Working in the (Music) Industry
A Personal Experience

Sertan Şentürk – Lead Data Scientist, R&D
Kobalt Music Group, London

UKRI Centre for Doctoral Training in Artificial Intelligence and Music
Queen Mary University of London
Tuesday, July 21, 2020

Kobalt

AI + MUSIC

Queen Mary
University of London

Outline

- Introduction
- About Me
- The (Music) Industry
- Working in a “Data” Role
- Wrap-Up
- Q&A

Disclaimers

- Based on observations, experiences, opinions of me & my echo chamber
- Mileage may vary
- Views of my own
- Try to be as honest as I can be
- Hypothetical examples (mostly)

How did I end up here?

About Me (in academia)

- BSc in Electrical and Electronics Eng, Classical Guitar Certificate from a Music School
- MSc @ **Georgia Tech Center for Music Technology** (2009 – 2011)
- PhD @ **Music Technology Group, UPF** (2011 – 2017)

- Part Time Data Science Consultant @ Music and Sound Cultures Research Group, **New York University, Abu Dhabi** (June 2019 – cont.)

Georgia Tech Center for Music Technology College of Architecture

upf. MTG Music Technology Group

MaSC جامعة نيويورك أبوظبي NYU ABU DHABI

About Me (circa 2017)

- Worked on many different topics and music traditions nothing really relevant to the industry
  - Audio-score alignment, makam/raga recognition, music corpus analysis and discovery, ontologies, music prediction and generation, interactive music, …
  - Turkish makam, Hindustani, Carnatic …
- Gained decent developer skills pretty desirable anywhere
  - Python, git, issue tracking, testing, continuous integration
- Advocated open source & reproducible research
  - Build my own toolbox, created/contributed to more than 10 corpora/datasets
  - Work available online at https://dunya.compmusic.upf.edu/ visibility helps
- Gave many talks and seminars

Switching to the industry

- Applied to ~5 postdocs
  - Got 1 offer
  - Feedback: Worked on a very specific music style / topic, have specific skills
- Applied to only 1 position in the industry
  - Got an offer
  - Feedback: Worked on diverse topics, good developer skills
- **Note:** Both feedback are meaningful; perspectives are different
  - **Postdoc:** Transferable work, research output, paper writing, supervision, …
  - **DS:** Transferable foundations, delivery record, presentation skills, mentoring, …

TIP: Both are just a job, and a job is not more than what it is

About Me (in the industry)

- Senior Data Scientist @ **SoundCloud** (summer 2017)
- Senior → Lead Data Scientist, R&D @ **Kobalt Music Group** (2018 –)

The Music Industry
(from a royalty perspective)

MUSIC INDUSTRY IN BRIEF (!)

