Heuristic and Evolutionary Algorithms Laboratory

Adaptive Analytic Prescriptions: Integrating Models, Data, and Optimization in Dynamic Environments

4th Hagenberg Prescriptive Analytics Days 2025

Bernhard Werth,
Michael Heckmann, Johannes Karder, Philipp Fleck, Stefan Habringer, Philipp Neuhauser, Florian Holzinger, Stefan Wagner, Michael Affenzeller

Why Adaptive Prescriptions?

> ***No plan of operations extends with any certainty beyond the first encounter with the main enemy forces.***
> *Only the layman believes that in the course of a campaign he sees the consistent implementation of an original thought that has been considered in advance in every detail and retained to the end. -- Helmuth von Moltke 1871 [1]*

- Many types of prescriptions:
  - Schedules, machine assignments, packings, …
- Many types of surprises:
  - Employee late, machine break down, product fail at QA, bad weather, manual error, new orders from management, stuck in traffic jam, IT outage, …
- “Just make & communicate a new plan”

Open-Ended Adaptive Prescriptions

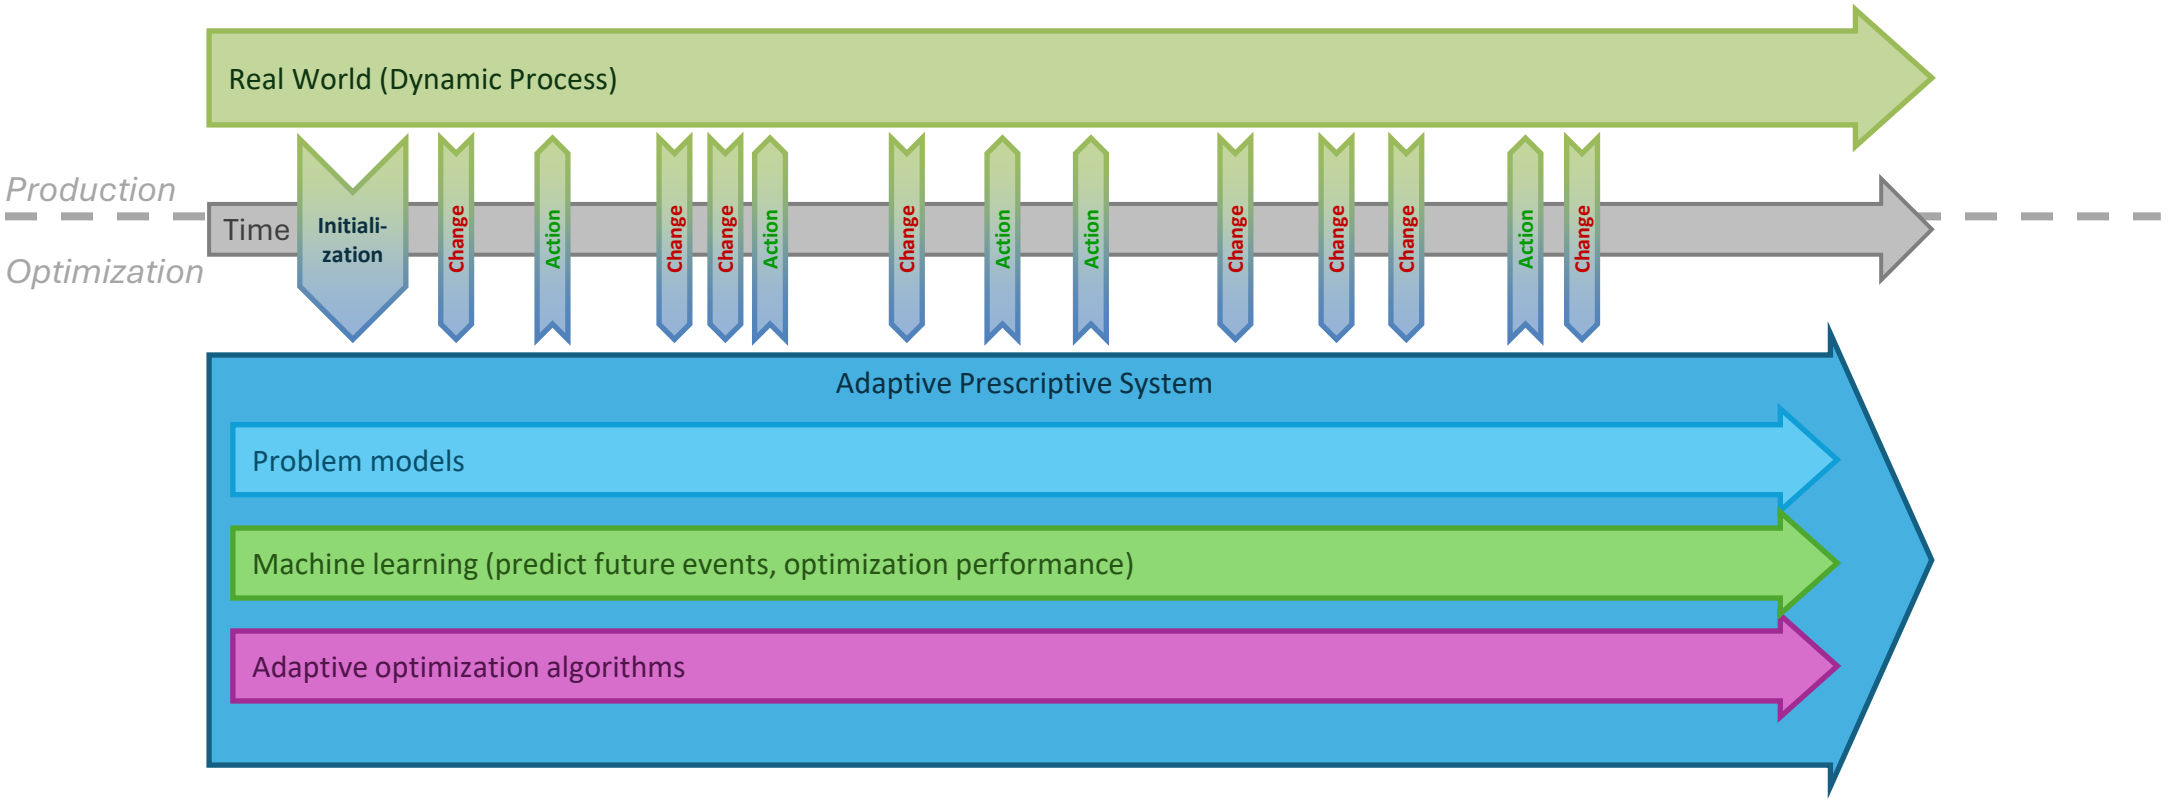

Where It Matters

- Complex environments with optimization potential
- Not easily solved via Reinforcement Learning or Exact Methods
- Uncertainties & dynamic influences
- Process is fast enough to prohibit repeated static solving

