Control Theory for Reinforcement Learning — A Tutorial Session

Anders Rantzer

Lund University, Sweden

Our Aim

- This is *not* a tutorial on how to use machine learning for control.
- Rather, it is an attempt to highlight some concepts and ideas that we believe have the potential to enrich control theory in coming years.

Our Aim

- This is *not* a tutorial on how to use machine learning for control.
- Rather, it is an attempt to highlight some concepts and ideas that we believe have the potential to enrich control theory in coming years.

Some Questions to Keep in Mind

- Does high-dimensional data require entirely different tools?
- Is there a conflict between transient and asymptotic performance?
- What aspects of control theory can make a difference in industry?

Some Questions to Keep in Mind

- Does high-dimensional data require entirely different tools?
- Is there a conflict between transient and asymptotic performance?
- What aspects of control theory can make a difference in industry?

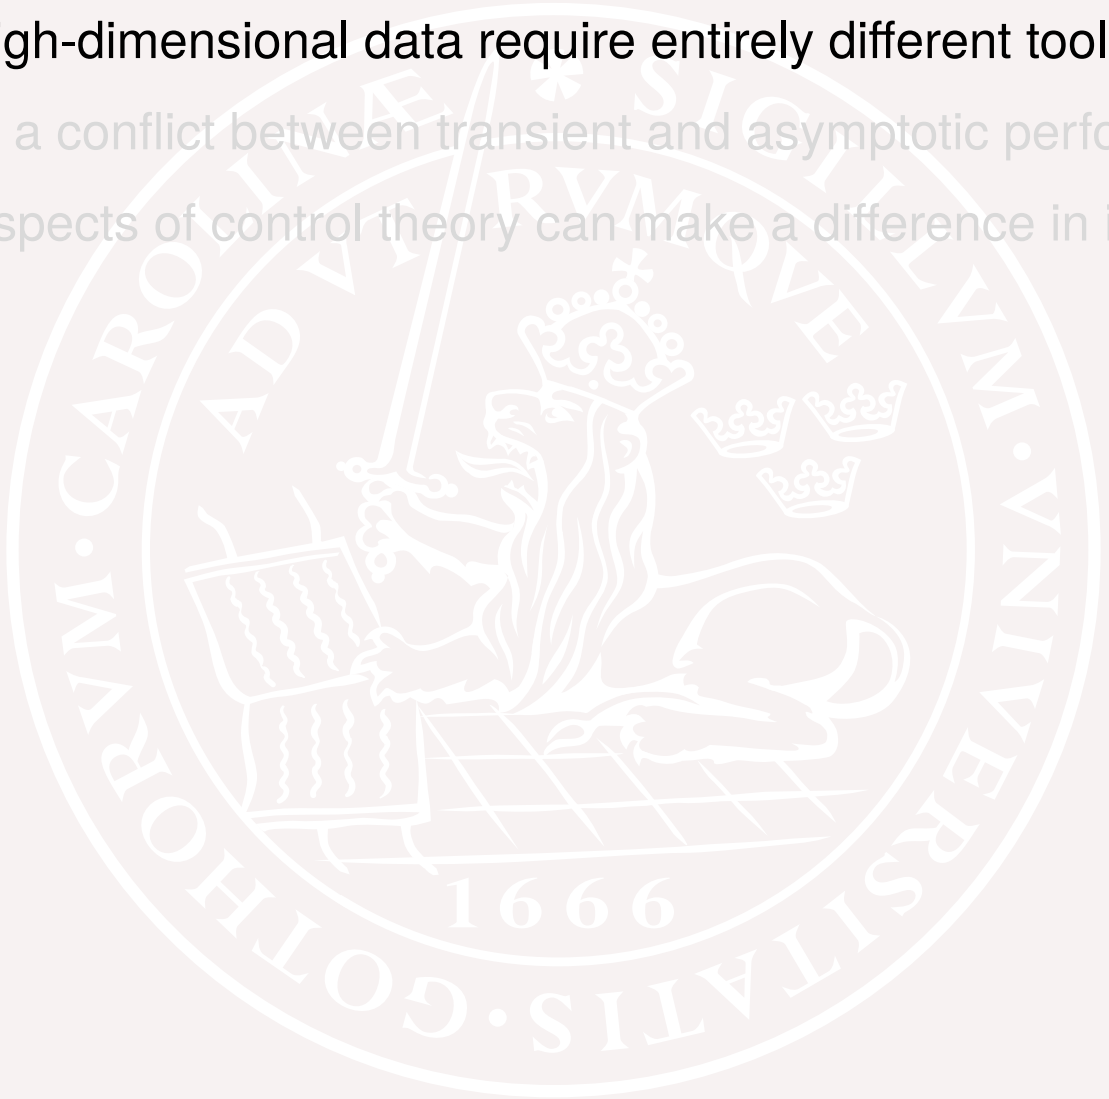

Some Questions to Keep in Mind

- Does high-dimensional data require entirely different tools?
- Is there a conflict between transient and asymptotic performance?
- What aspects of control theory can make a difference in industry?

