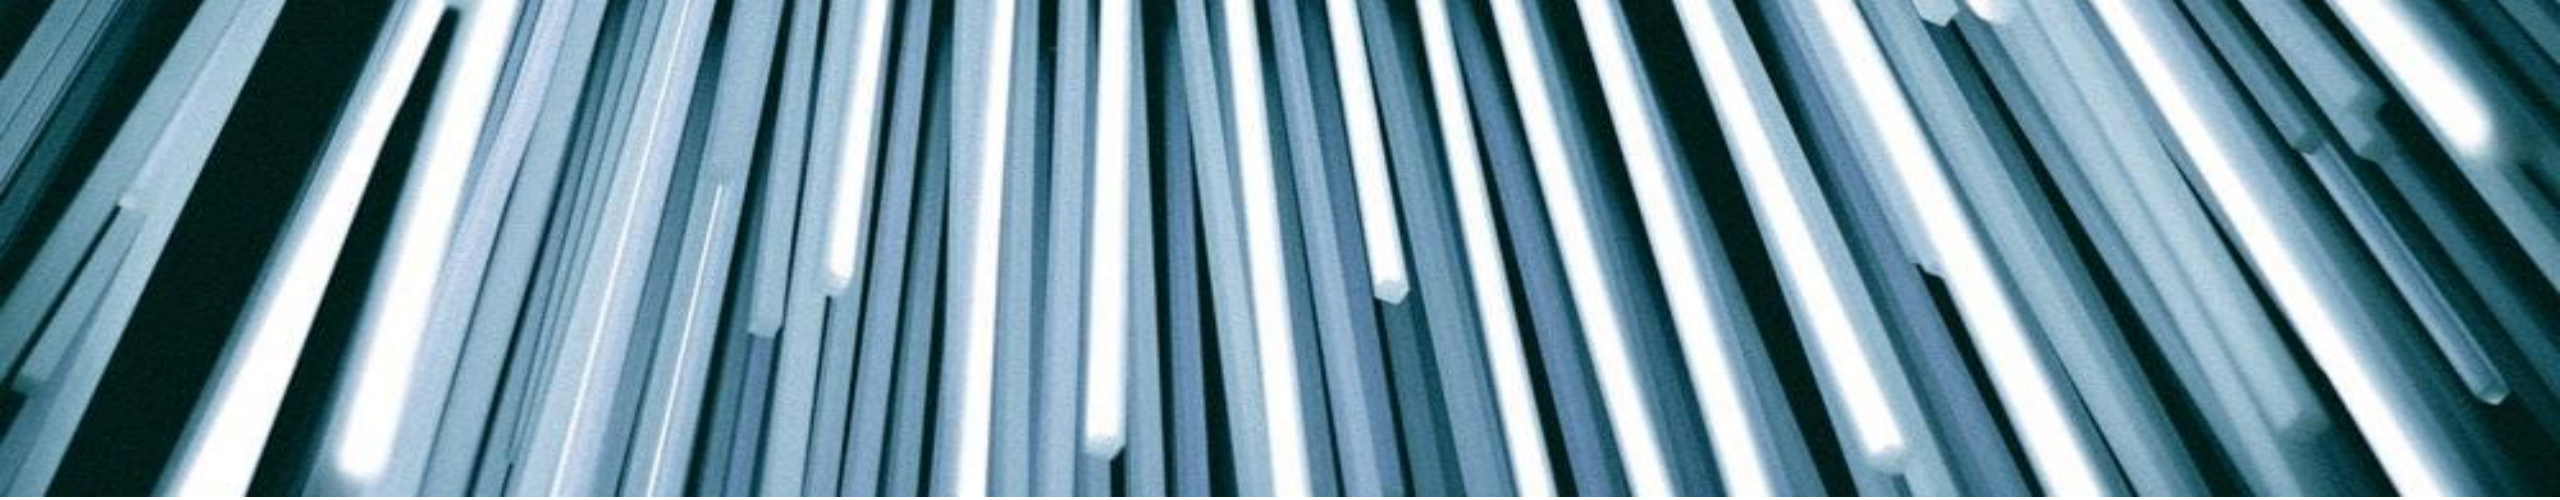

Reciprocal Learning

Julian Rodemann[1], Christoph Jansen[2], Georg Schollmeyer[1]
[1]Department of Statistics, LMU München
[2]School of Computing and Communications, Lancaster University Leipzig

