



Deep Learning in Production

Konstantinos Bitsakos

VAIX

ABOUT VAIX

- VAIX is a self-funded startup (2.5 years old) with the mission to support internet businesses utilize the power of AI for their customers and operations.
- VAIX's DNA **merges** deep **business** expertise (operations, strategy, product) with **AI expertise** at PhD levels.
- Transform businesses putting AI-powered personalization at their core.



[John O'Malia](#) (London, Paris, Gibraltar)

- 3x Serial Entrepreneur
- 1 x private-equity funded transaction
- 15 yrs Gaming & Betting
- 3 yrs Board Director
- 2 yrs PLC C-level



[Andreas Hartmann](#) (SF, Gibraltar, London)

- 17 yrs Internet Product Mgt
- 10 yrs Online Gaming
- 5 yrs Operations
- 4 Patents
- 3 yrs C-level



[Kostas Bitsakos](#) (Maryland, Detroit, Athens)

- 2010 - PhD Univ. Maryland - Computer Science
- 12 yrs SW Development
- 8 major Journal papers
- 5 yrs Machine Learning Research
- 5 yrs Technical Mgt



[Dimitris Stefanidis](#) (Philadelphia, Athens)

- 16 yrs Systems Engineering
- 9 yrs Full Stack
- 4 yrs NLP
- 4 yrs Deep Learning
- 2 yrs Game AI

- Offices in
 - London, United Kingdom
 - Athens, Greece
- Team of 16 and growing
 - Technical Team (Software and Data Engineers, ML and Backend Developers, DevOps): 11

- Recommender Systems
- Natural Language processing
 - Chat log analysis
 - Sentiment Analysis
- Lifetime customer value Modeling
- Customer Classification
 - Fraud detection/Wise Guy detection
 - VIP detection
- Image Processing
 - Object Identification/Object Segmentation
 - Action detection



Deep Learning: Hype or Panacea?

Personal Experience

- Domains
 - Image/Video
 - Voice
 - Natural Language Processing
 - Time-Series Data
 - Prediction
 - Regression
- Cons
 - Requires LOTS of data
 - Complex
 - Difficult to explain/interpret

Video Processing for security

Identify and count entities of interest (humans, vehicles, animals etc) in video camera feeds or single images.
Applicable in the security, marketing, surveillance domains amongst others.