Map of Contents

(I get one 'funny' image per presentation and I WILL USE IT)

Data Mismanagement

Path to academic neverland

Company partner bringing new requirements in the last months of the project

Beware spAIders

System Overview

- Optimizer and Real World exchange solutions and data
- Optimizer decides on solutions by simulation
- Simulation queries ML-Models
- (Not drawn) Optimizer uses other ML-models for meta-purposes

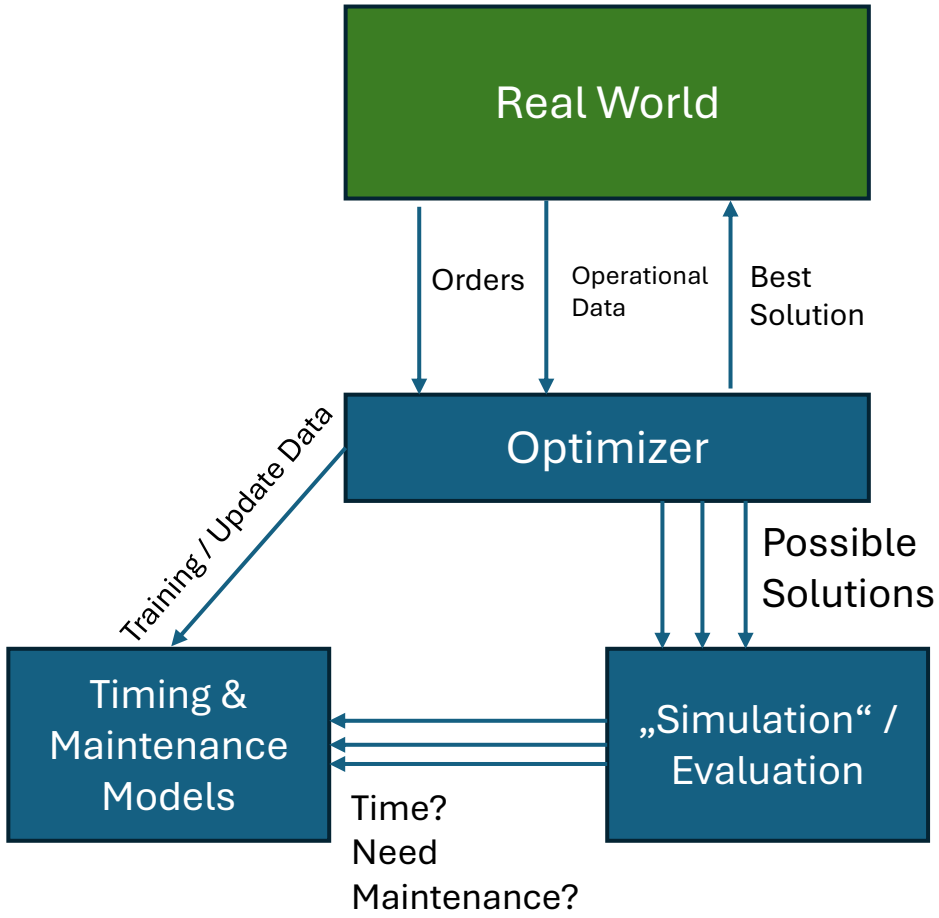

Real World Data Streams

- Machine Data
  - Status Updates
  - Sensor Signals (electrical, thermal, vibration,…)
- Operational Data
  - Orders
  - Cancellations
  - Remakes/Reworks
  - Restrictions
    (“not this place”, “freeze this part of the solution”)
- *Other Optimizers*

Kafka / MQTT Messaging

- Thou shalt NOT query the production database all the time
- Dedicated message brokers
- MQTT / Kafka
  - MQTT: IoT messaging, best effort, fire and forget
  - Kafka: Slightly more backend, broker persists messages
  - You will most likely meet both

