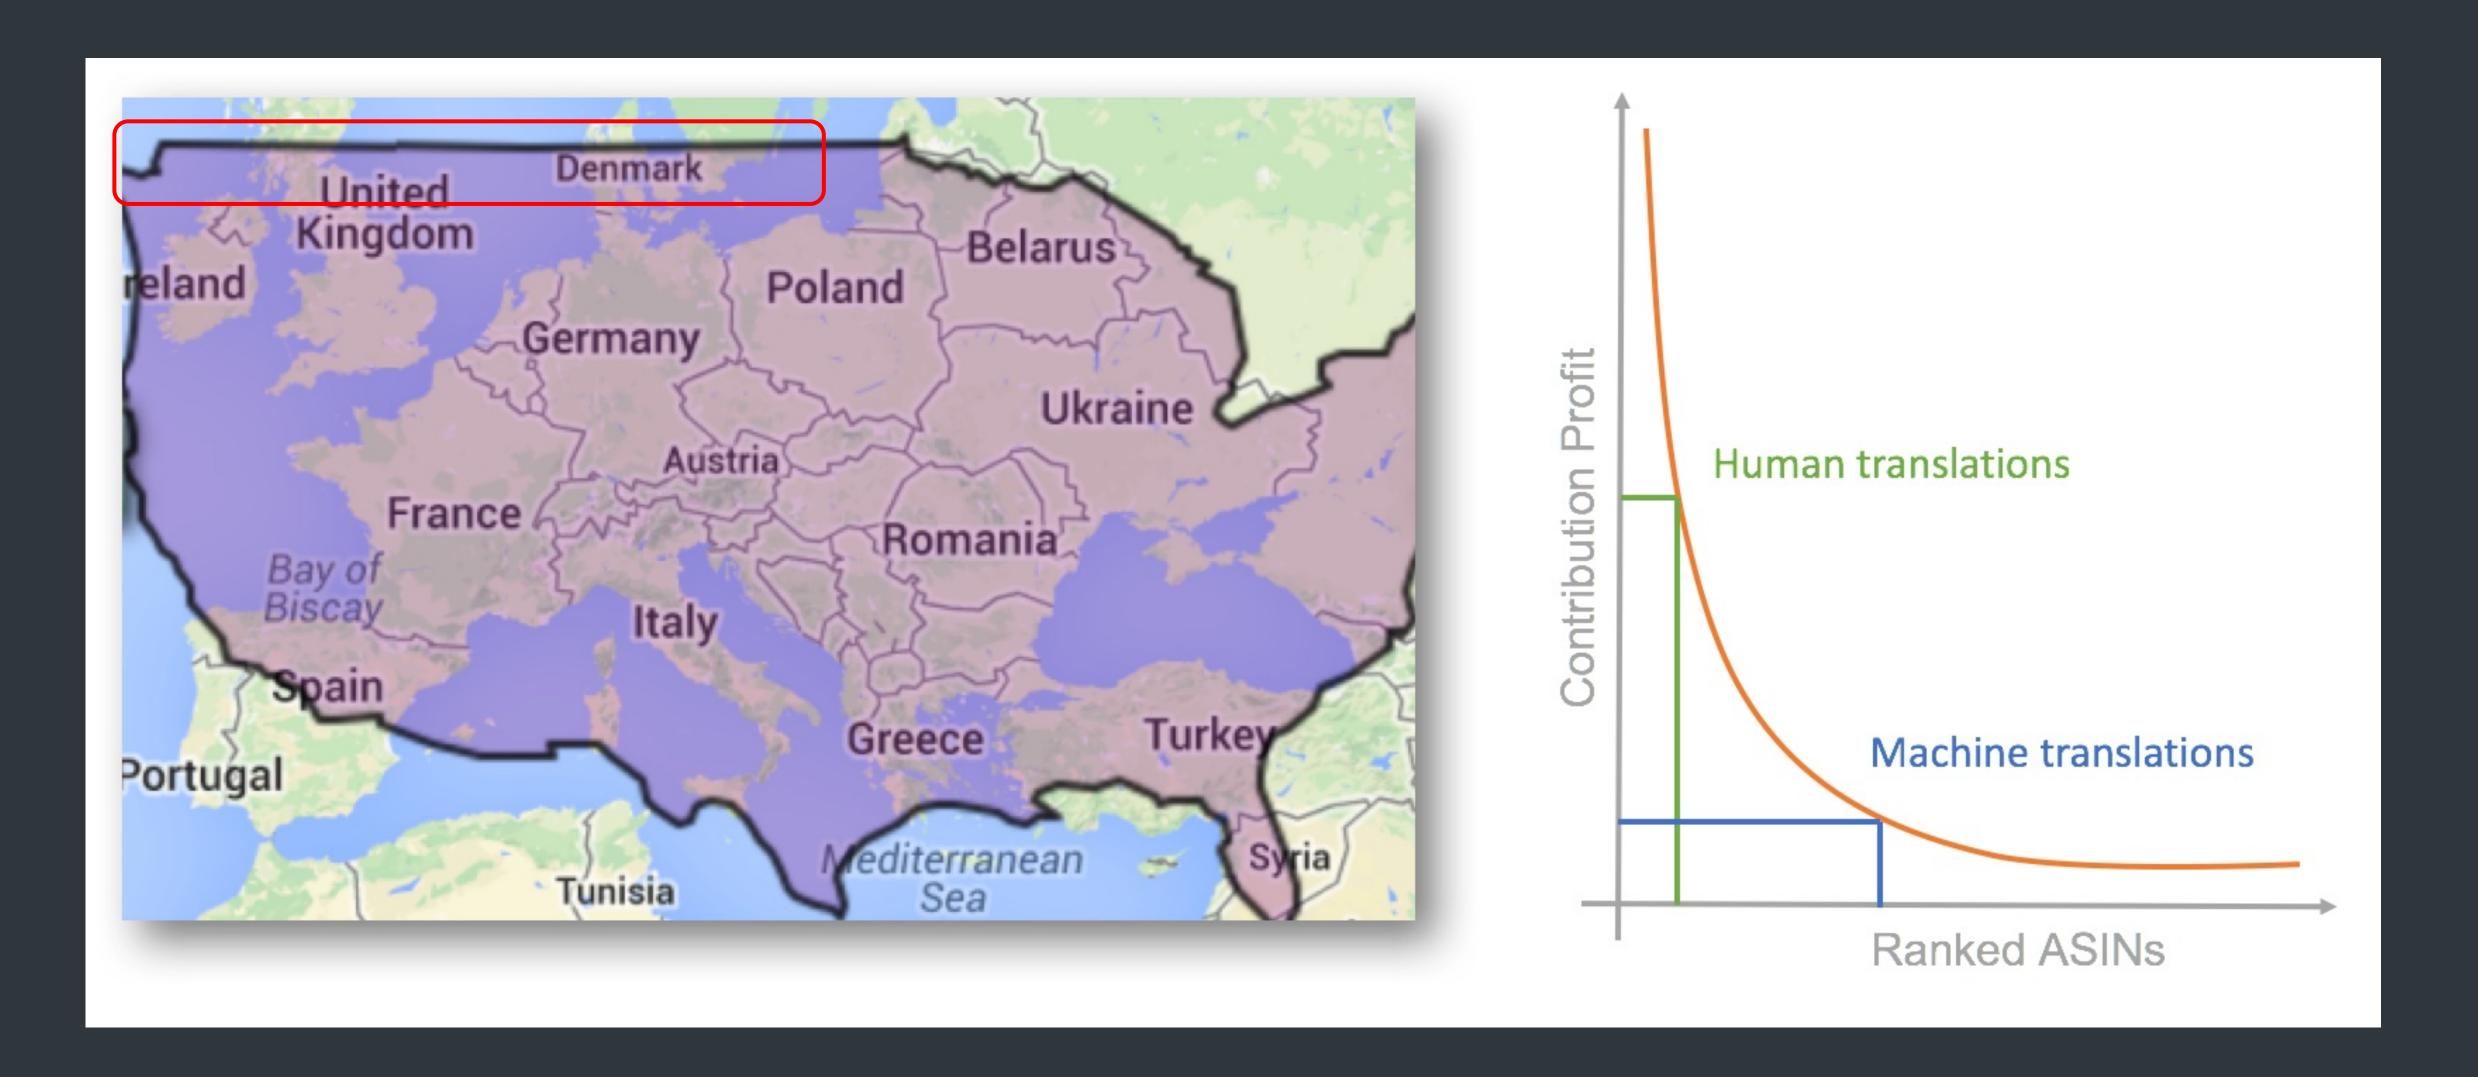


(Archambeau and Bach, NIPS 2008)

Shallow Latent Variable Models



Machine Translated Detail Pages



Neural Machine Translation (NMT) with Sockeye

github.com/awslabs/sockeye

- Open-sourced toolkit for sequence-tosequence modeling in MXNet
- Implements encoder-decoder models with attention (Bahdanau, et al., 2014)
- Supports different attention models (Luong, et al., 2015)
- Applicable to Named Entity Recognition,
 Semantic Parsing, ...



(Image credit: Washington Department of Fish & Wildlife.)

Recurrent Neural Network Language Model

Language model without Markov assumption:

$$P(\boldsymbol{v}) = \prod_{t=1}^{n} P(v_t | \boldsymbol{v}_{1:t-1})$$

Embedding layer:

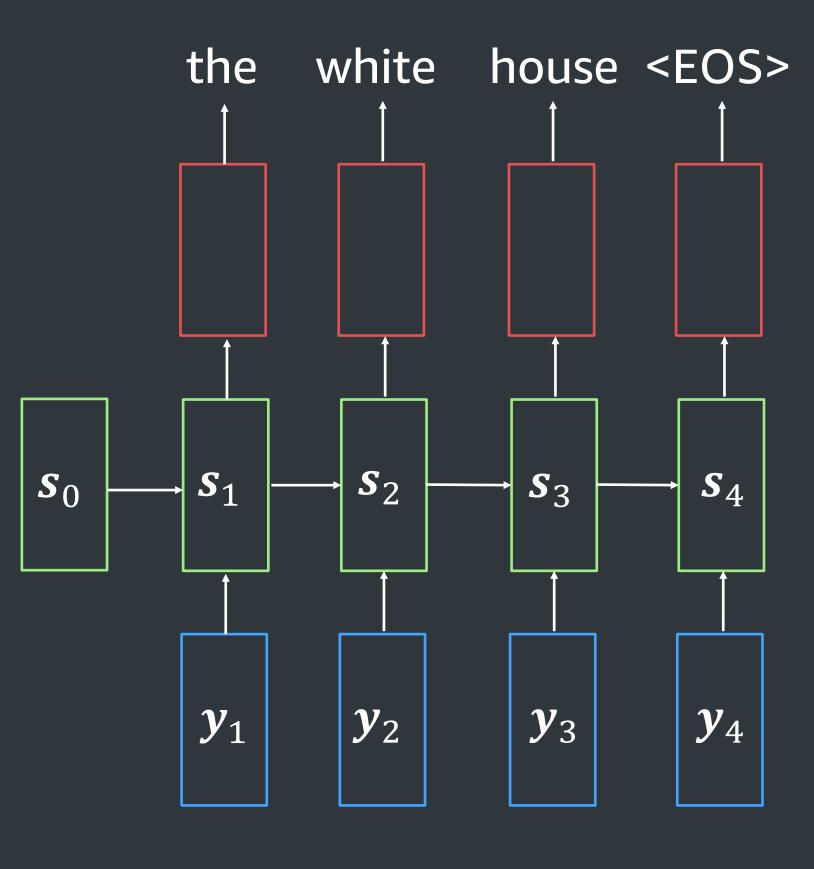
$$\mathbf{y}_t = \mathbf{W}_E \mathbf{v}_t$$

