ML Infrastructure and Operations - Overview

ML Lifecycle, Intro to ML Systems and Infrastructure

Endri Deliu: endri.deliu@univie.ac.at

EURO AUSTRIA

About Me

Focus in *controllable and safe* AI, *modular* alternatives to LLM-s

- Distinguished Fellow, ML, Indeed - San Francisco, USA
- AI Architect, Salesforce AI - San Francisco, USA
- Principal Architect Recommendations, Playstation - San Francisco, USA
- Architect Search, Autodesk - San Francisco, USA

endri.deliu@univie.ac.at

EURO AUSTRIA

ML Infrastructure and Operations

- Set of processes, architectures, infrastructure and tools to ensure, reproducible, scalable, robust, and observable ML *lifecycle* development and **deployments** in production (offline/streaming/online)

ML Canonical Lifecycle - Simple

- From Data to Deployment and Beyond

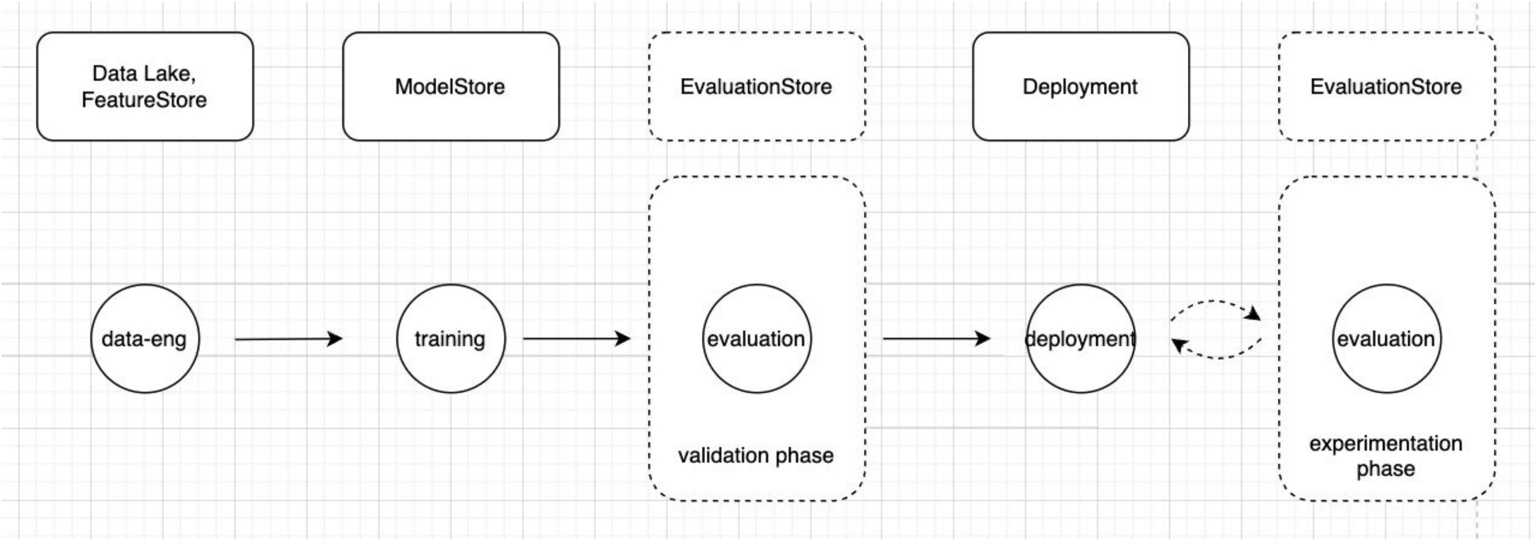

ML Canonical Lifecycle - Simple

- From Data to Deployment and Beyond

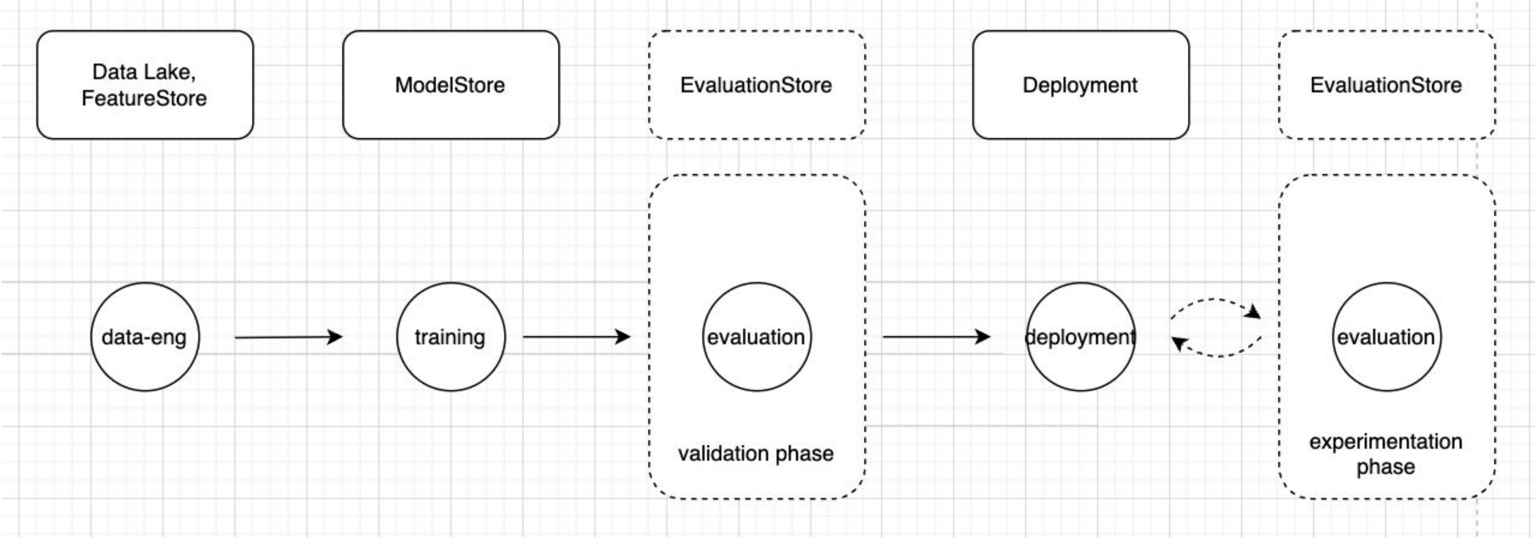

ML Lifecycle - Simple Pipeline

- Lifecycle expressed as flow/pipeline(s)

