Why data-driven methods will shape the future of relevance search

Pedro Balage
Lead Data Scientist @ Farfetch
Berlin Buzzwords 2019

Is relevance search an art?

“The Birth Of Venus”
Sandro Botticelli

How Search teams are created?

- Search capabilities are created by demand
- Initially, teams focused only on technical aspects
- Attention to search quality comes much later

“Planning the Voyage”
Max Gaisser

Pain problems

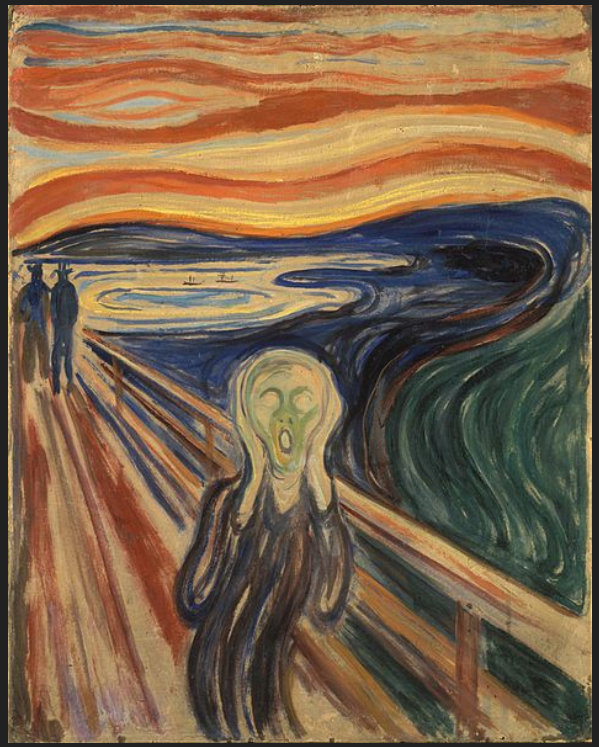

"The Scream"
Edvard Munch

- The lack of planning ahead leads to restricted growth
- You have a good team of engineers, but you lack expertise in search quality
- The complexity of search relevance tasks increase exponentially over time

Understanding relevance search

- What should not happen
  - You are retrieving items that are not relevant
  - You are missing relevant items
  - You are not showing the most relevant items first

Some steps into relevance search

- TF-IDF, Okapi BM25
- multifield matching, boosting
- phrase matching, query matching, fuzzy matching
- text normalizations, stemming
- synonyms, taxonomies, ontologies
- rules and exceptions!

“Janissary and a Merchant in Cairo”
Emile Prisse d'Avennes

How easy is it to maintain control over all these parameters?

Relevance Search is a complex task!

“Convergence”
Jackson Pollock

How is your organization prepared for the future?

- Disruptive Digital Innovation
- New Products
- New Categories
- New Markets
- Internationalization
- Fast adaptation!

“Love is in the Bin”
Banksy

Can relevance search be more science?

A Data Science?

Using Data to Leverage your Search Results

“The Tax Collectors”
Quentin Massys

- Search logs
- Click-through logs
- Navigation logs
- User preferences
- AB Testing
- External sources

Machine learning is just a mathematical function!