Real-time architecture with Apache Kafka

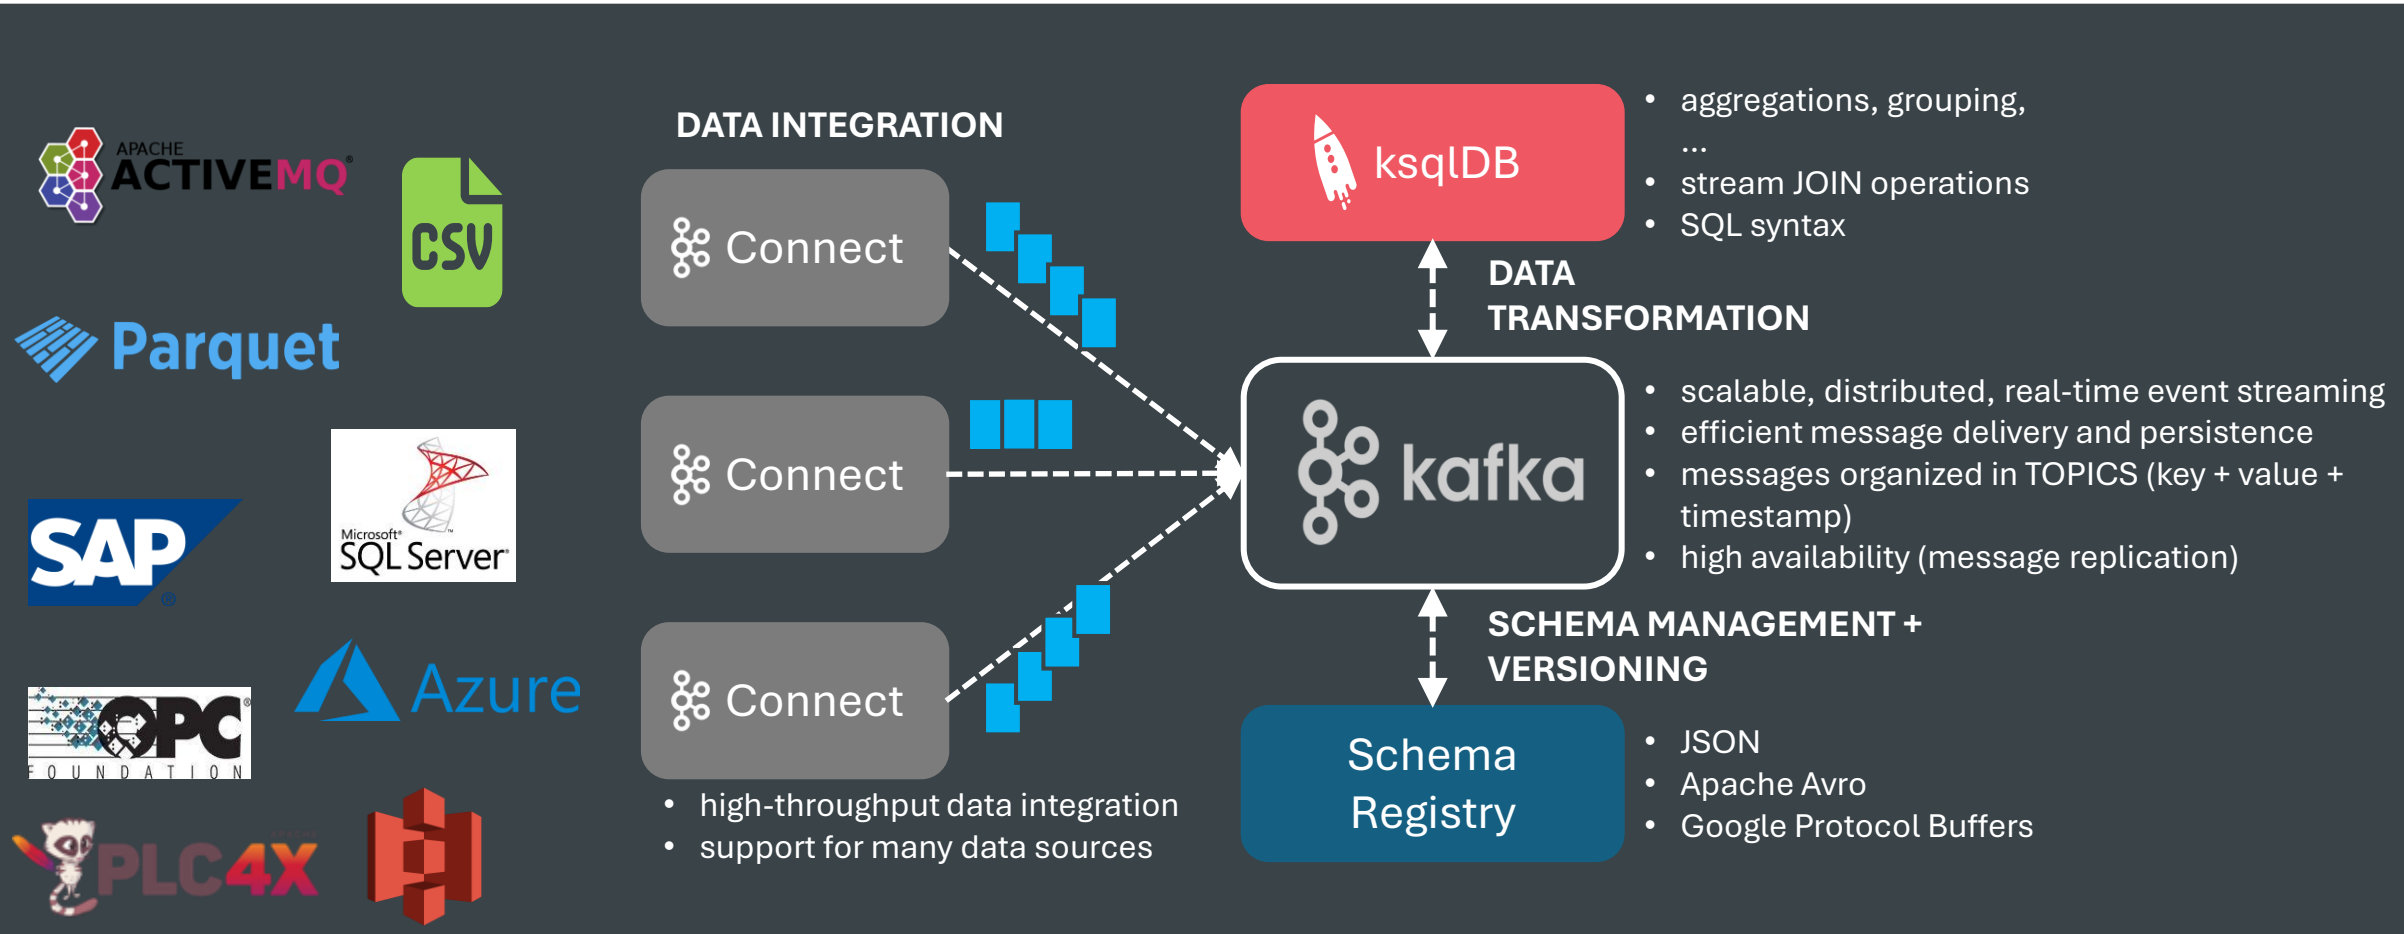

Why Evolutionary Algorithms?