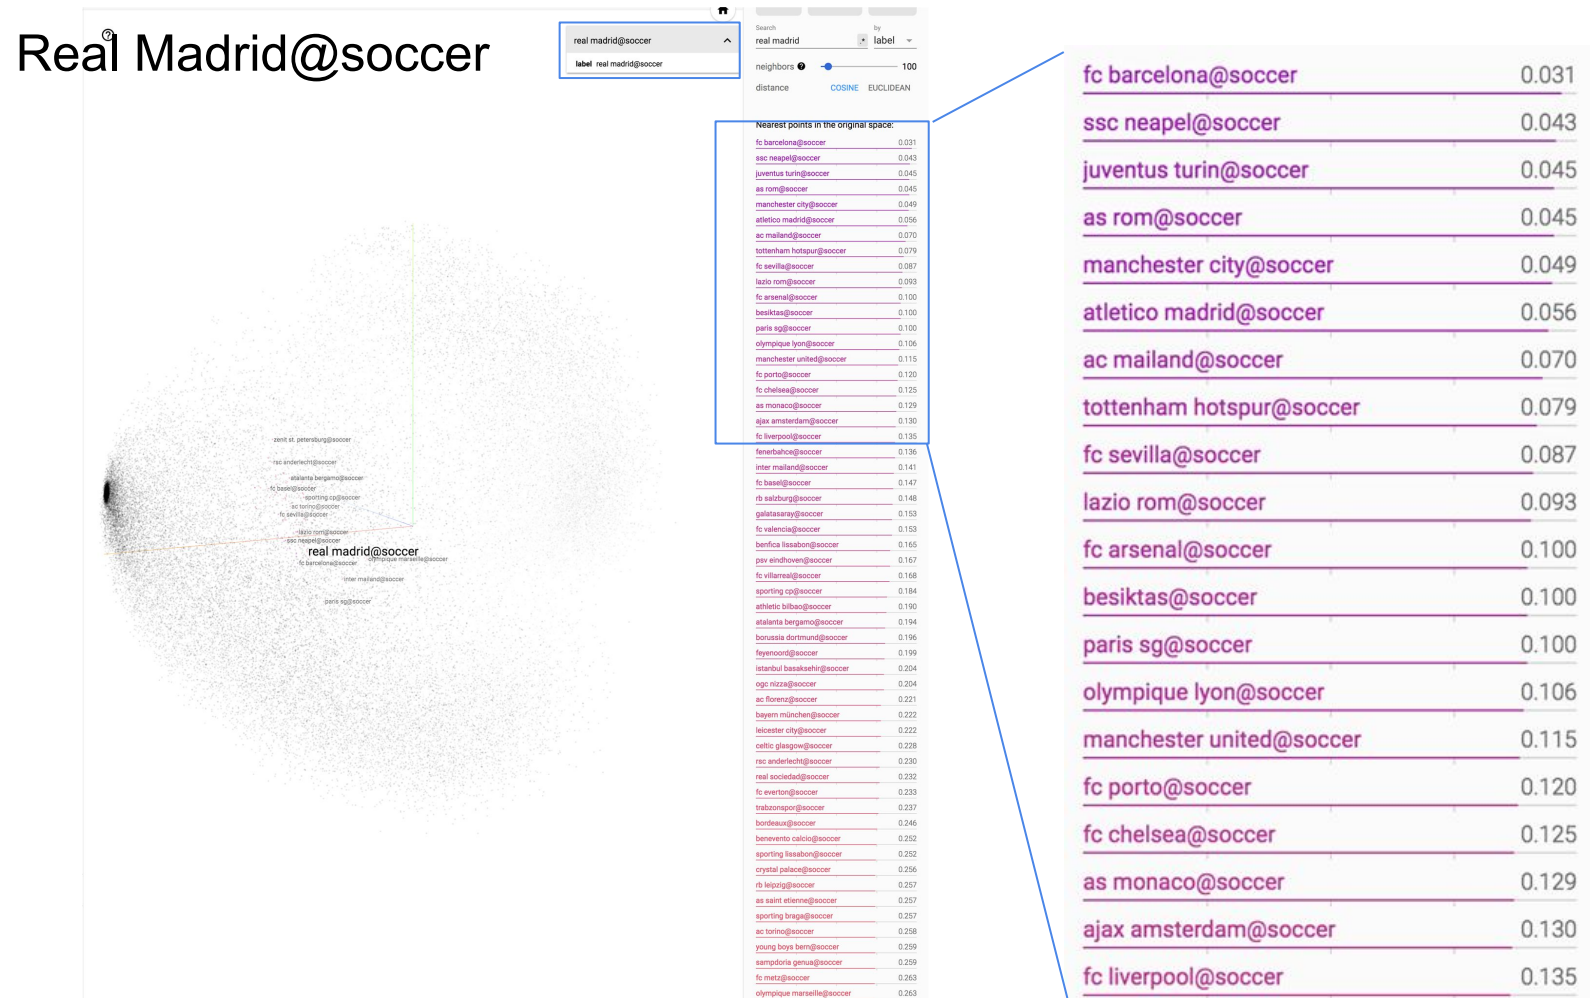
Card & Action Detection in Blackjack

Identify mis-pays in Blackjack tables by recognizing and grouping the cards of each player and detecting the actions during play



Recommender: Sport Team Embeddings

The algorithm learns on its own that the closest entity to Real Madrid is Barcelona without any explicit domain inputs. This is all learned end to end by just watching the customer purchasing patterns and activity.