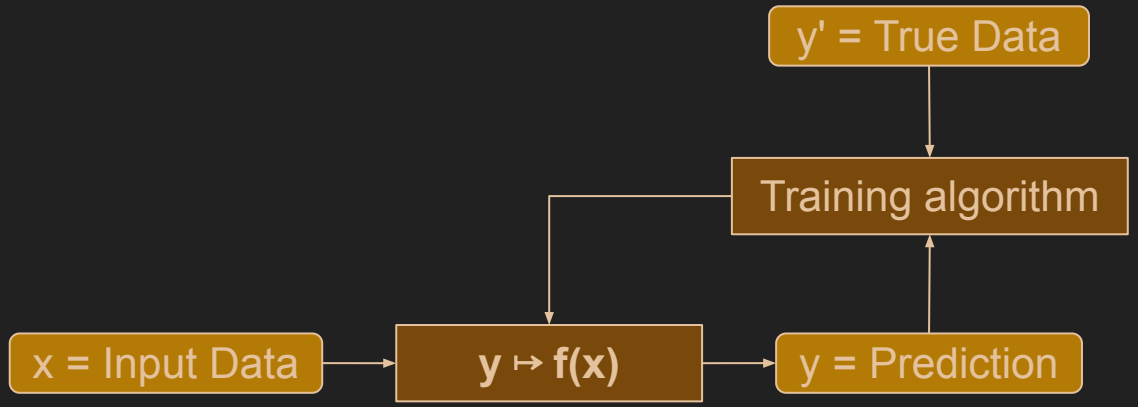

“Portrait of Leonhard Euler”
Jakob Emanuel Handmann

Using machine learning to build better search engines

- Traditional Approach
  - Autocomplete
  - Query Understanding
  - Query Expansion
  - Relevance Re-scorer
  - Ranking

