Improving Trust and Efficiency via Human-Centered AutoML

Marius Lindauer @
Summer School for Responsible AI PhD Programm

based on a lecture series at the European Summer School on AI ‘23 with Katharina Eggensperger

Machine learning is this …

“Machine learning is the science of getting computers to act without being explicitly programmed.”

by Andrew Ng
(probably inspired by Arthur Samuels)

… and also this

source: XKDC

Design Decisions

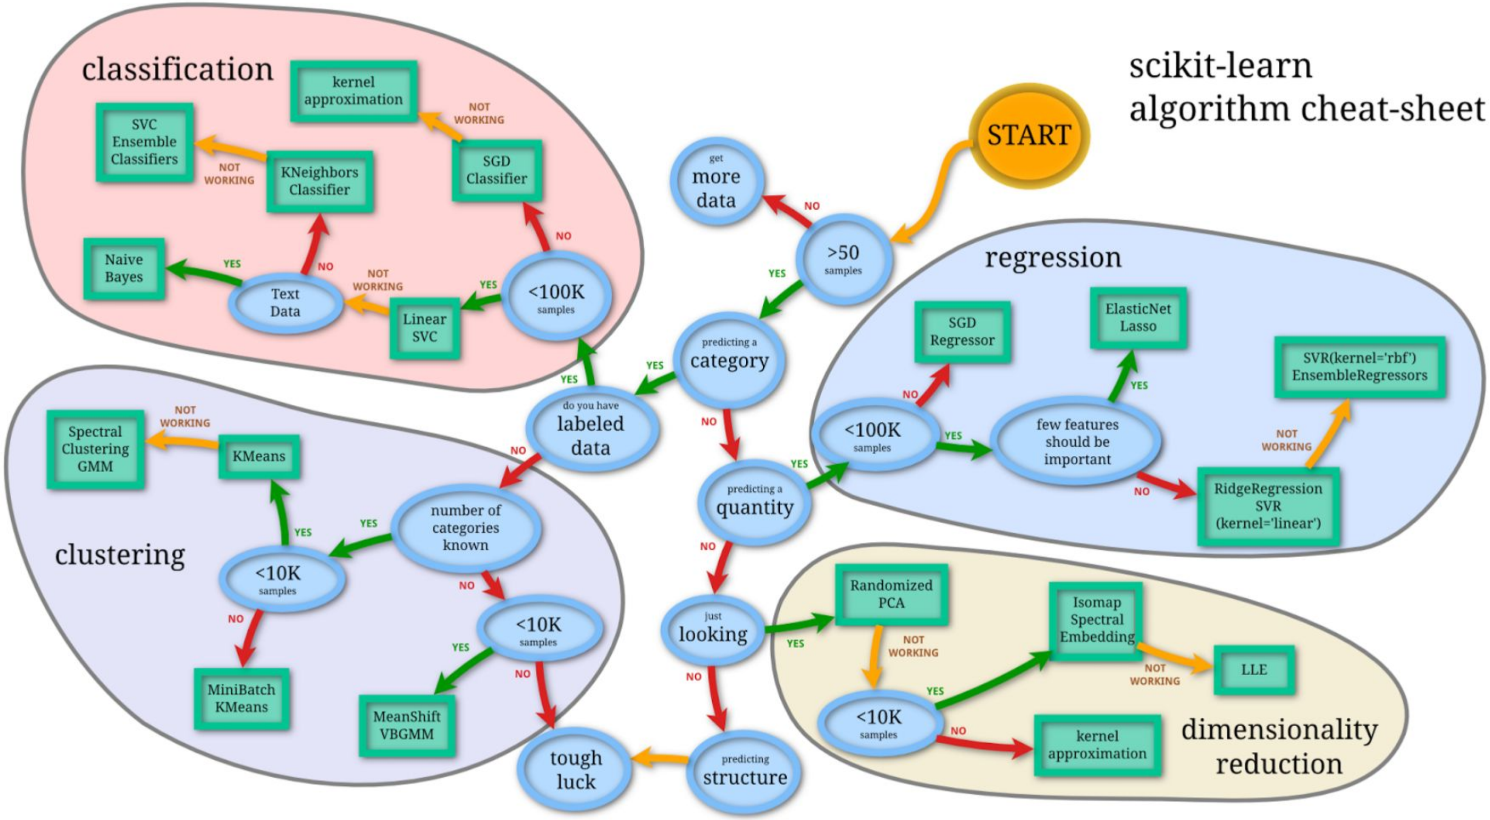

source: https://scikit-learn.org/stable/tutorial/machine_learning_map/index.html

Challenges in Applying AI/ML these days

Required expertise in ML and AI

Long development time for new AI applications

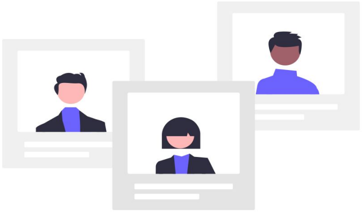

Few experts are available on the job market

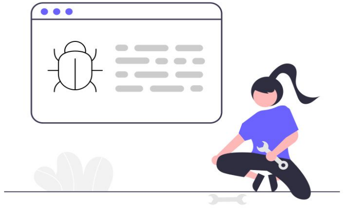

Unstructured and error-prone development of AI application

Why does ML development take a lot of time?

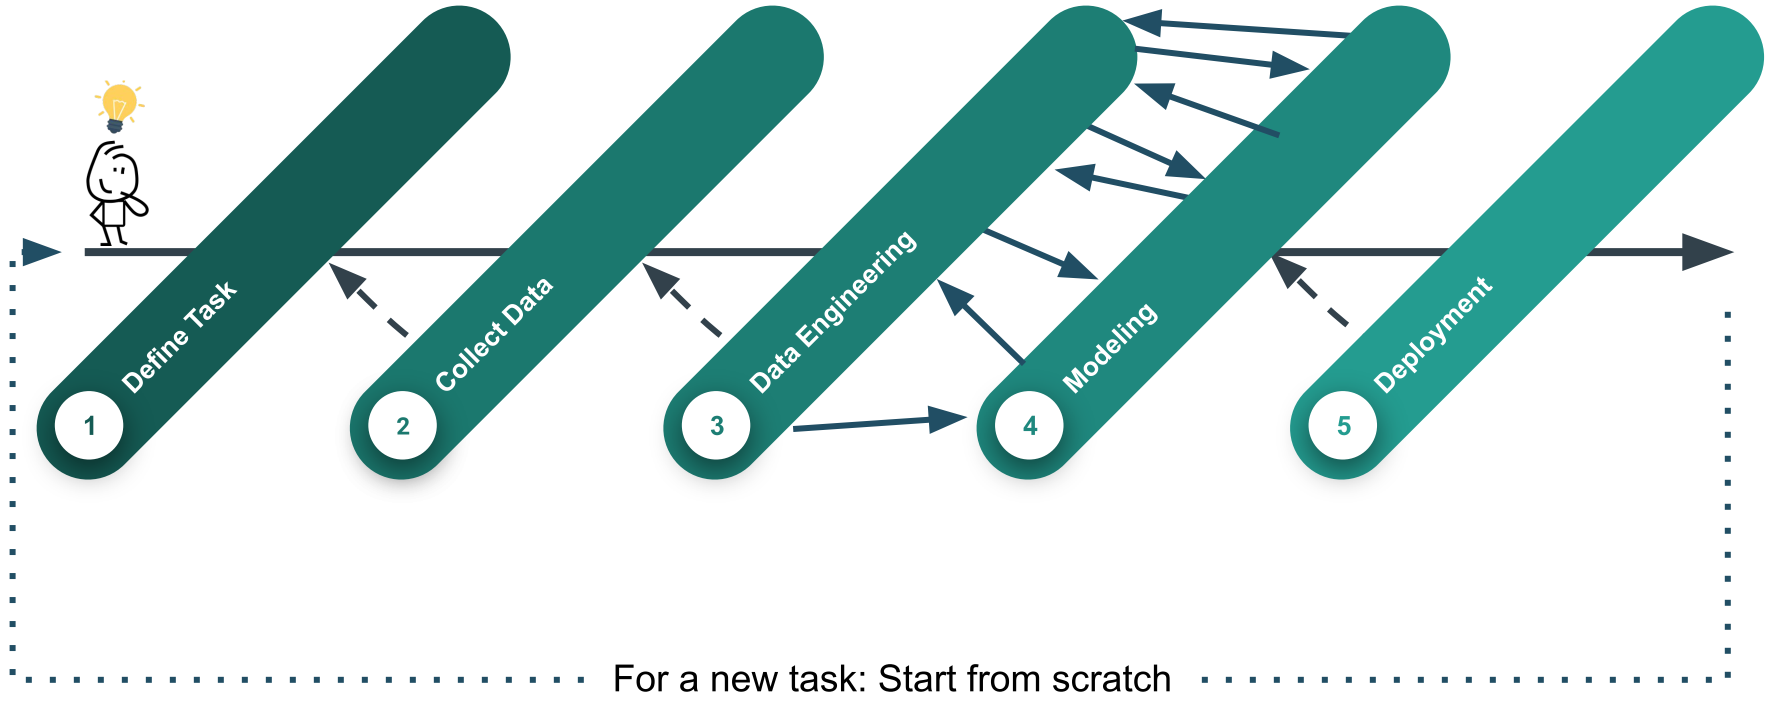

Toy Example: kNN

- *k*-nearest neighbors (kNN) is one of the **simplest ML algorithms**
- Size of neighbourhood (*k*) **is very important for its performance**
- The performance function depending on *k* is **quite complex** (not at all convex)

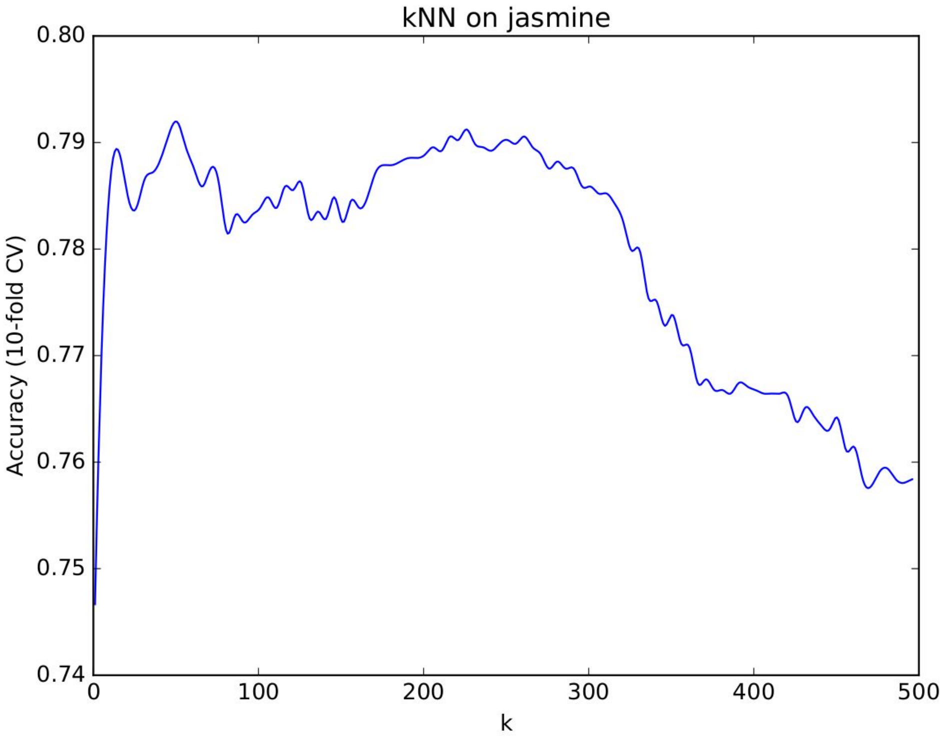

Goals of AutoML

Goal: Progressively automate all parts of machine learning (as needed) to support users efficiently building their ML-applications.