NEURAL INFORMATION PROCESSING SYSTEMS

Machine Learning – A Visual Perspective

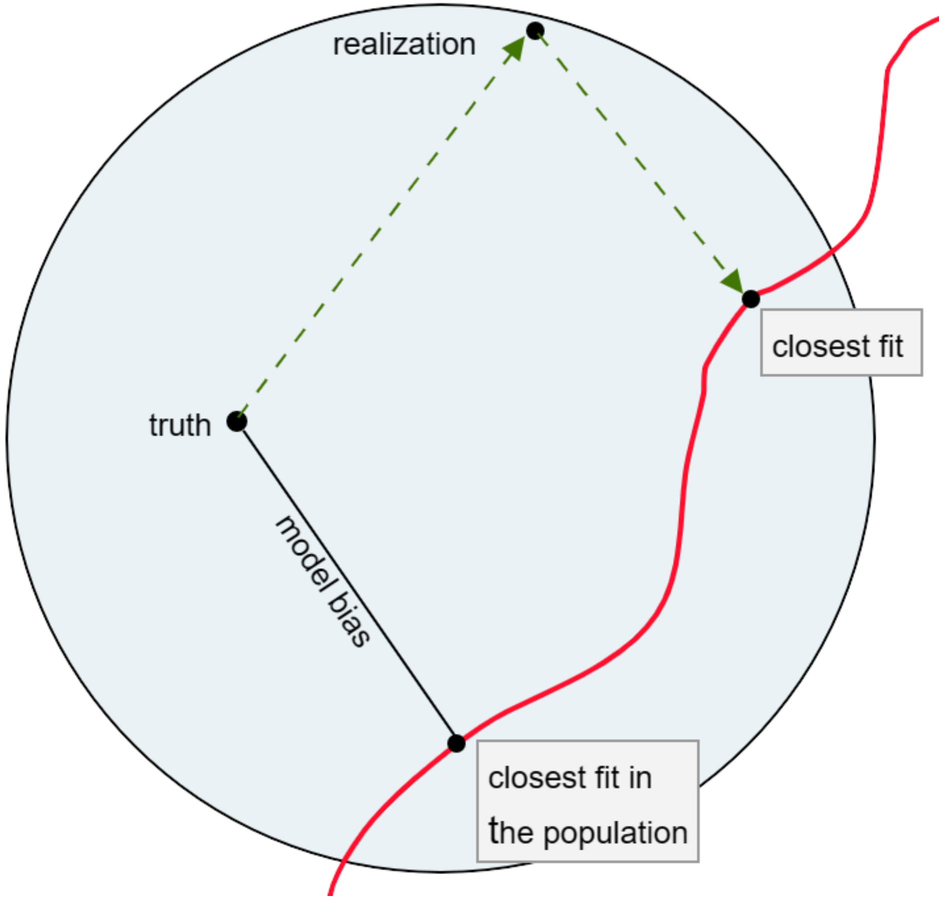

- Red: Model space
- Blue-grey: Sample space
- Machine Learning: Find closest model fit to realized sample

Figure replicated from Jerome H. Friedman, Robert Tibshirani und Trevor Hastie: "The Elements of Statistical Learning"