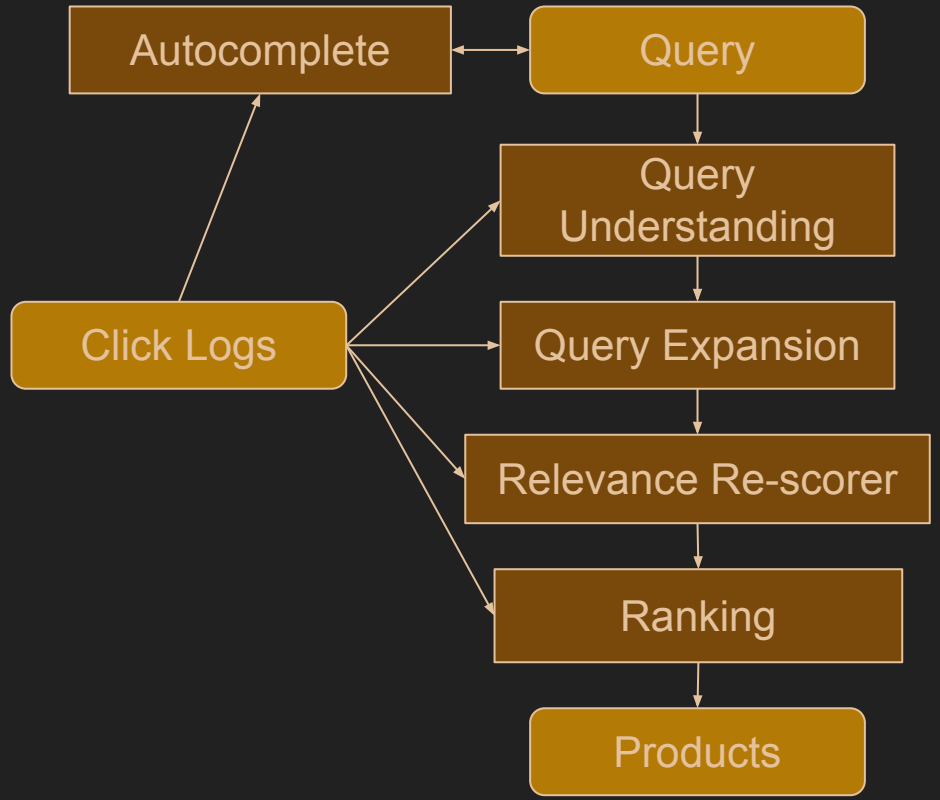

Using machine learning to build better search engines

- Modern Approach
  - End-to-End Product Search with Deep Learning

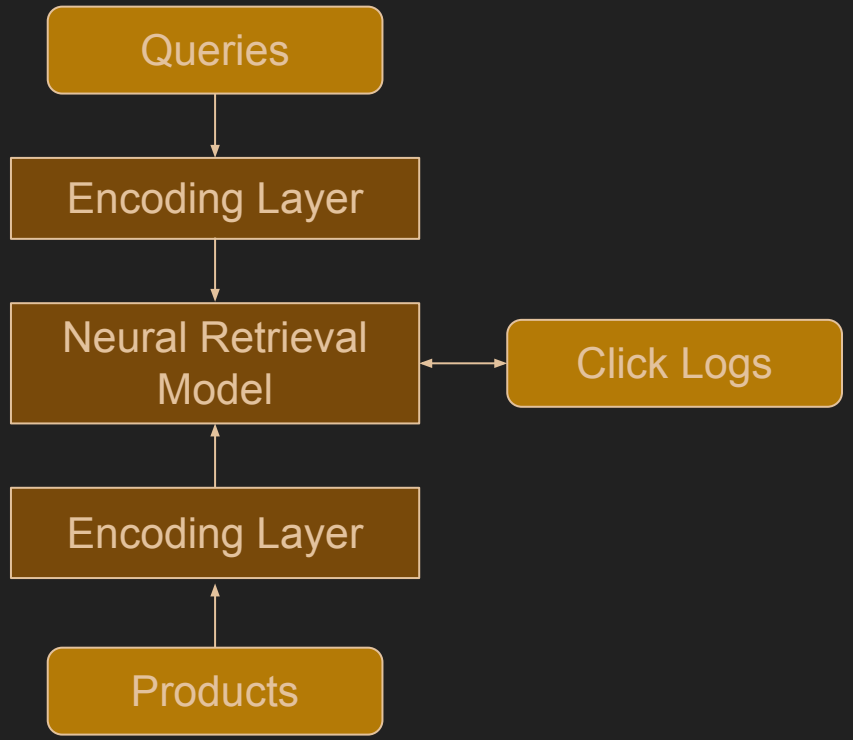

Why should we favor a more traditional architecture?

Iron and Coal
“William Bell Scott”

- Traditional architecture for search still very solid and tested in the industry
- This architecture is easily interpreted and changed to support business needs.

The path to a data-driven search engine

Autocomplete suggestions

- Field-based Autocomplete
- Phrases database Autocomplete
- Language Modeling
  - Hidden Markov Models
  - Neural Language Modeling