Informal Definition: AutoML System

Given

- A **dataset**,
- a **task** (e.g. supervised classification),
- a **cost metric** (e.g., accuracy or RMSE),
- (optional) a **budget**

an AutoML System automatically determines the approach that performs best for this application.

ML vs AutoML

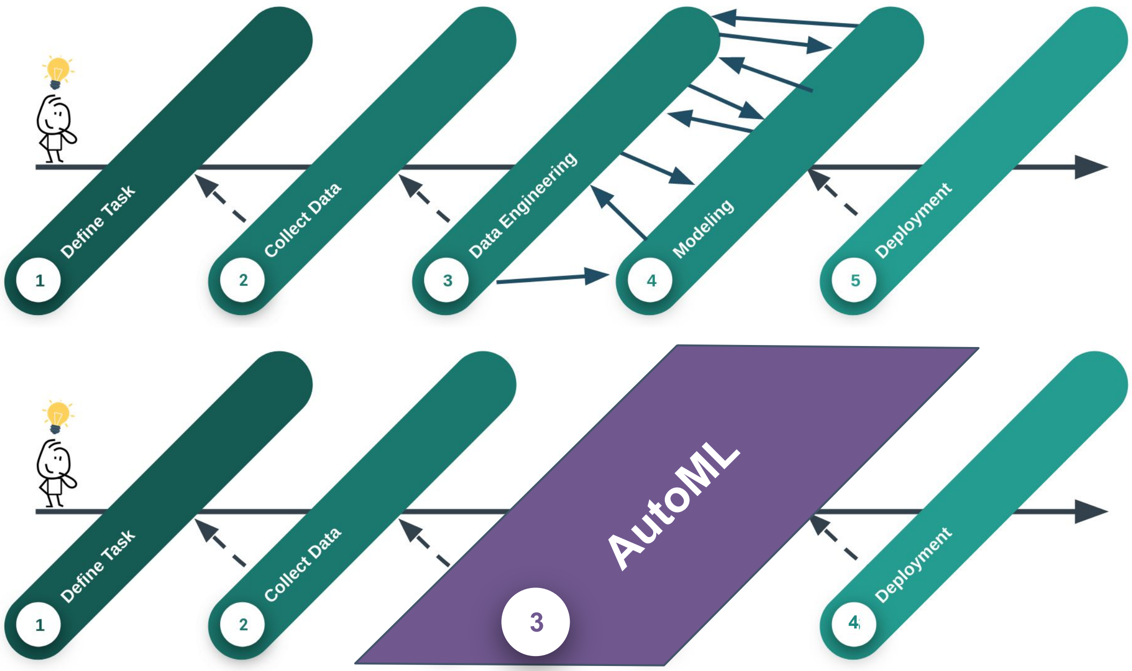

Define Task
1
Collect Data
2
Data Engineering
3
Modeling
4
Deployment
5
Define Task
1
Collect Data
2
AutoML
3
Deployment
4

Motivating Example: Shape Error Prediction in Milling Process

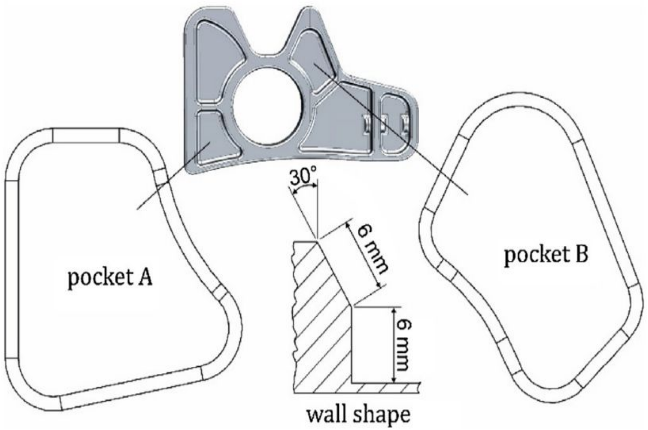

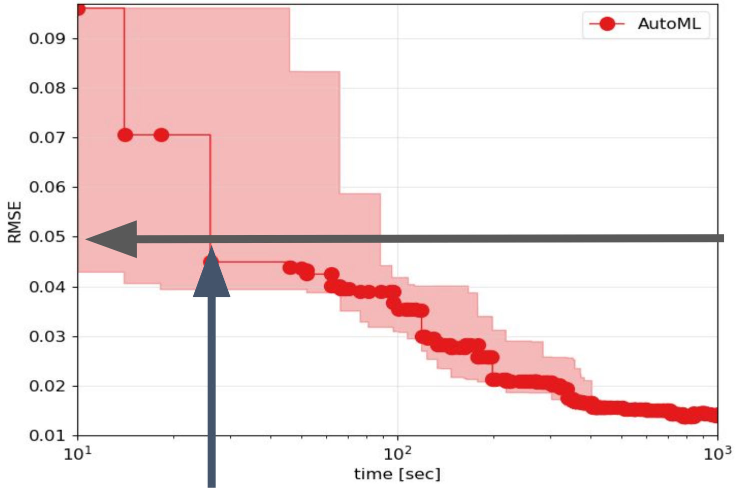

State of the art by human domain expert

Outperforming human domain expert after ~30sec
(+ some time to write a parser for the data)

[Denkena et al. 2020]

Advantages

AutoML enables

- 🚀 More **efficient** research (and development of ML applications)
  - → AutoML has been shown to outperform humans on subproblems
- 🧮 More **systematic** research (and development of ML applications)
  - → no (human) bias or unsystematic evaluation
- 📋 More **reproducible** research
  - → since it is systematic!
- 🦾 **Broader use** of ML methods
  - → less required ML expert knowledge
  - → not only limited to computer scientists

Challenges

But, it is not that easy, because

- 🔁 Each dataset potentially requires **different optimal ML-designs**
  - → Design decisions have to be made for each dataset again
- ⌛ Training of a single ML model can be **quite expensive**
  - →We can not try many configurations
- ❓ Mathematical **relation** between design and performance is (often) **unknown**
  - → Gradient-based optimization not easily possible
- 🎛️ 🎚️ Optimization in **highly complex spaces**
  - → including categorical, continuous and conditional dependencies