- Exact methods too expensive
- Implicit & Explicit Memory
- Robust to encoding changes
  - Solutions get longer or shorter or change parts
  - Constraints change
- No pre-training

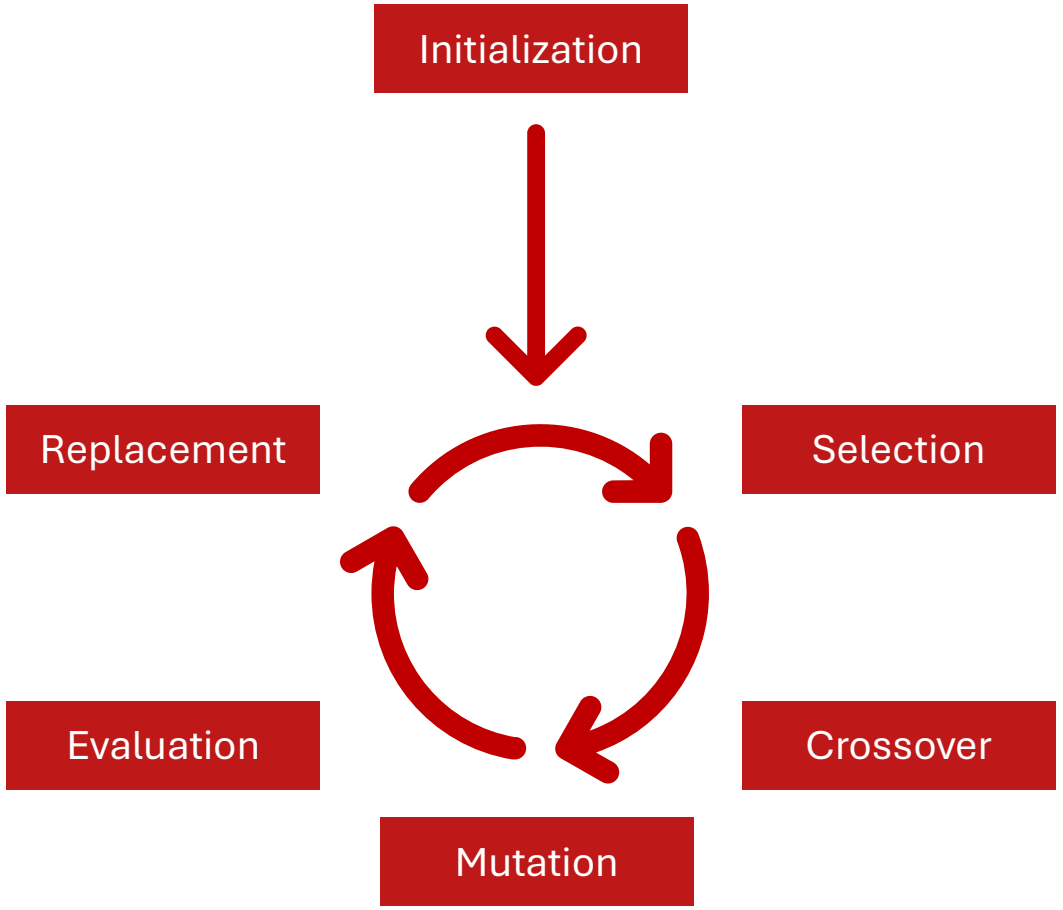

Open-Ended Evolutionary Algorithms

- Convergence:
  - Solutions get „more similar“ to each other
  - Usually by design
  - Helps with „fine-tuning“ of solutions
- Can brick your optimizer

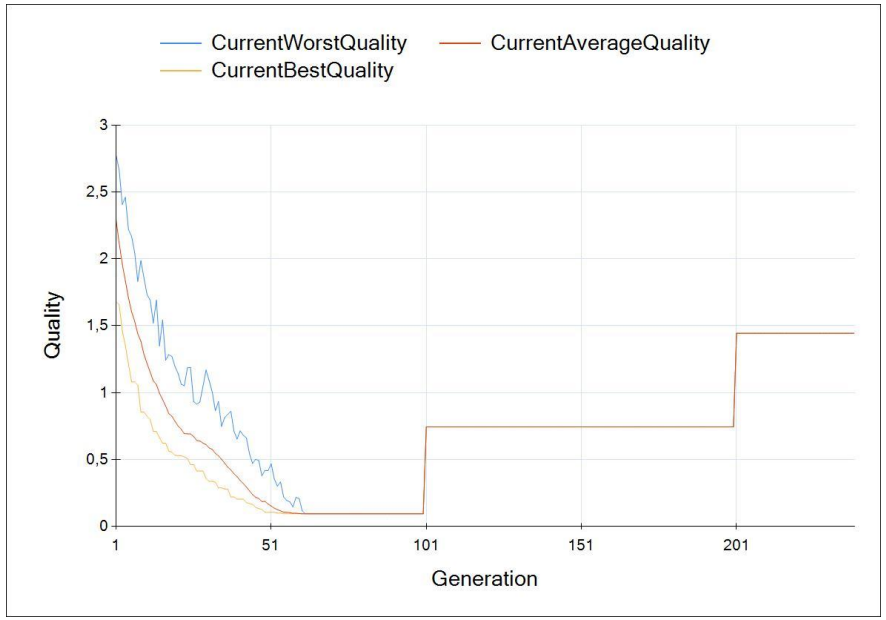

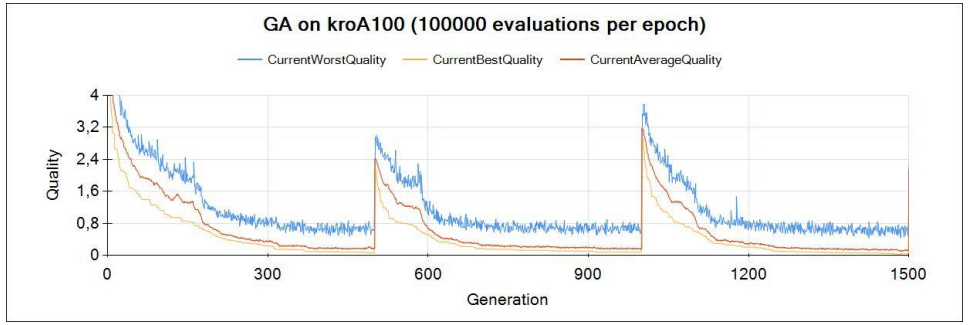

Open-Ended Genetic Algorithms

- Adds three new steps to evolutionary cycle

1. Synchronization
   - Applies update procedures
   - Applies repair procedures
   - Communicates best solution
2. Self-Adaptation
   - Switches between algorithm modes
   - Adjusts parameters
3. Population Reseeding
   - No relevant offspring found