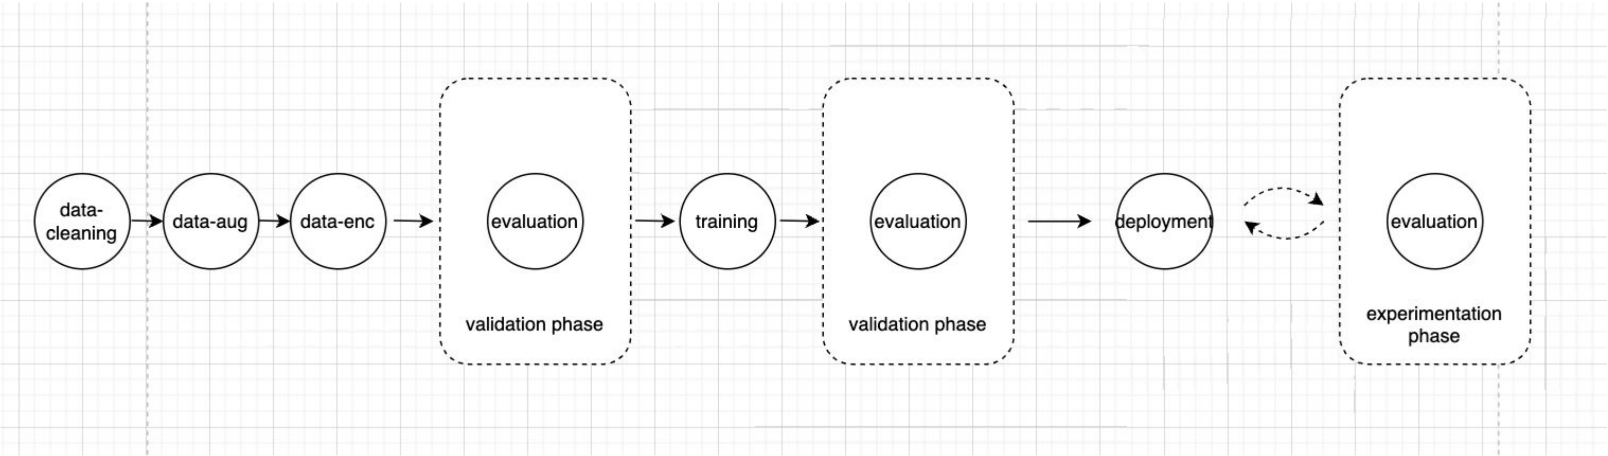

Systems - Orchestration and Lifecycle Management

- Paradigm on popular orchestrator systems:
  - process centric (Airflow, Argo, ….) - lifecycle stages are coupled with (comp.) process
  - event centric (Step Functions, ..) - manages lifecycle stage transitions as events
- Focus on managing **lifecycle(s),** - evtl. lots of them (...millions)

Lifecycle - What about Realtime/Online?

- Lifecycle of request…
- Recommender system example
- Realtime pipeline systems (internal to big companies...)