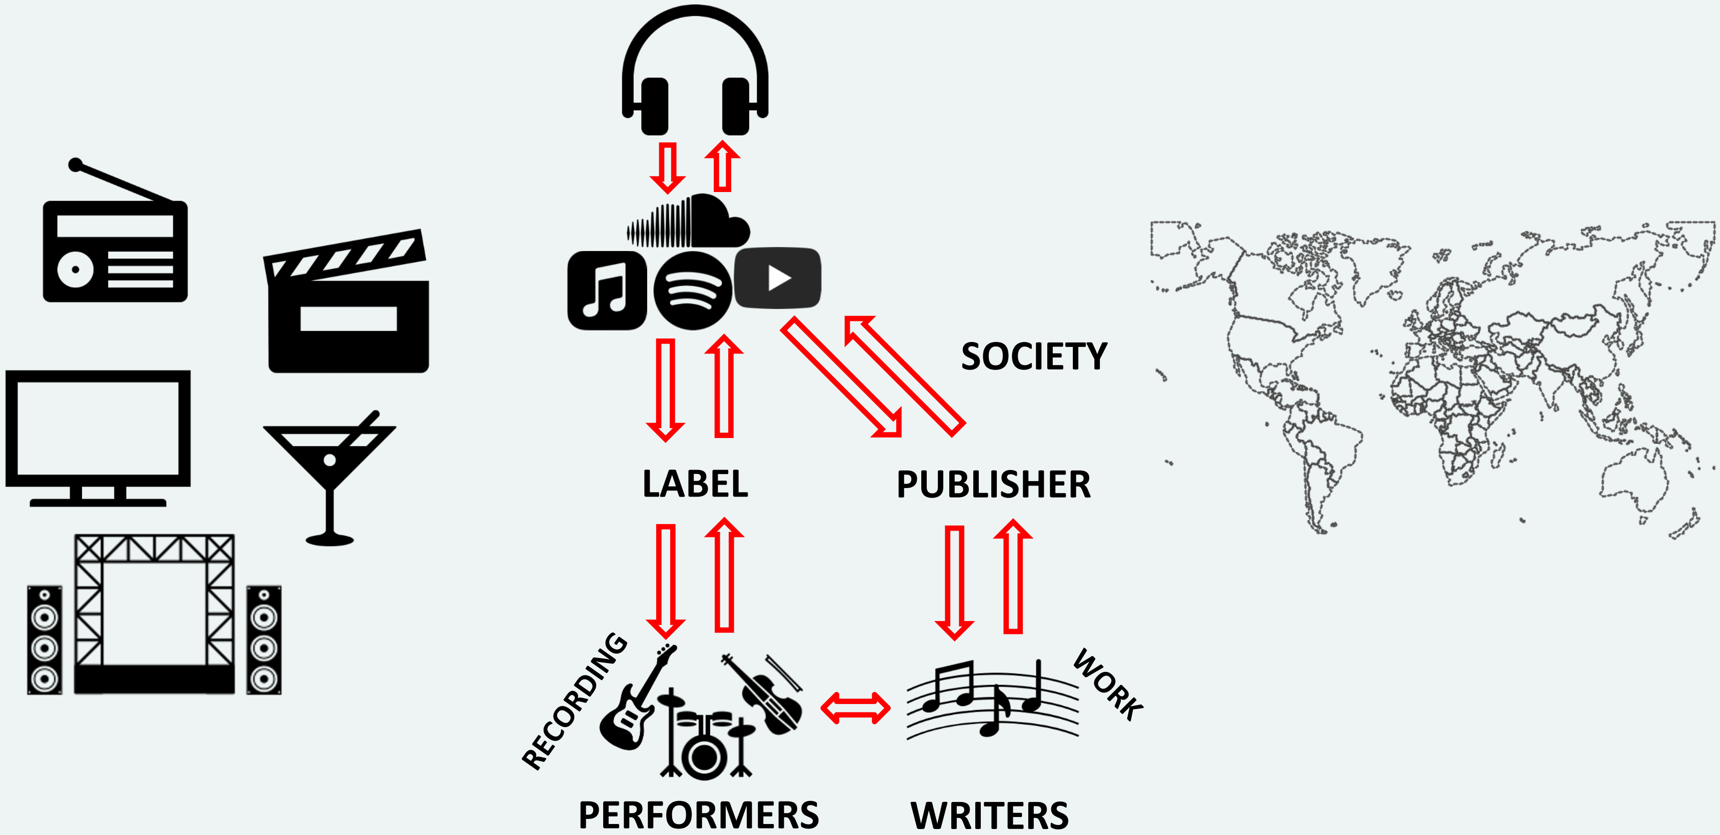

COMPLEX ECOSYSTEM

CONSUMER	100 000 000's…
RETAILER	1000's…
SOCIETY	100's…
PUBLISHER / RECORD LABEL	1000's…
CREATOR	1 000 000's…