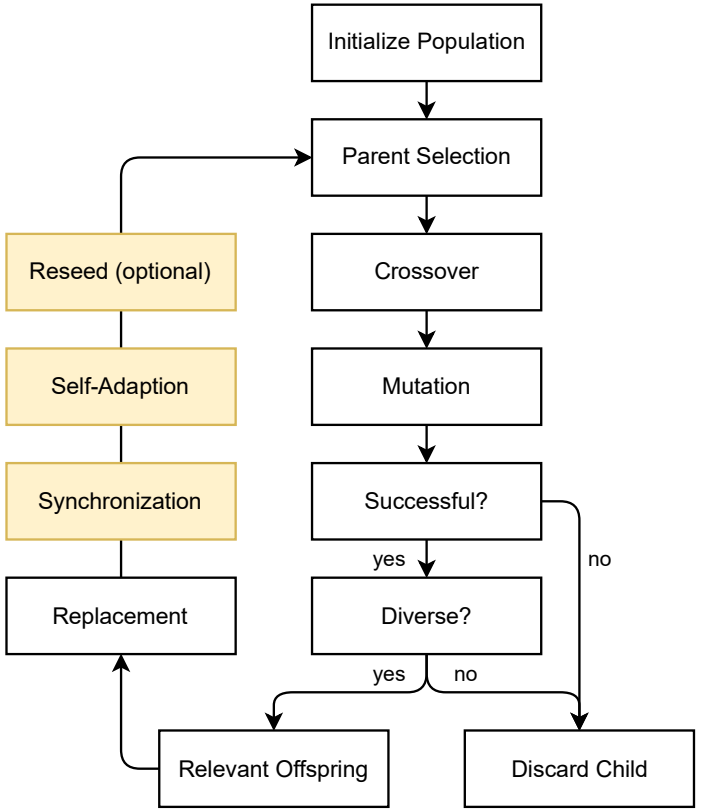

Open-Ended Relevant Allele Preserving Genetic Algorithm

- Robust
- Self-restarting
- Faster as Convergence approaches
- Longest continuous run: 2 months

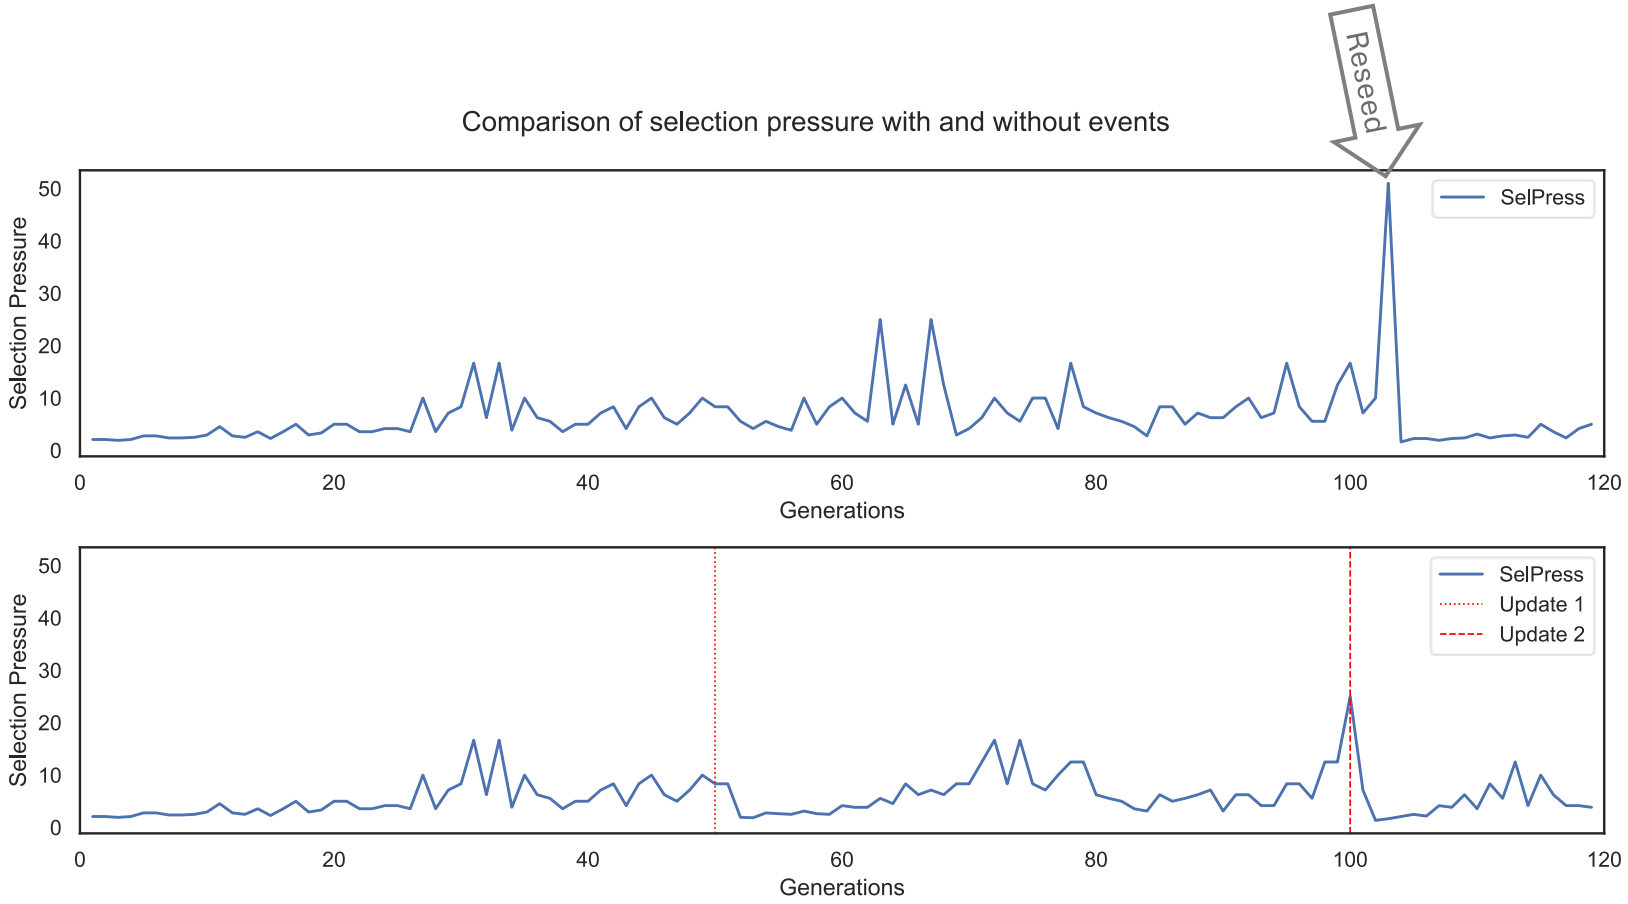

Role of Machine Learning in Adaptive Prescriptions

- Algorithm performance prediction
  - How good is the solution quality going to become?
  - Trade off: Time vs Quality
- Simulation features (very often timings)
  - How good is this solution?
- Surrogate models
  - How good is this solution? (let's skip the simulation)
- Solution generators
  - Oh, great and eldritch gods of reinforcement learning, what plan do I need to follow?
- Problem development
  - Consequences of decisions compound
  - Problem may become unmanageable
  - System needs to escalate to the user