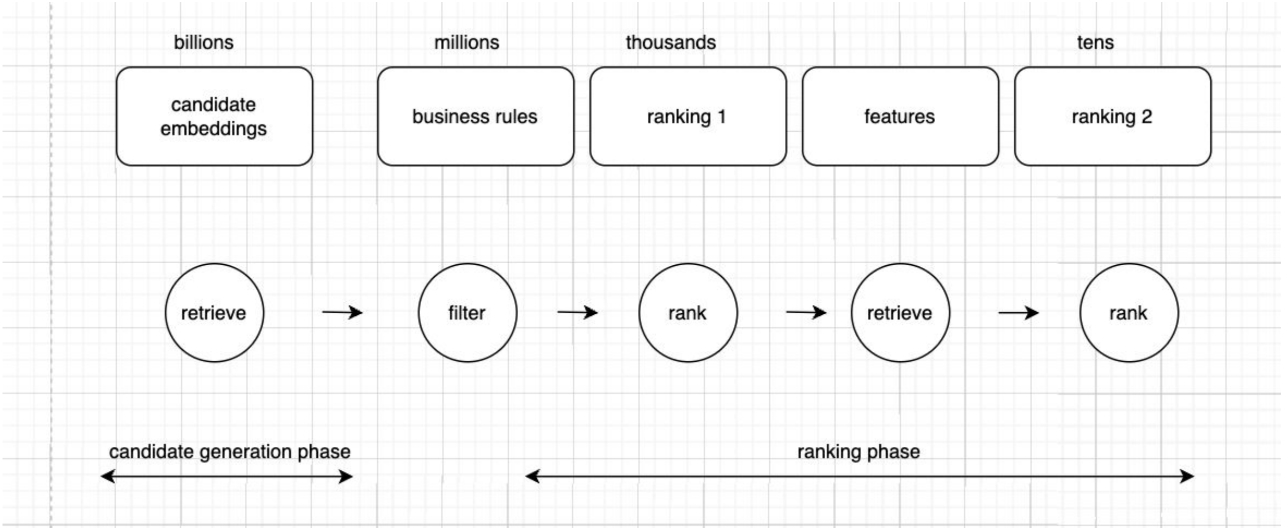

Data Systems in ML Operations

- Discover, Store and Reuse data for high scale ML
- **Data Warehouses**: large tables, curated data: used for analytics and history
  - Requires Query Processing: AWS Redshift + Tableau, Google BigQuery + Looker
- **Data Lakes**:
  - structured/unstructured, large-scale, offline, data of all company, analytics, ML, etc. cheap(ish).
  - simple* metadata systems for schemas, versions and raw data
  - Requires processing (typically **Spark**)
- **Feature Store(s)**:
  - Specialized for ML features, offline **and** online, not cheap, for curated and reusable data
  - Online data in DB, manages offline online skew, realtime features via streaming

Data Systems in ML Operations

- **Data Warehouses**:
- **Data Lakes**:
- **Feature Store(s)**:
- **Build vs Buy**

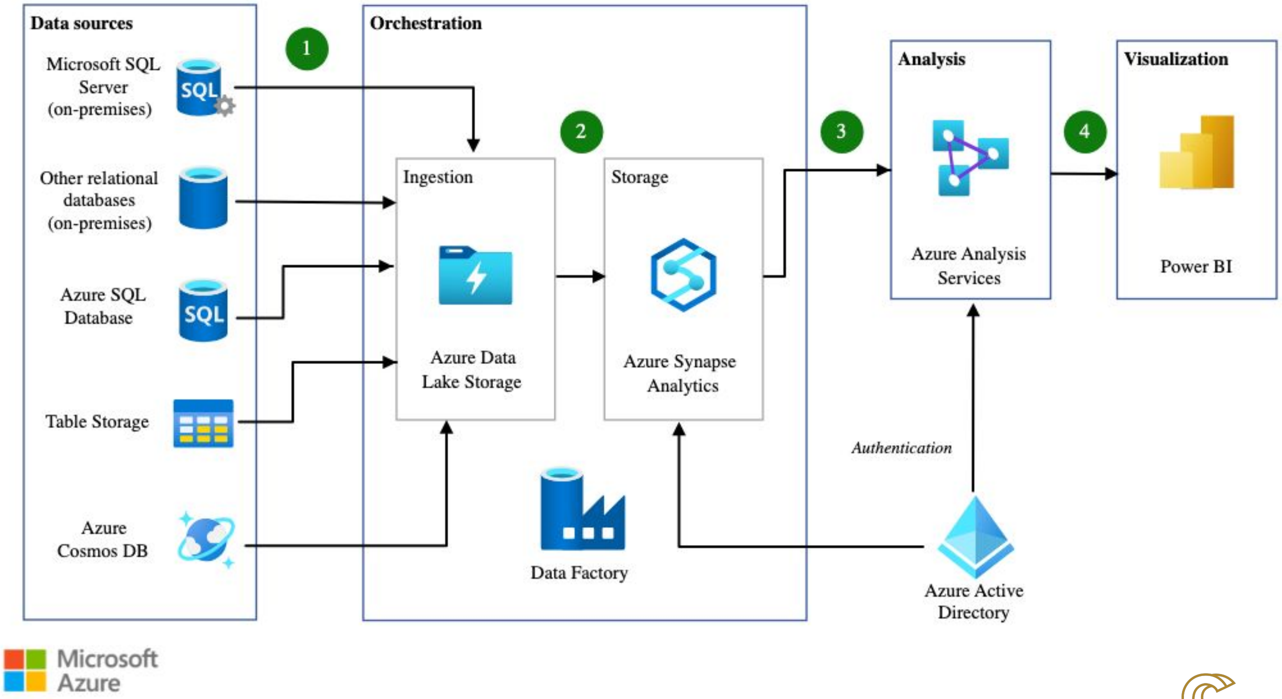

Model Infra and Systems

- Model Training Systems/Infrastructure:
  - Large **Spark** Cluster(s), or **K8**, or **Ray, Cloud Vendors** … Support for distributed training (Decision trees, Boosted Models, DNN etc.). Scheduling, Multi-tenancy, Rate limiting, Billing, …
- Inference Systems/Infrastructure:
  - Trained model != deployed model i.e. compilation
- ModelStores:
  - Stores models, provides versioning, and model *metadata*,
  - Versions, tracks code/lib dependencies, model lineage, input/output schemas, model cards,... checkpoints, cadence of retraining, App specific tags, …
- Build vs Buy