How everything is connected

LUH|AI Leibniz Universität Hannover

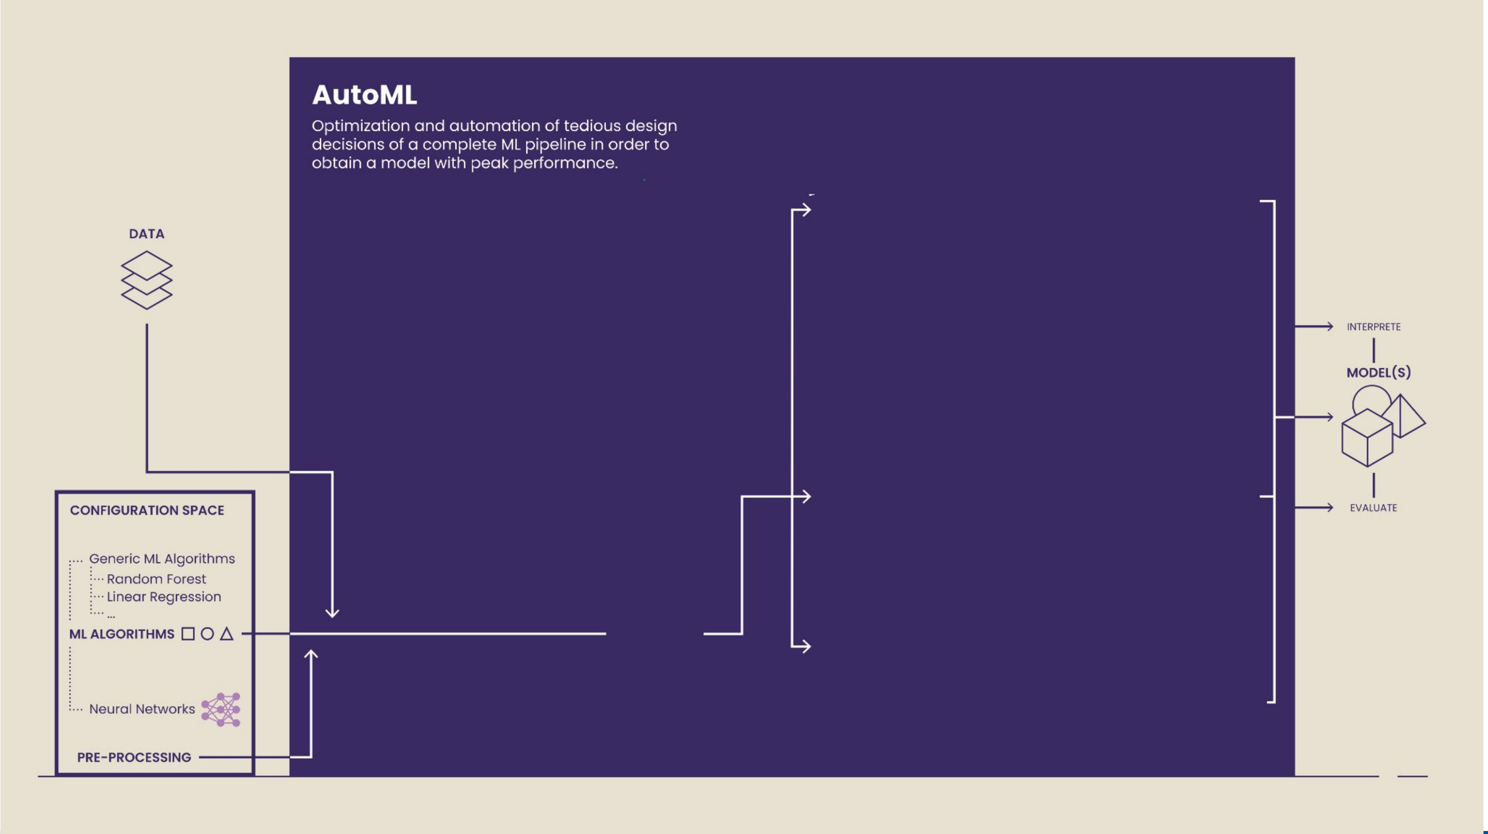

How everything is connected

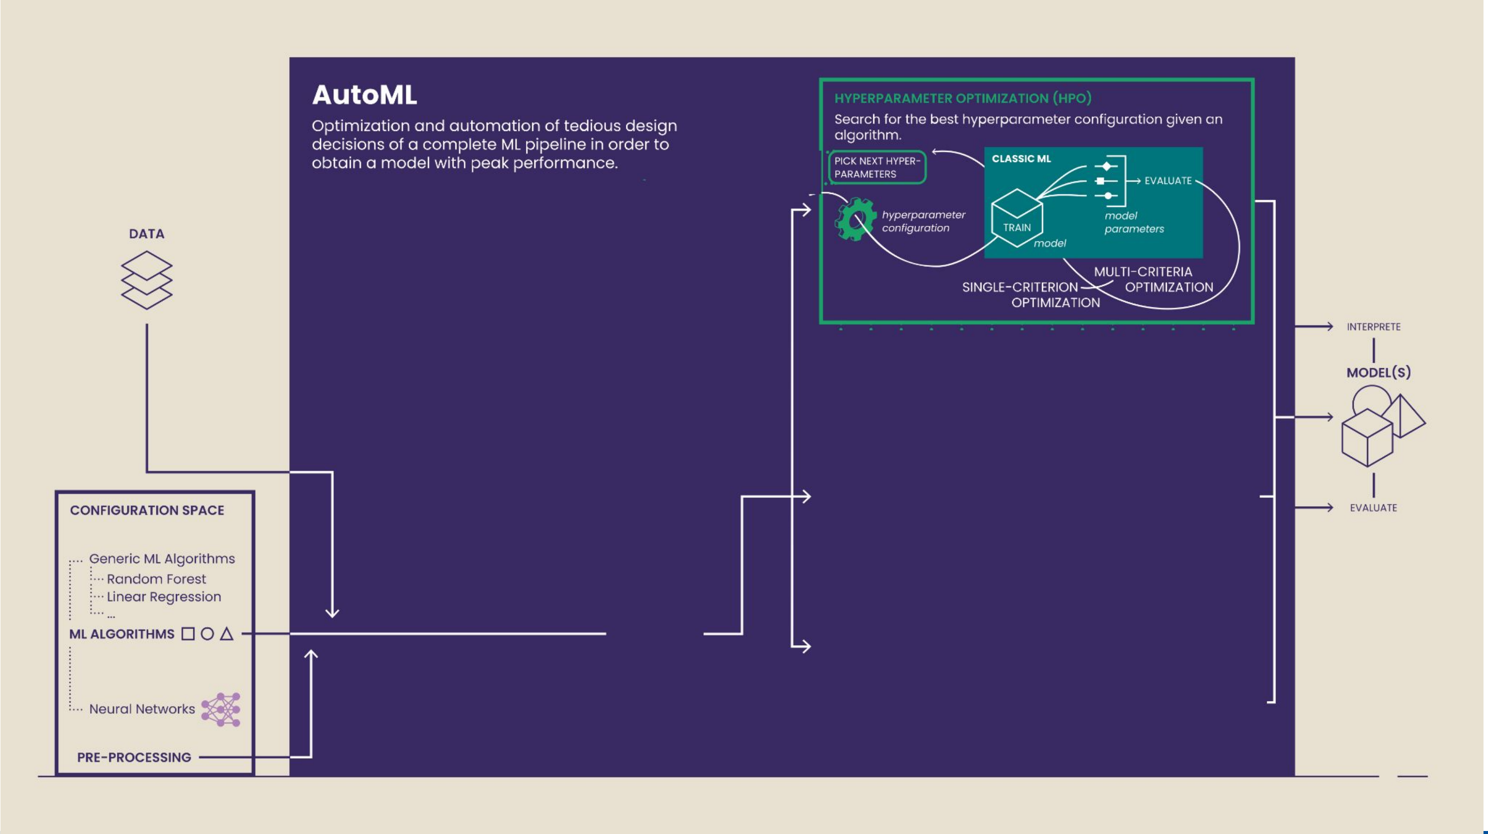

How everything is connected

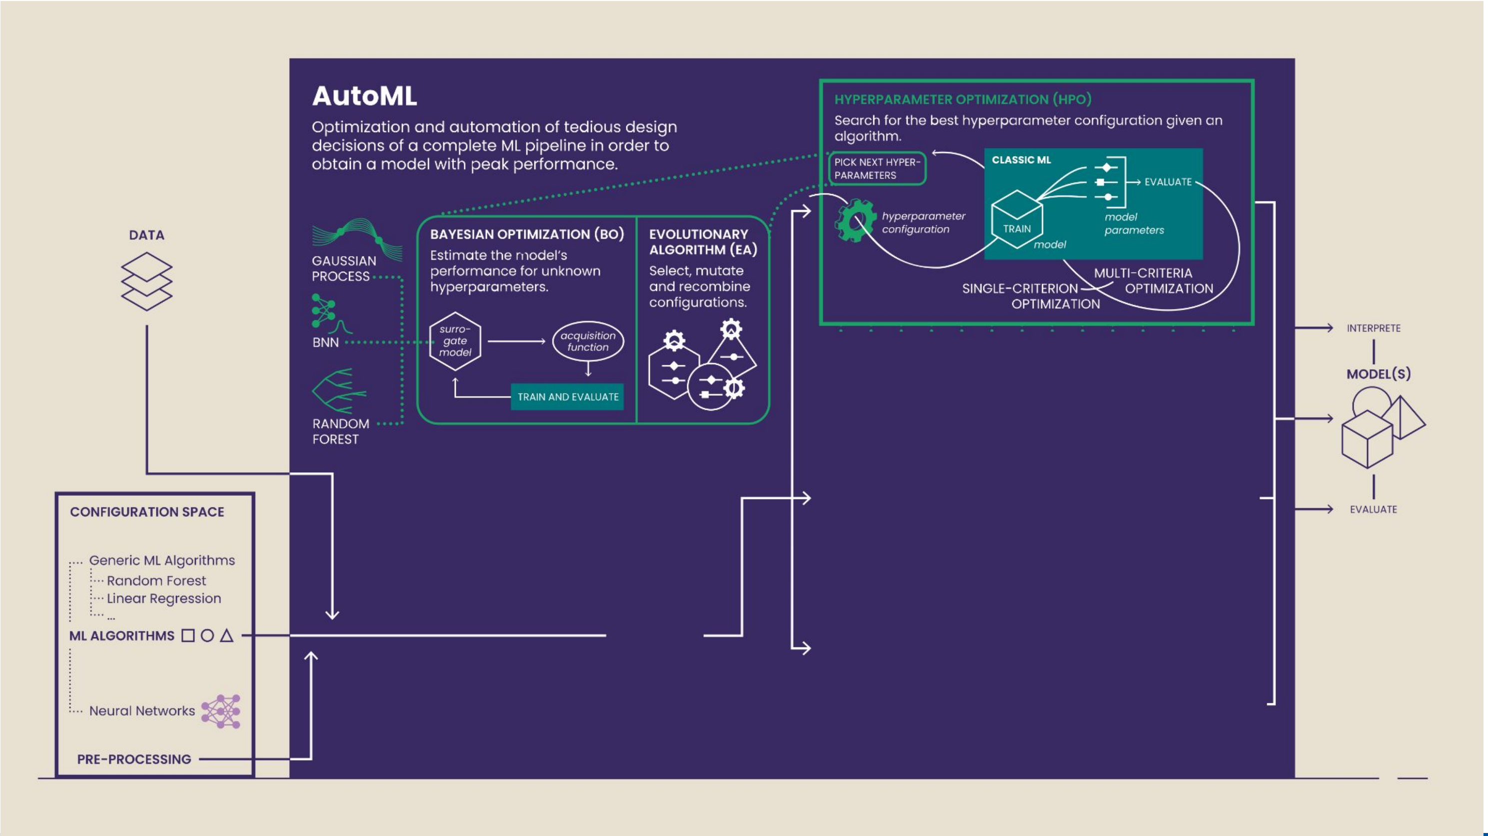

How everything is connected

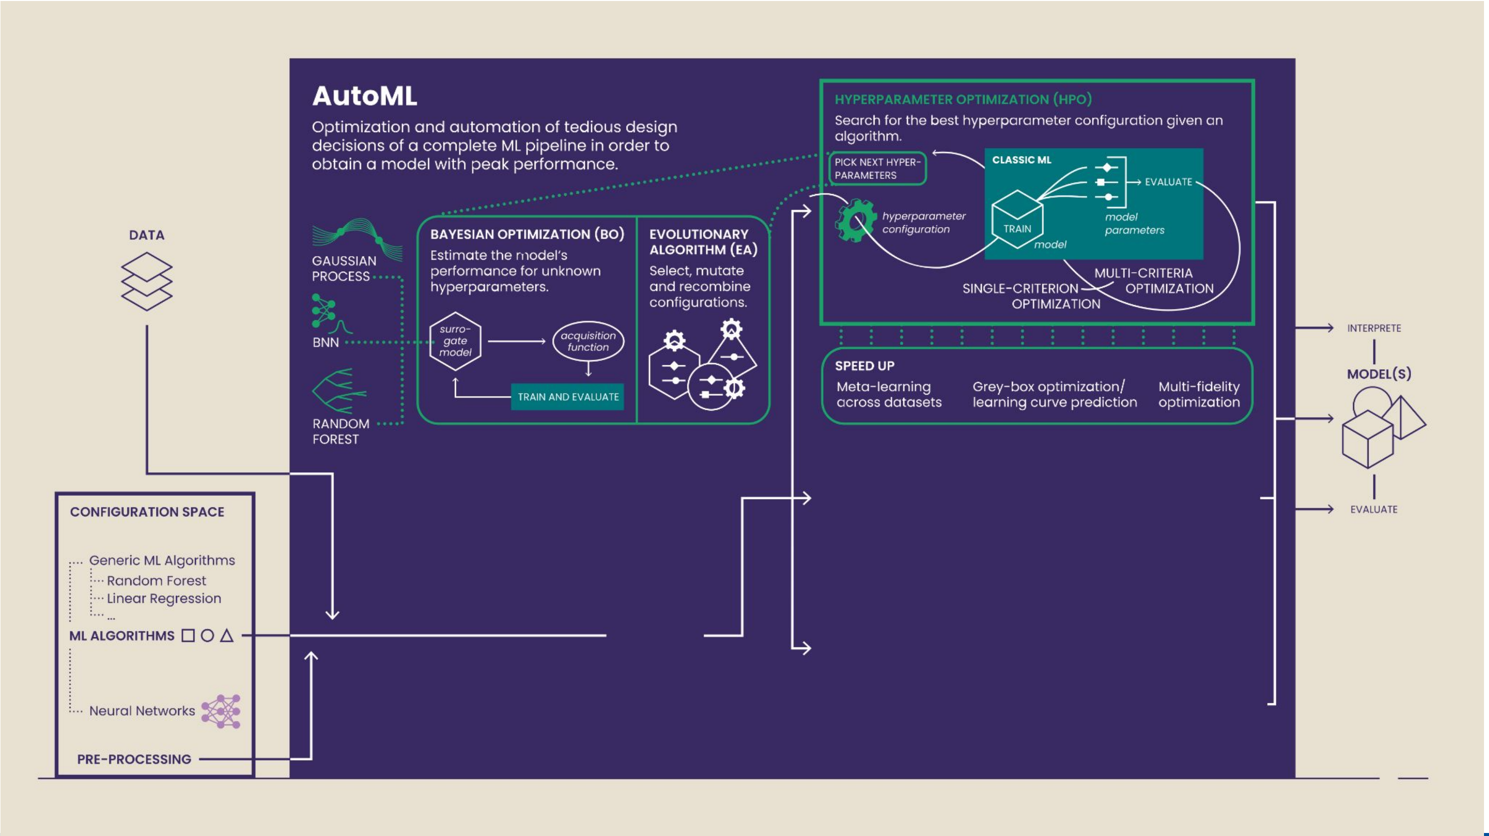

How everything is connected

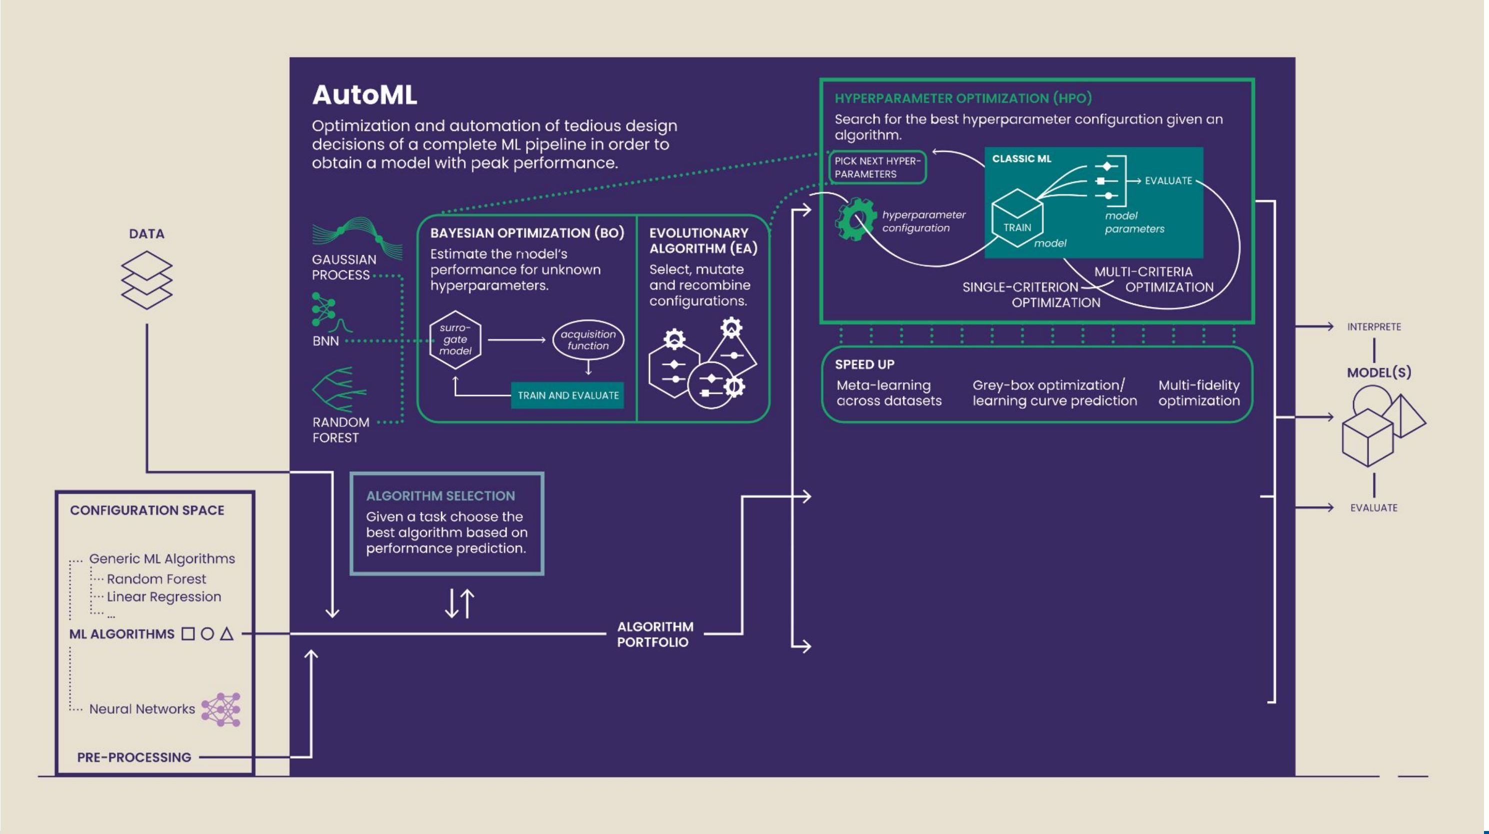

How everything is connected

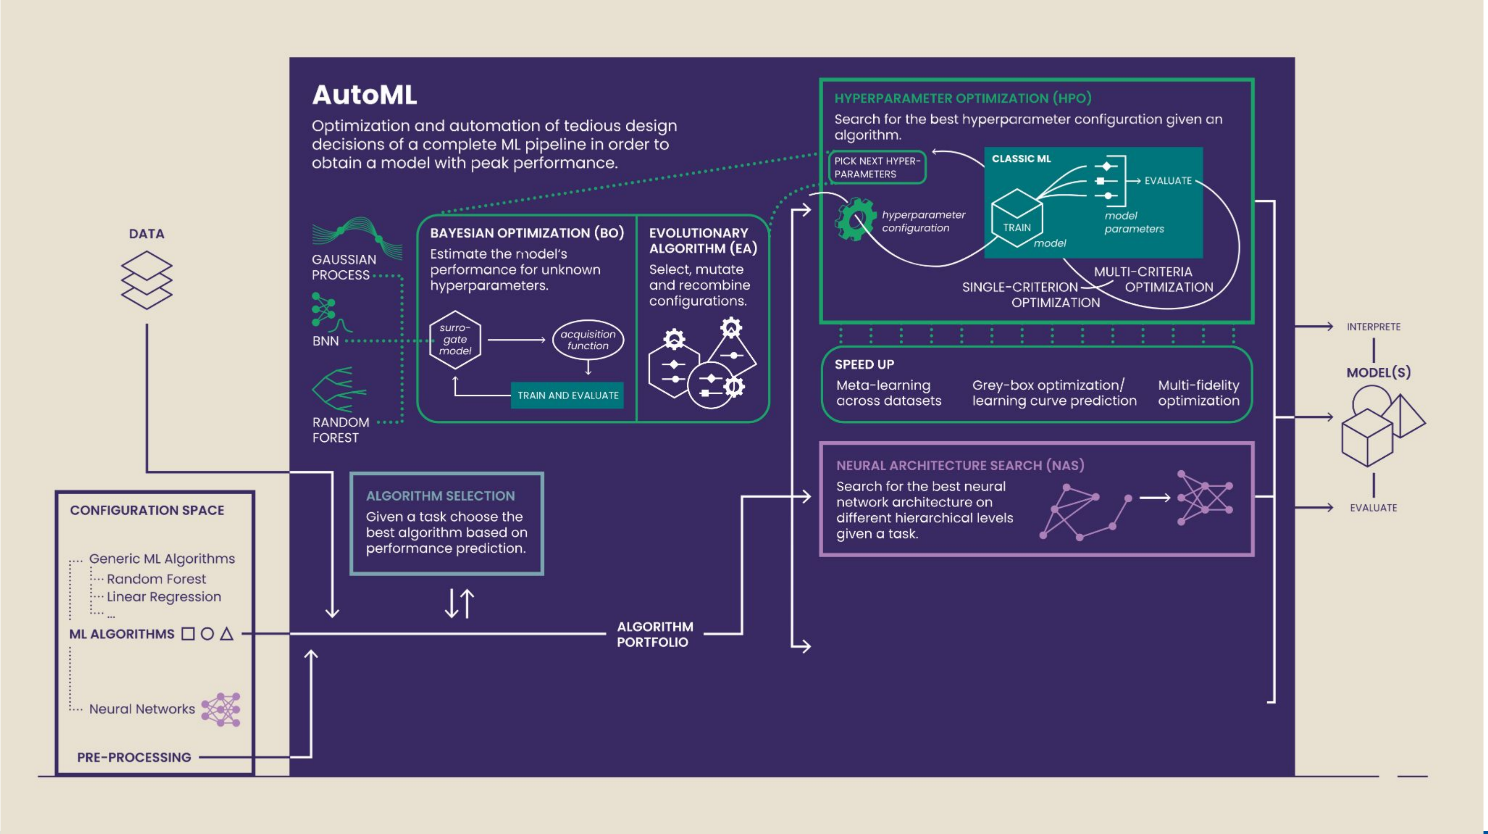

How everything is connected

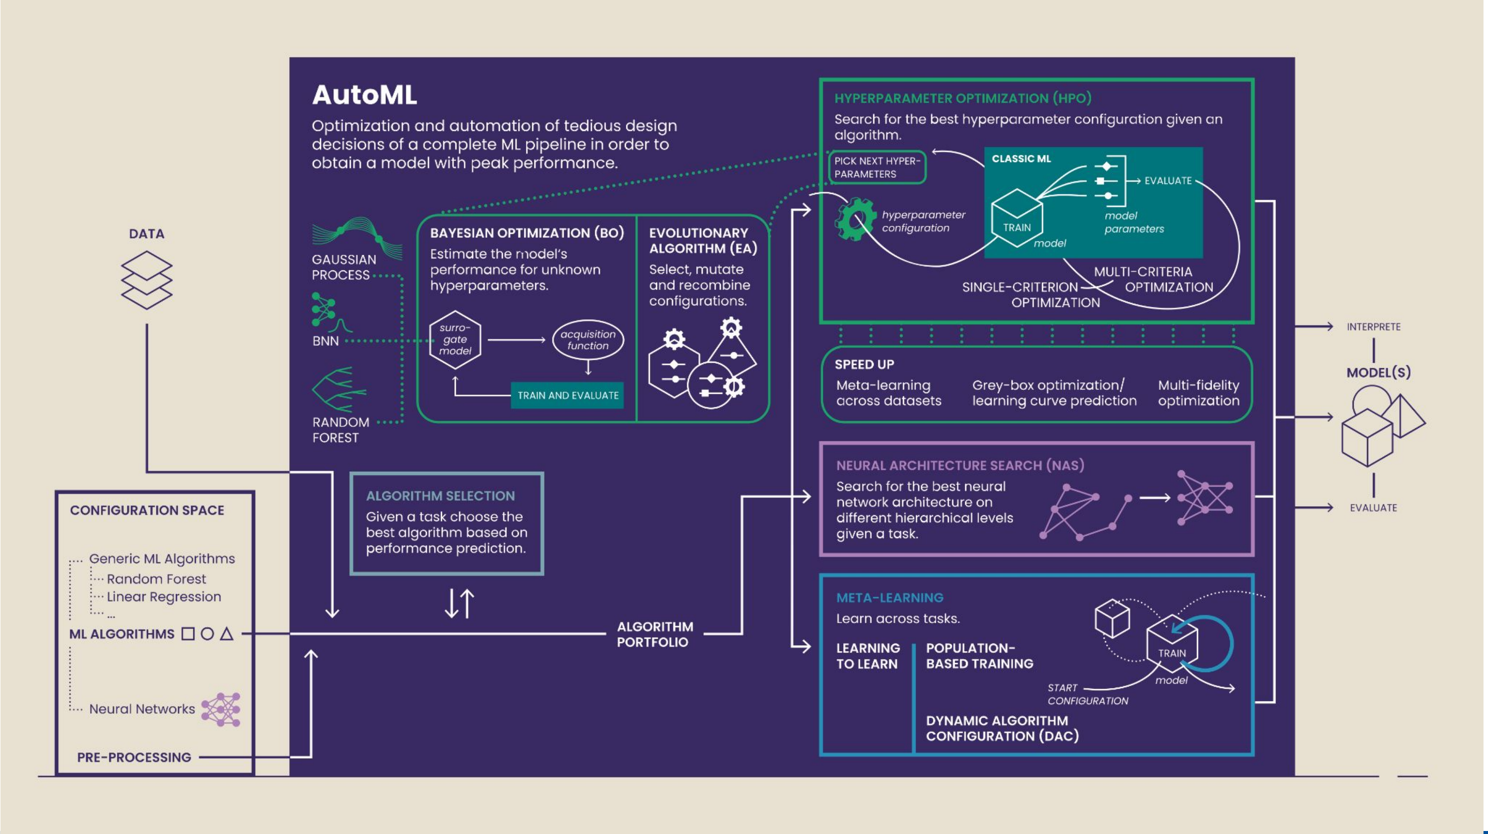

Human-Centered AutoML

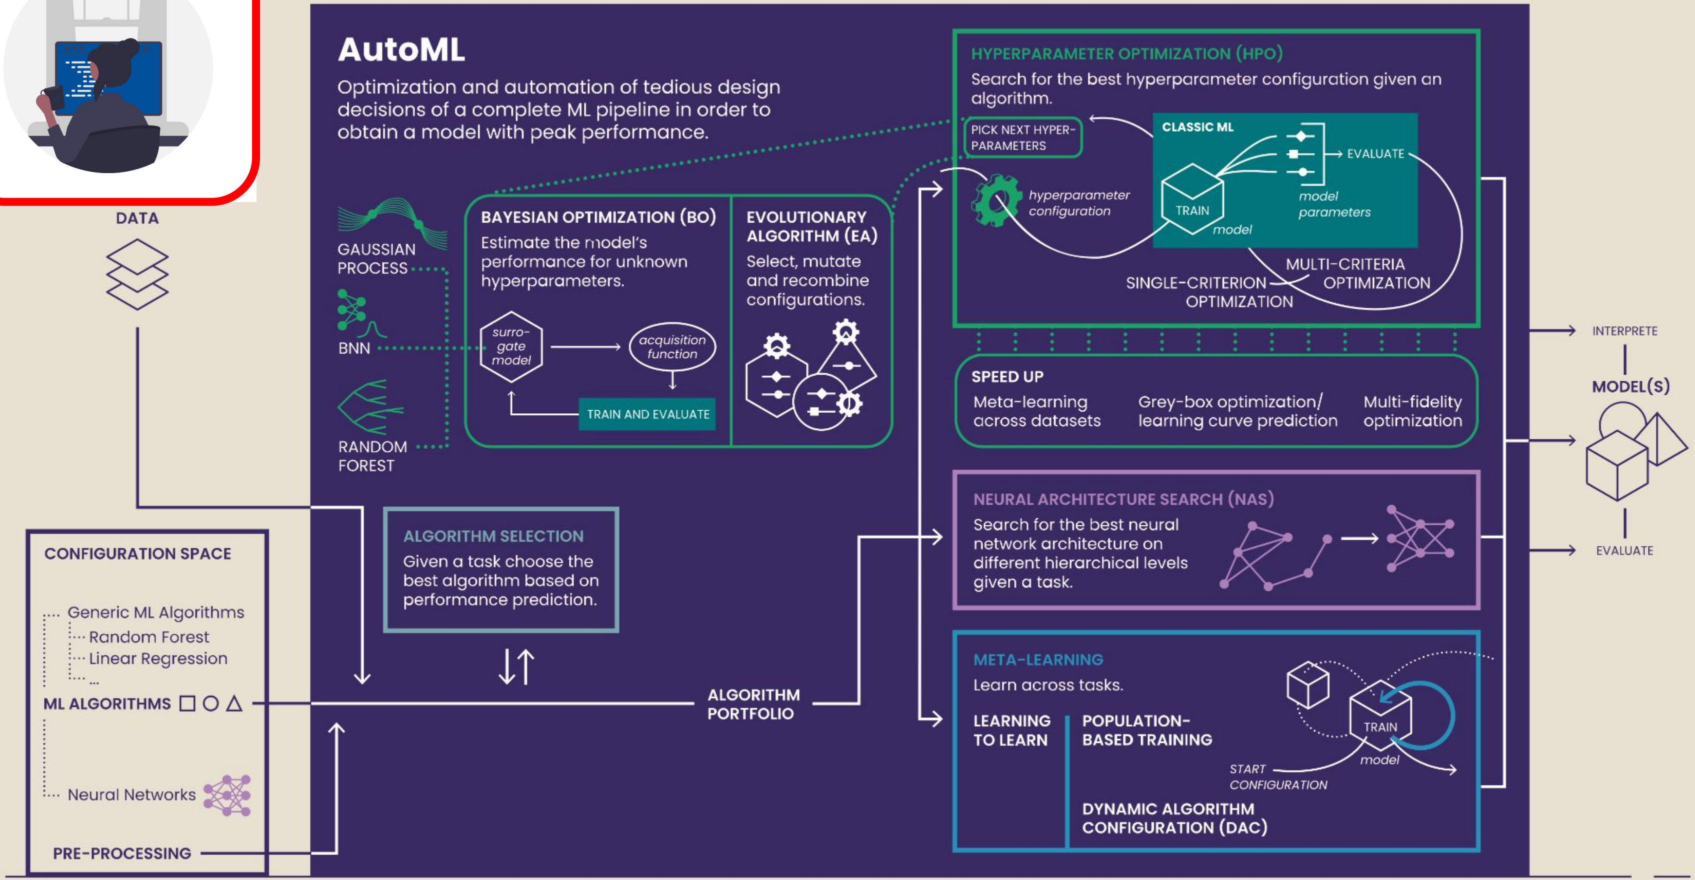
AutoML
Optimization and automation of tedious design decisions of a complete ML pipeline in order to obtain a model with peak performance.