Online Learning

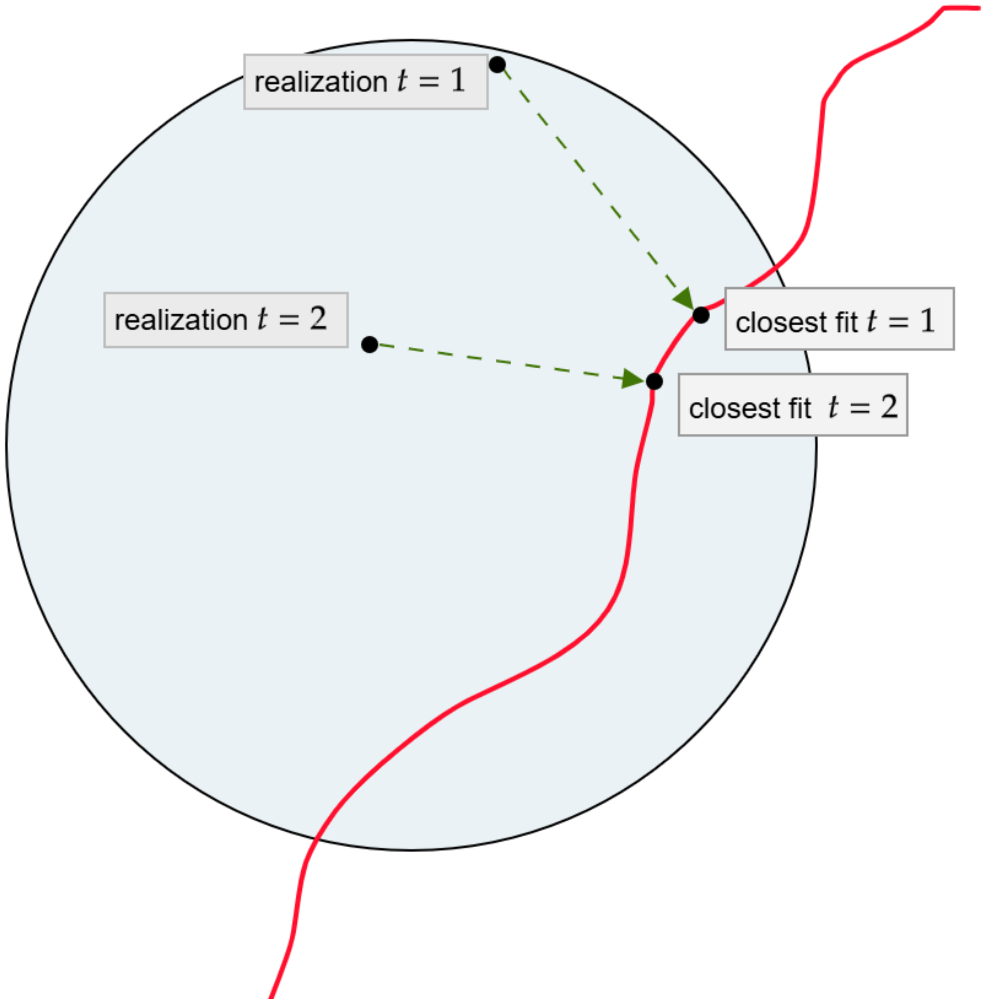

- Online Learning: Realized sample changes, and so does the model fit

NEURAL INFORMATION PROCESSING SYSTEMS

Reciprocal Learning

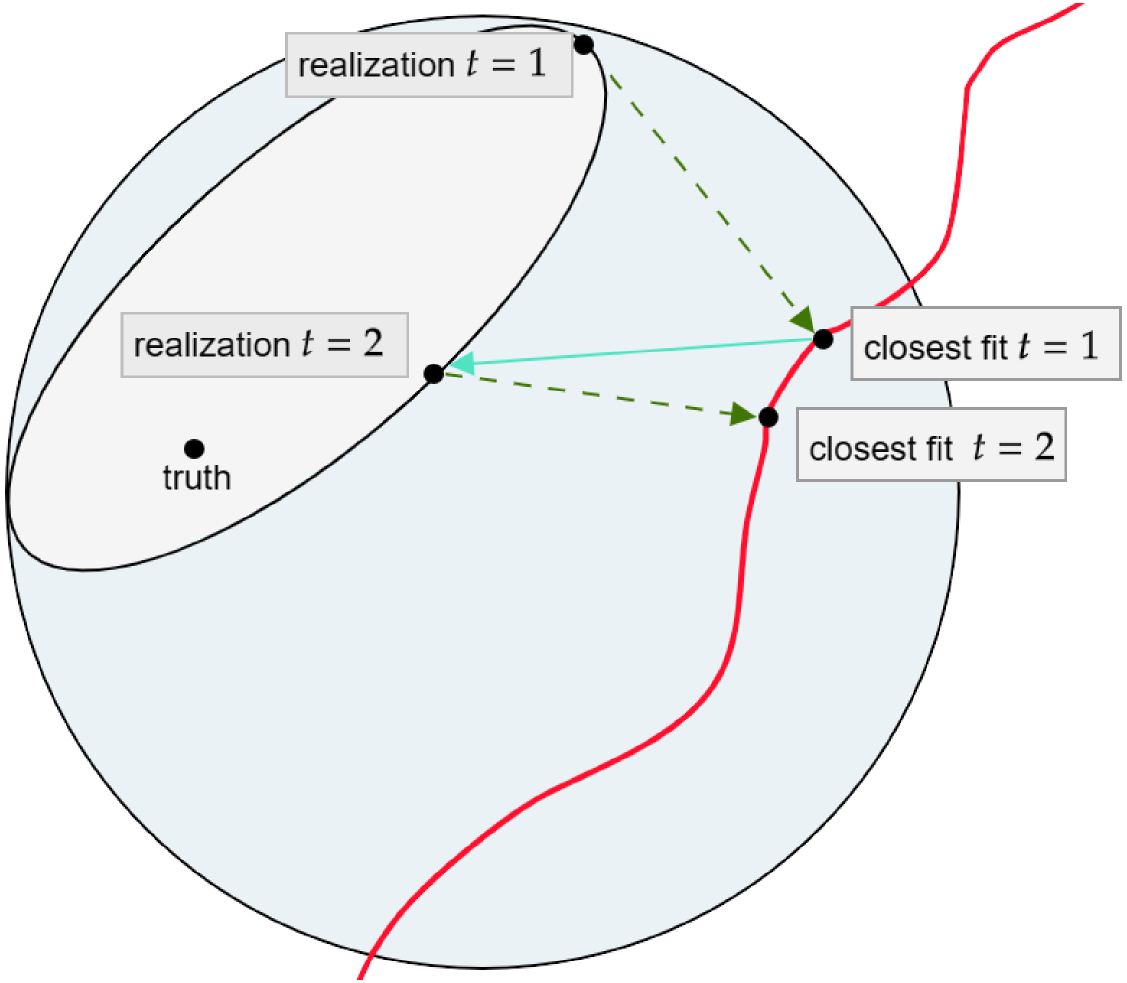

- Observation: Many algorithms change the sample **themselves**
- Sample changes **in response** to the fit
- Grey ellipse: restriction of sample space in t through realization in t−1
- Sample in t depends on model in t−1 and sample in t−1.