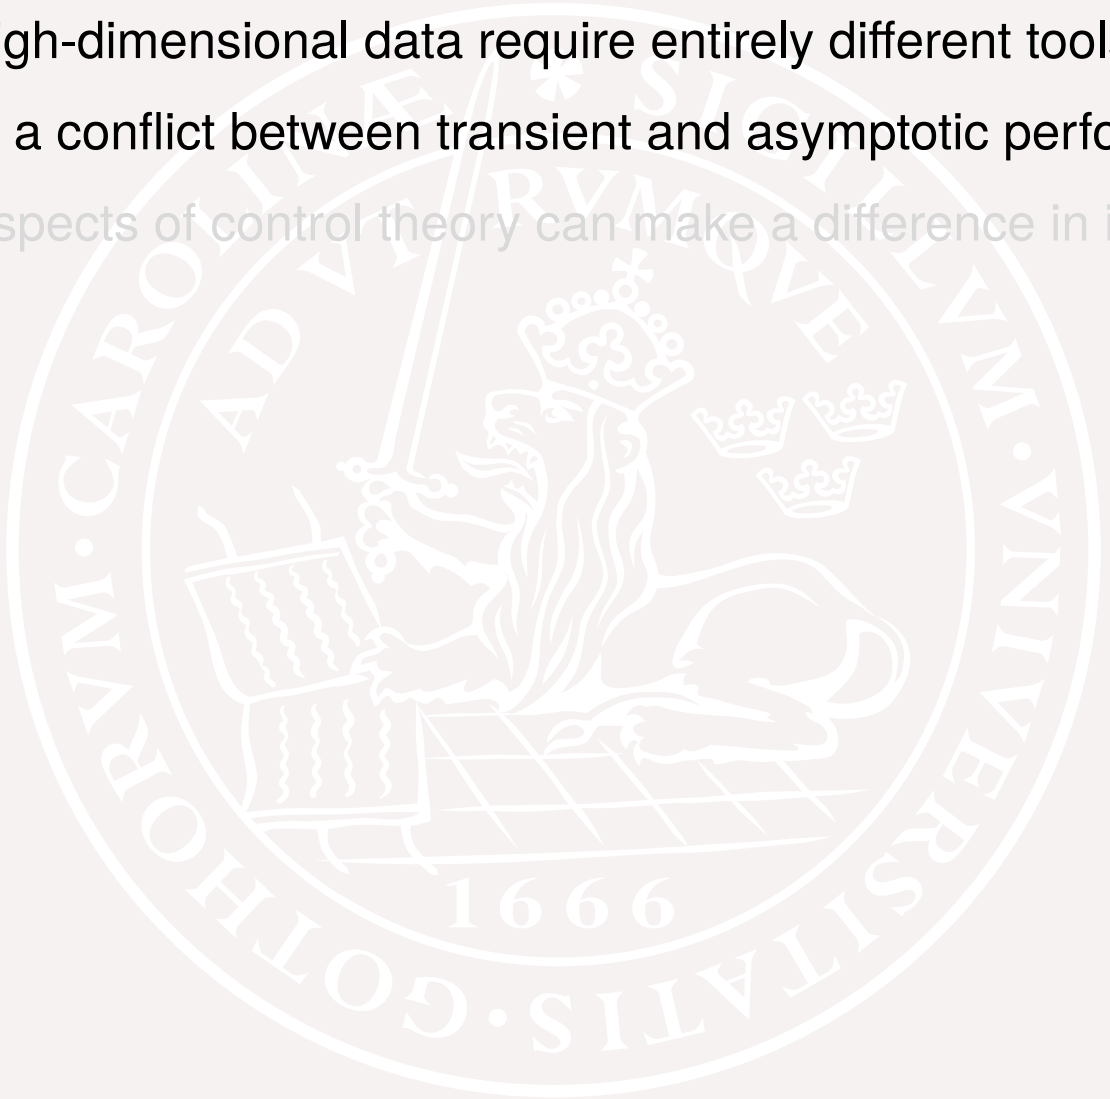

Some Questions to Keep in Mind

- Does high-dimensional data require entirely different tools?
- Is there a conflict between transient and asymptotic performance?
- What aspects of control theory can make a difference in industry?