DATA
GAUSSIAN PROCESS
BNN
RANDOM FOREST
BAYESIAN OPTIMIZATION (BO)
Estimate the model's performance for unknown hyperparameters.
surro-gate model
acquisition function
TRAIN AND EVALUATE
EVOLUTIONARY ALGORITHM (EA)
Select, mutate and recombine configurations.
HYPERPARAMETER OPTIMIZATION (HPO)
Search for the best hyperparameter configuration given an algorithm.
PICK NEXT HYPER-PARAMETERS
CLASSIC ML
EVALUATE
hyperparameter configuration
TRAIN
model
model parameters
MULTI-CRITERIA OPTIMIZATION
SINGLE-CRITERION OPTIMIZATION
SPEED UP
Meta-learning across datasets
Grey-box optimization/ learning curve prediction
Multi-fidelity optimization
NEURAL ARCHITECTURE SEARCH (NAS)
Search for the best neural network architecture on different hierarchical levels given a task.
META-LEARNING
Learn across tasks.
LEARNING TO LEARN
POPULATION-BASED TRAINING
TRAIN
model
START CONFIGURATION
DYNAMIC ALGORITHM CONFIGURATION (DAC)
ALGORITHM SELECTION
Given a task choose the best algorithm based on performance prediction.
ALGORITHM PORTFOLIO
CONFIGURATION SPACE
Generic ML Algorithms
Random Forest
Linear Regression
...
ML ALGORITHMS
Neural Networks
PRE-PROCESSING
INTERPRETE
MODEL(S)
EVALUATE

AutoML aims at

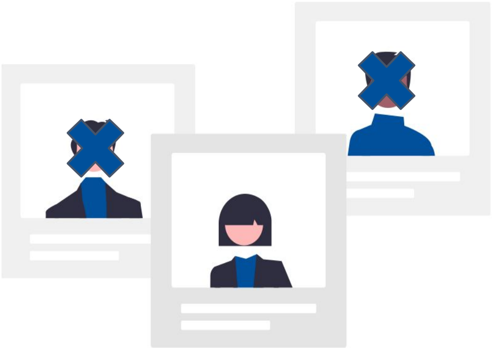

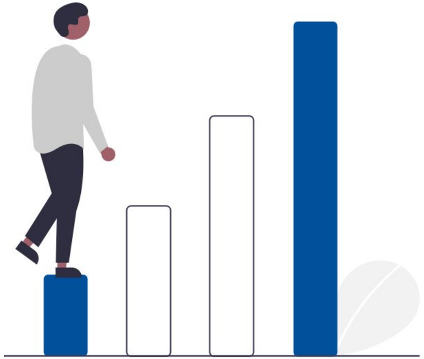

Automating workflows of ML development

Reduce required expert knowledge

Scaling up

But what if ...