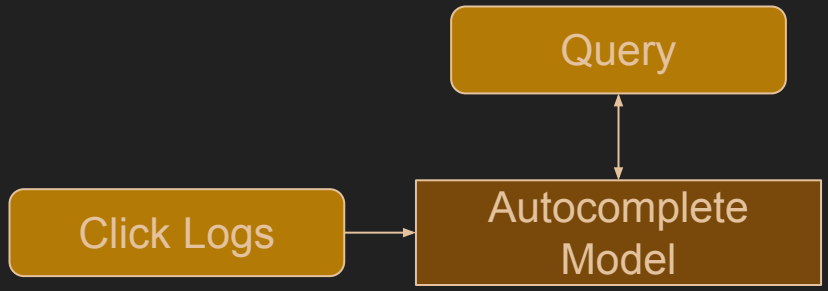

Neural language modeling

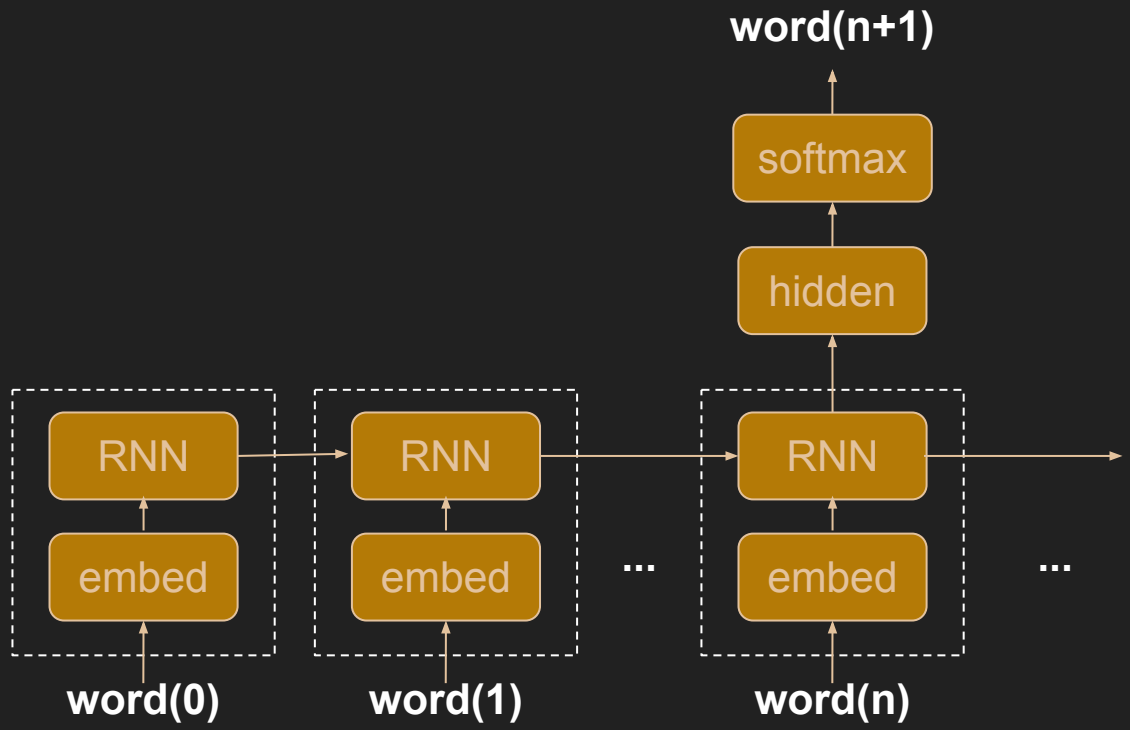

The same approach could also work on character level.

Neural language modeling

- Vector representation of words
  - Word2vec
  - Query2Vec [1]
- Unbounded vocabulary and sentence sizes
- Possible adaptation to deal with spellings and different words [2]

“Bedroom in Arles”
Vincent van Gogh

1. Dongyeop Kang. 2016. Query2Vec: Learning Deep Intentions from Heterogenous Search Logs.
2. Mostafa Dehghani et al. 2017. Learning to Attend, Copy, and Generate for Session-Based query Suggestion. CIKM’17.

Query understanding

- Percolate Query
- Named-entity Recognition
  - Conditional Random Fields (CRF)
  - Neural named-entity recognition