Recurrent hidden layer (e.g., RNN, LSTM, GRU):

$$s_t = \tanh(Us_{t-1} + Wy_t)$$

Output layer:

 $P(\text{house} \mid \langle BOS \rangle, \text{ the, white}) = \text{softmax}(W_O s_3 + b_O)$



<BOS> the white house



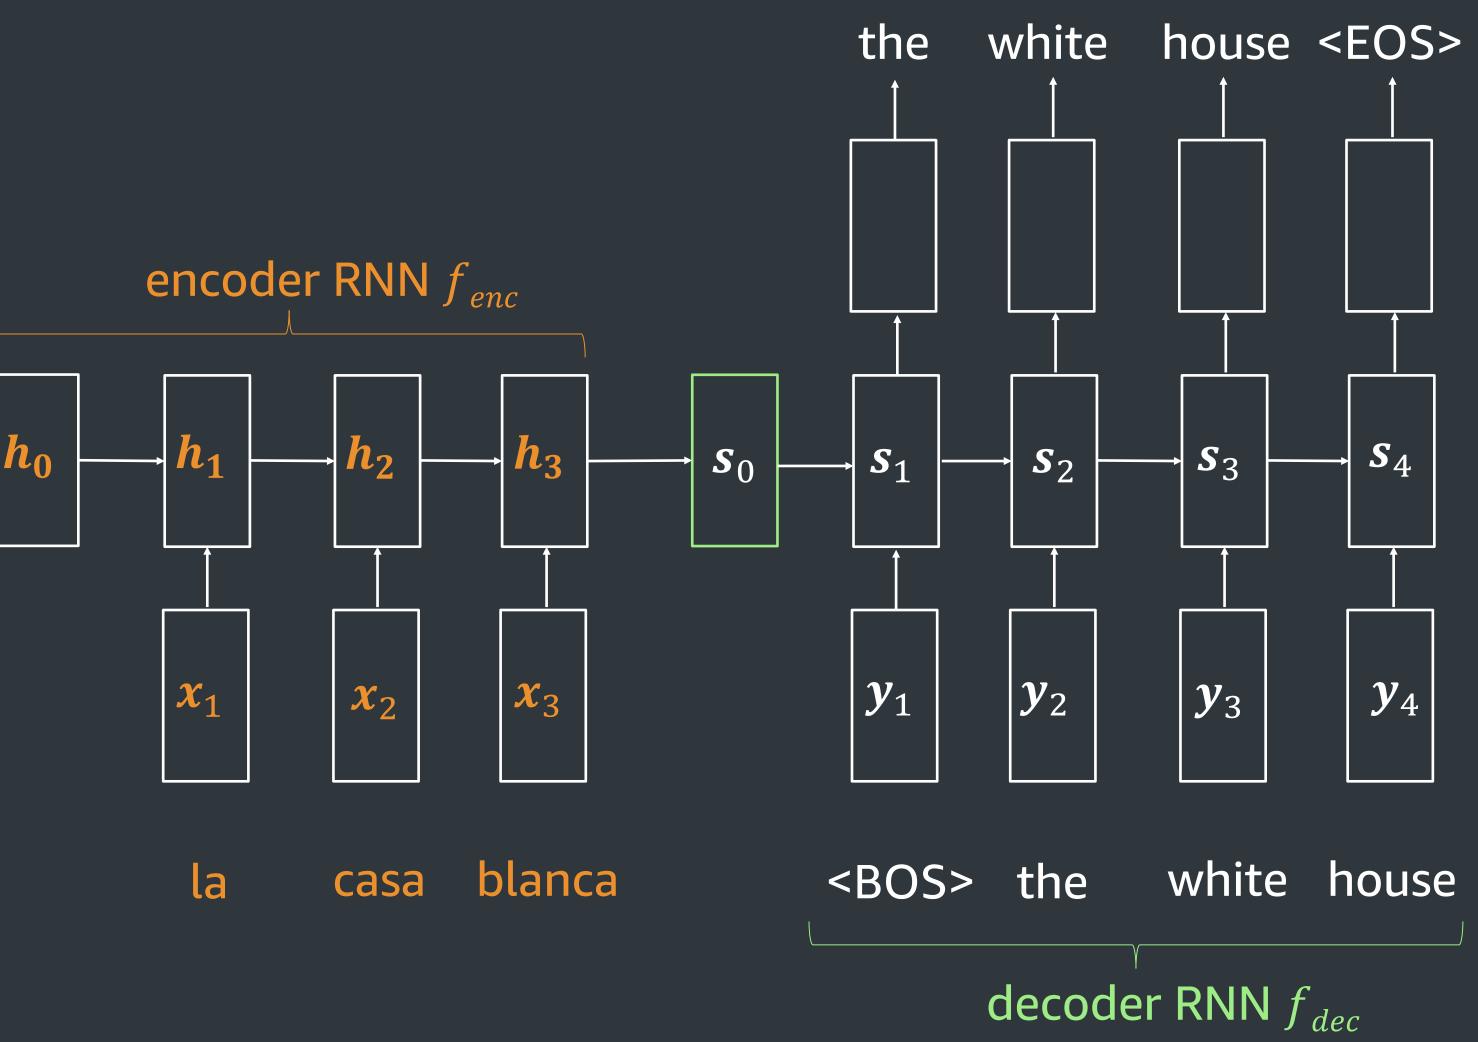
Sequence-to-Sequence Model (Sutskever, et al., 2014)

Language model conditioned on the source sentence:

$$P(\boldsymbol{v}|\boldsymbol{x}) = \prod_{t=1}^{n} P(v_t|\boldsymbol{v}_{1:t-1},\boldsymbol{x})$$

Encoded source sentence initializes decoder RNN:

$$s_0 = \tanh(W_I h_m + b_I)$$



Sequence Decoding with Attention (Bahdanau et al., 2014)

Decoder consumes an attention vector:

$$s_t = \tanh(Us_{t-1} + W[y_t, \overline{s}_{t-1}])$$

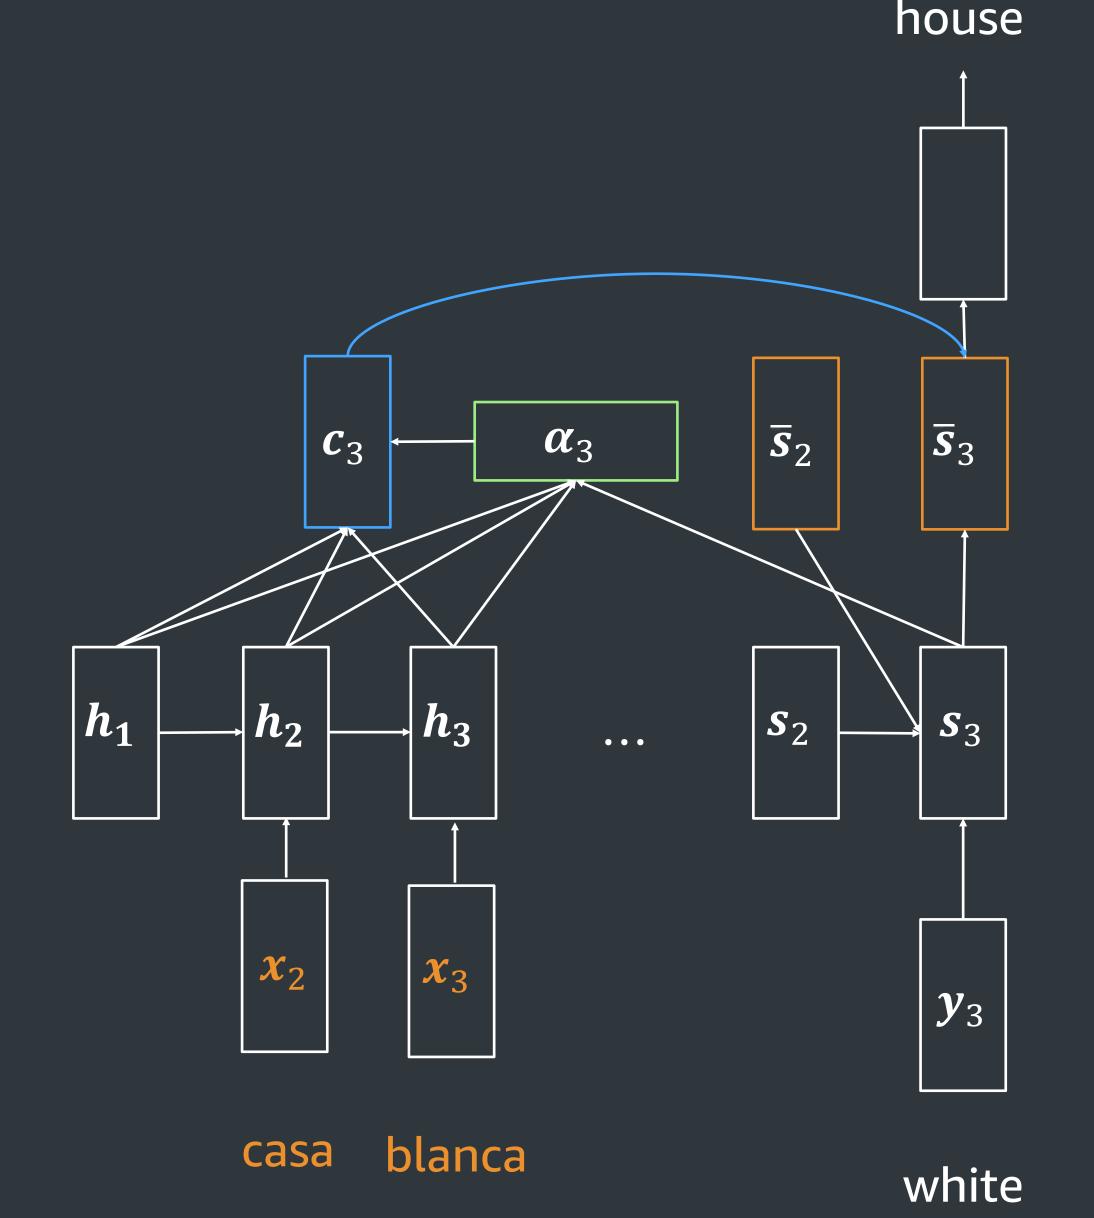
The attention vector consumes a context vector:

$$\overline{s}_t = \tanh(W_s[s_t, c_t])$$

The context vector is a linear combination source states:

$$c_t = \sum_{i=1}^m \alpha_{ti} h_i$$

where $\alpha_{ti} = softmax(score(s_t, h_i))$.