NEURAL INFORMATION PROCESSING SYSTEMS

Example: Self-Training

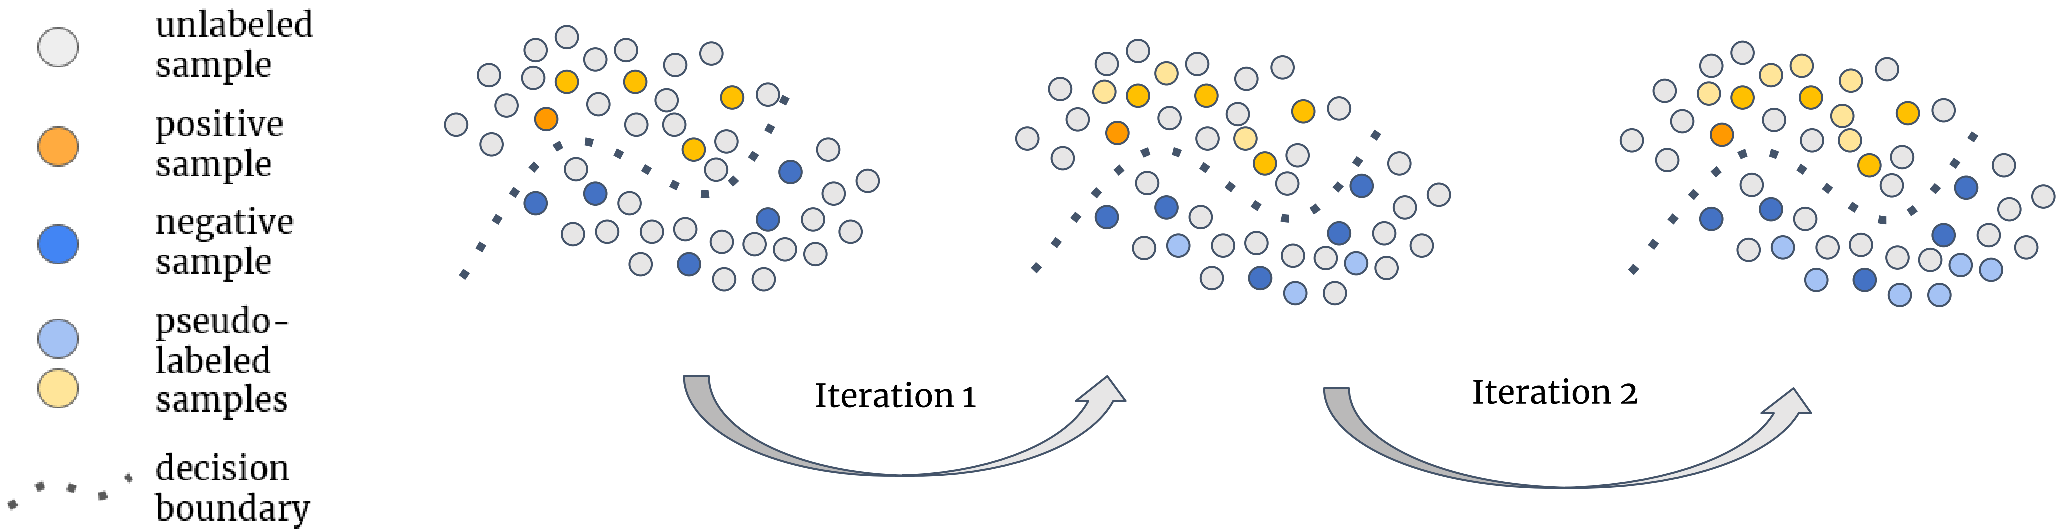

Sketch of Self-Training in Semi-Supervised