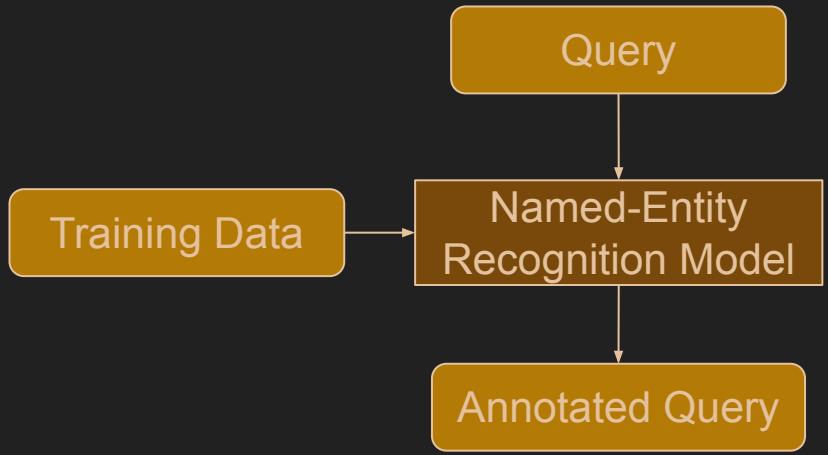

Neural named-entity recognition

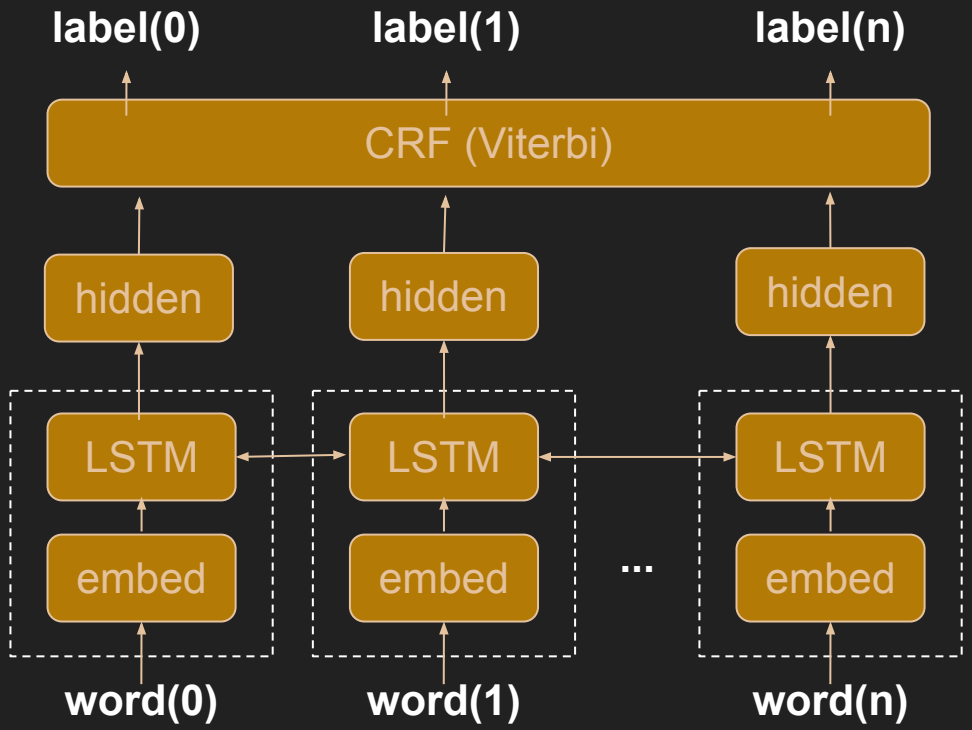

Zhiheng Huang et al. 2015. Bidirectional LSTM-CRF Models for Sequence Tagging. arXiv:1508.01991

Neural named-entity recognition

“Woman with a Dog”
Pablo Picasso

- State-of-the-art
- Open-source libraries available
- Flexible to word variations and misspellings
- Usually requires manually labeled training data

Query expansion and retrieval

- Additional knowledge
  - Facets
  - Synonyms, similar words
  - Taxonomy, ontologies, knowledge graphs
  - Graph-based techniques