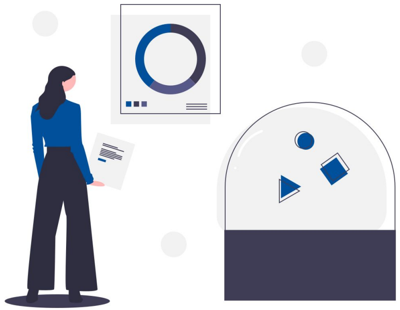

insights into the AutoML black box are crucial?

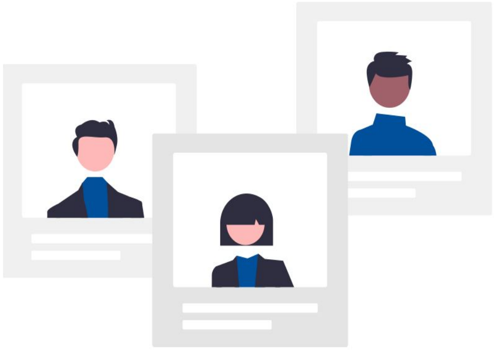

we have expert knowledge?

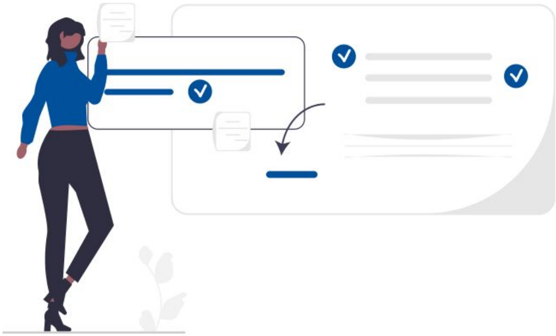

we want to learn from AutoML?

AutoML for Who?

Domain experts with little to no ML expertise

ML practitioners and researchers

Goals in Using HPO [Hasebrook et al. 2023]

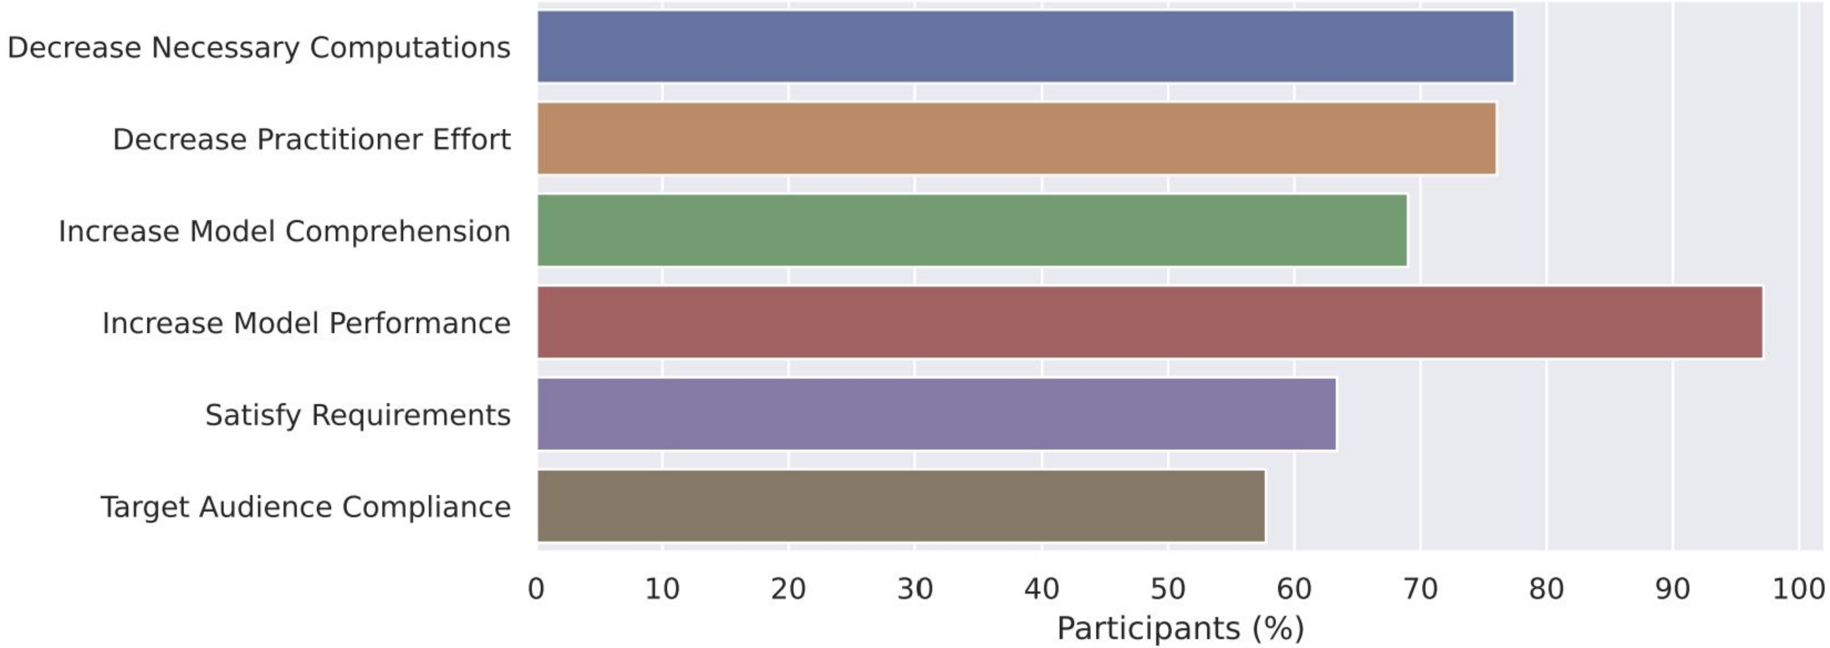

Survey on HPO Use [Hasebrook et al. 2023]

What do you use typically?