Learning for Binary Classification

Reciprocal Learning: Outline of Paper

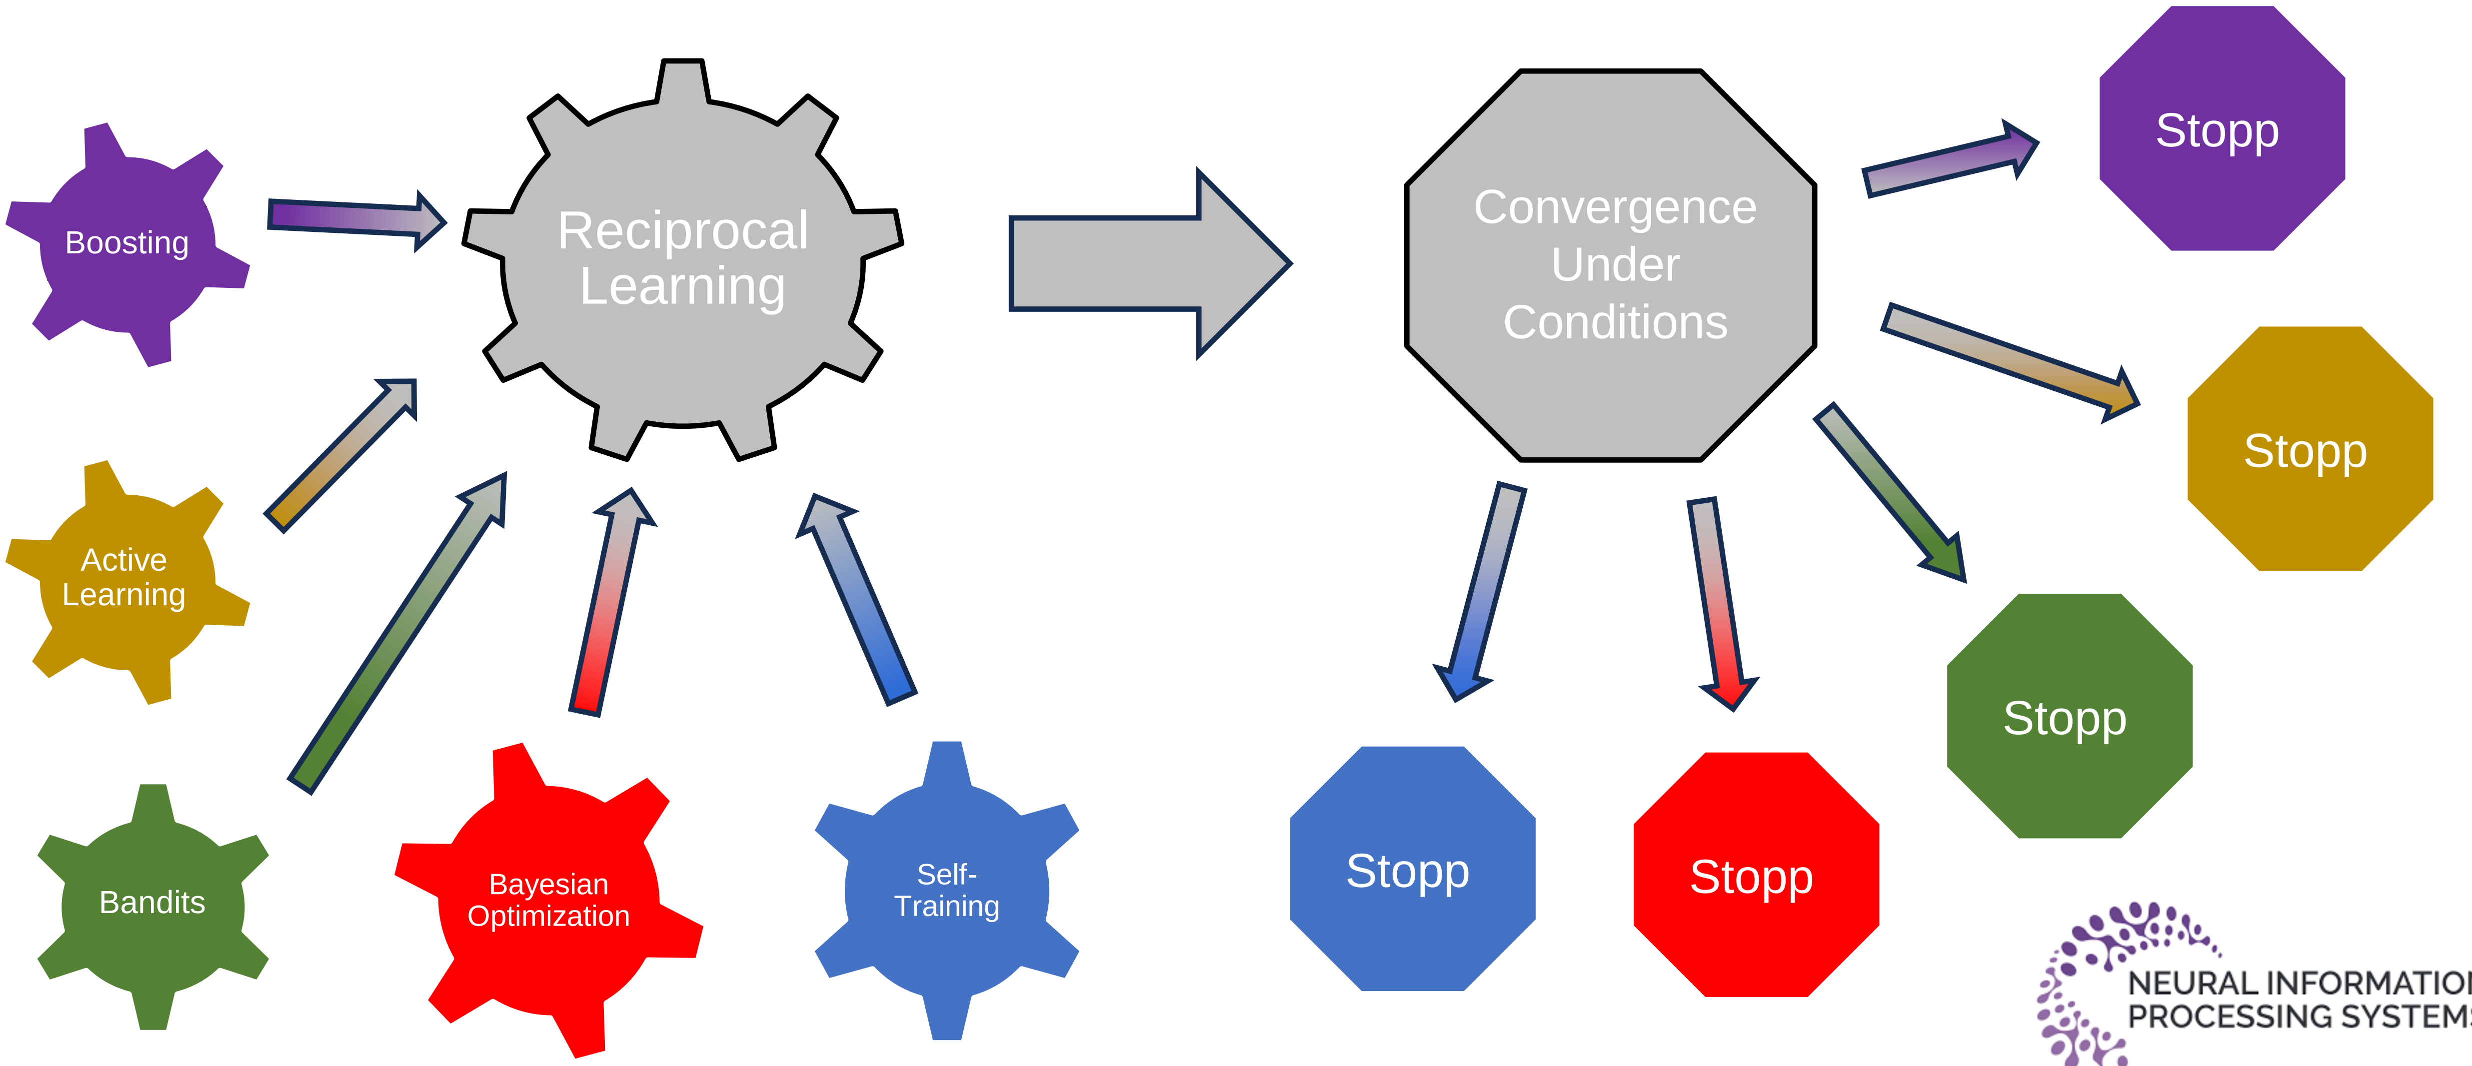

Convergence: Lipschitz Is All You Need