Game Recommender Results: Overall*

VAIX.ai

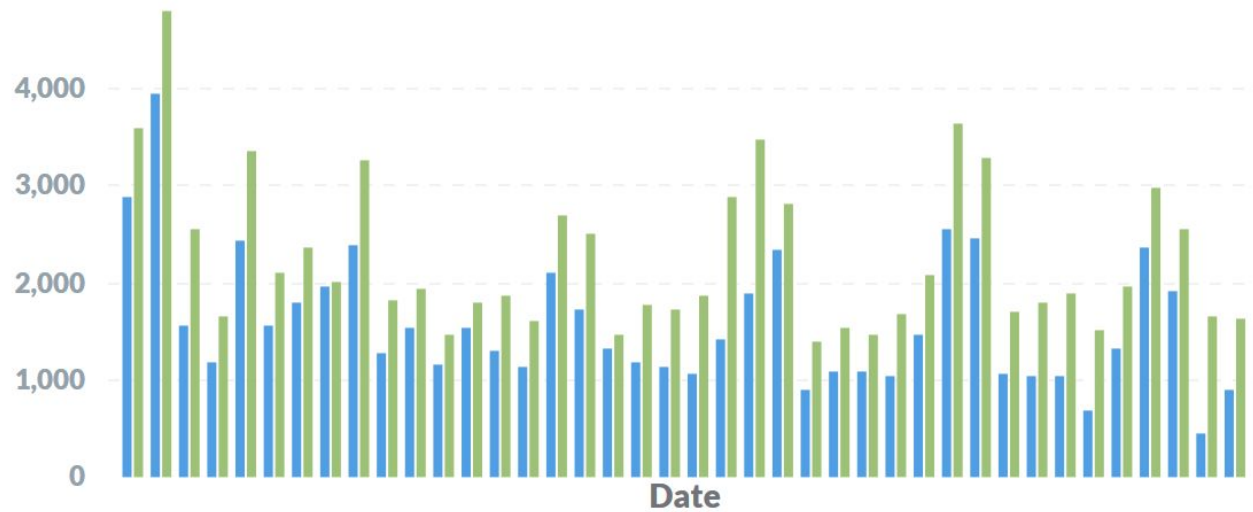
- 20% click-through rate uplift
- 10% revenue uplift
- 50% more visits in the recommendations page
- 200+% more recommended games played

* A/B test vs New+Most Popular Games

Recommender Results: Visits on recommendations page

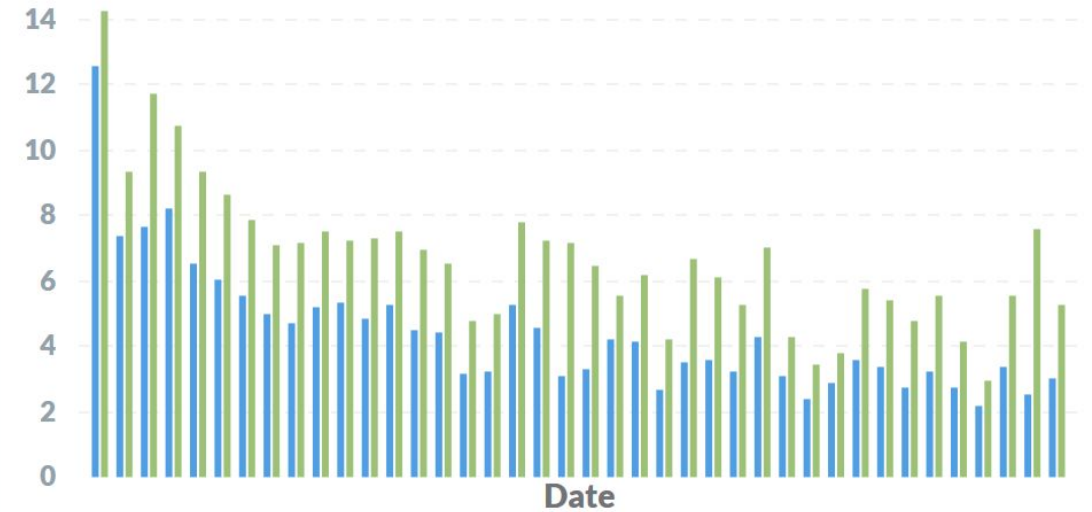
Total Clicks (Control vs Test)