Model Infra and Systems

- Model Inference and Systems:

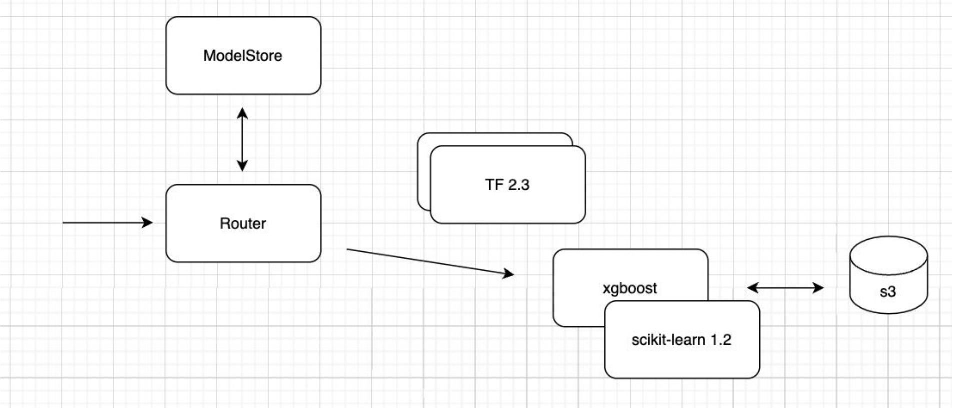

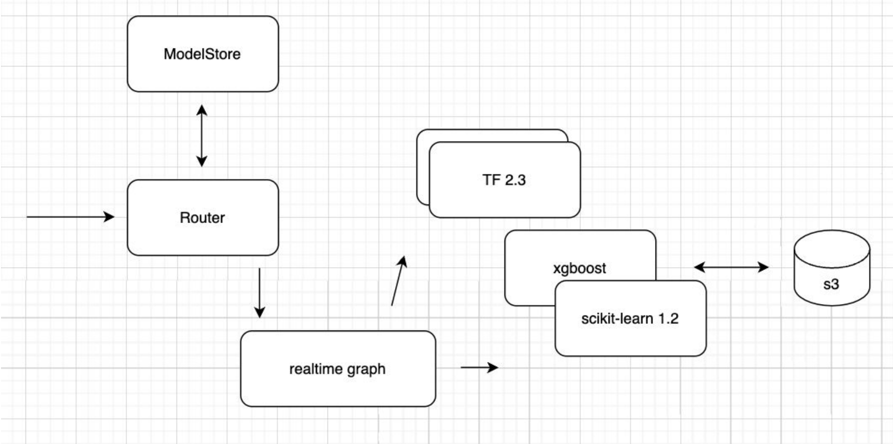

Testing in AI

- Current Situation in Evaluation Approach
  - Not very principled: manual, ad-hoc, blinders on narrow performance aspects (i.e. accuracy)
  - Metric centric
- Quality in ML/AI Context:
  - Quality is about validating **behavioral scenarios**
  - Clear **pass/fail** outcome, similar to **software eng. testing** (unit, integration…)
  - Metrics are just part of story, they represent *data*
  - Talk about Quality Assurance
  - Shift from Metric Centric to **Test-Centric Paradigm**

EURO
AUSTRIA

Testing and Metrics in AI

- Metric Systems: required and various providers
- Metrics sourced/calculated by sql engines: Presto, Athena, Trino, …
- OpenTSDB + Graphana
- Elastic Search + Kibana
- Vendors: emerging ecosystem - few companies, Arize AI, Evidently AI
- Build vs Buy