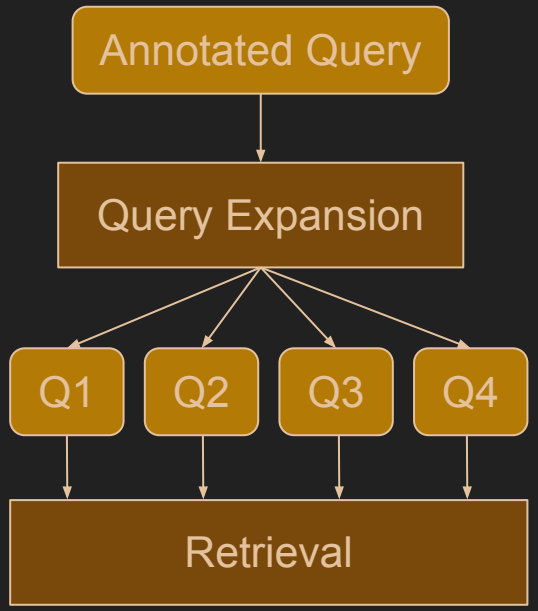

Graph-based techniques

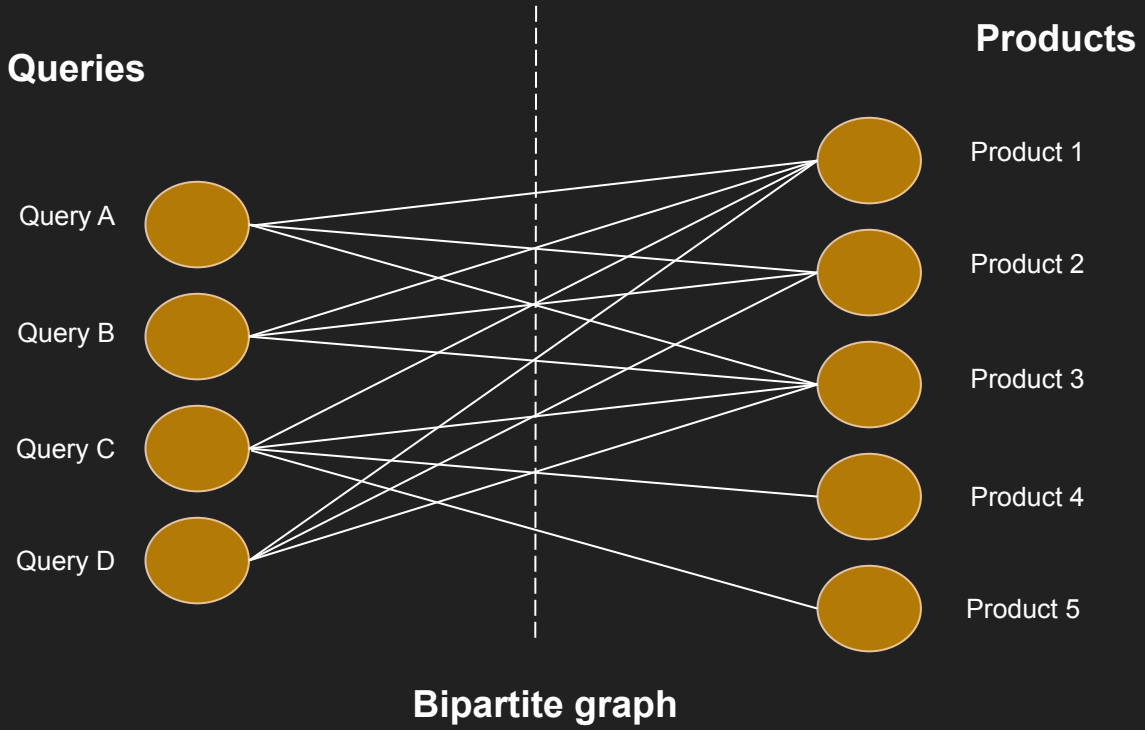

Graph-based techniques

- Similar queries
- Similar products
- Vector propagation [1]
- Graph embedding [2]

“Moses by the red sea”
Michael Tingsgård

1. Shan Jiang et al. 2016. Learning Query and Document Relevance from a Web-scale Click Graph. SIGIR '16.
2. Ming Gao et al. 2018. BiNE: Bipartite Network Embedding. SIGIR '18.

Relevance Re-scorer

- Score function is a parameterized model
- One may use click logs to train the parameters
- A threshold cut may also be learned by machine learning

Ranking

- Learning to Rank algorithms
  - RankSVM
  - LambdaRank
  - LambdaMART
  - ConvRankNet

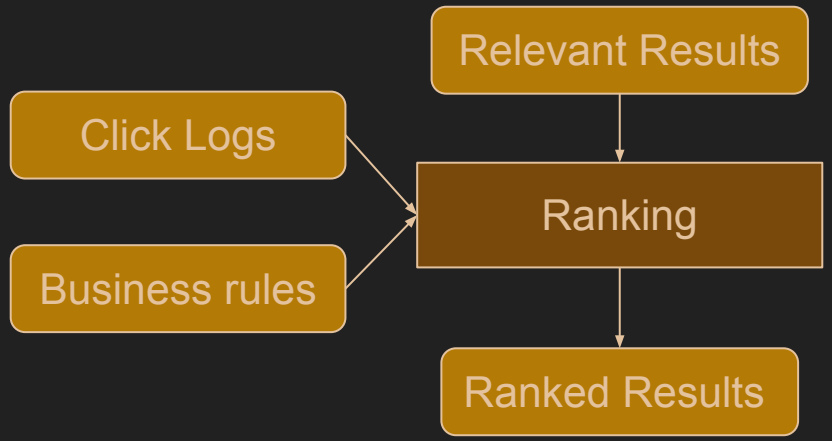

Relevance search as a mathematical function

- Every component of our search architecture is a function
- A composition of functions is also a function