Attention Models in Sockeye

Name	score(s, h)	Available in Sockeye
mlp (Bahdanau, et al., 2014)	$\mathbf{v}_a^{\top} \tanh(\mathbf{W}_u \mathbf{s} + \mathbf{W}_v \mathbf{h})$	
concat (Luong, et al., 2015)	$\mathbf{v}_a^\top[\mathbf{s};\mathbf{h}]$	
dot (Luong, et al., 2015)	$\mathbf{s}^{\top}\mathbf{h}$	
location (Luong, et al., 2015)	$\mathbf{v}_{at}^{ op}\mathbf{s}_{t}$	
bilinear (Luong, et al., 2015)	$\mathbf{s}_t^\top \mathbf{W}_a \mathbf{h}$	
coverage (Tu, et al., 2015)	$\mathbf{v}_a^{\top} \tanh(\mathbf{W}_u \mathbf{s} + \mathbf{W}_v \mathbf{h} + \mathbf{W}_c \mathcal{C})$	



Artificial Intelligence at Amazon

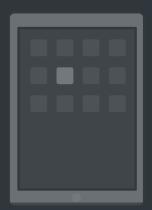
Thousands of employees across the company focused on machine learning & Al



Discovery & Search



Fulfilment & Logistics



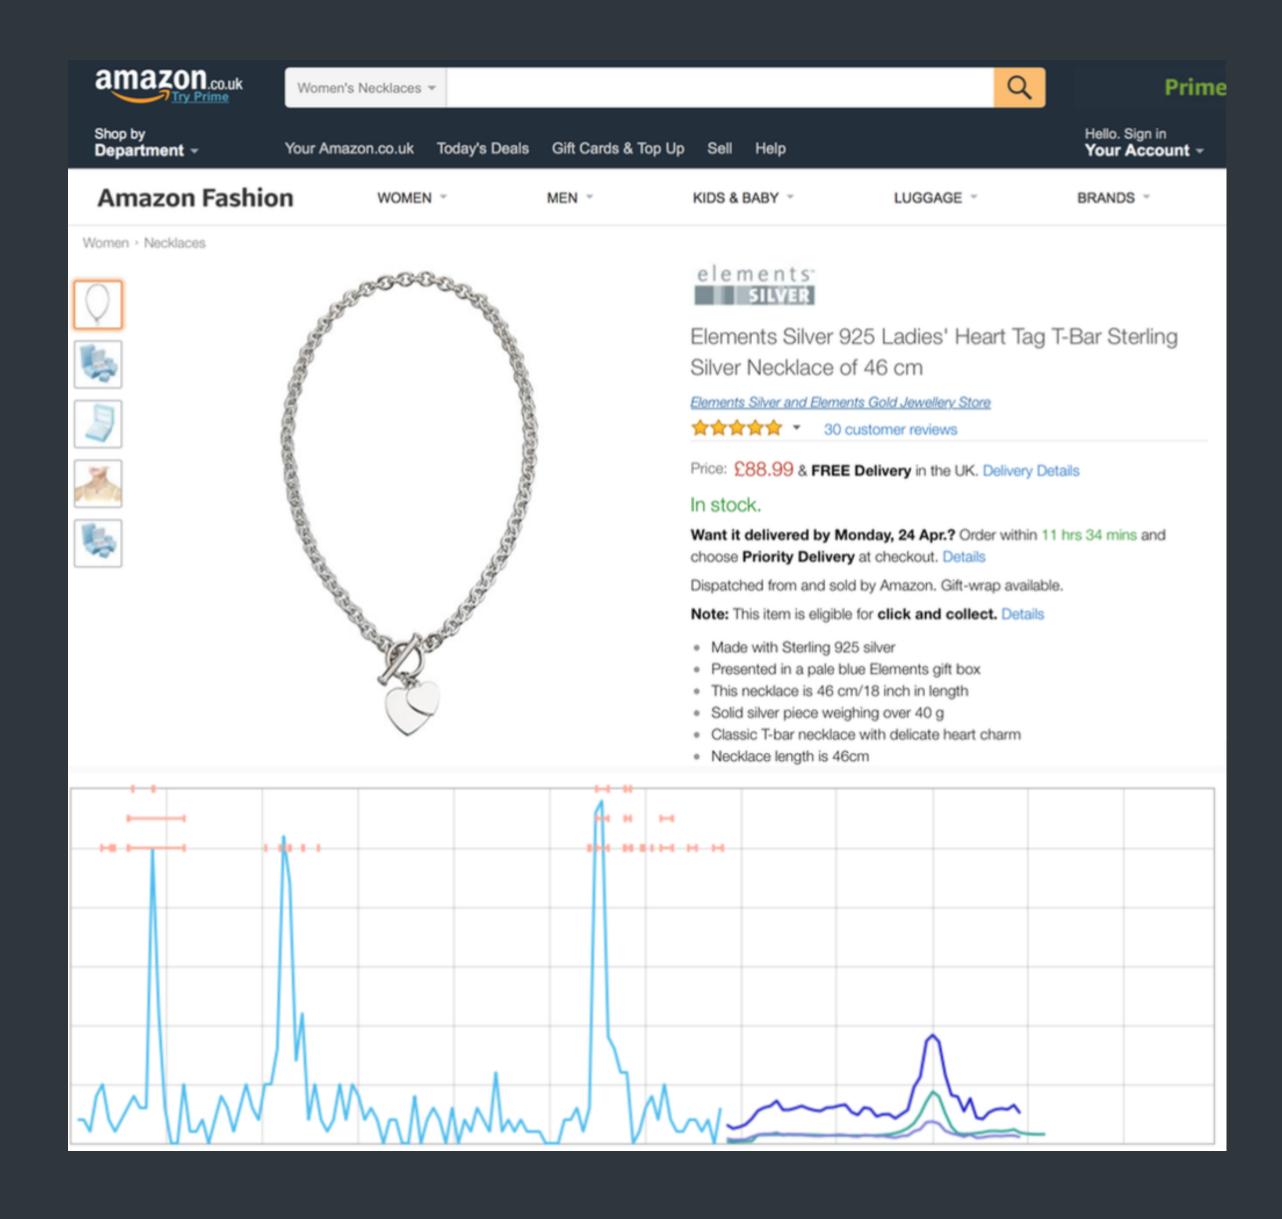
Enhance & Create Products



Bring Machine Learning to All



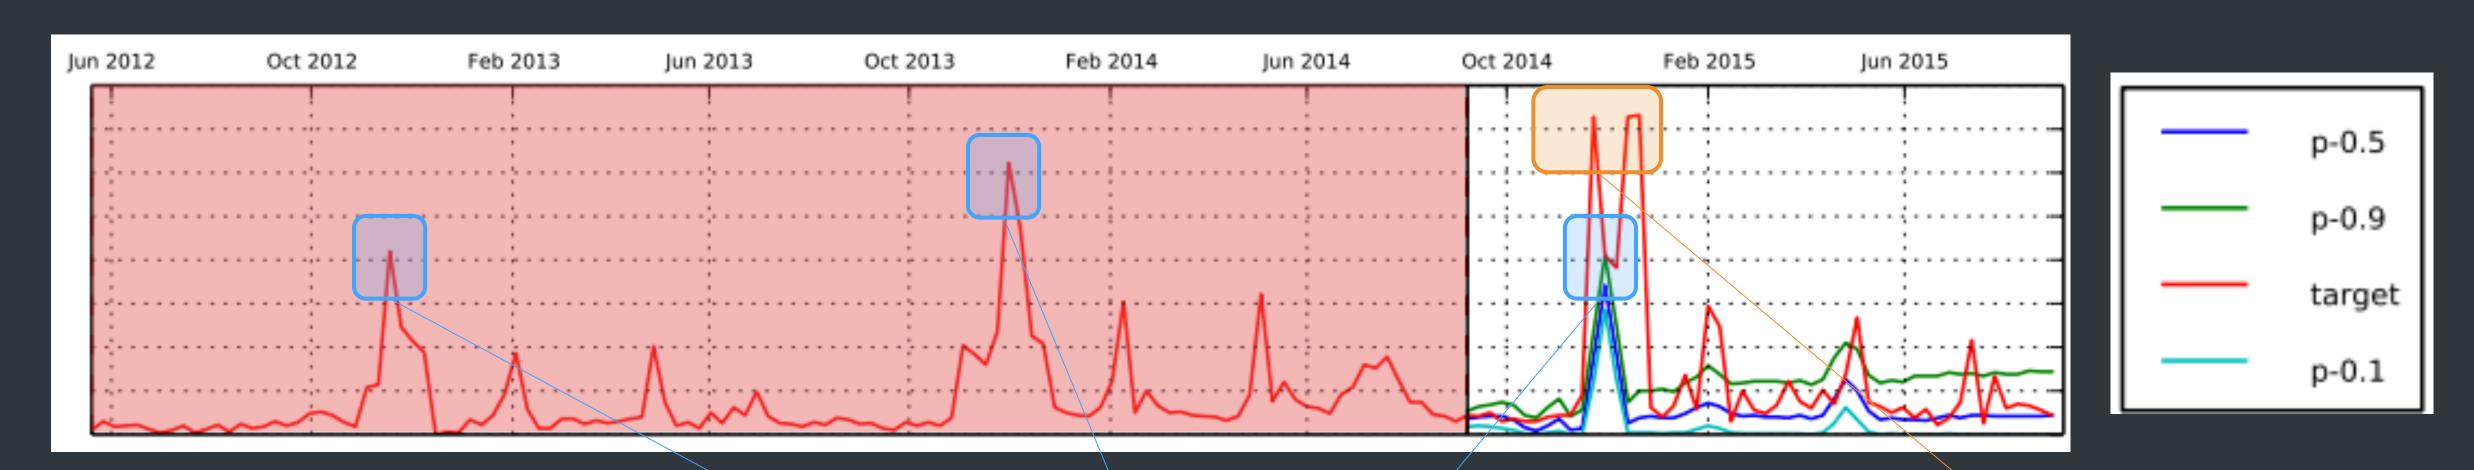
Demand Forecasting



Scale 20M+ products fulfilled by Amazon alone! Sparsity Many product sell very infrequently Regionalised 100+ FCs worldwide New products No past demand



Seasonality and External Events



Training Range

Non-fashion items have long(er) training ranges.