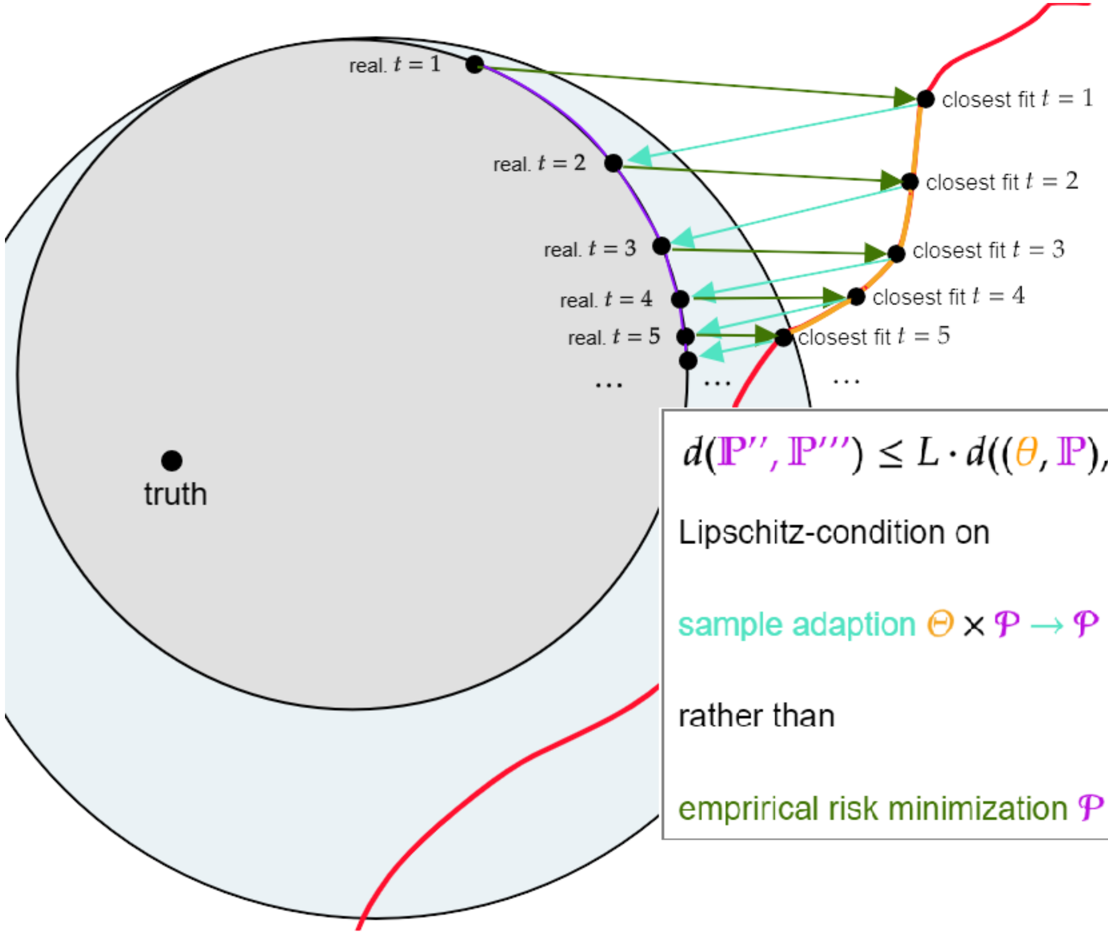

$$d(\mathbb{P}'', \mathbb{P}''') \leq L \cdot d((\theta, \mathbb{P}), (\theta', \mathbb{P}'))$$