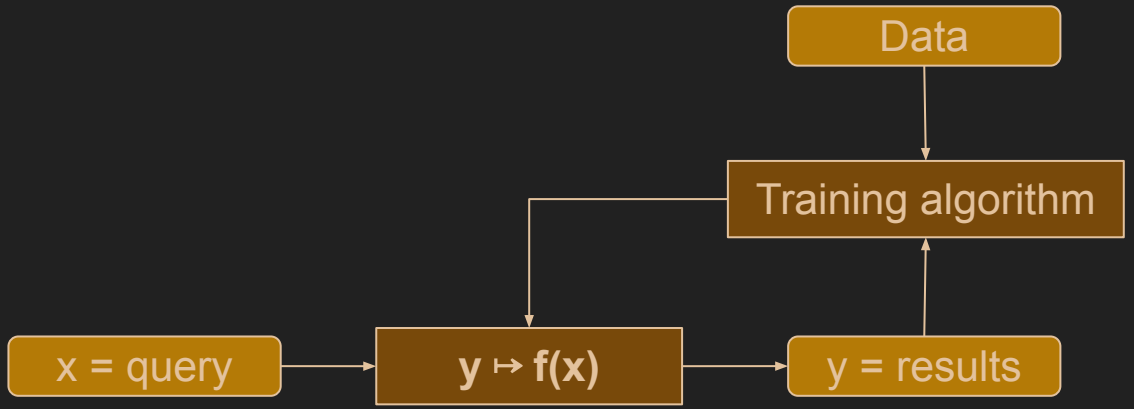

Adapting to new conditions

- The data used for training our relevance search could be conditioned in any particular interest
  - User segment / Specific users
  - Marketing campaigns
  - Regionality