Online Machine Learning (River)

- Special Requirements
  - Adding / forgetting data points
  - Constraints on model behavior (similarity between versions)
  - Prediction speed (non-batch)
  - Bounded memory footprint
  - Retraining/update speed (Many models scale poorly with large training datasets)
  - Cold start: pre-training is not always possible/sensible
  - Robust, robust, robust

- Practical Information:
  - Very little software available
  - https://github.com/online-ml/river
  - Wide array of models / preprocessors
  - (not affiliated / this is not an ad)

Training Data

Challenge #1: Time Series

Challenge #2: Dynamic Data

Challenge #3: Dynamic Schema

ID	Mat A	Mat B	Mat C	...	Temperature	Pressure	Electric Voltage	...	Quality
1	10.3	22.0			[20, 25, 30, ..., 25]	[1.5, 1.4, 1.8, ..., 1.2]			**5.0**
2	12.0	23.0			[18, 22, 23, ..., 28]	[1.6, 1.3, 1.7, ..., 1.2]			**5.2**
3	11.5	22.4			[20, 21, 28, ..., 30]	[1.6, 1.6, 1.7, ..., 1.3]			**5.3**
4	11.3	23.1			[20, 22, 25, ..., 31]	[1.4, 1.5, 1.4, ..., 1.4]			**5.3**
5	12.1	22.6			[19, 22, 23, ..., 27]	[1.7, 1.3, 1.6, ..., 1.3]			**5.4**
6	10.7	-	14.9		[20, 23, 24, ..., 29]	[1.4, 1.6, 1.6, ..., 1.7]	[2100, 2800, ..., 4200]		**5.2**
7	11.7	-	14.5		[20, 18, 22, ..., 24]	[1.3, 1.7, 1.8, ..., 1.5]	[2200, 2900, ..., 3900]		**5.5**
n	...	-	...		...	...	...		**...**

Expression Trees and Genetic Operators

Representation

$$c = \sqrt{a^2 + b^2}$$

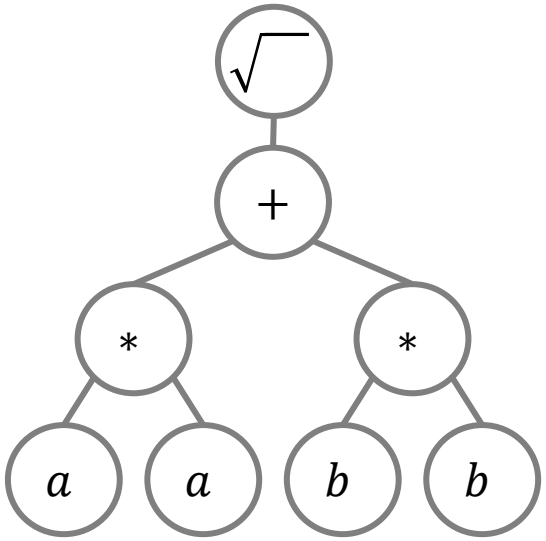

Crossover

$$c = \sqrt{a^2 - 1/b}$$

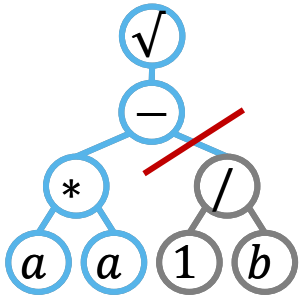

$$c = b^2/(a+5)$$

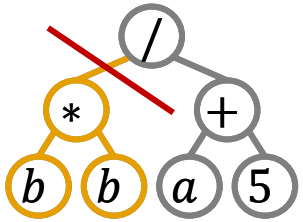

$$c = \sqrt{a^2 - b^2}$$

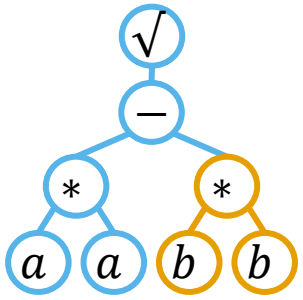

Mutation

$$c = \sqrt{a^2 - b^2}$$

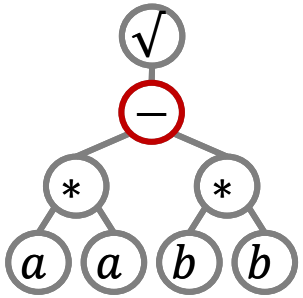

$$c = \sqrt{a^2 + b^2}$$

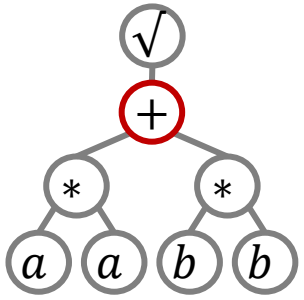

Symbolic Regression on Time Series

- Special extensions of “classic” equation finding
- Vectorial operators in the grammar