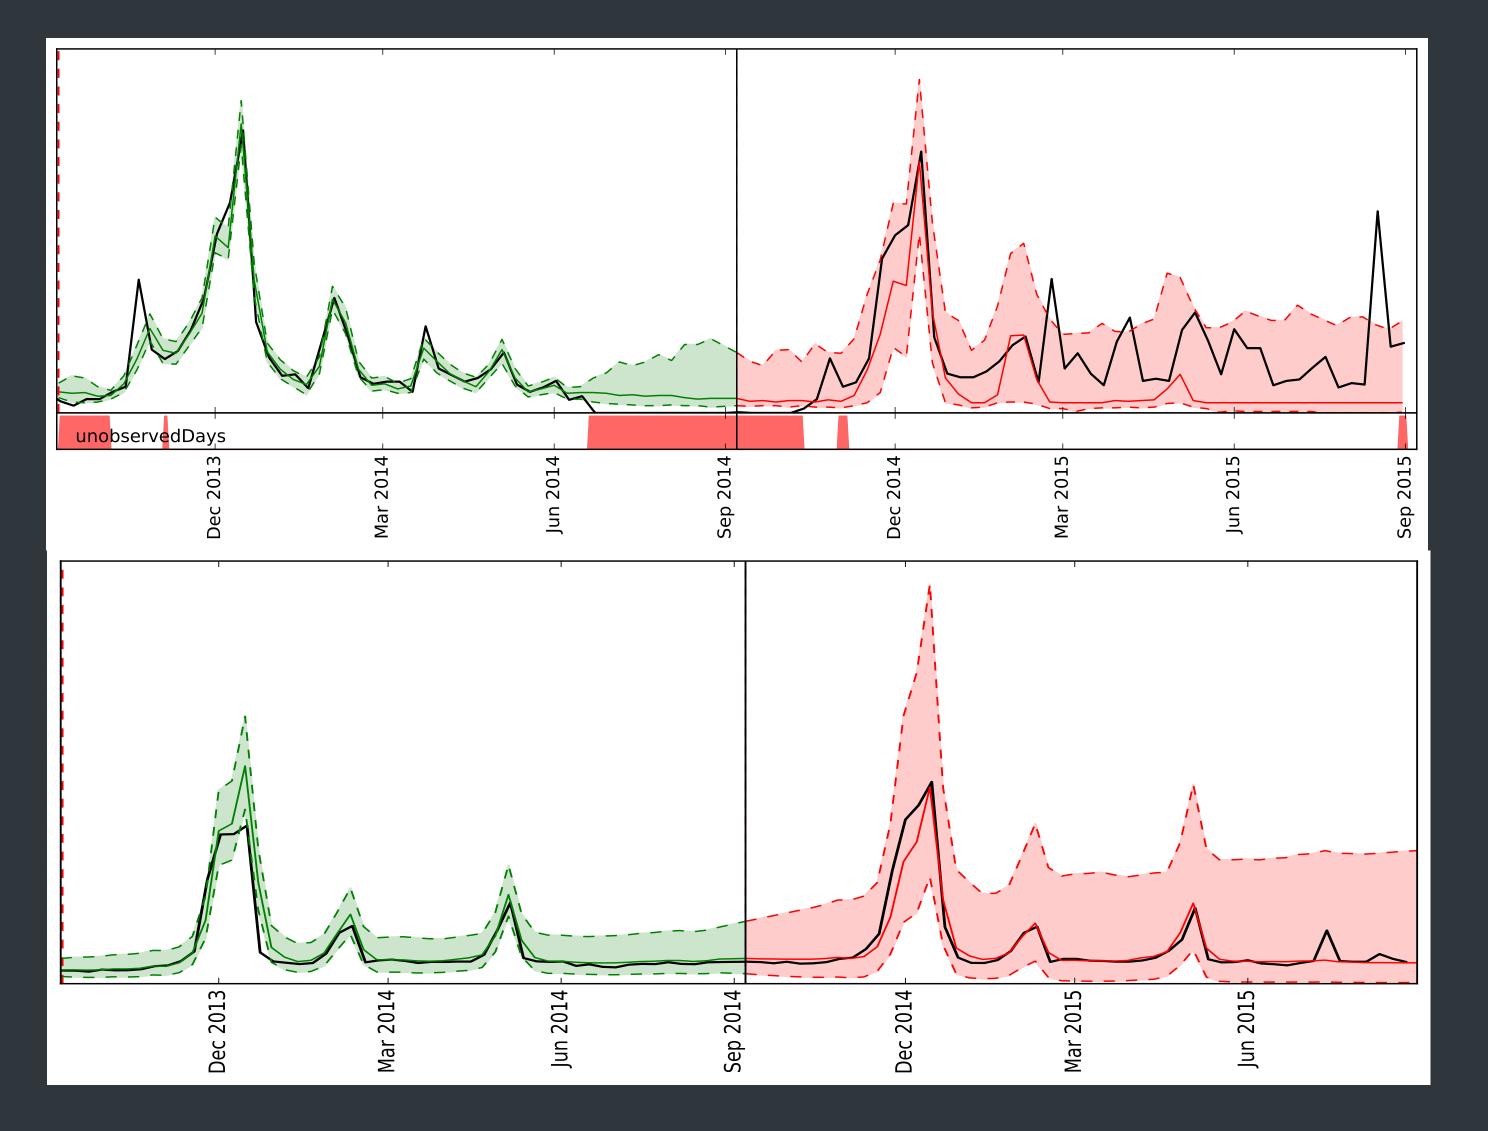
Seasonality

This item has Christmas seasonality with higher growth over time.

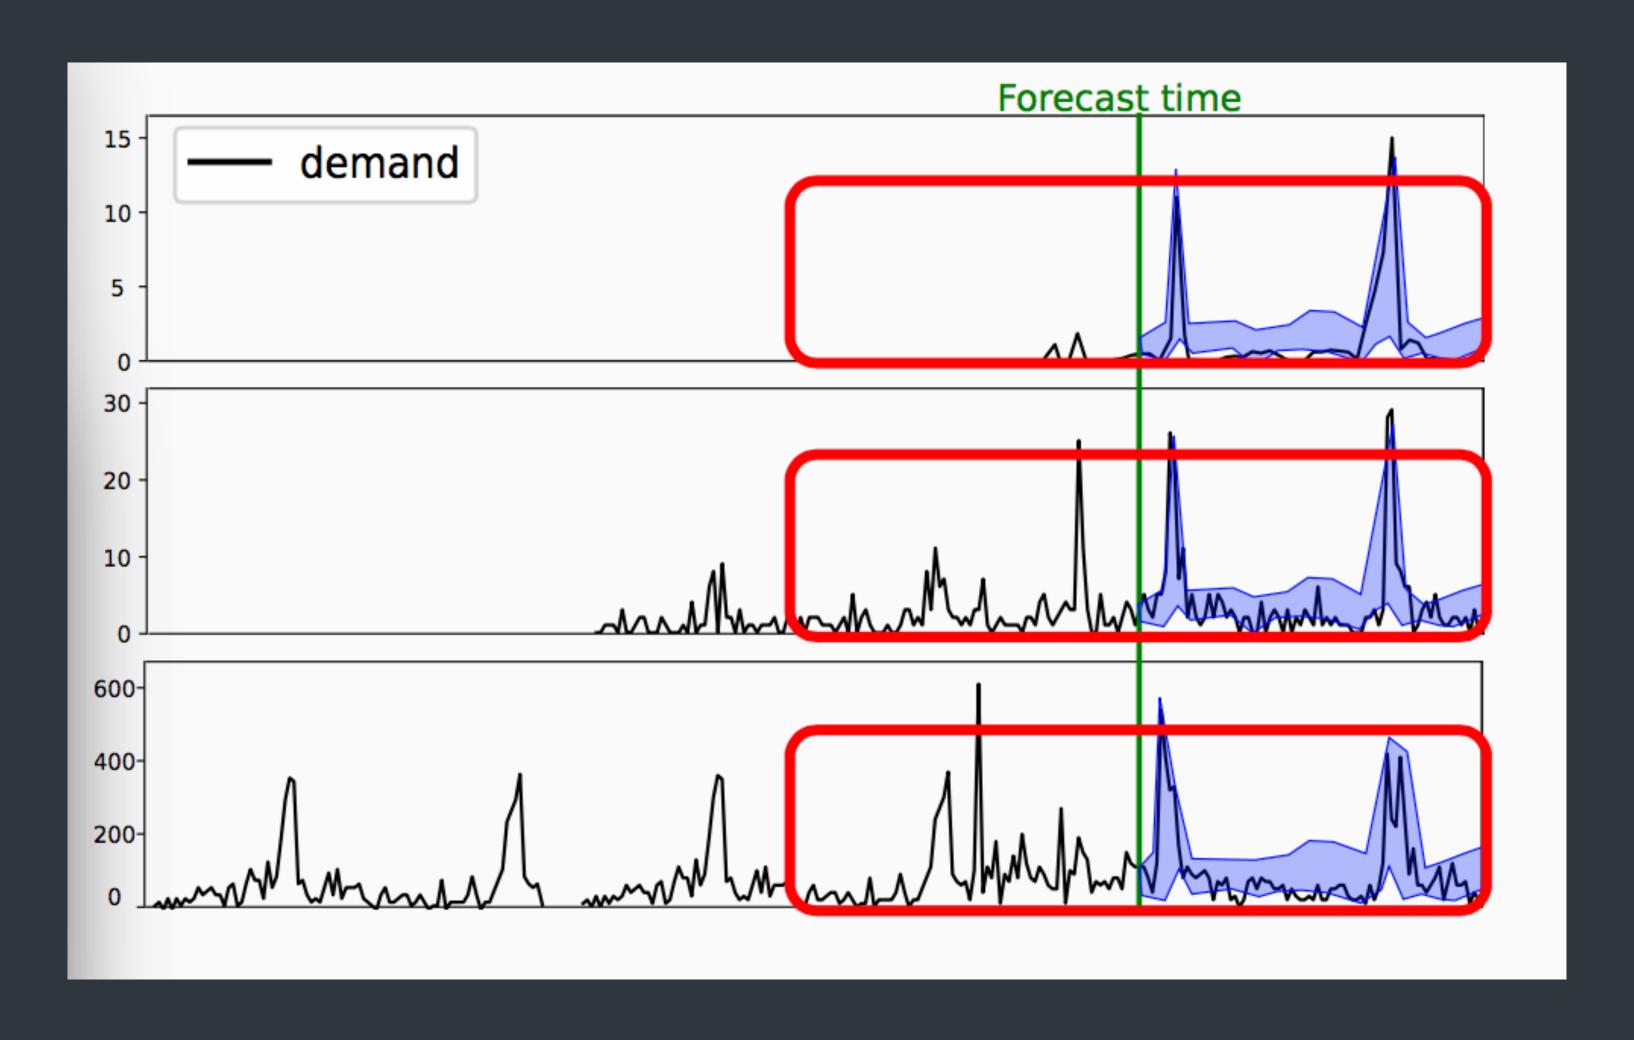
Missing Features/Inputs
Unexplained spikes
in the demand.