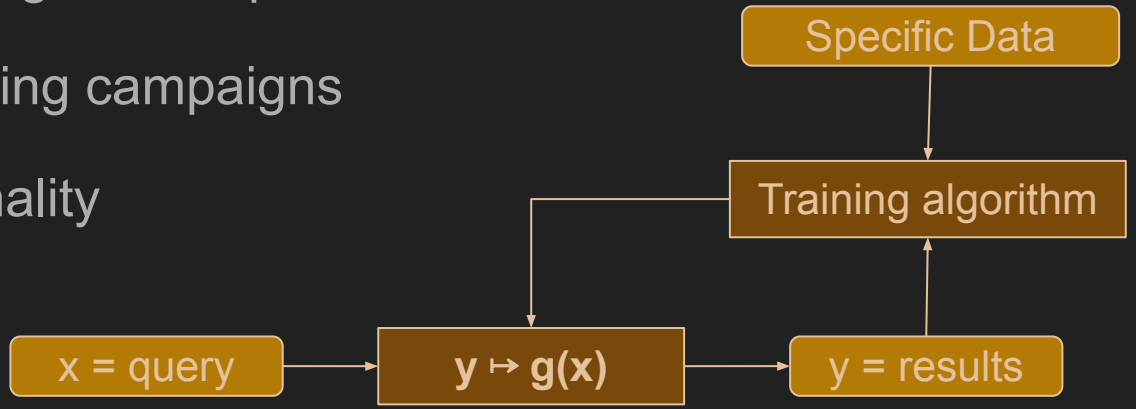

Multiple models

- One could also aggregate multiple search functions

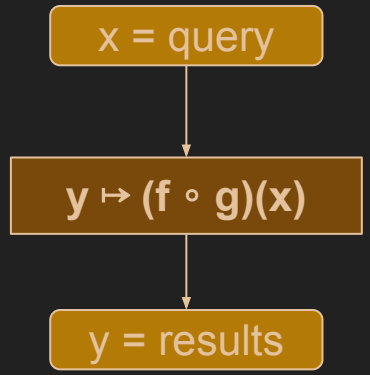

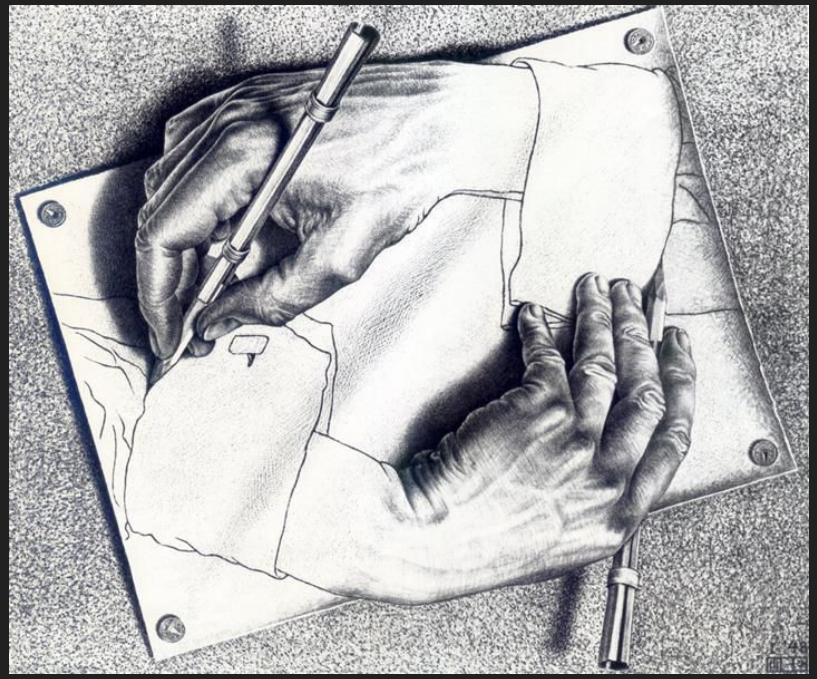

“Drawing Hands”
M.C. Escher

Wrap Up

- Companies with significant volume of Search use should focus more on data-driven search methods as a way of scaling and opening new markets
- Data Scientists are important to provide data-driven approaches to your Search
- Relevance Engineers are still important to track your Search quality and provide search adjustments quickly.

Relevance search is a science and an art!

Data-driven methods are the future of relevance search

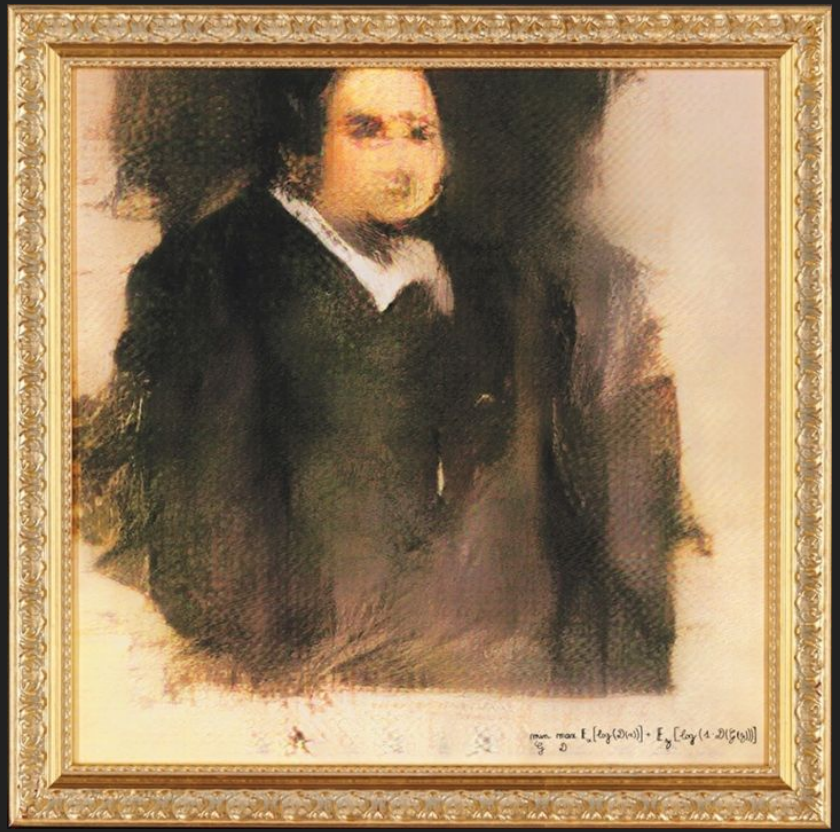

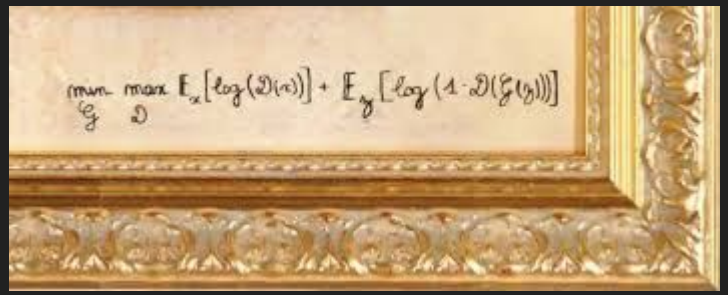

“Portrait of Edmond de Belamy”
AI generated
Sold in 2018 for over $400,000

Thank You!

Pedro Balage

@pedrobalage

pedrobalage.com