MULTI-LAYERED

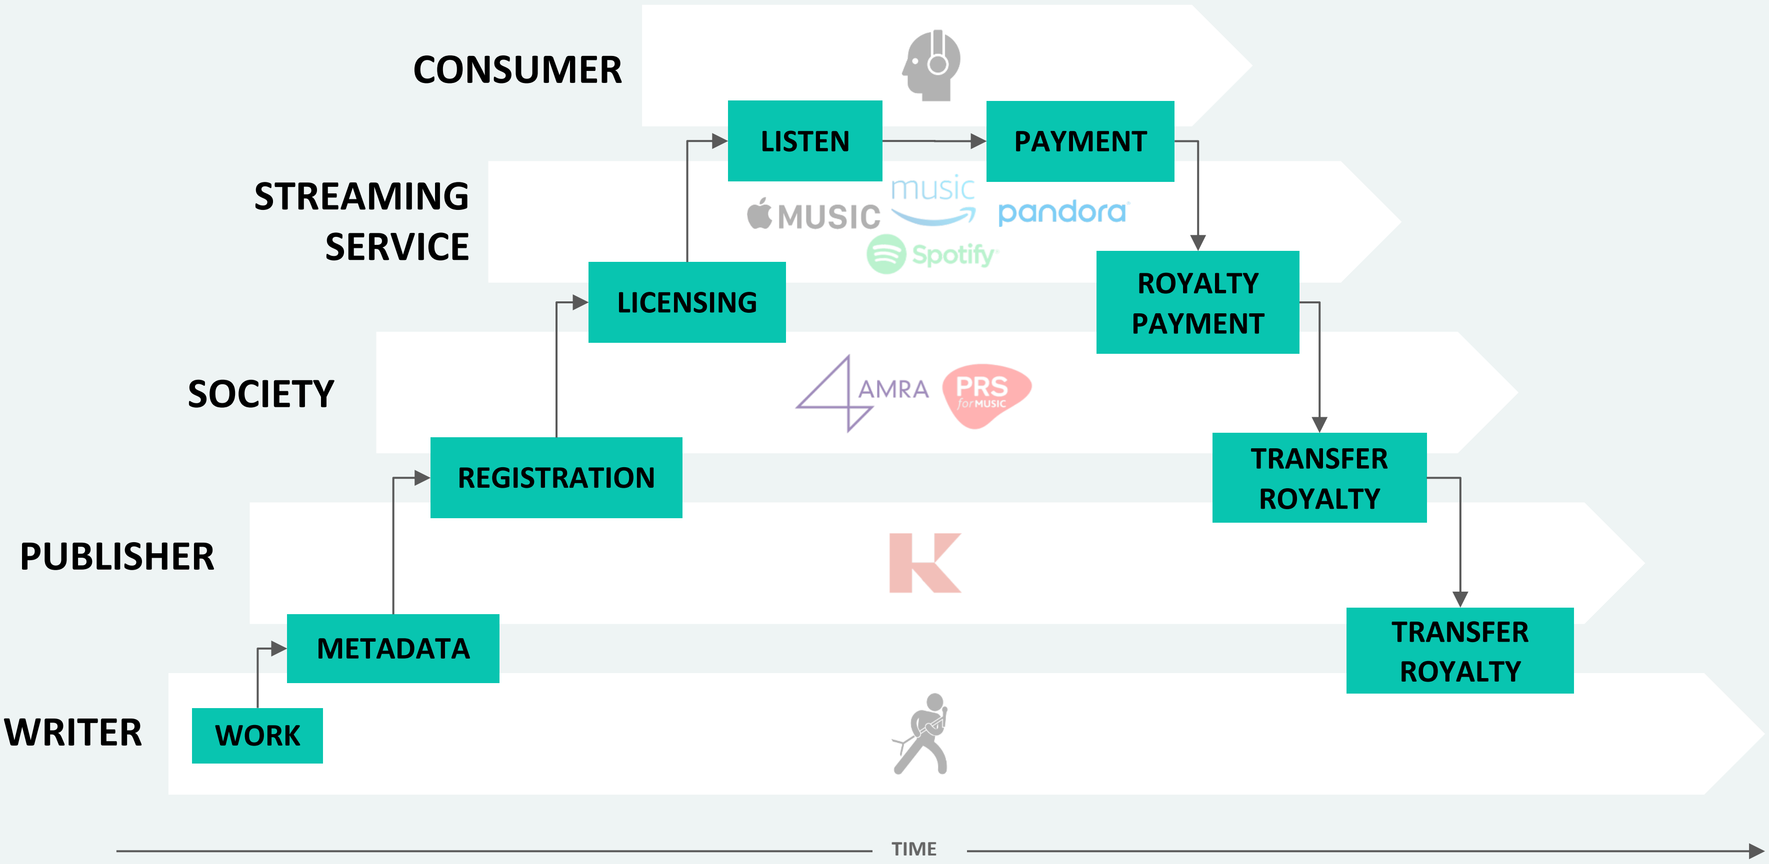

OPAQUENESS

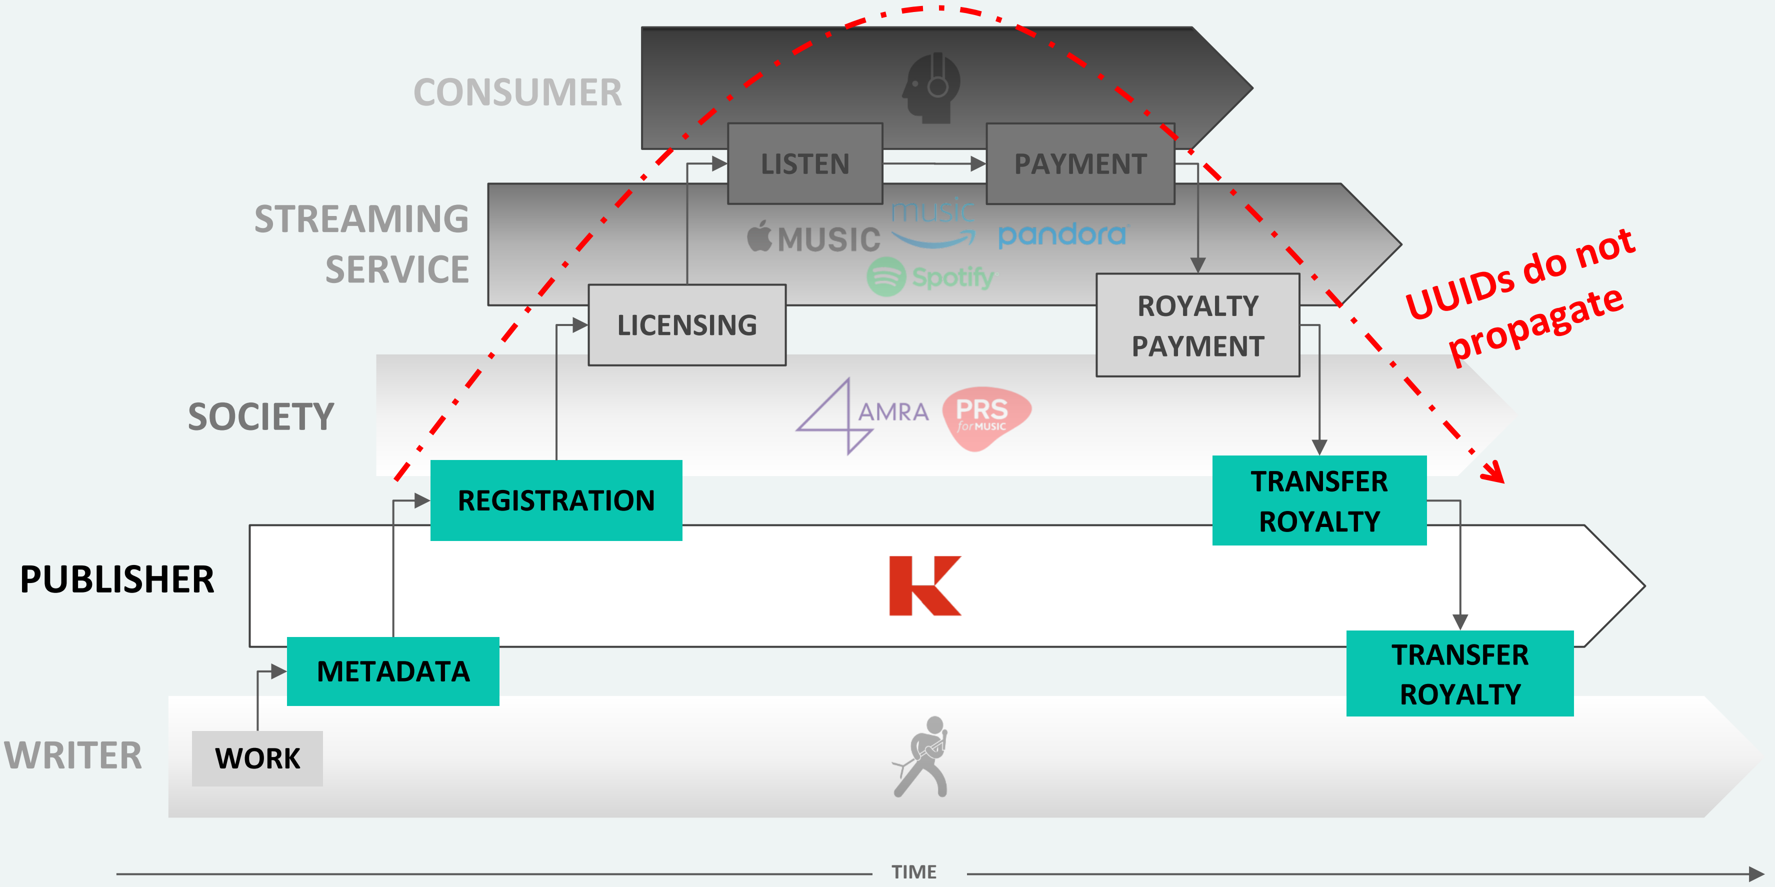
CONSUMER
LISTEN
PAYMENT
STREAMING SERVICE
MUSIC
music
pandora
Spotify
LICENSING
ROYALTY PAYMENT
UUIDs do not propagate
SOCIETY
AMRA
PRS for MUSIC
REGISTRATION
TRANSFER ROYALTY
PUBLISHER
METADATA
TRANSFER ROYALTY
WRITER
WORK
TIME