EURO
AUSTRIA

Testing in AI

- Beyond Accuracy:
  - Holistically validate diff. **behavioral scenarios**
- Quality has multiple dimensions
  - *Performance (accuracy, rmse etc.)*
  - *Robustness (perturb inputs and check changes)*
  - *Privacy (check for leaking private info)*
  - *Security (red teams, attack own ML system)*
  - *Fairness (segments, under/overrepresented..)*
  - *Cost (inference latency, overall $ cost)*

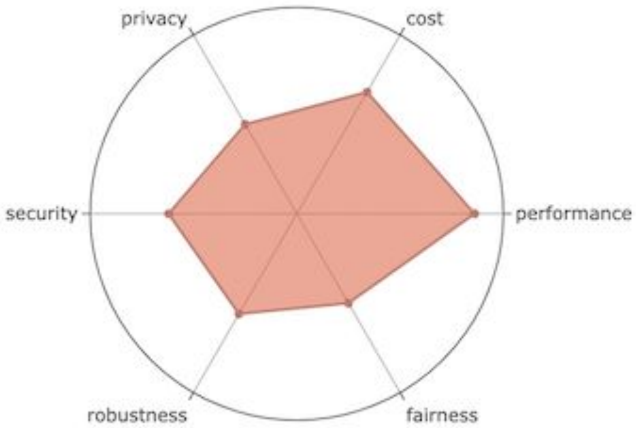

Testing and Deployment - An Interplay

- Lifecycle - see as continuous journey to check/ensure **quality**:
- Deployments - (long) **Processes,** not Events