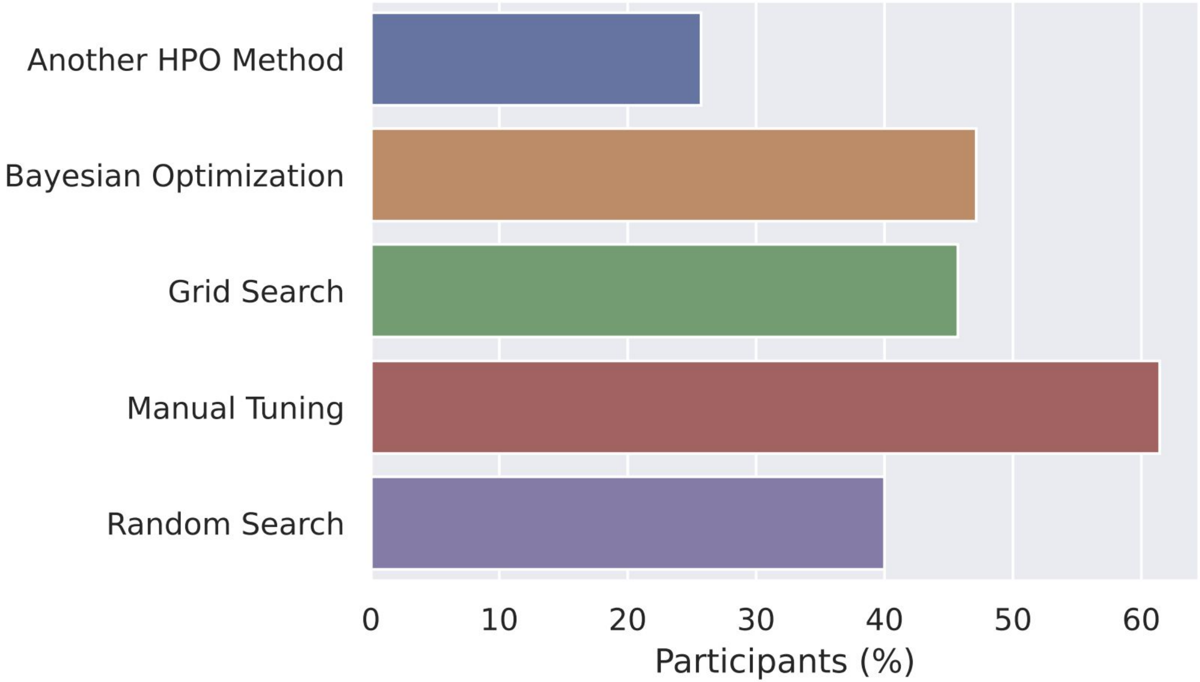

Explainable AutoML

interactive and explainable AutoML

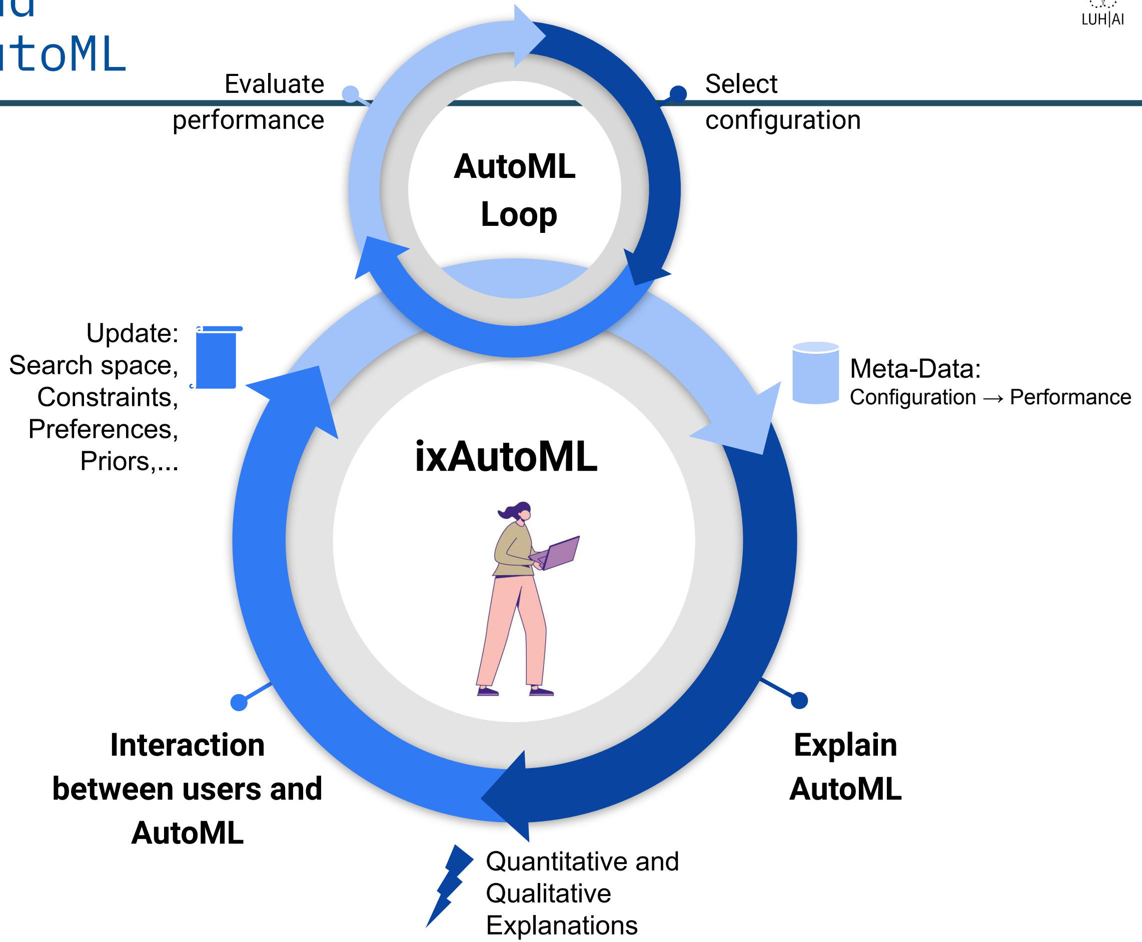

Explaining HPO with PDPs [Moosbauer et al. NeurIPS'21]

Ground truth
PDP
incumbent

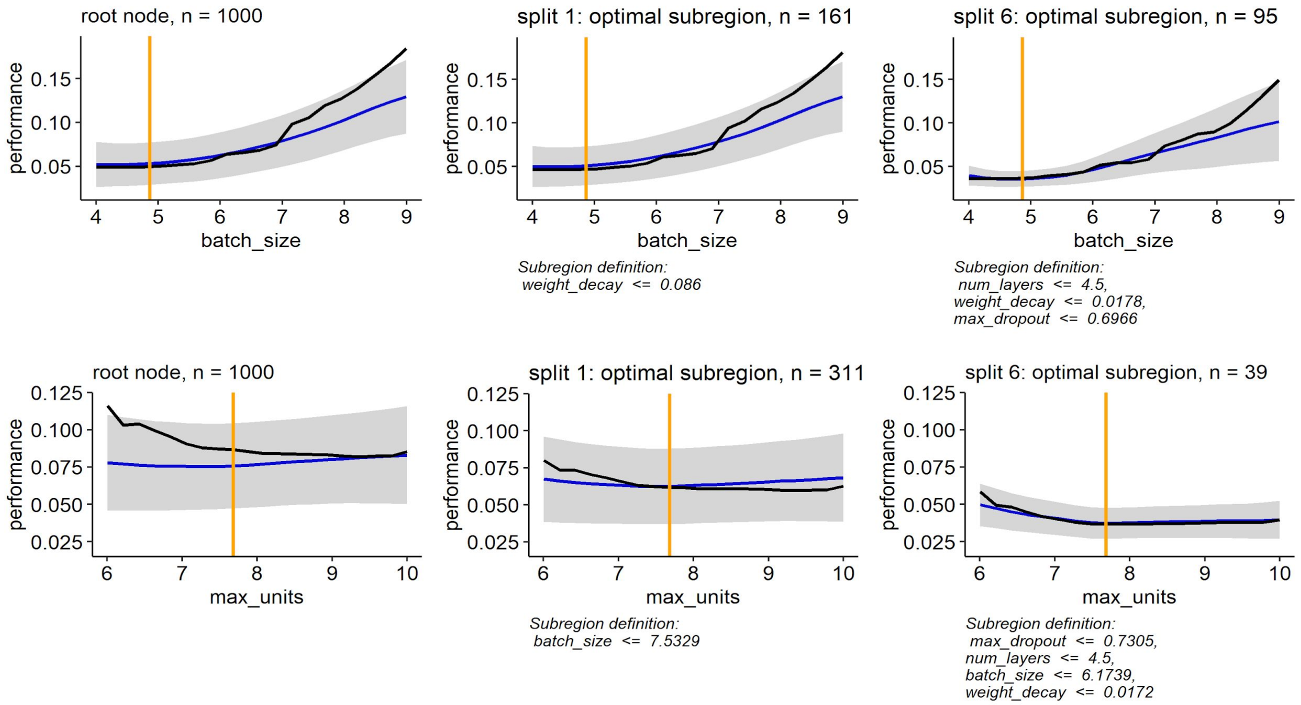

Problem: Biased Sampling (example on PDPs)

- Partial Dependence Plots (PDPs) assume that the data is independently, identically distributed (iid)
- Obviously not the case for efficient AutoML tools with a focus on high-performance regions

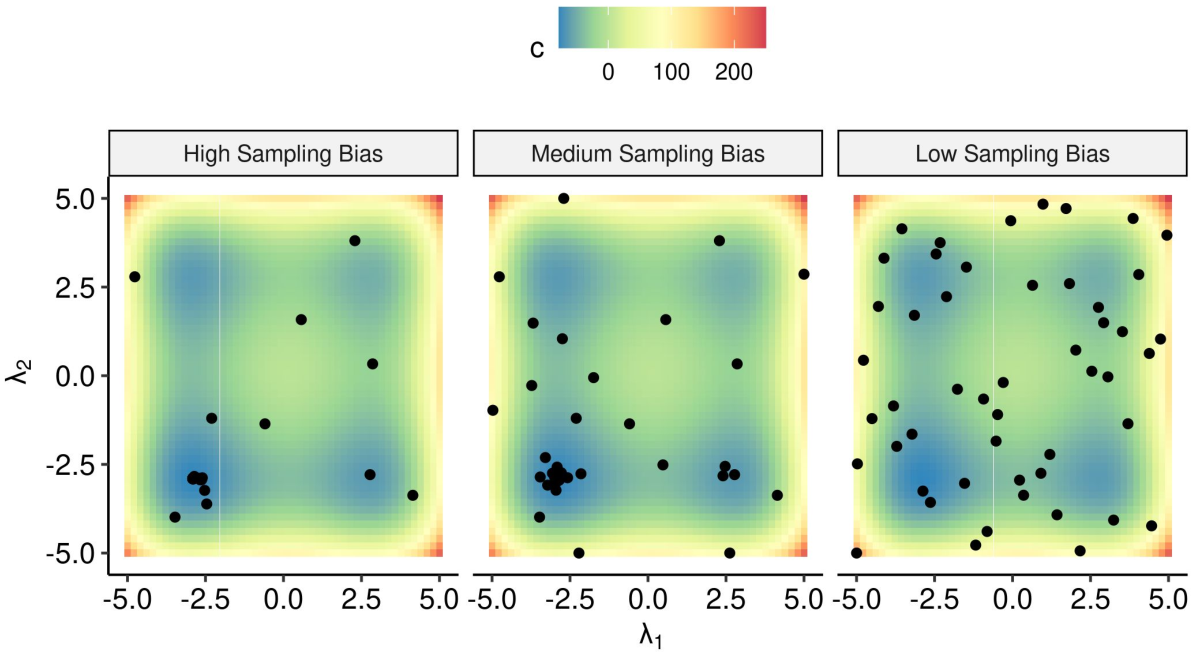

Impact of the Sampling Bias

- Simply using all observations from AutoML tools might lead to misleading PDPs
- Uncertainty estimates help to quantify the poor fits

→ Sampling bias is wanted and a solution to this problem should not change the sampling behavior

⇒ **Adapt explanation techniques or develop new sampling techniques**

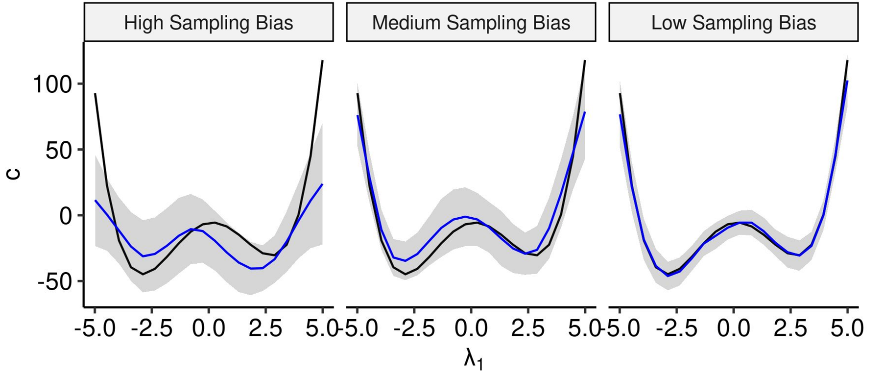

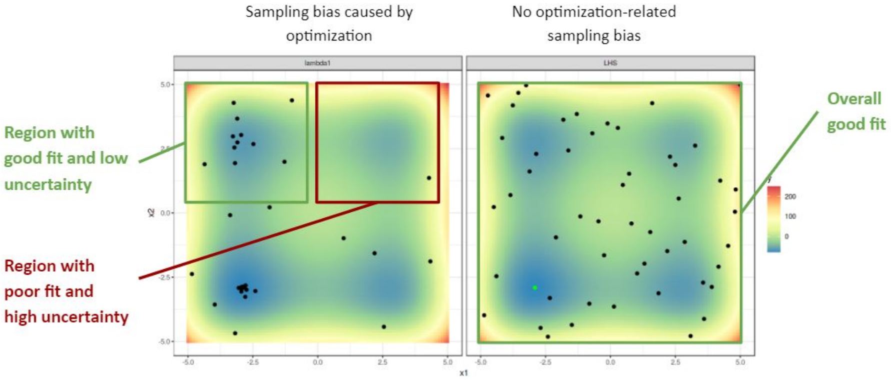

Fair AutoML

AutoML x Fairness [Weerts et al. 2022]

One of many examples

"During the coronavirus crisis, students had to take exams at home. Universities used anti-cheat software to prevent fraud. Among other things, the software had to recognize the student's faces. But it couldn't recognize the student in question, Robin Pocornie. It wasn't until she pointed an extra light at her face that the surveillance software Proctorio finally recognized her. And in the meantime, she had a lot of extra stress to deal with. She feels discriminated against. " [NL Times, 15.07.2023, Webcam exam software "discriminatory," doesn't recognize darker skin tones, says student]

→ Could've AutoML helped here?