$$Quality = \text{avg}\left(\frac{Temp.}{Press.}\right) + Mat\ A$$

Aggregation-function

Component-wise operator

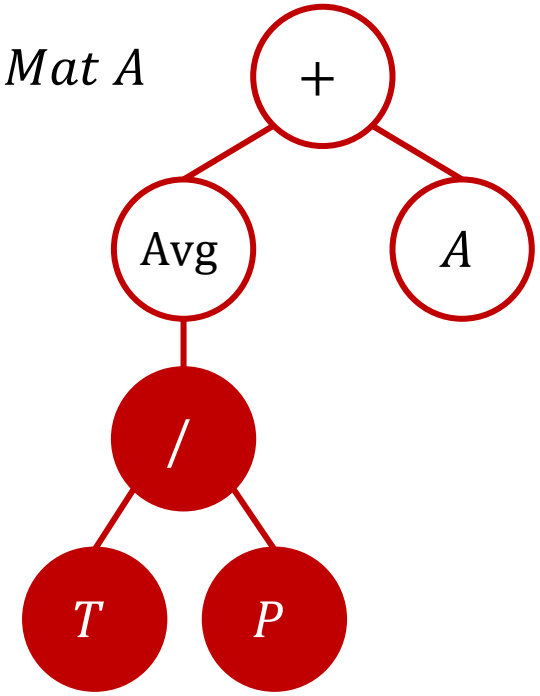

Why Symbolic ML Helps

- Can be trained/updated via dynamic evolutionary algorithm
- Resistant to schema changes
- Fast single predictions (with dynamic compilation)
- “Verifiable“
- Accommodates all kinds of shape and output constraints
  - Must be quadratic with x
  - Must be at/near zero for large y

Practical Considerations

- Subject of the model often stateful
- *What will happen if I do A and then B and then C….?*

- Will be queried a lot
- And in parallel

- Target
  - Being accurate is good
  - Guiding the optimizer is better

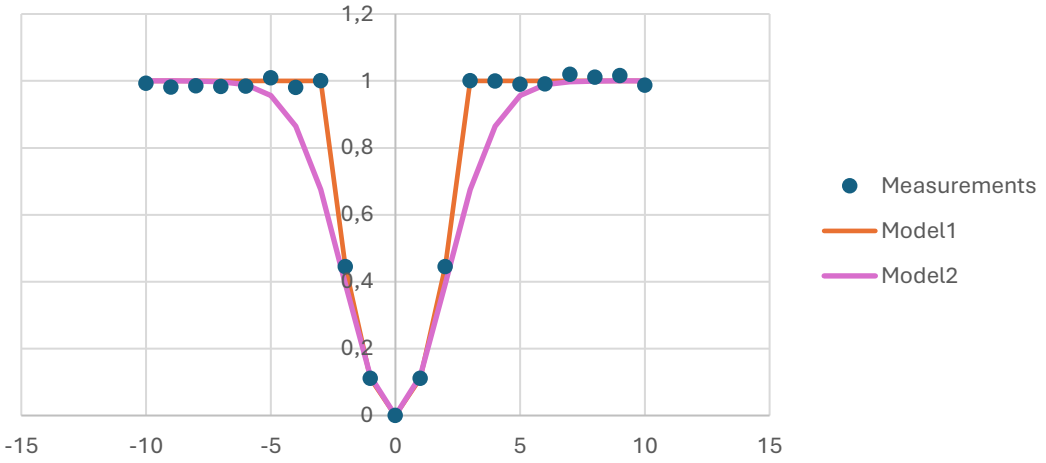

Machine Learning – Optimization Feedback Loop

- Models learn from the Real World*
- Models guide the Optimizer
- The Optimizer influences the Real World

- Training data is coupled to previous model iterations
  - Heavily skewed
  - Self-approving

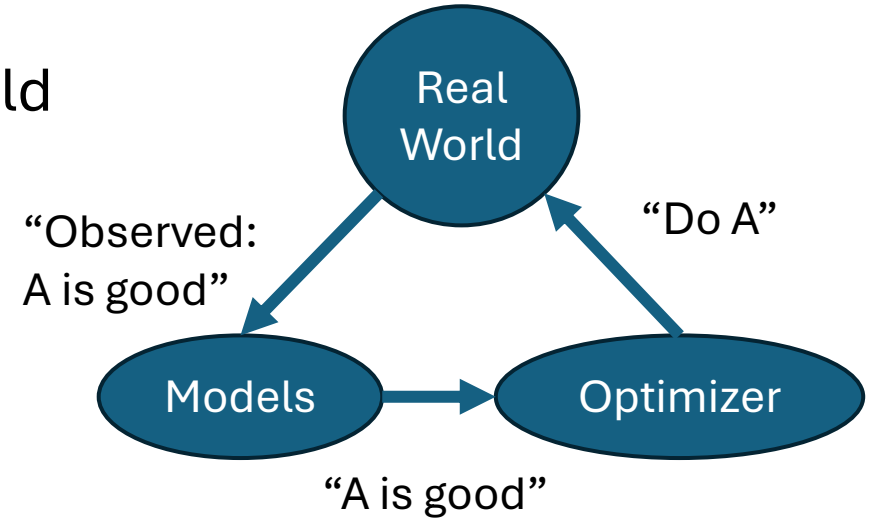

"B would have been better"

Case Studies

- Industrial applications
- Non-fungible material
- Success factors:
  - “Industrial Digitalization“ & data collection
  - Many many meetings…
- Often very simple ML-models
- As in every real-world case:
  - Constraints galore
  - Objectives as well

Coil Warehouse

- Crane scheduling
- Main difficulty: Other cranes are in the way

- Open-Ended Genetic Algorithm

Hot-Storage

- Stacking, Scheduling & Delivery Lot Building
- Main difficulty: Stack Management

- Tree-search
  - Fast enough
  - Repeated static optimization

Glass Production

- Scheduling and Assignment
- Main difficulty:
  - Buffer limits
  - Many unknowns
- Open-ended Genetic Algorithm

Battery Production

Lead Plates

Paletts

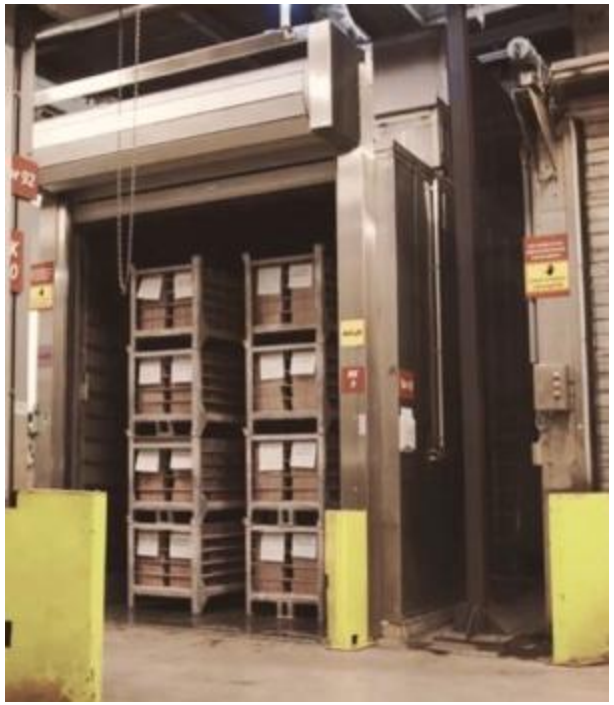

Curing Chamber

- Material Batching
- Main difficulty: Utilization
- Dynamic Tree Search

5th Dynstack competition 2026

https://dynstack.adaptop.at/

- Current ranking (– ish):
  1. Tree-search + custom heuristics
  2. Hand-written rules
  3. General tree-searches
  4. Reinforcement learning (no working entry…)