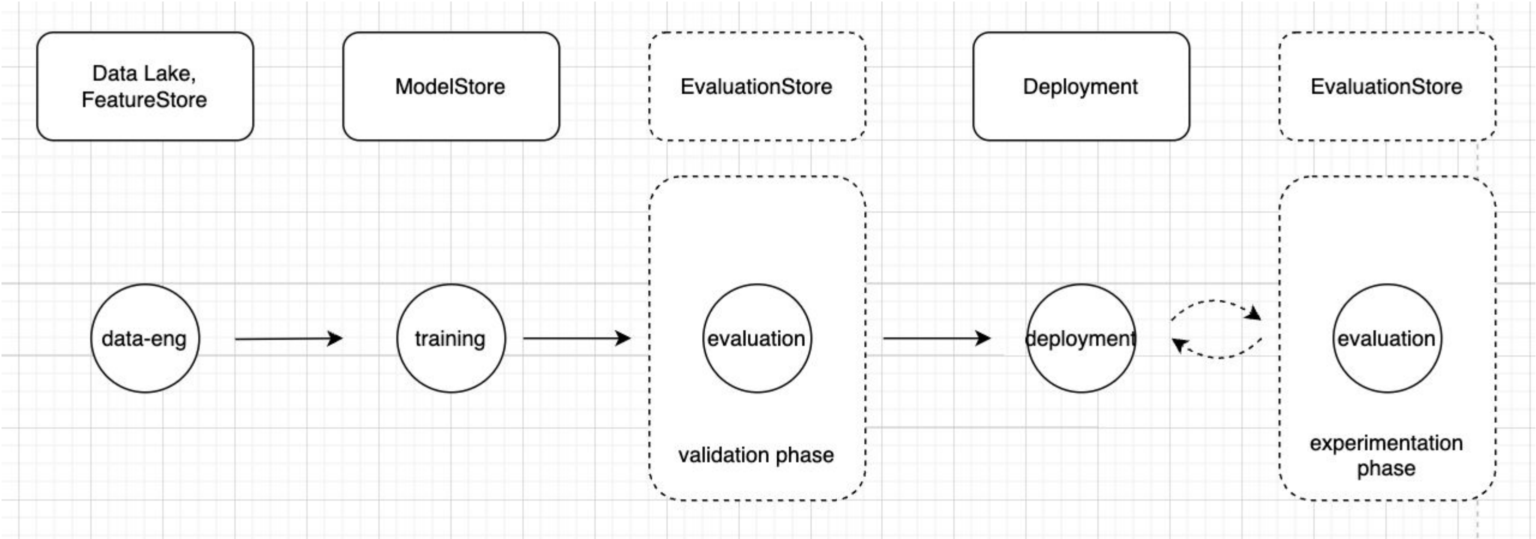

Testing and Deployment - An Interplay

- **Validation**, **Experimentation** and **Monitoring** via ML Testing

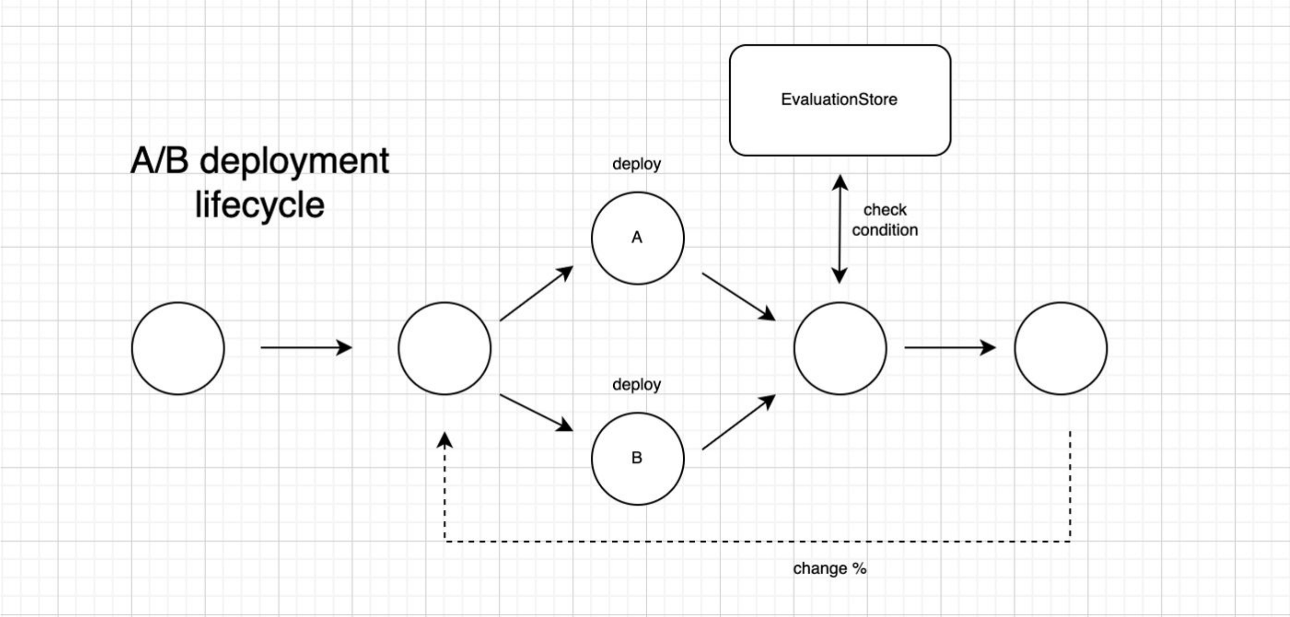

Testing and Deployment - An Interplay

- Many Deployment types - many *winning* versions (>> # models)
- a/b, multi-arm and contextual bandits, …
- Single model vs many models