SLOWNESS

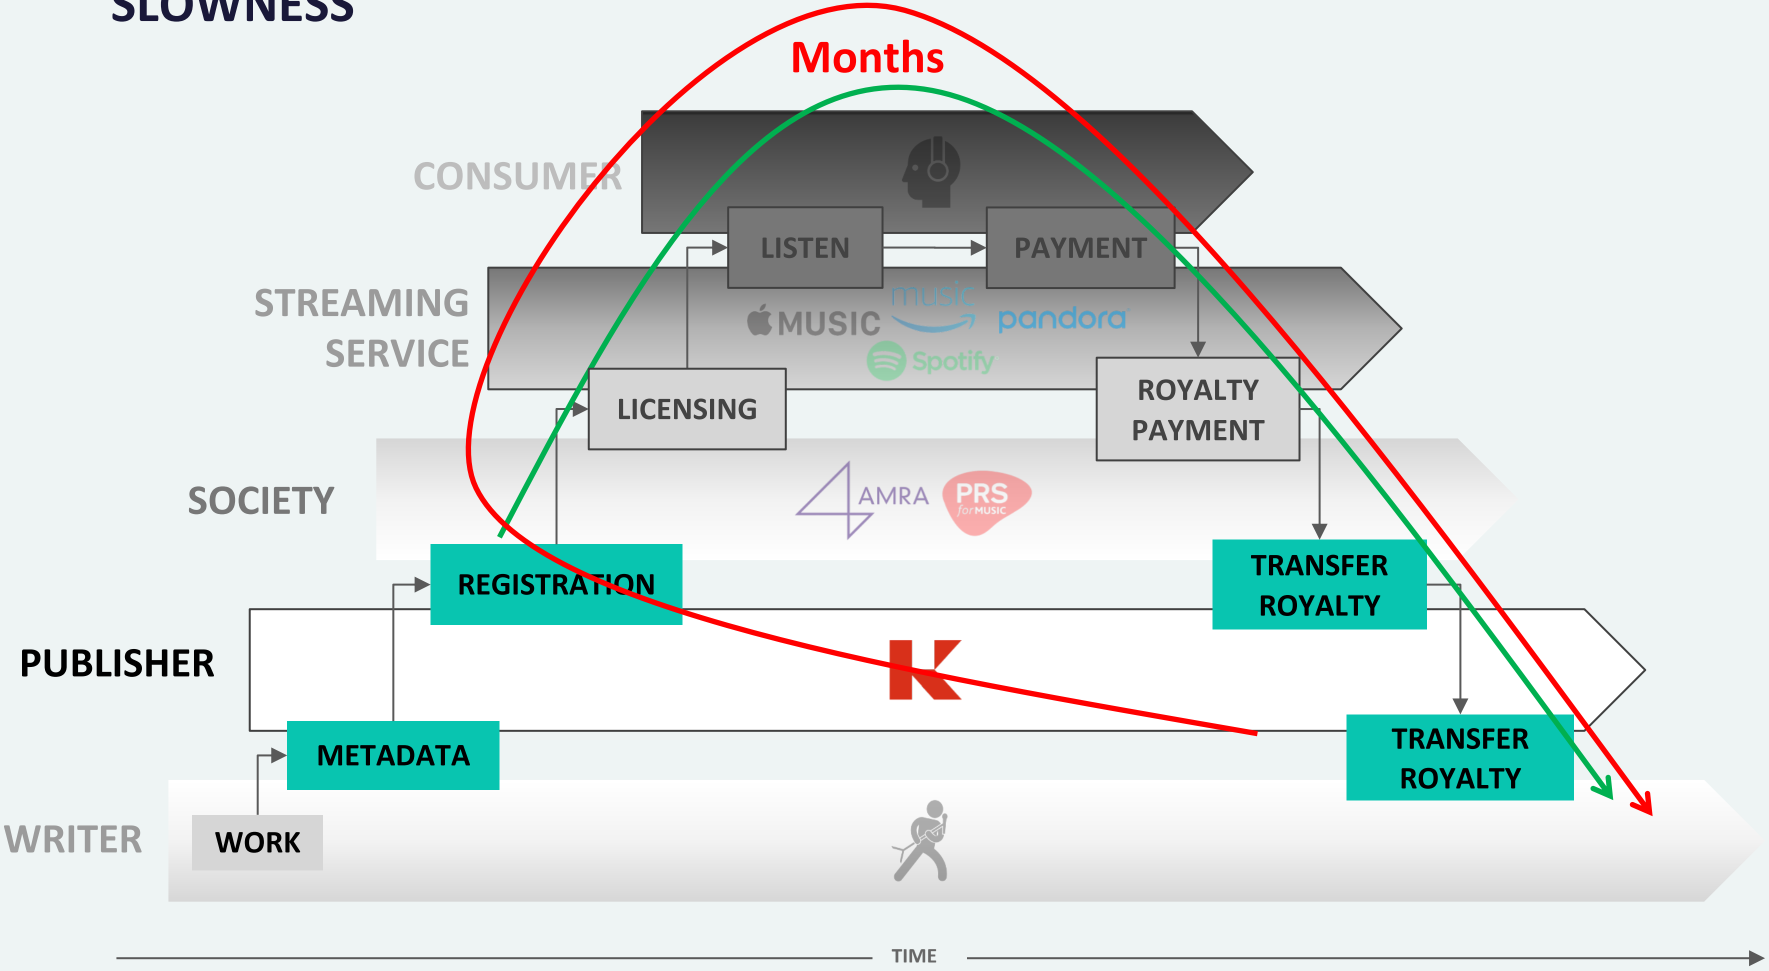
Months
CONSUMER
LISTEN
PAYMENT
STREAMING SERVICE
MUSIC
music
pandora
Spotify
LICENSING
ROYALTY PAYMENT
SOCIETY
AMRA
PRS for MUSIC
REGISTRATION
TRANSFER ROYALTY
PUBLISHER
METADATA
TRANSFER ROYALTY
WRITER
WORK
TIME

Potential data-driven tasks in copyright companies

- Entity Recognition (Metadata Matching)
- Knowledge Graphs
- Music Autotagging
- Audio fingerprinting and version identification
- Music Similarity
- Predictive analytics (for artists, copyright orgs etc.)
- Recommendations (to creators, marketing etc.)
- Music generation

How does the industry work?