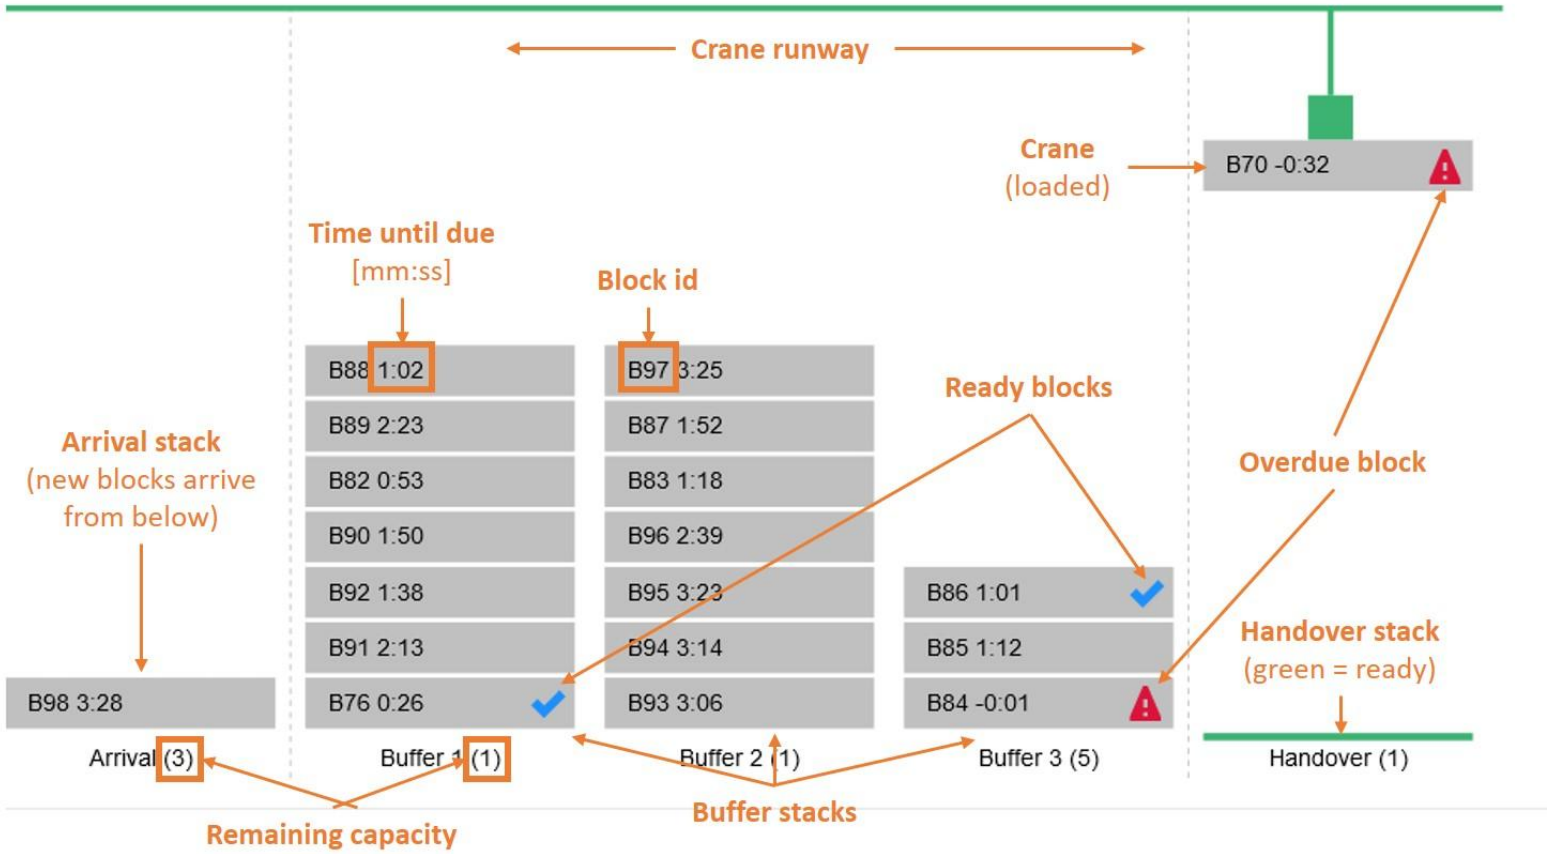

Conclusion

- Dynamic Prescriptive Analytics is messy but powerful
  - Many simple sub-problems *(good for practitioner, bad for publications)*
  - Requires a lot of live data about the system *(bad for practitioner, trivial for publications)*
  - Needs to be effective on longer time scales *(very bad for publications)*
  - Needs to deal with all kinds of distribution and asynchrony *(messy all around)*
- Many different facets improved by ML/AI
  - Many simulations => Training data to feed the models
  - Different model uses:
    (Auto-)Regression, Clustering, Detection?
    Dynamic Prescriptive Analytics probably has a place for you

References

- [1] 1900, Moltkes Militärische Werke: II. Die Thätigkeit als Chef des Generalstabes der Armee im Frieden. (Moltke's Military Works: II. Activity as Chief of the Army General Staff in Peacetime) Zweiter Theil (Second Part), Aufsatz vom Jahre 1871 Ueber Strategie (Article from 1871 on strategy), Start Page 287, Quote Page 291, Publisher: Ernst Siegfried Mittler und Sohn, Berlin, Germany.