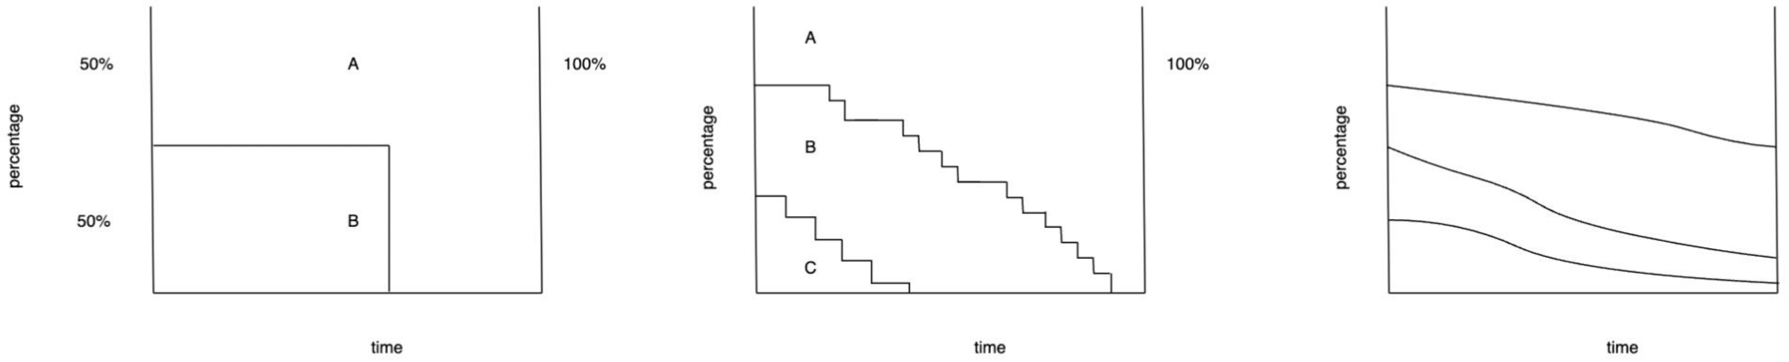

Testing and Deployment - An Interplay

- Live Experimentation Infrastructure
- Build vs Buy

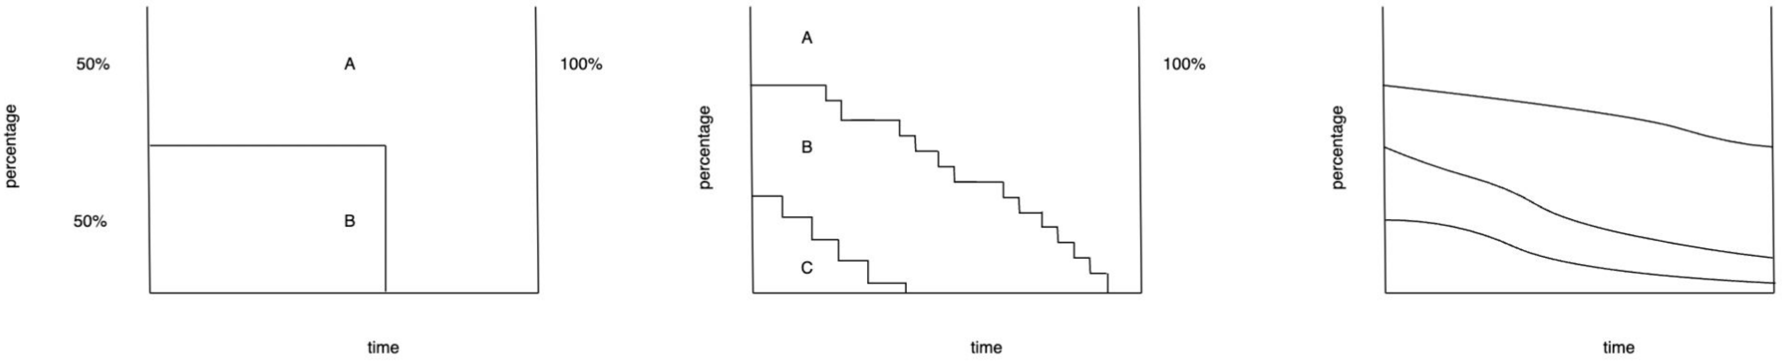

AI Org - Inception to Excellence

Infrastructure & Platform	Applications	Safety, Quality & Governance	Training & Talent Dev.	Research & Collaboration	Venture & Acquisitions

EURO AUSTRIA

Be Bold, Be Hungry, Be Fearless - Thank You

- Questions