● Control Group Clicks ● Test Group Clicks



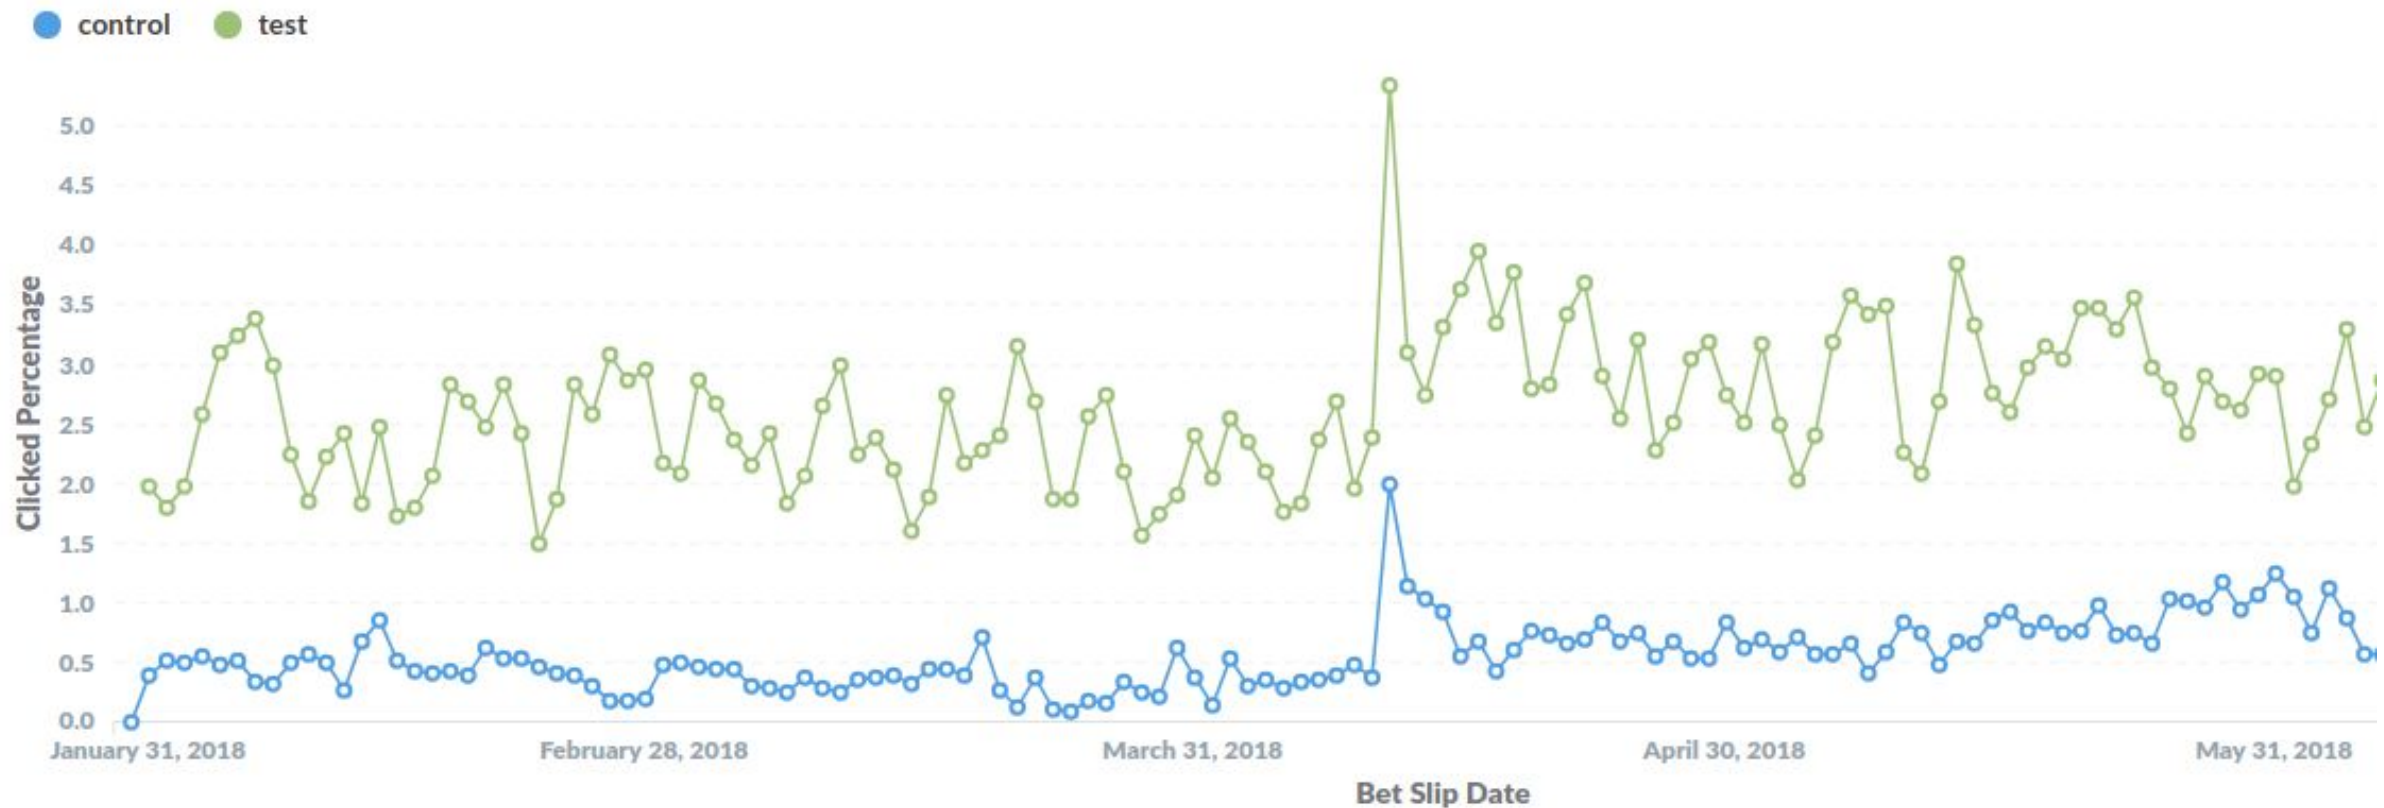
Clicks Percentage (Control vs Test)