We produce probabilistic forecasts



Deep Autoregressive Recurrent Networks





Same day and early morning home delivery of grocery



Defect Detection is a Challenging Task



OK



Some Damage



Serious Damage



Decay





Artificial Intelligence at Amazon

Thousands of employees across the company focused on machine learning & Al



Discovery & Search



Fulfilment & Logistics



Enhance & Create Products



Bring Machine Learning to All



Amazon Al



Apache MXNet

Deep learning engine









Polly

Text-to-Speech

Rekognition

Image Analysis

Lex

ASR & NLU

ML Applications



Apache MXNet is the deep learning framework of choice for Amazon



Why MXNet?

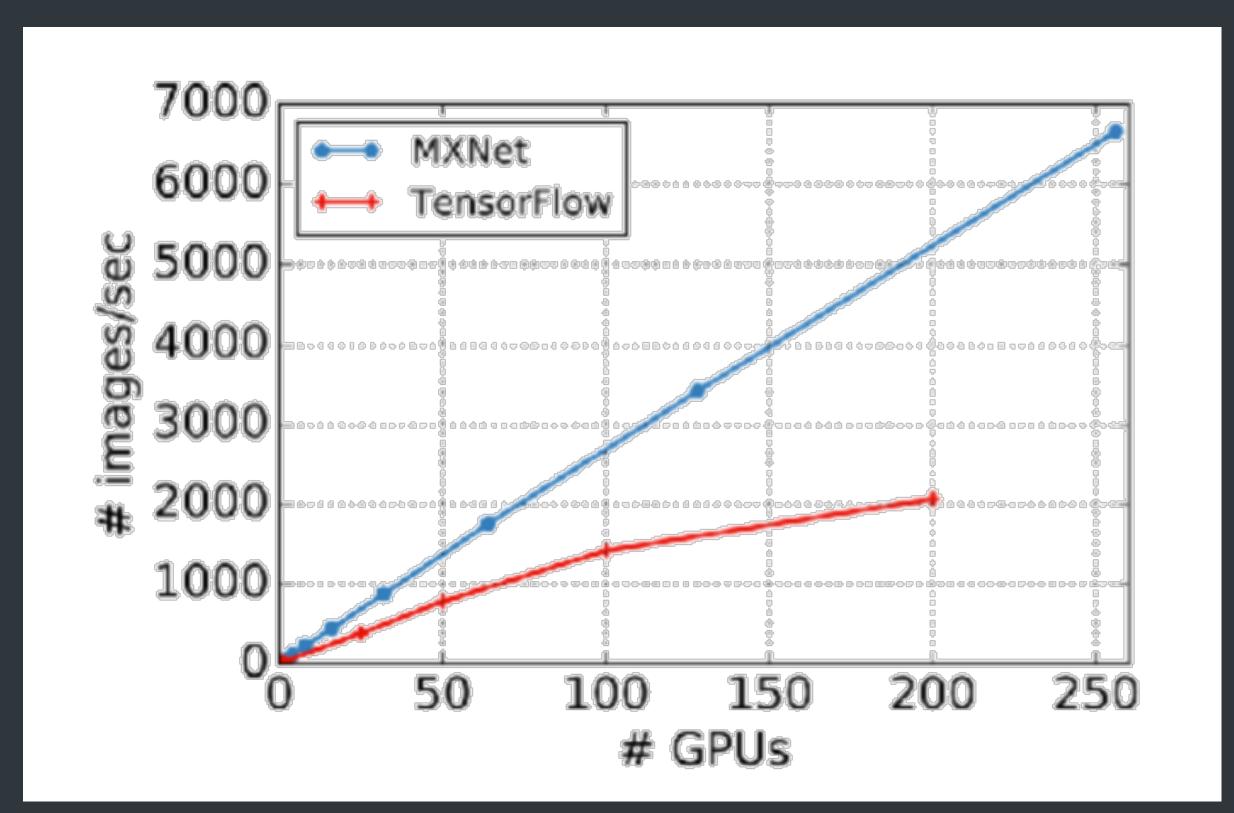
Flexible programing model:

- Symbolic API
- Imperative API

Bindings for Python, C++, Scala, R, Julia, Perl.

Portable: GPU and CPU

Open source project.



Google Inception v3 (image recognition)



Amazon Polly





Powers Alexa



47 voices



24 languages



Low latency, real time

A Focus On Voice Quality & Pronunciation

1. Automatic, Accurate Text Processing



"Today in Seattle, WA, it's 11°F"



"We live for the music live from the Madison Square Garden."



A Focus On Voice Quality & Pronunciation

1. Automatic, Accurate Text Processing



2. Intelligible and Easy to Understand



A Focus On Voice Quality & Pronunciation

- 1. Automatic, Accurate Text Processing
- 2. Intelligible and Easy to Understand
- 3. Add Semantic Meaning to Text



"Richard's number is 2122341237"



"Richard's number is 2122341237"

Telephone Number



A Focus On Voice Quality & Pronunciation

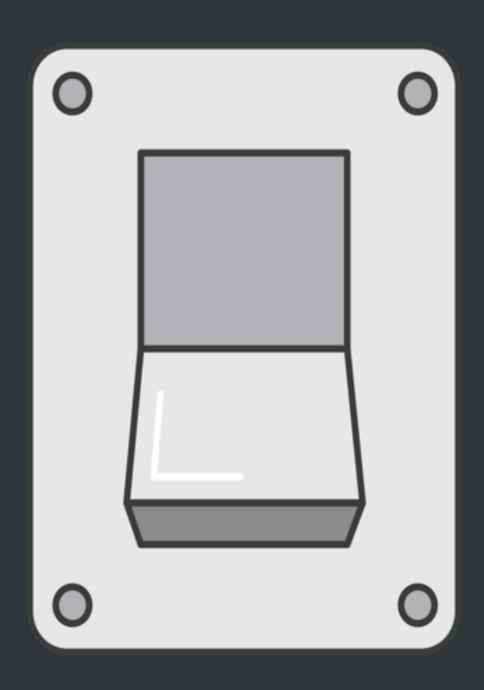
- 1. Automatic, Accurate Text Processing
- 2. Intelligible and Easy to Understand
- 3. Add Semantic Meaning to Text
- 4. Customized Pronunciation

"My daughter's name is Kaja."

"My daughter's name is Kaja."



Democratising Machine Learning

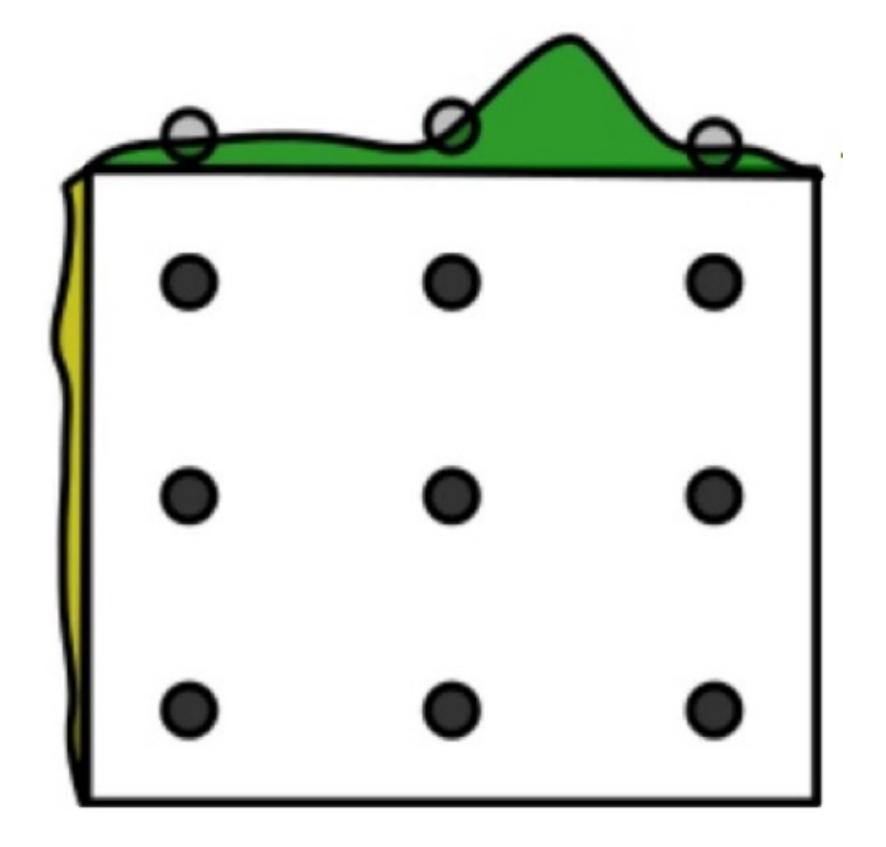


- Focus on the problem at hand!
- Abstract away learning algorithms
- Abstract away feature engineering

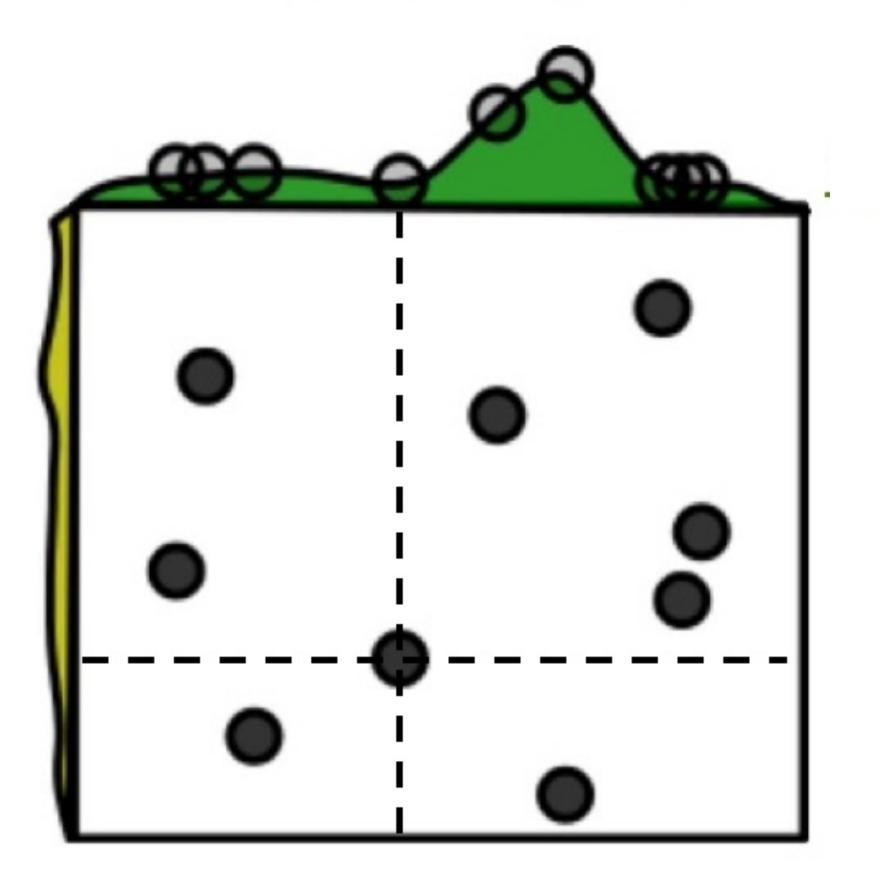
Requires to automate hyperparameter tuning!