→ Can we automate fairness?

Photo by cottonbro studio

Based on https://www.automl.org/can-fairness-be-automated/ and [Weerts et al. 2022]

Fairness Considerations in the ML Workflow

[Weerts et al. 2022]

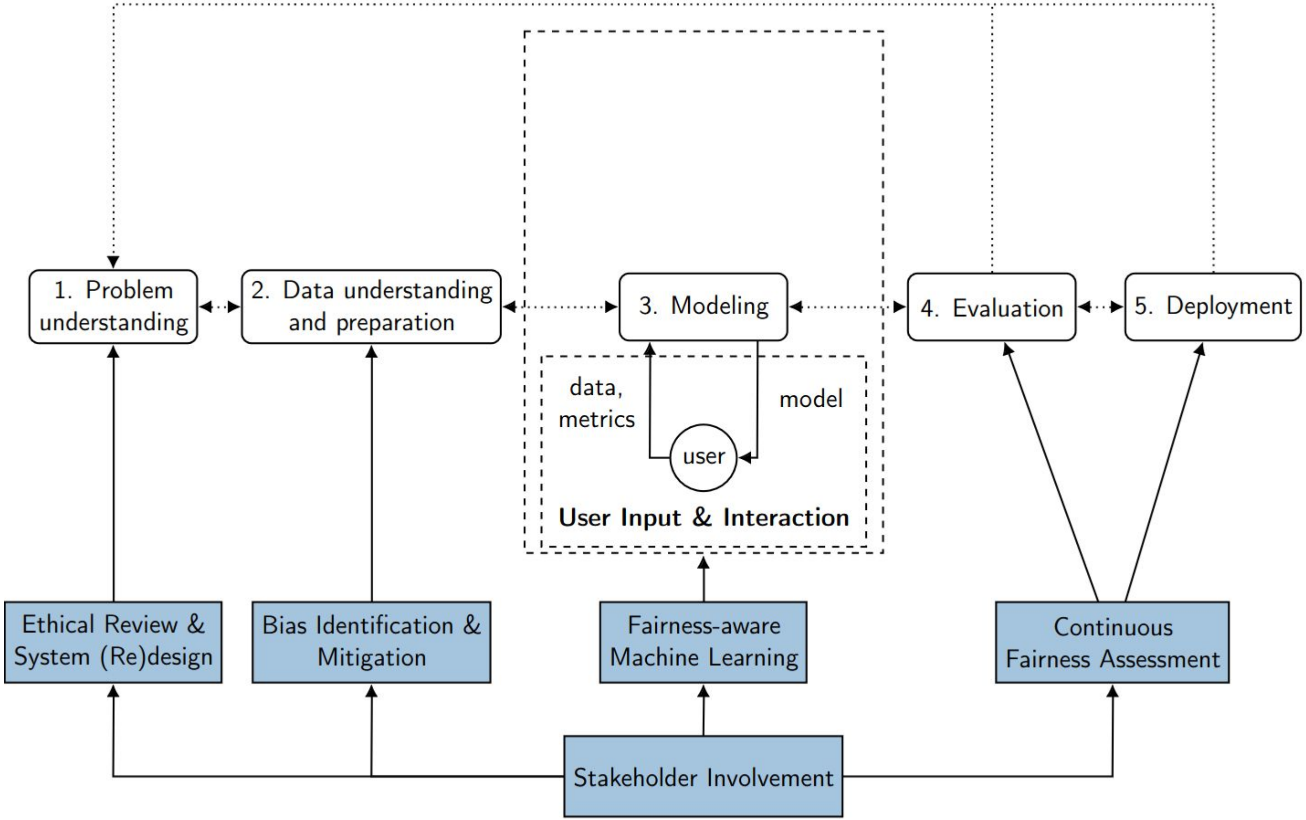

Opportunities for fairness-aware AutoML

[Weerts et al. 2022]

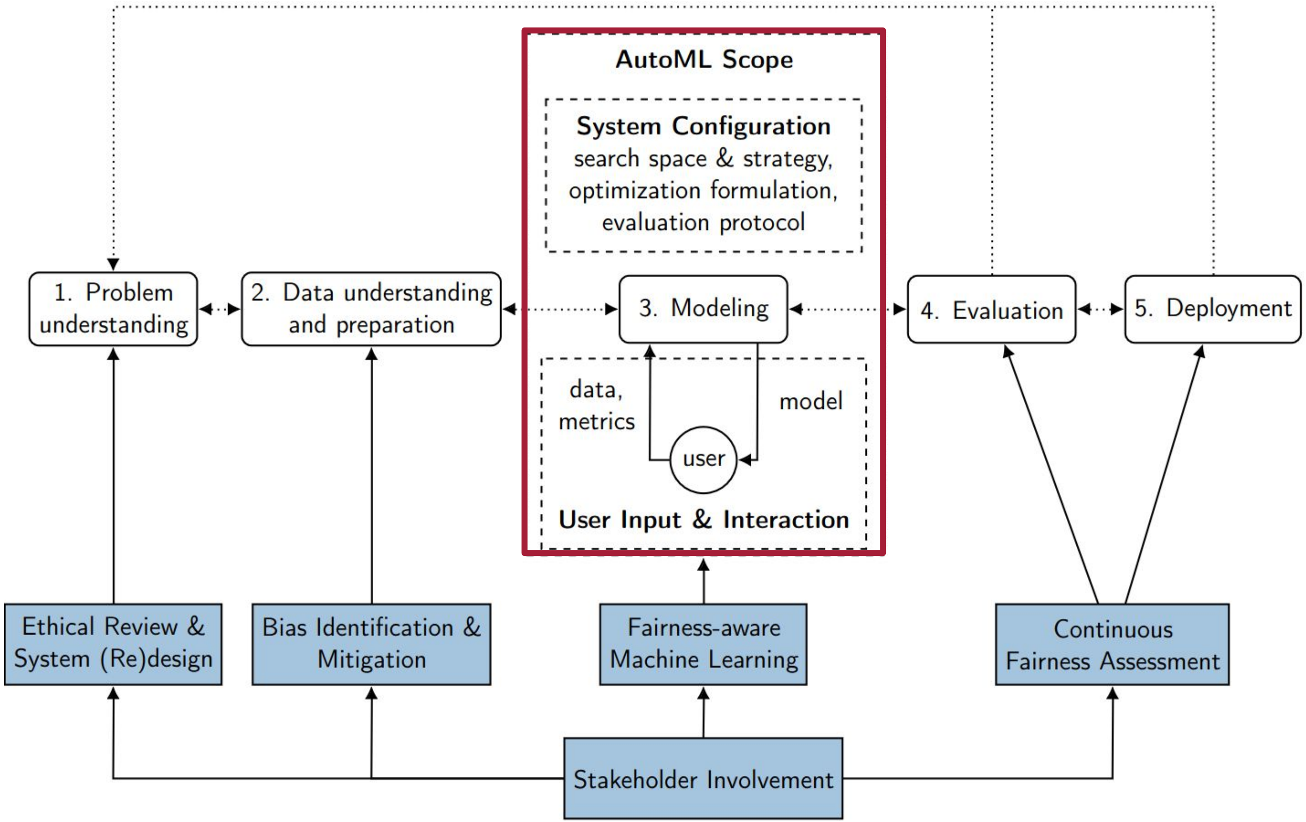

What can we do? Opportunities?

- Codifying best practices
- Better Multi-objective/Constrained optimization
- Better (contextualized) benchmarks
- Better interpretability/explainability
- Better reporting

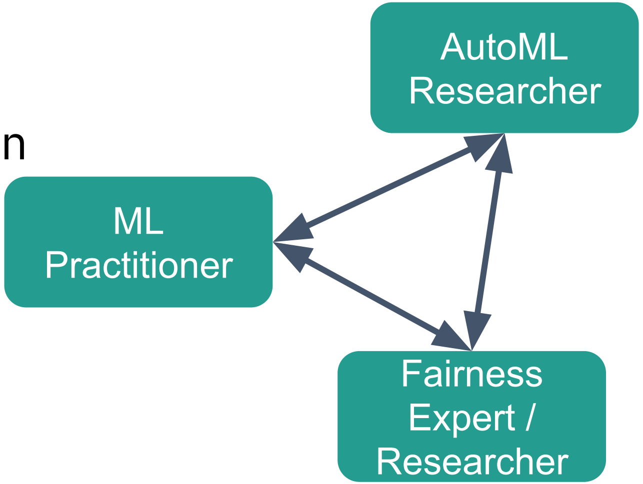

Technical interventions are **not the sole tool for addressing unfairness!**

→ No, we can <u>not fully</u> automate fairness!

→ But AutoML can allow the user to **spend more time on aspects where a human in the loop is essential**

Find Us

AutoML.org

- @AutoML_org
- automl
- @AutoML_org

LUH|AI

- @AIHannover
- LUH-AI
- @luh-ai

Funded by:

European Research Council (Established by the European Commission) · DFG Deutsche Forschungsgemeinschaft · Federal Ministry of Education and Research · Federal Ministry for Economic Affairs and Energy · Federal Ministry for the Environment, Nature Conservation, Nuclear Safety and Consumer Protection · Niedersachsen · DAAD Deutscher Akademischer Austauschdienst German Academic Exchange Service