● Control Group Percent ● Test Group Percent



Recommender Results: Actual bets placed through recommendations

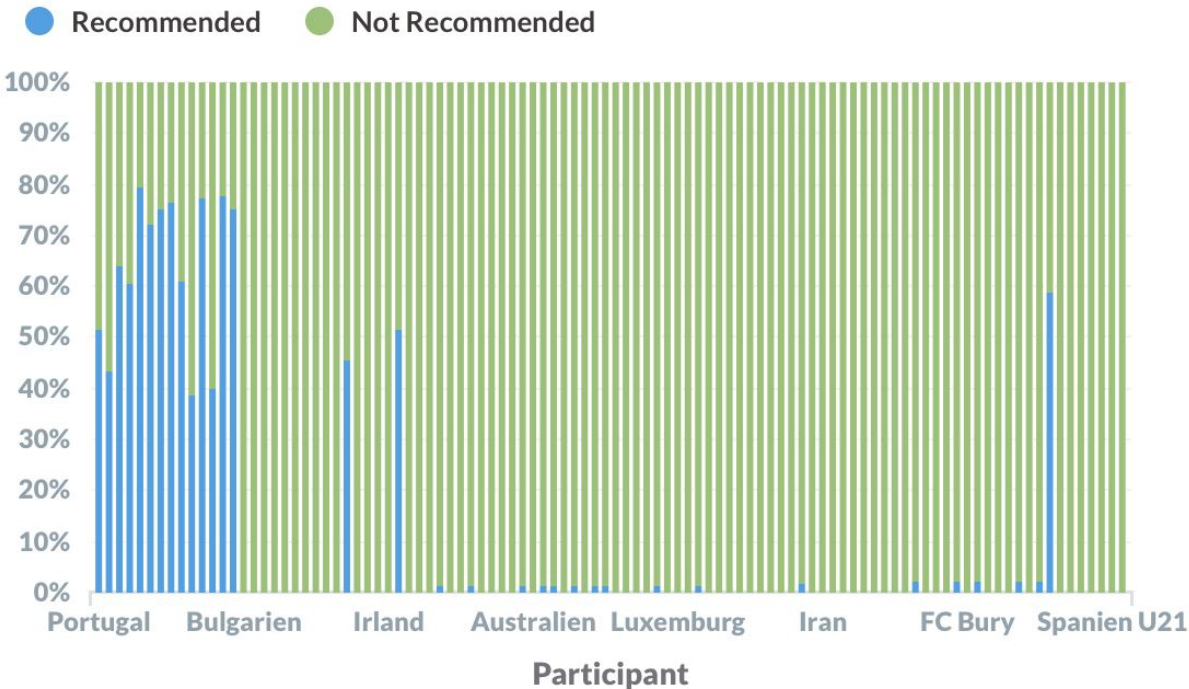
% of total bets placed



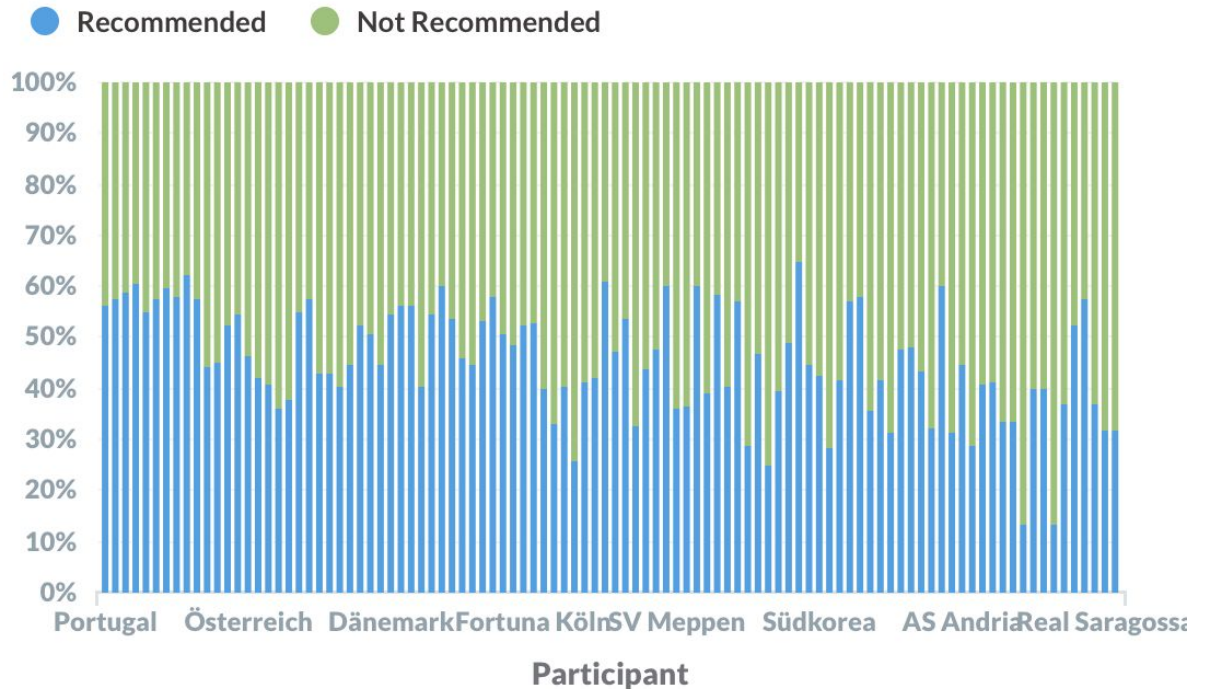
Recommender Results: Catalogue coverage

Much wider variety of products recommended via the VAIX recommendations page

Control group - Top participants (Recommended vs Not Recommended)



Test group - Top participants (Recommended vs Not Recommended)



- Use case: Casino operator data
 - training: 150K users, first 7 days of activity
 - test: 5K users
- Predicting the number of active days of a user within a 90 day window:
 - 5 classes
 - **87.1% accuracy VAIX vs 68% accuracy BI**
- Predicting customer turnover within a 90 day window:
 - ~10 classes
 - **92.8% accuracy VAIX vs 56 % accuracy BI**
- Predicting customer turnover vs acquisition cost (CPA):
 - 3 classes (below/ at/ over cost)
 - **89.7% accuracy VAIX vs 56% accuracy BI**

* Benchmark created by a CRM/BI expert using the same data set.

** The expert's model cannot be applied to users with low activity (~20% of users)

- Use case: Casino operator data (15K users, 1 year/16M transactions)
- Predicting the number of active days of a user within a 30 day window:
 - 4 classes
 - **73% accuracy VAIX vs 35% accuracy BI**
- Predicting customer turnover within a 30 day window:
 - 3 classes (below/equal/above user specific past turnover)
 - **57% accuracy vs 50 % accuracy**
- Predicting activity span (=last day of transaction):
 - 6 classes
 - 55% accuracy. Could not be predicted by the expert

* Benchmark created by a CRM/BI expert using the same data set.