Grid Layout

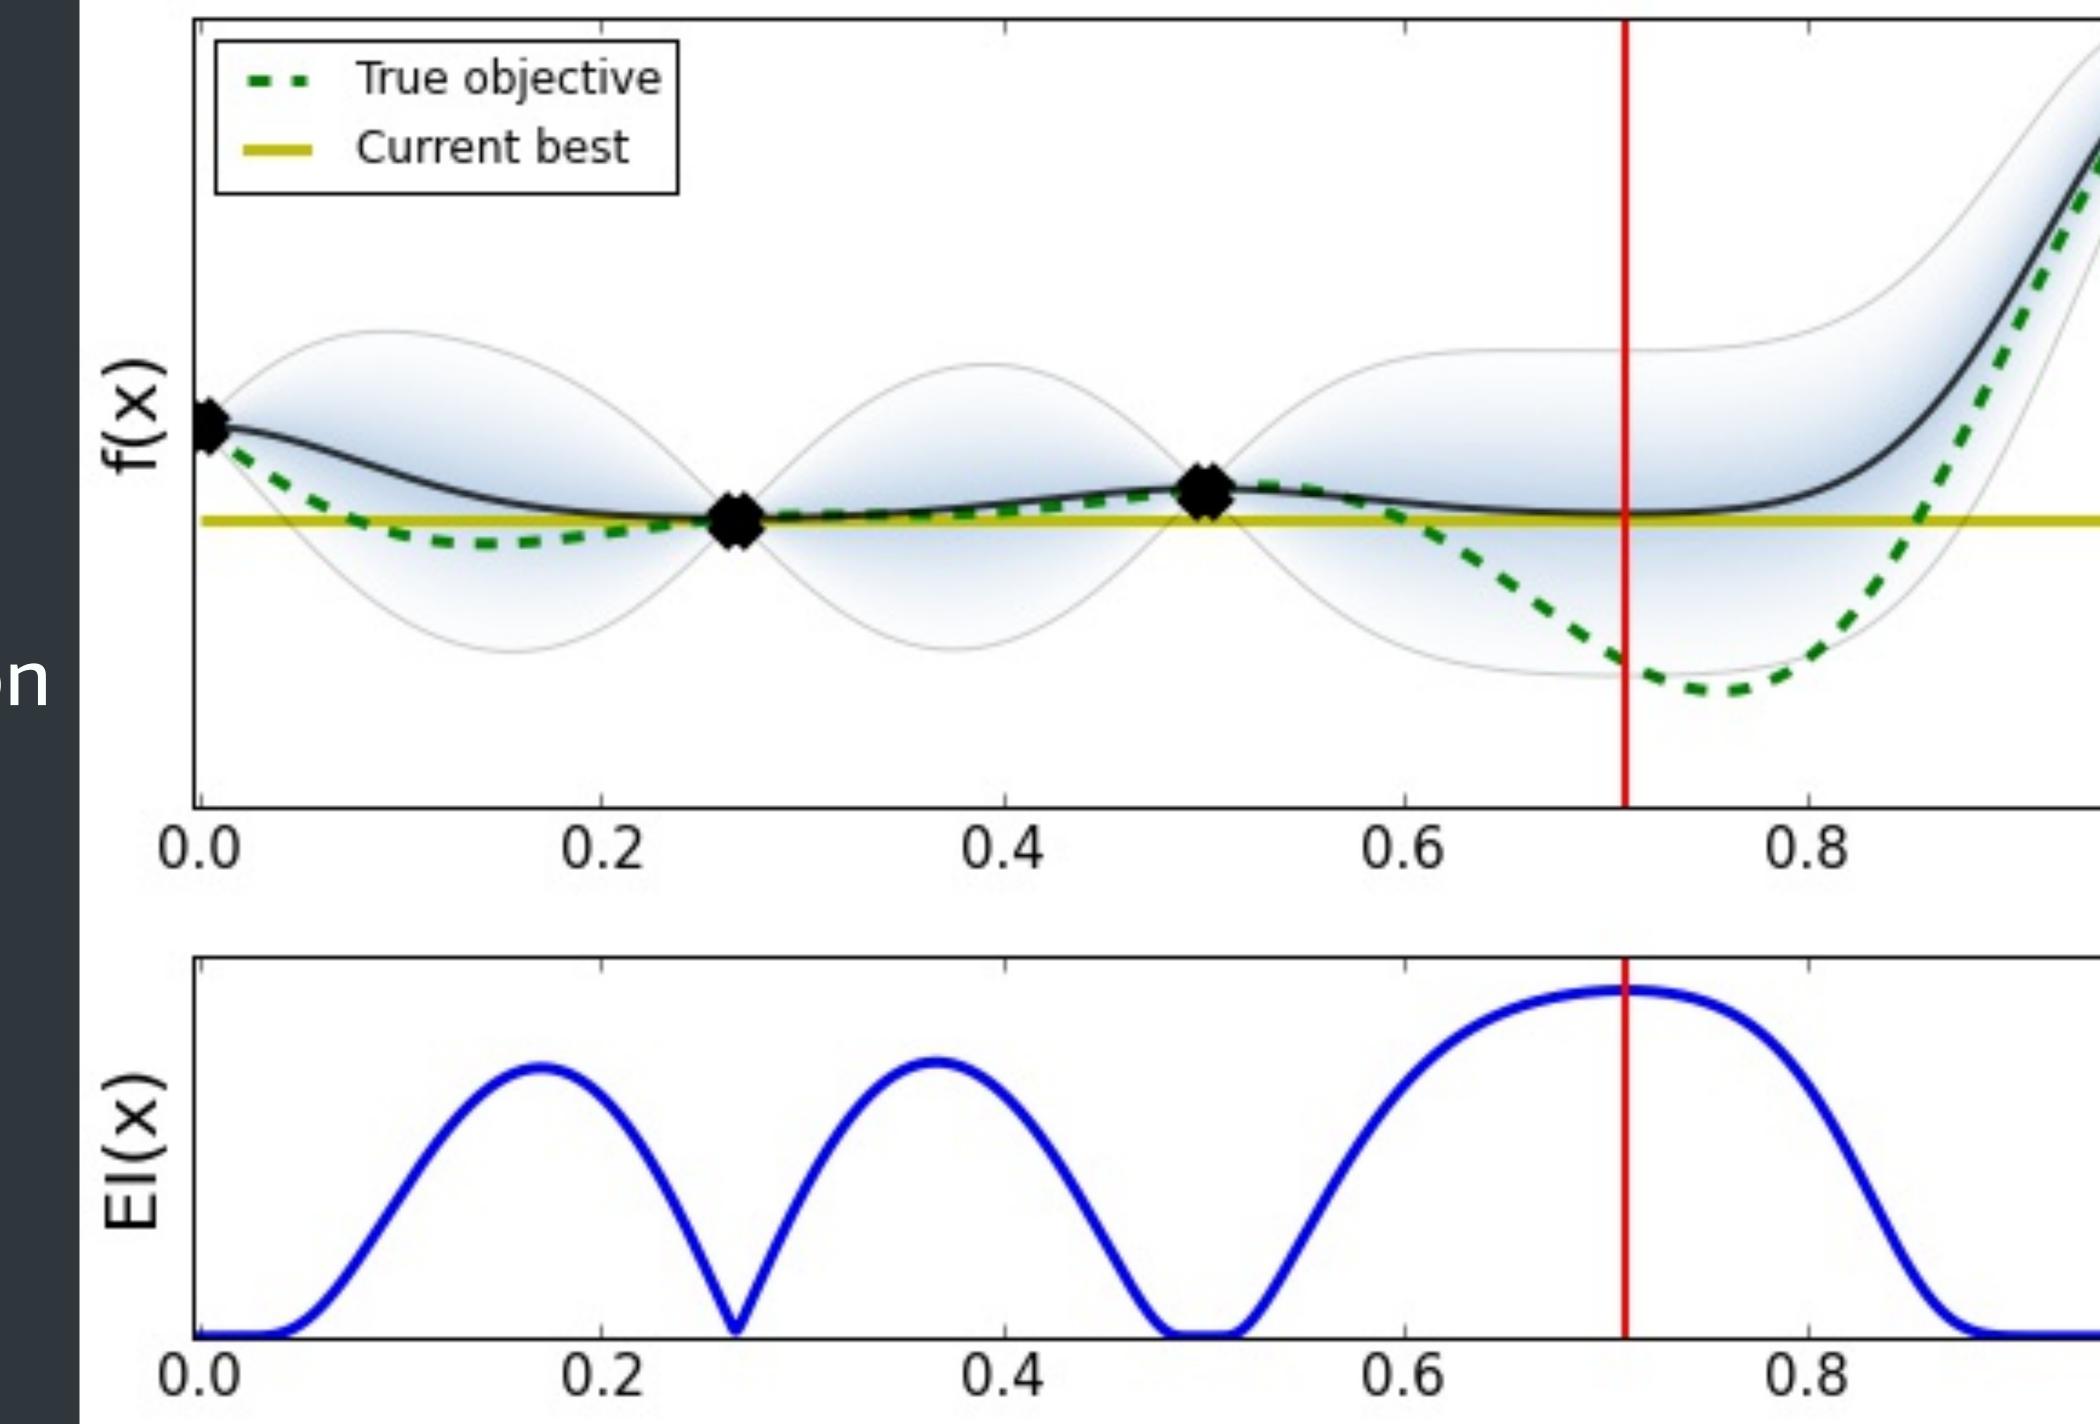


Random Layout





Can we do better?



Bayesian Optimisation



A Knight of the Seven Kingdoms (A Song of Ice and Fire)

In Stock. Ships from and sold by Amazon.com. Gift-wrap available. \$20.43



****** I quickly became absorbed in the tales of "Dunk and Egg" and the ancestors of the great houses of

Westeros

By Amazon Customer on November 24, 2015

Format: Kindle Edition Verified Purchase

After reading Martin's other series, I was eager to find any and all related materials. This story is set about 100 years before the main action in Westeros and introduces some new characters and fills in some blanks on ones who were referred to in the later story. I quickly became absorbed in the tales of "Dunk and Egg" and the ancestors of the great houses of Westeros. I loved the angle of Egg traveling around living as a regular child instead of a prince of the realm. This book holds three short tales of adventures they have together and different lessons they both learn. My only issue with it was that it was too short! I wanted more; I wanted to see how Dunk developed as a person because he had a lot to learn about how noble power players might use hapless knights such as he. I hope there are plans to continue this series because I'd like to see how Egg learned from his experiences living among the people and how that changed the man he would become.