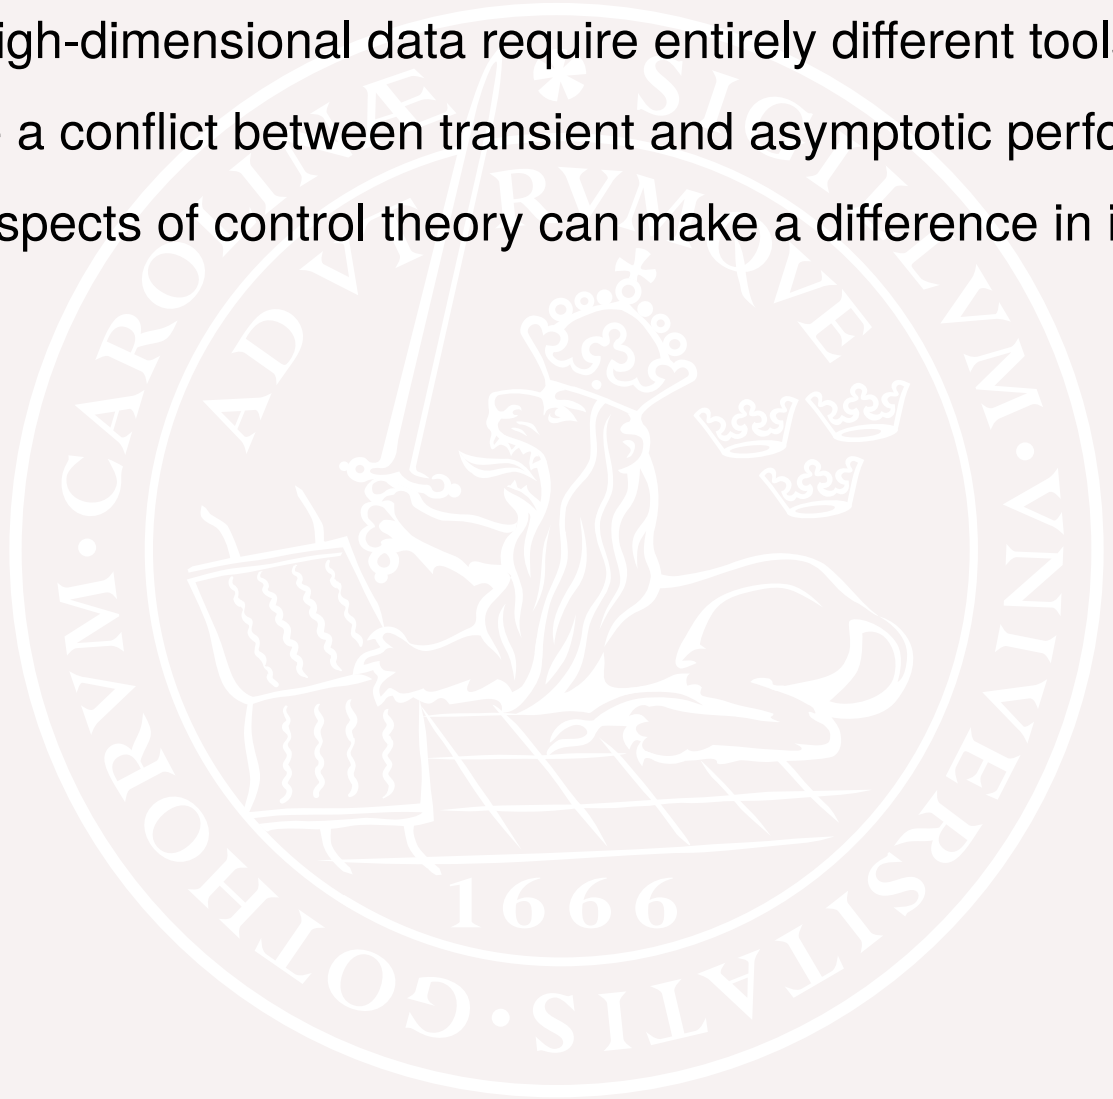

Some Questions to Keep in Mind

- Does high-dimensional data require entirely different tools?
- Is there a conflict between transient analysis and asymptotic bounds?
- What aspects of control theory can make a difference in industry?