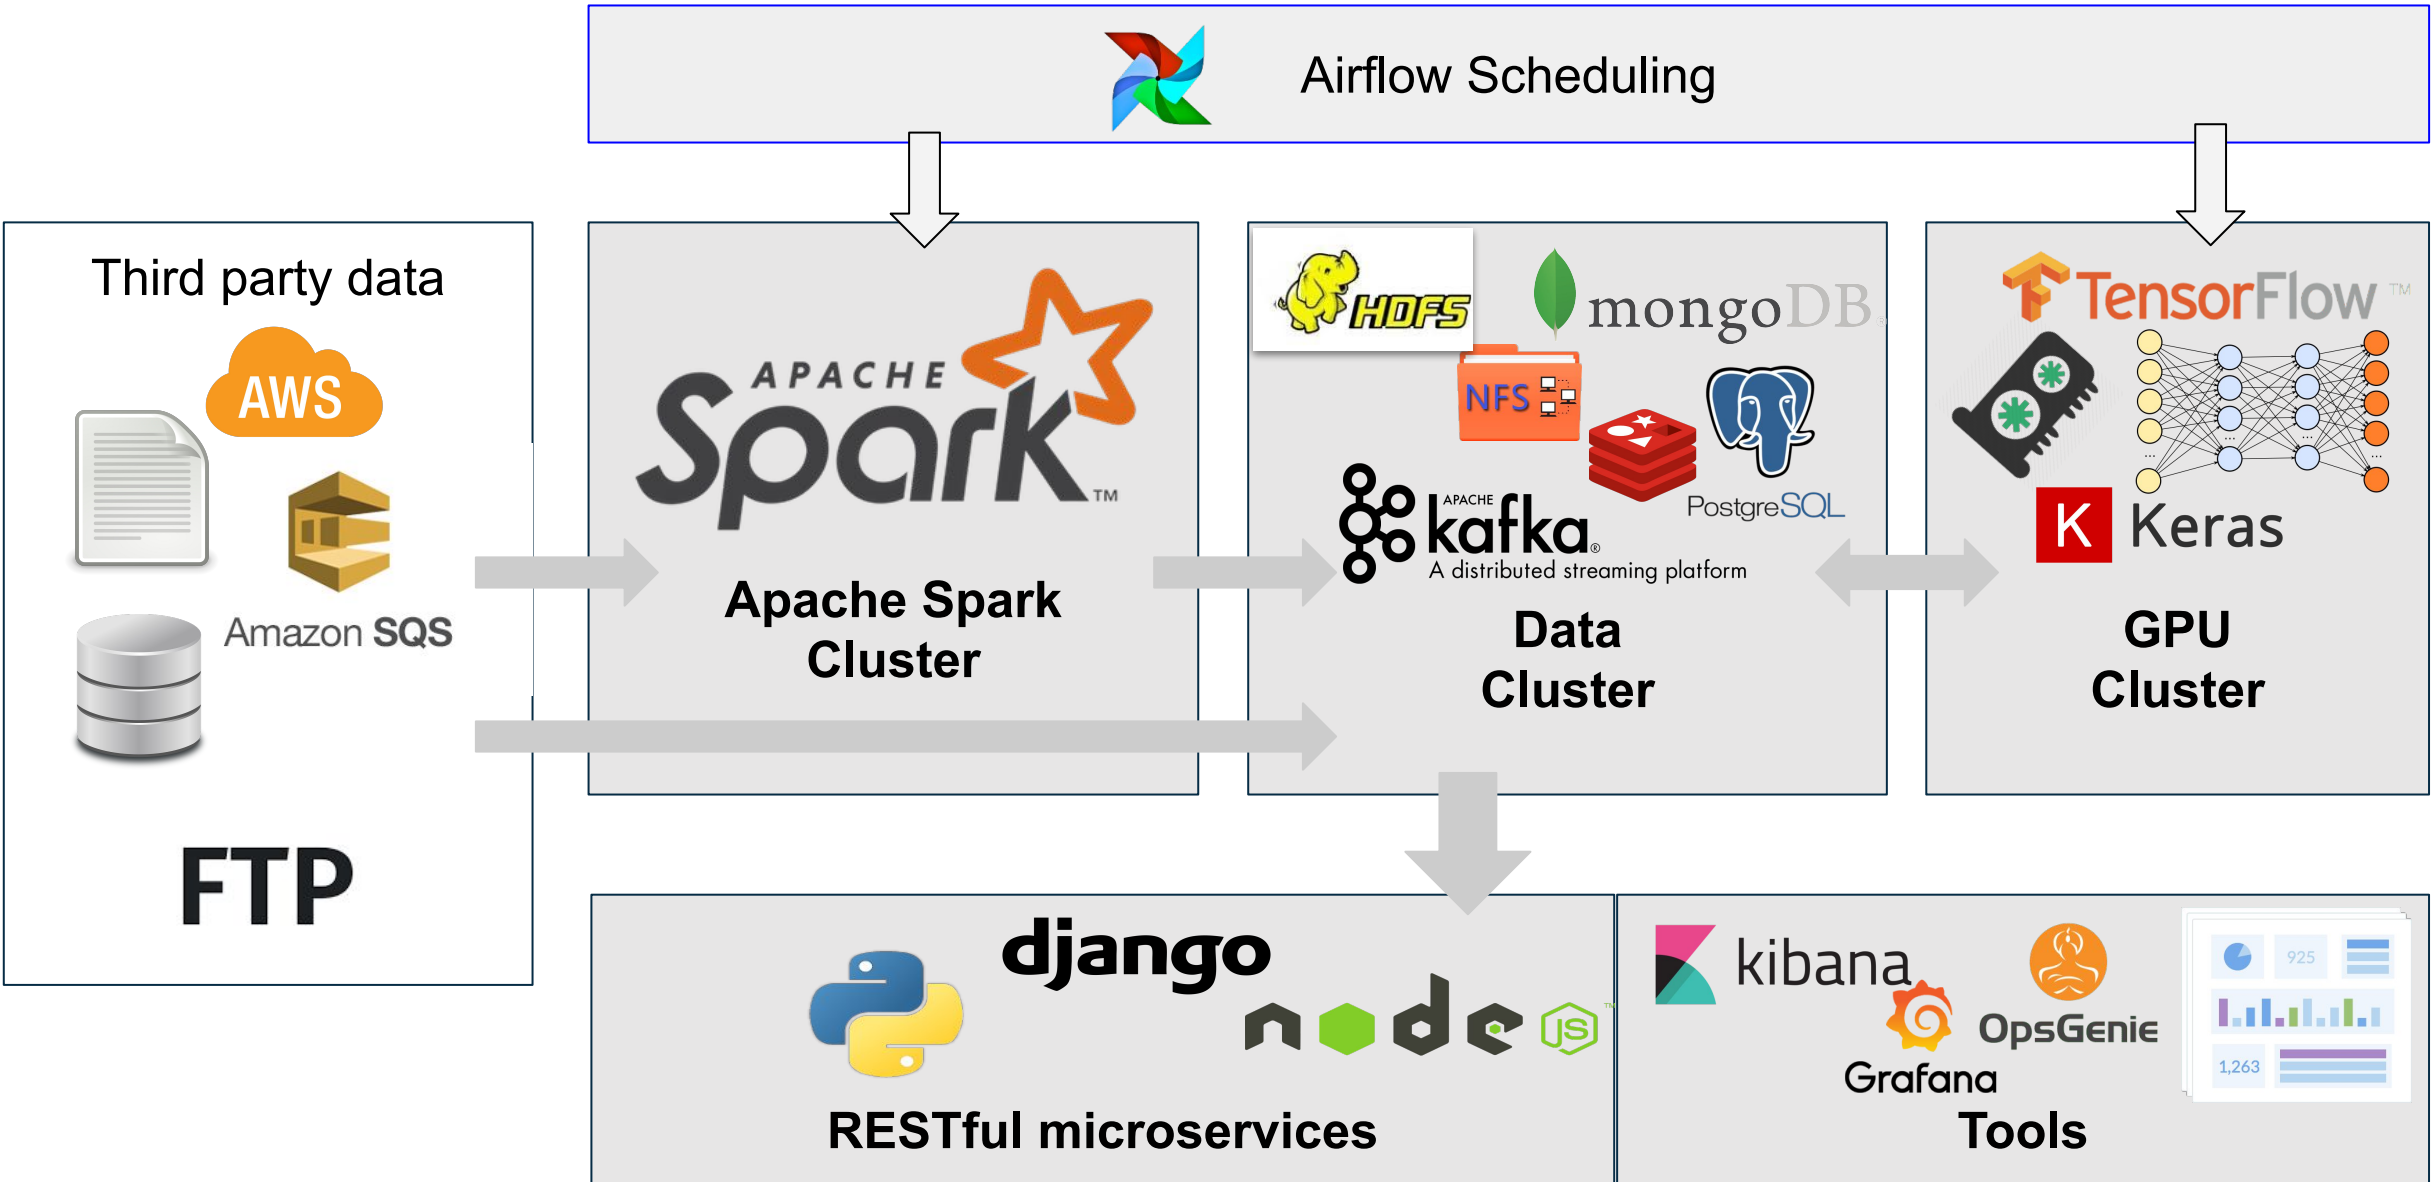
** The expert's model cannot be applied to users with low activity (~20% of users)

Deep Learning Engineering

- Data Acquisition from Operator
 - Real-time or batch (e.g. daily)
 - Anonymized data
- Data transformation
- Data storage
- ML Model training (periodic e.g. daily)
- ML Model deployment
 - Offer API endpoints to customers
- Monitor Performance & Analytics
 - 24/7 Operation
 - Automatic notification in case of potential issue

- Data Storage in a format suitable for model training
- Periodic Model training/Incremental training
- Deployment of DL models in production
 - Model testing
 - Real-time inference
 - Seamless transition to new models
- A/B Testing of DL models

Infrastructure Overview





Thank you!

Konstantinos Bitsakos

kb@vaix.ai