Comment

3 people found this helpful. Was this review helpful to you?





*** A taste of game of thrones before 6th book!

By Steven M McLaughlin on December 5, 2015

Format: Kindle Edition Verified Purchase

I tore through the game of thrones series and have been waiting and waiting for the latest book. An associate told me about this book and I was psyched. I was traveling and was thankful to download onto my kindle for a long flight home. It was entertaining but I got lost with all the characters and couldn't really keep up with who was doing what. Might need to go back and read again slowly to try to comprehend what happened! Didn't have this problem with the other game of thrones books...

Comment 2 people found this helpful. Was this review helpful to you?

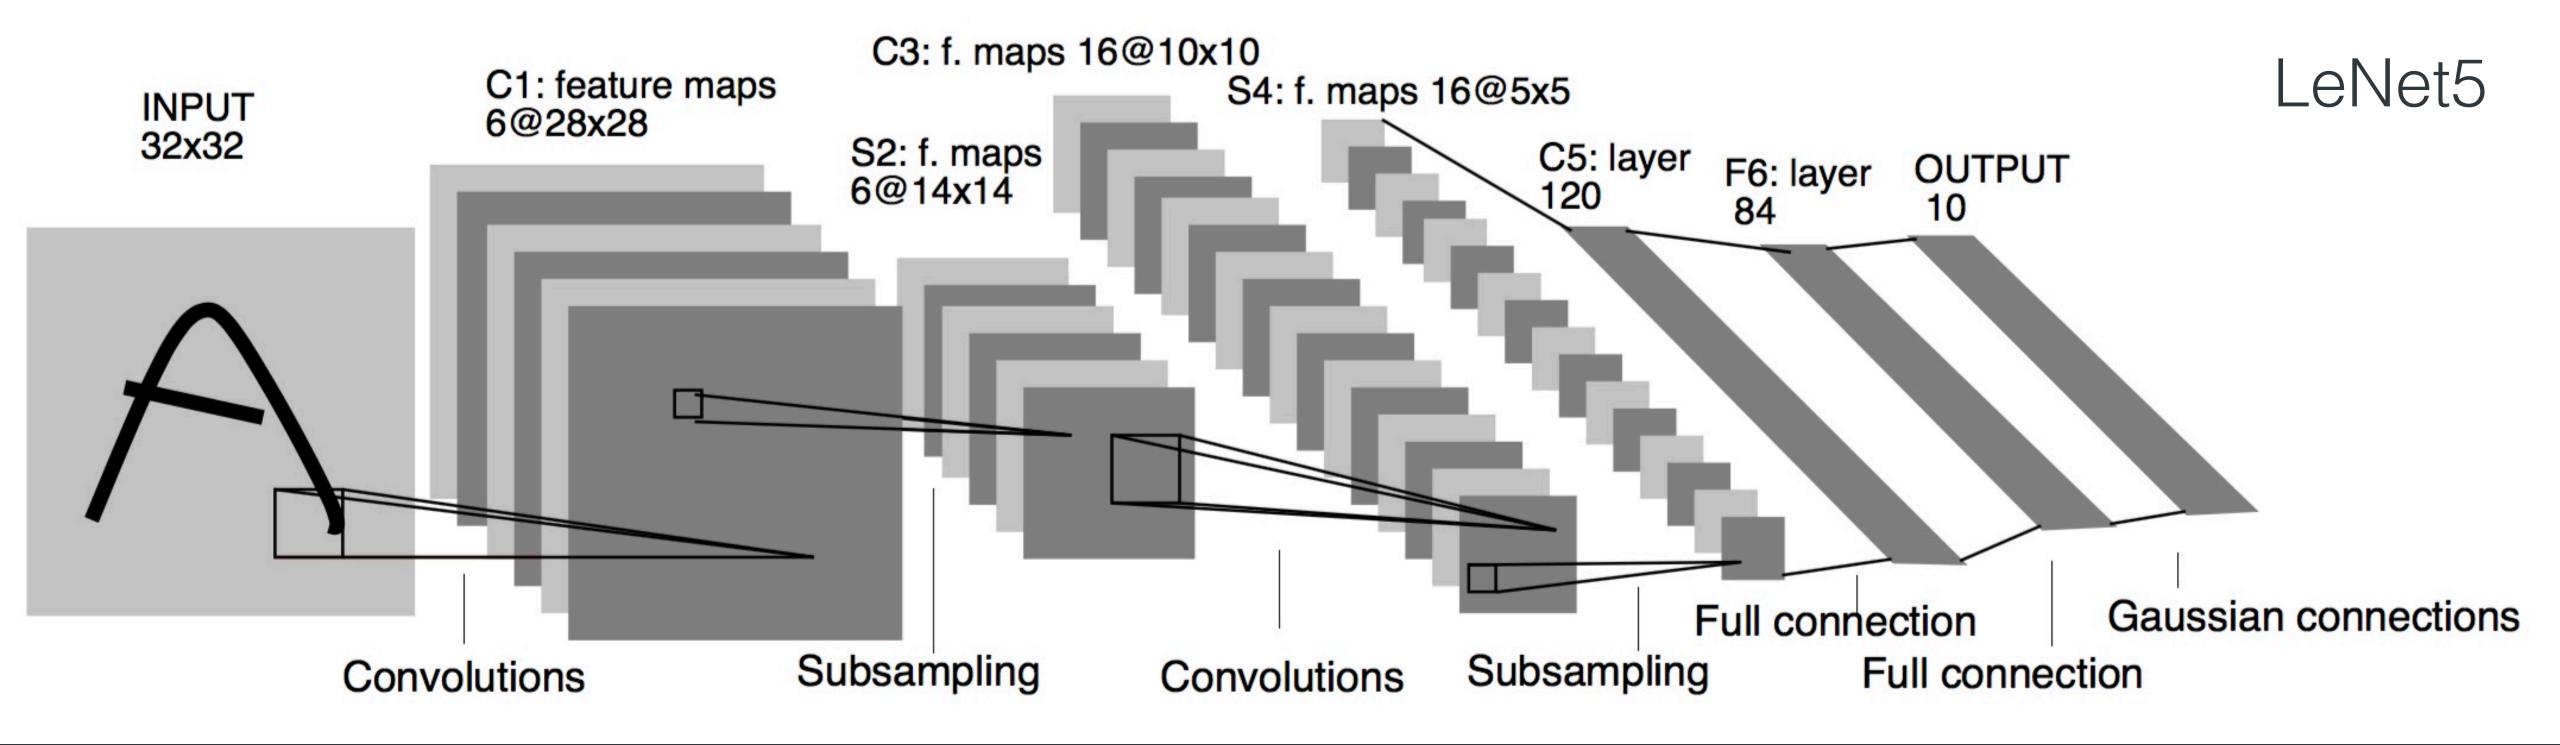


Revisiting Sentiment Analysis

Hyperparameter	Values
n_{min}	$\{1, 2, 3\}$
n_{max}	$\{n_{min},\ldots,3\}$
weighting scheme	{tf, tf-idf, binary}
remove stop words?	{True, False}
regularization	$\{\ell_1,\ell_2\}$
regularization strength	$[10^{-5}, 10^5]$
convergence tolerance	$[10^{-5}, 10^{-3}]$