Organization in High-Level

- **Business**
  - e.g. in Kobalt: Label, publisher, collection society, synch, neighbouring rights, …
- **Product**
  - Identifies problems (beyond symptoms) and new opportunities
  - Captures requirements, sets strategy and planning to bring a solution as well as features

    i.e. **why, what,** and **when** an engineering team builds of a product
  - Communicates between the business and engineering
- **Tech**
  - Builders: Backend, frontend, QA, data engineer, research scientist, …
  - Ideally autonomous in choosing the product solutions
- Delivery, design, …

Organization Structure

- Can be grouped into business units, functions, and/or roles
  - Publishing Tech, “Creators”, Data
- Different approaches
  - Pyramid, Matrix, “Spotify” model, …
- Very fluid from org to org, with organization size, and in time
  - “Re-orgs” happen all the time

- In our level, teams are more interesting:
  - Units dealing with a tangible product, task etc.
  - Personalized recommendations, audio fingerprinting, cloud infrastructure, …
  - Mostly (or at least aimed to be) self-sufficient

Ways of Working in Tech Teams

- **Agile** (http://agilemanifesto.org/)
  - Iterative changes, constant feedback, responsive to change, …
  - Scrum vs Kanban
  - For research teams it gets vague…
- **Vision – Mission**
- **Long Term Plans**
  - Roadmap, end of year success criteria
- **Objectives and key results (OKR)**
  - Monthly, quarterly etc…
  - Planning, design and product management intensive

    "stakeholder" management, requirements capturing, user story mapping, tech/design review…
  - Approval by upper management
  - Agreed by stakeholders before the cycle starts
  - Progress scored in shorter intervals (e.g. two weeks) and communicated to the stakeholders

OKR Example

Objective: Improve the audio auto-tagging pipeline

KR1: Handle loads up to 10M recordings daily

KR2: Reduce the training time from 10 days to 8 hours

KR3: Monthly cost reduced by at least 30%

KR4: Improve Precision by 0.05

KR5: Reduce the manual ops from 3 full-time equivalent (FTE) to 2.5

OKR Feedback

Objective: Improve the audio auto-tagging pipeline

we don’t know how much of improvements is enough, could be more precise

why do we need to improve it?

what’s wrong with the current solution?

KR1: Handle loads up to 10M recordings daily

KR2: Reduce the training time from 10 days to 8 hours

KR3: Monthly cost reduced by at least 30%

KR4: Improve Precision by 0.05

KR5: Reduce the manual ops from 3 full-time equivalent (FTE) to 2.5

3 types of KR here - scale, cost, quality; maybe there should be several objectives

OKR Recap

Objective1: Unlock the growth of audio auto-tagging service into new markets

- We will reach the capacity of the solution (300k recordings per day) in 6 months. It will hinder our growth to new markets, reduce customer satisfaction, leading to churn

- Current model takes 10 days to train, which slows down our iterations, has a high risk of failure (3 out of 11 jobs last quarter), and makes it difficult to test and deploy. It takes up to 2 full days of an engineer to monitor a training jobs.

KR1: Handle loads up to 10M recordings daily

KR2: Reduce the training time from 10 days to 8 hours

…

Why all the fuss?

- **Maximize Return on Investment (ROI)**
  - Bring revenue
  - Lower costs
- **Minimize Opportunity Cost**
  - "the loss of potential gain from other alternatives when one alternative is chosen"

Working in a “Data” Role

Data-related roles

- Data engineer
- Data analyst
- Data scientist
- ML engineer
- Research scientist
- ML scientist
- …
- Product manager
- Software Engineer

- Overlapping functionalities
- Meanings differ from org to org or even by job specs
- New titles emerge all the time
- Traditional roles also start having ML deployment as a "good-to-have"

DS Ways of Working

- Pairing with/Getting embedded into teams
  - Help (consultancy, training) another team to achieve data-driven objectives
  - Impact/influence dependent on the attitude and trust
  - Communication/convincing skills are crucial

- Cross business team
  - Serving other parts of the organization (data lakes, internal ML tracking API)
  - Operational (e.g. run predictions on a one-off tasks)
  - Great to capture the inner workings and the overall picture

DS Ways of Working (2)

- Team building a product
  - Very focused, dynamic, hands-on
  - or (depending on the point of view) narrow-scoped, less exploration
  - Depending on the team and expectations, it can be amazing or horrendous
- Research & Development
  - Create or bring novel technologies & methods into the organization
  - Most independent, but you should still have impact
  - Typically there are other teams who productionize proof of concepts