Lipschitz-condition on

sample adaption $\Theta \times \mathcal{P} \to \mathcal{P}$

rather than

emprirical risk minimization $\mathcal{P} \to \Theta$

NEURAL INFORMATION PROCESSING SYSTEMS

Data Regularization

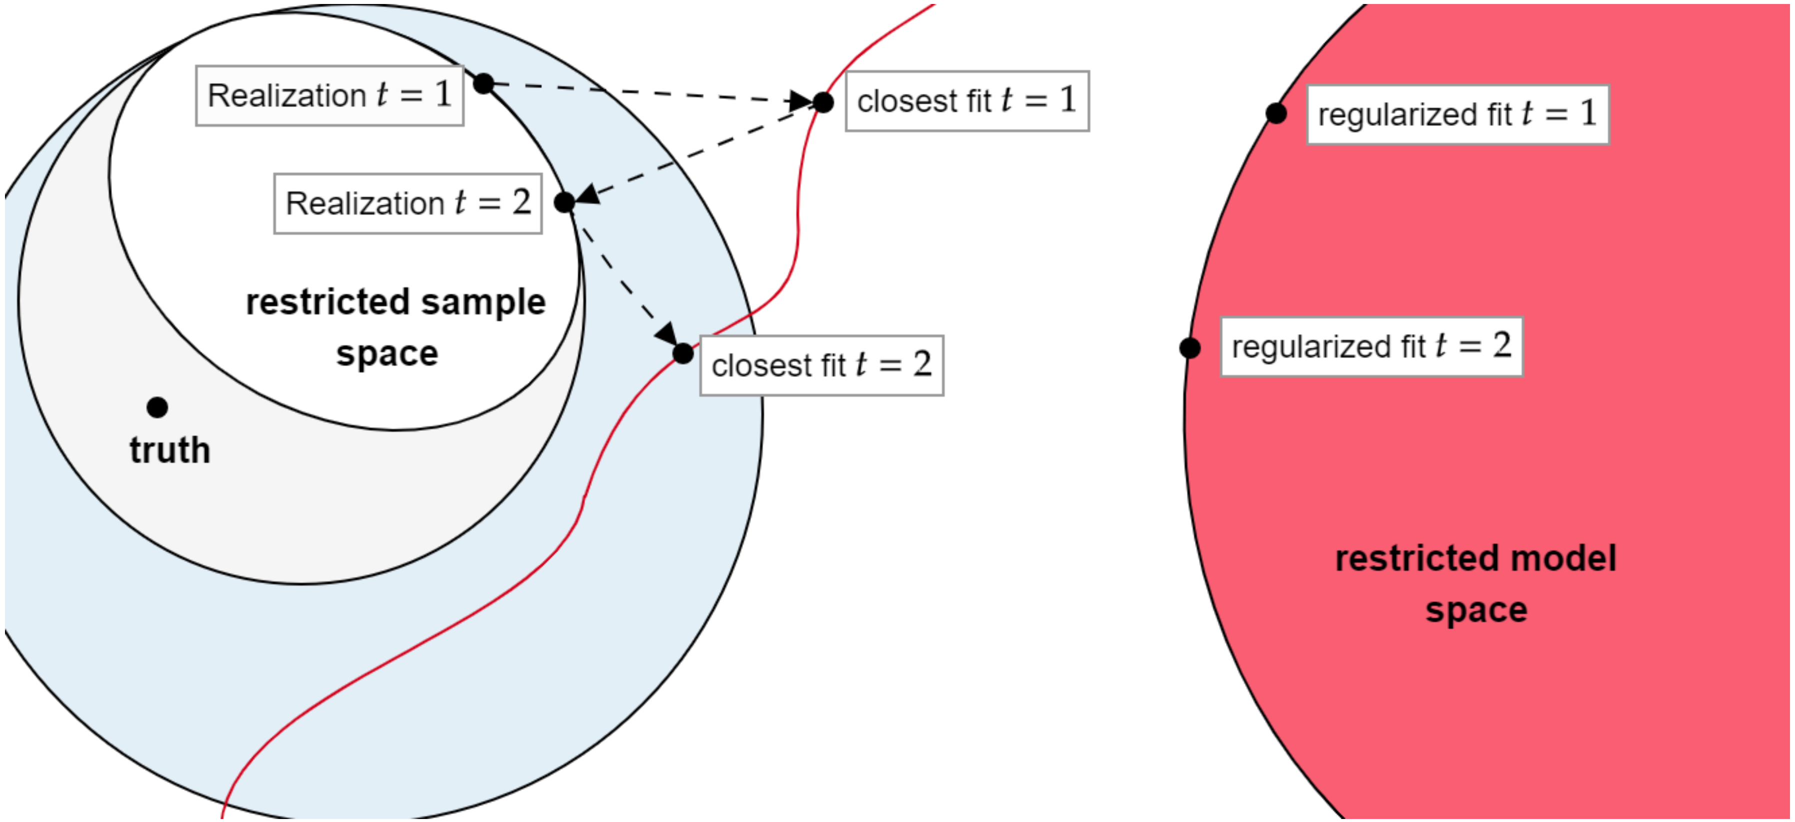

NEURAL INFORMATION PROCESSING SYSTEMS

Thank You for Your Attention!

Reciprocal Learning

Julian Rodemann
Department of Statistics
LMU Munich
j.rodemann@lmu.de

Christoph Jansen
Computing & Communications
Lancaster University Leipzig
c.jansen@lancaster.ac.uk

Georg Schollmeyer
Department of Statistics
LMU Munich
g.schollmeyer@lmu.de

Abstract

We demonstrate that numerous machine learning algorithms are specific instances of one single paradigm: *reciprocal learning*. These instances range from active learning over multi-armed bandits to self-training. We show that all these algorithms not only learn parameters from data but also vice versa: They iteratively alter training data in a way that depends on the current model fit. We introduce reciprocal learning as a generalization of these algorithms using the language of decision theory. This allows us to study under what conditions they converge. The key is to guarantee that reciprocal learning contracts such that the Banach fixed-point theorem applies. In this way, we find that reciprocal learning converges at linear rates to an approximately optimal model under some assumptions on the loss function, if their predictions are probabilistic and the sample adaption is both non-greedy and either randomized or regularized. We interpret these findings and provide corollaries that relate them to active learning, self-training, and bandits.

1 Introduction

The era of data abundance is drawing to a close. While GPT-3 [9] still had to make do with 300 billion tokens, Llama 3 [102] was trained on 15 trillion. With the stock of high-quality data growing at a much smaller rate [67], adequate training data might run out within this decade [58, 107]. Generally

NEURAL INFORMATION PROCESSING SYSTEMS