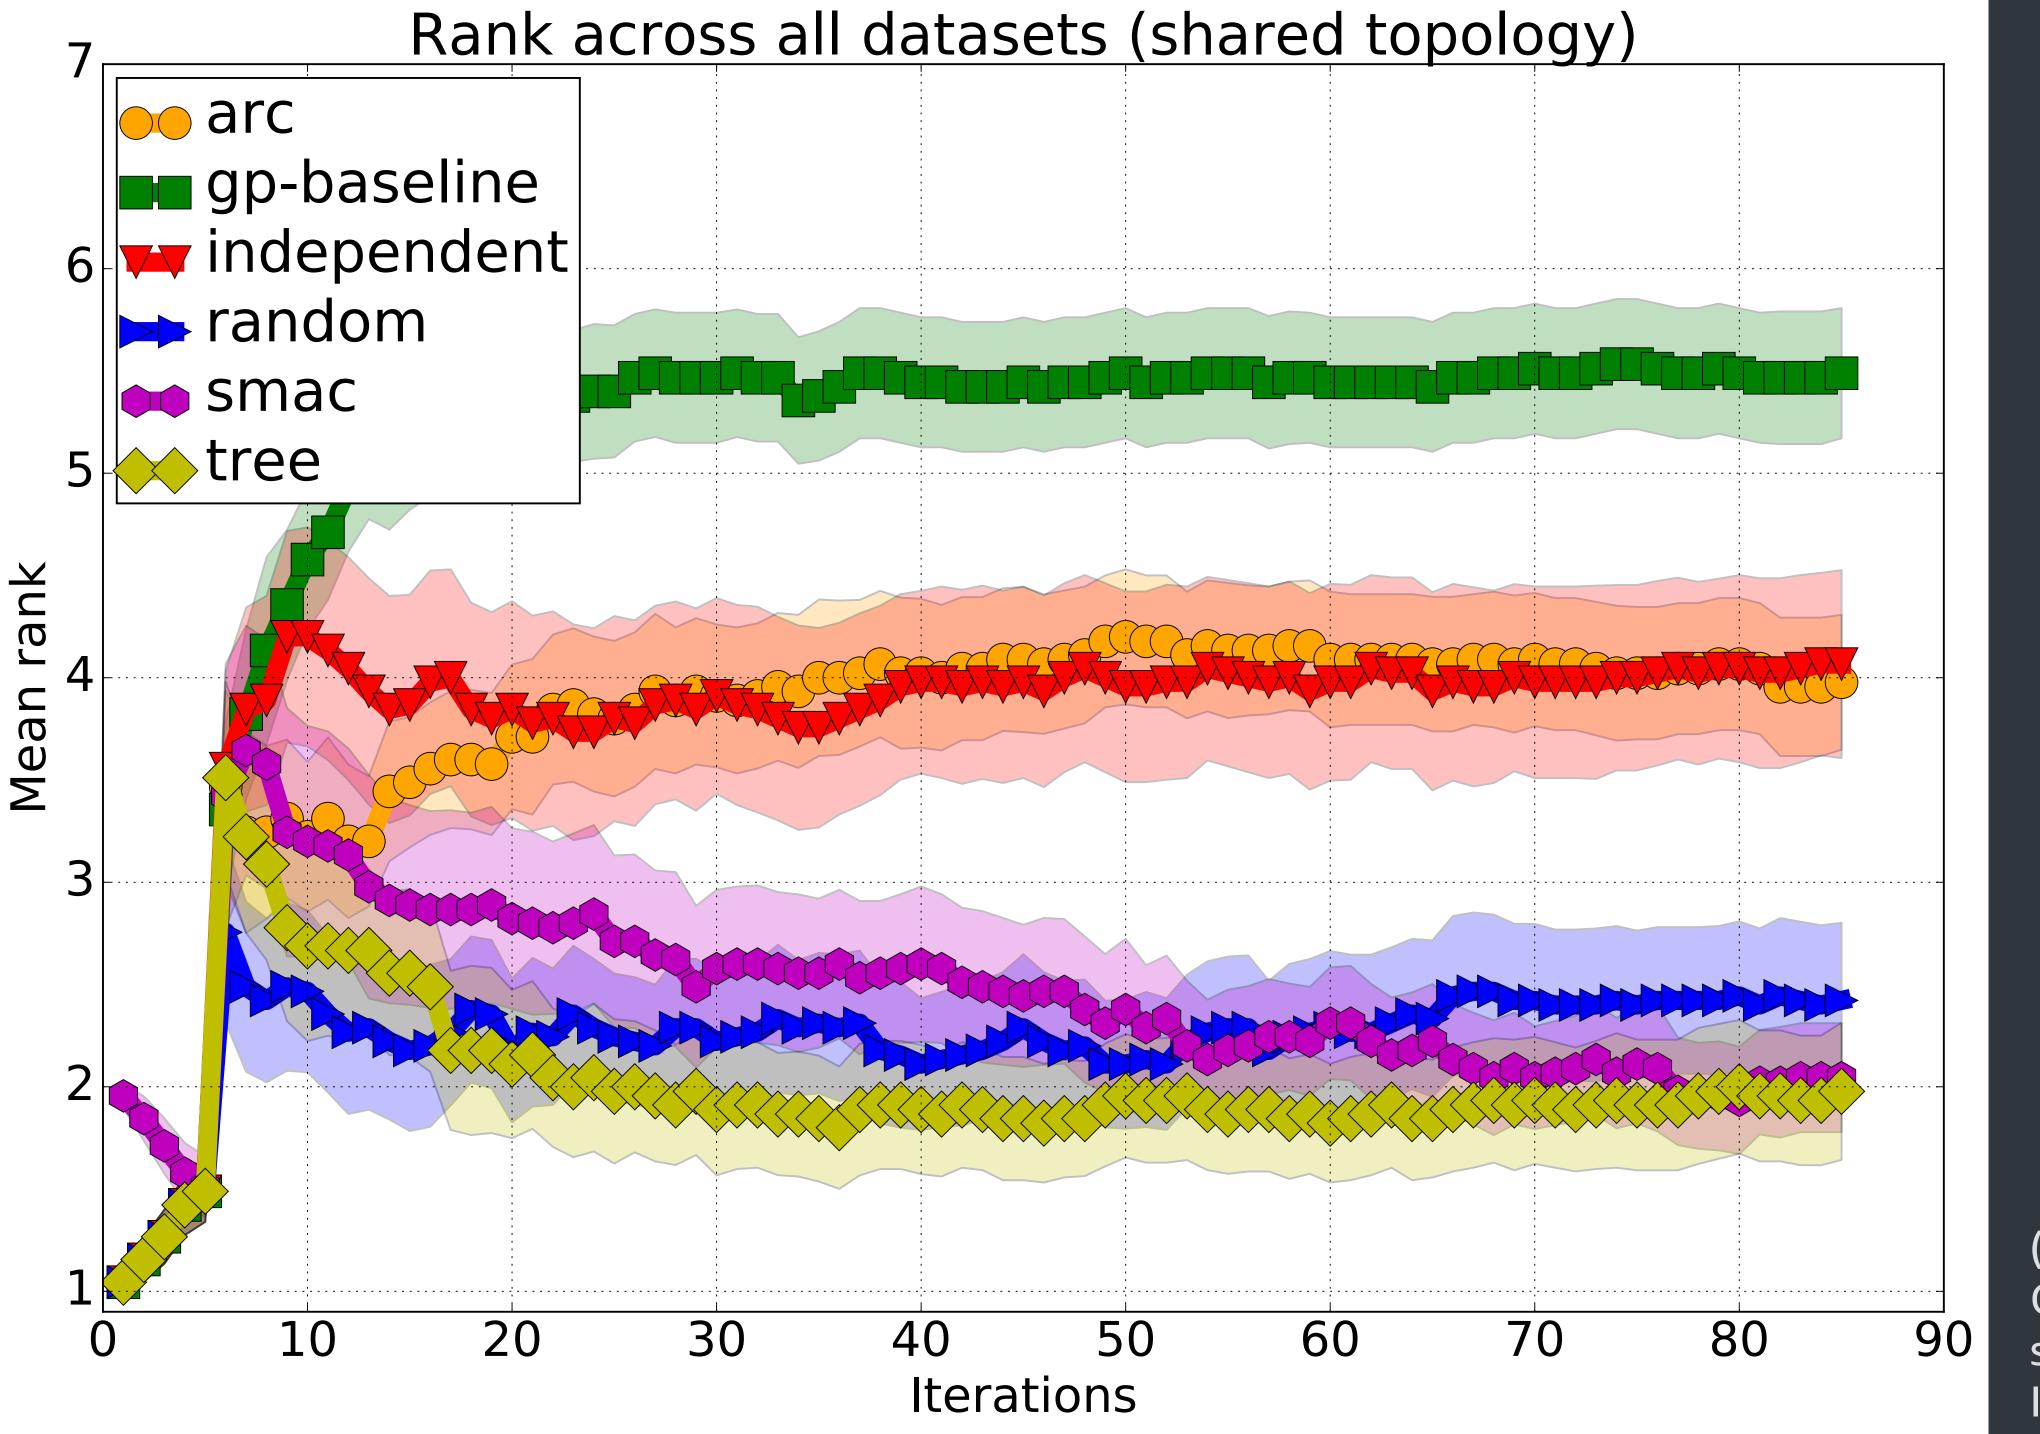
Method	Acc.
SVM-unigrams	88.62
SVM- $\{1, 2\}$ -grams	90.70
SVM- $\{1, 2, 3\}$ -grams	90.68
NN-unigrams	88.94
$NN-\{1,2\}$ -grams	91.10
NN- $\{1, 2, 3\}$ -grams	91.24
LR (this work)	91.56
Bag of words CNN	91.58
Sequential CNN	92.22

Table 5: Comparisons on the Amazon electronics dataset. Scores are as reported by Johnson and Zhang (2014).



- Evaluating a neural networks takes time
- Search space is large:
 - Architecture: #layers, #units per layer, activation function
 - Model complexity: regularization, drop-out, ...
 - Optimisation parameters: learning rate, momentum, ...





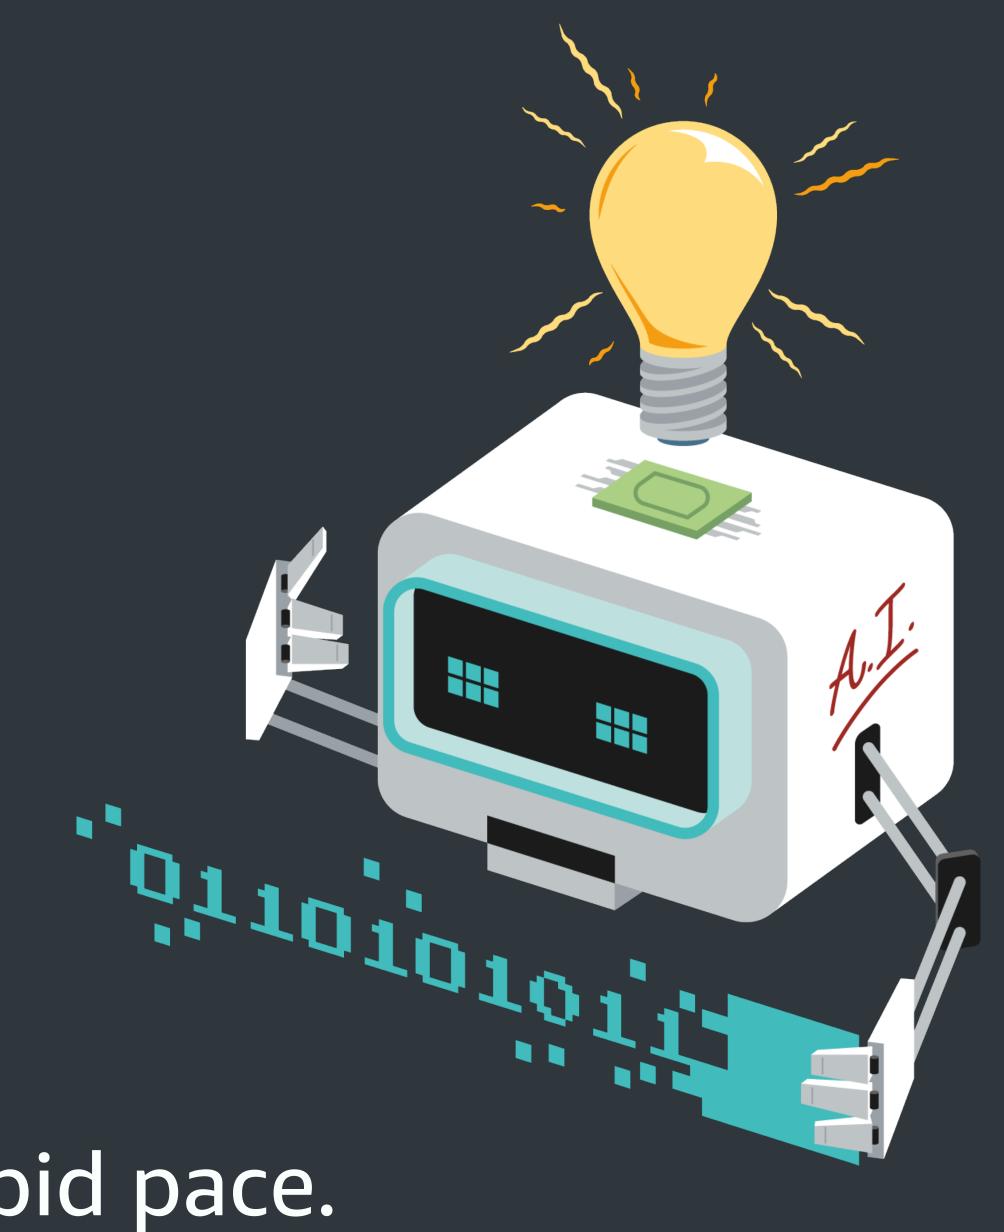
(Jenatton, et al.: Bayesian Optimization with Treestructured dependencies. ICML 2017.)

Take away

Machine learning & Al are transformational.

Machine learning & Al are ubiquitous.

The tools for building intelligent applications are maturing at a rapid pace.



amazon 3

cedrica@amazon.com

aws.amazon.com/amazon-ai mxnet.incubator.apache.org github.com/awslabs/sockeye aws.amazon.com/blogs/ai