Data Science Cycle

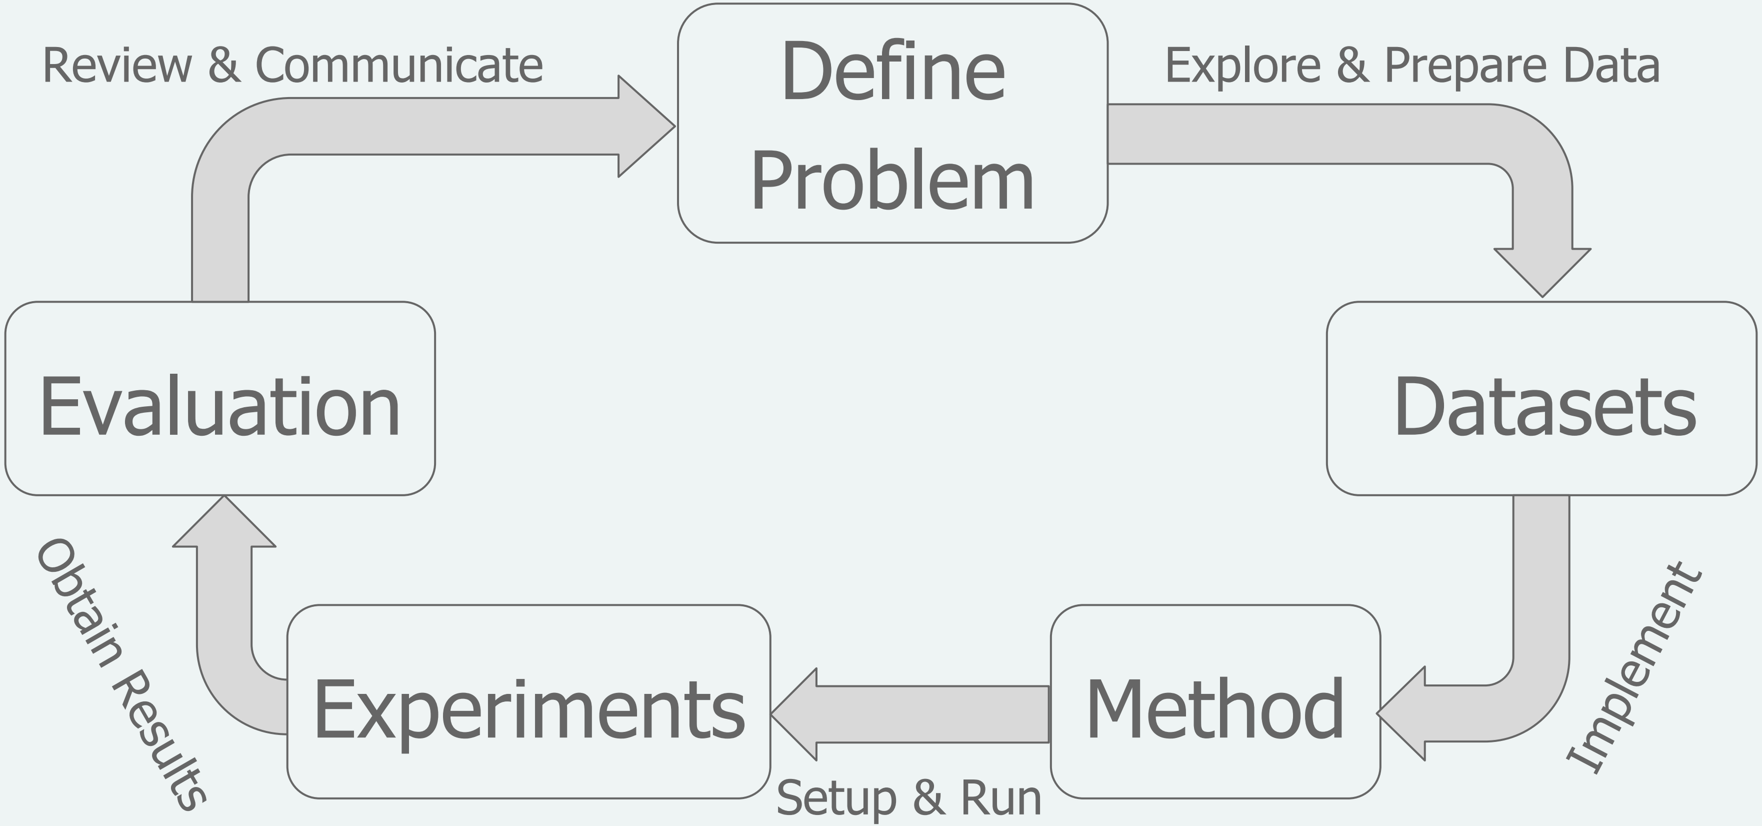

Example ML Production Cycle

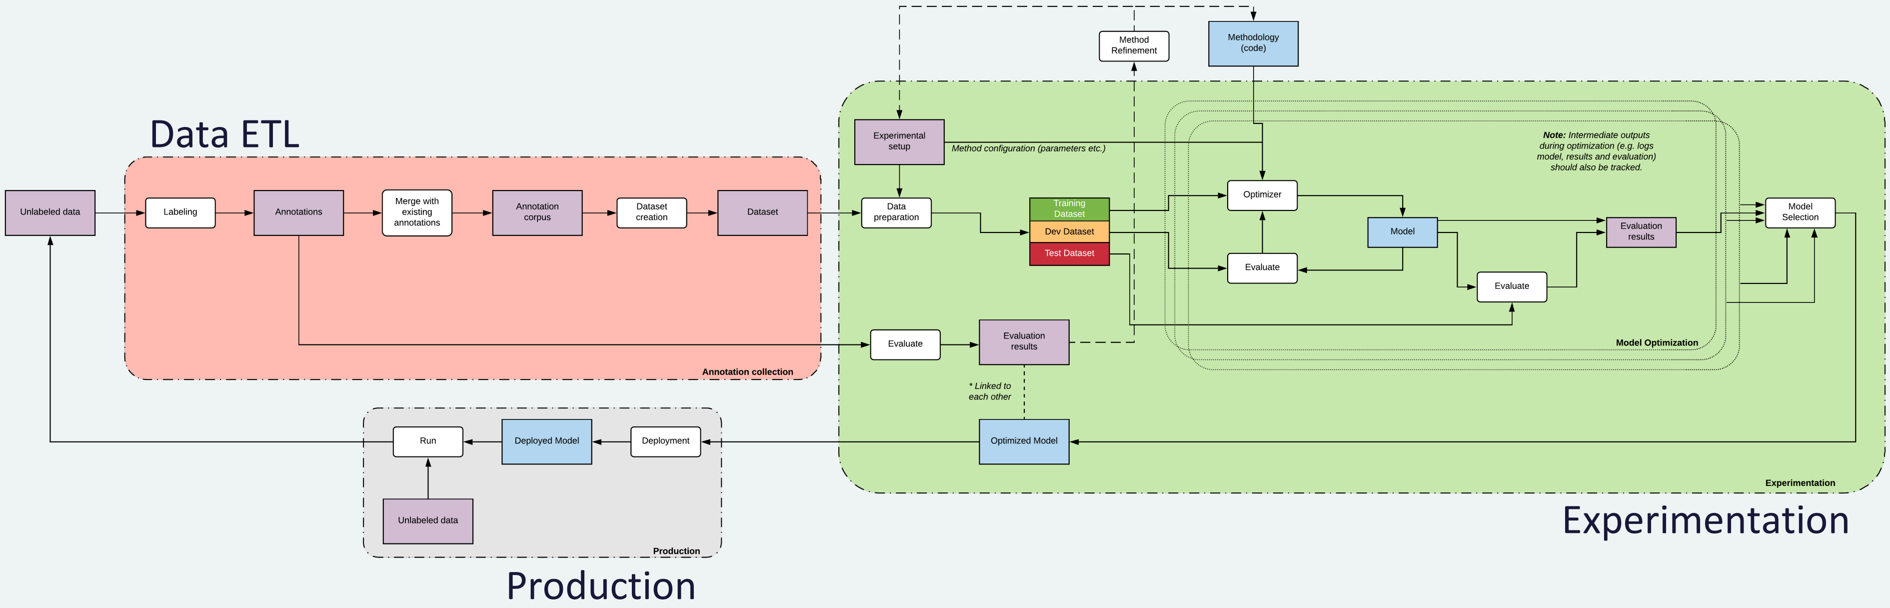

Tech Stack & Skills

- Cloud Computing
  - AWS, Google Cloud, Azure, …
- Microservices
  - Docker – one functionality/task/job per container
  - API-driven
- Big Data
  - AWS EMR, Google Big Query, EC2 with GPU attached etc. …
  - Scala/Spark
  - Large (but not infinite) resources
- Data Science Ecosystem
  - Very similar to academia (Jupyter, pandas, tensorflow, Keras, pytorch, seaborn …)
  - 3rd party tools (Amazon Sagemaker, Glue, Databricks Platform, ElasticSearch, …)
- Software Developer/Engineer Skills
  - Unit tests, PRs, Code Style, Continuous Integration/Deployment, Containerization, Infrastructure as Code, Databases, …

Wrap-Up

- Your skills are transferable! Don't make anyone tell you otherwise!
- Your work cycle may not change drastically
- But ways of working and responsibilities will change (occasionally)
- Communication is sometimes the most important part
- Good manager/team is preferable than a well-known company/area
- In many cases, **Data > ML models** in the industry
- Learn about software development (before), cloud computing and production processes (on the job)

Resources

- Chip Huyen (2020). Machine Learning Production Pipeline. ICML
- Allen AI (2018). Writing Code for NLP Research. EMNLP
- Ed Newton-Rex (2019). Six Principles Of Applied Research. Medium Post
- Martin Zinkevich (2018). Best Practices for ML Engineering. Google Developers
- Erik Bernhardsson (2020). Never attribute to stupidity that which is adequately explained by opportunity cost. Personal Blog
- Agile Alliance (2001). Manifesto for Agile Software Development.
- Andrej Karpathy (2018). Programming the 2.0 Stack. Full Stack Deep Learning
- Jeremiah Lee (2020). Failed #SquadGoals. Personal Blog

QUESTIONS?

Sertan Şentürk – Lead Data Scientist, R&D
Kobalt Music Group, London

@sertansenturk
contact@sertansenturk.com
https://www.sertansenturk.com