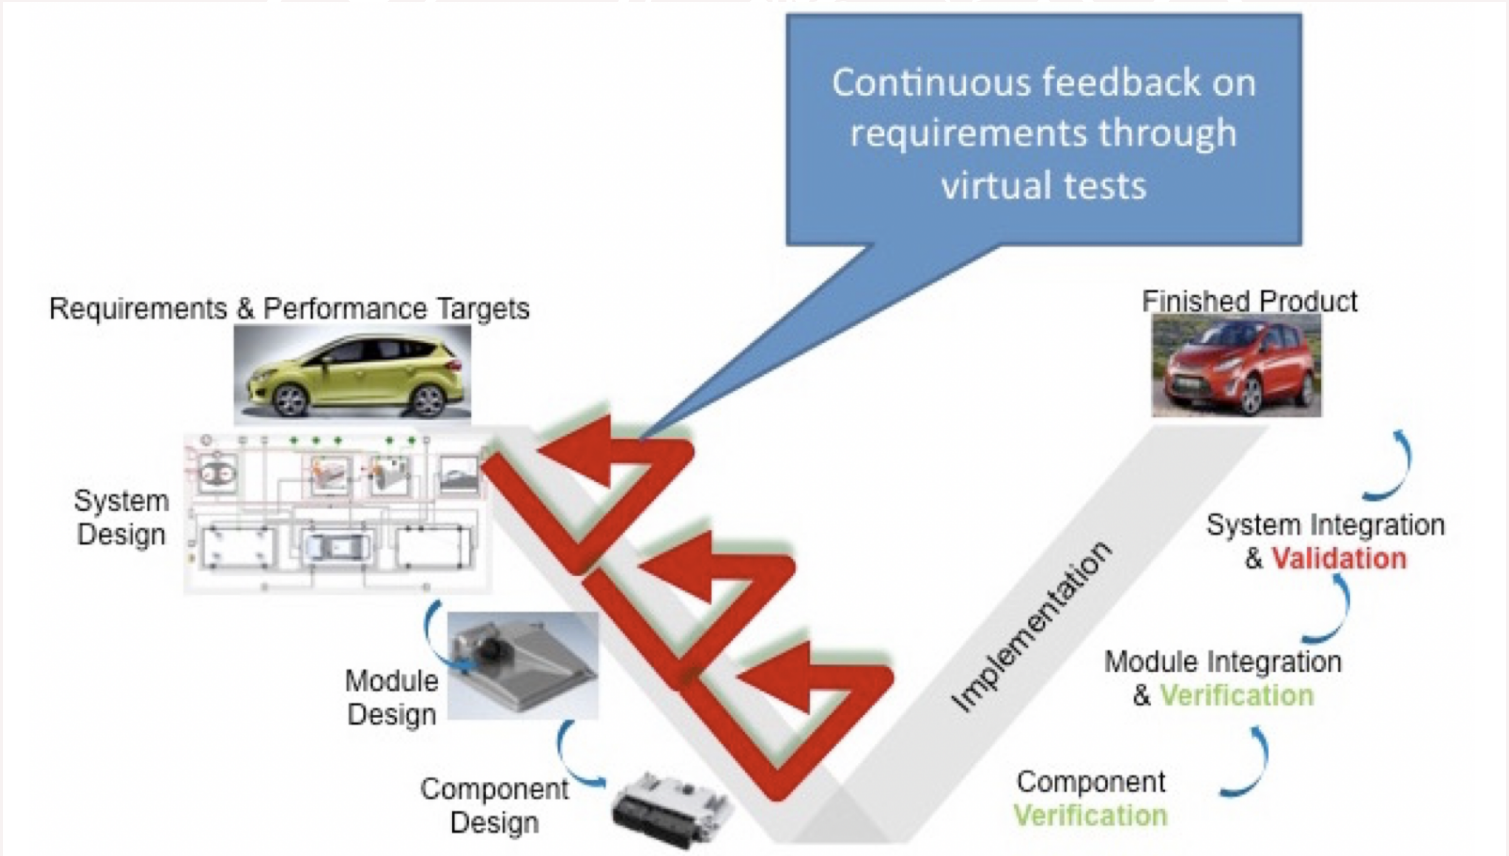

Schedule

10:05
Nikolai Matni, University of Pennsylvania
Self-tuning regulators and reinforcement learning

10:45
Alexandre Proutiere, KTH, Stockholm
Controlling Unknown Discrete Systems

11:10
Anders Rantzer, Lund University
Optimization based exploration/exploitation

11:35
Stephen Tu, Google Brain, New York City
Concentration Bounds for System Identification

12:00 End