A Next Step in Industrial Digitalization

Dan Smith, PhD
Business Director
BTG Group

The Digital Journey – Driving Performance

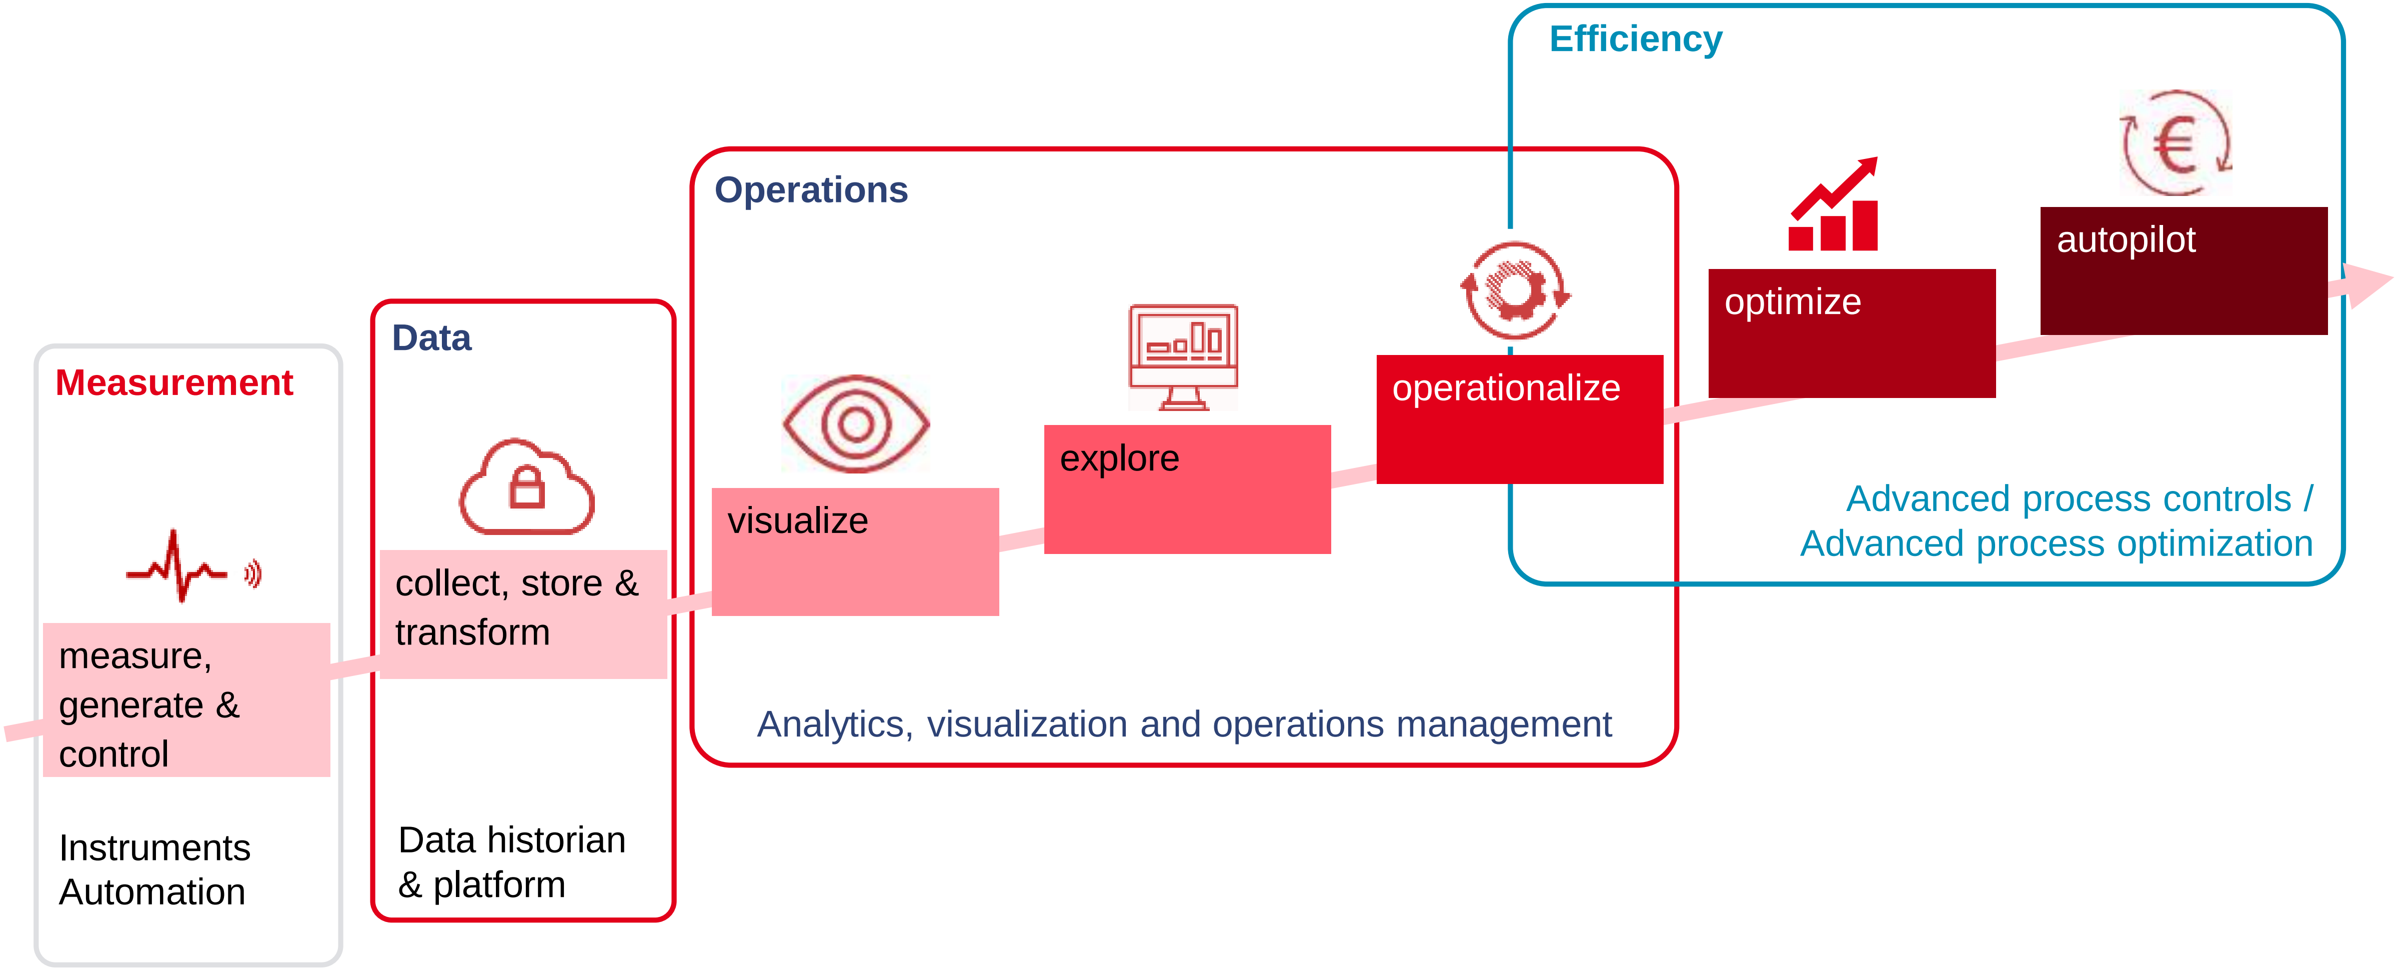

A Brief History of Industrial AI – What Didn't Work

1980 – Expert Systems – "*Solving All Problems*"

Software that mimics human expertise using a knowledge base and inference engine to solve domain-specific problems.

Example: Recovery boiler advisor

Adoption: 65% of Fortune 500 companies used expert systems in the 1980s for diagnosis, configuration, and decision support

Decline in the 2000s:

- **Knowledge acquisition issues**
- **Scalability limits**
- **Verification challenges**
- **Overhyped expectations**

1990 – Neural Networks – “*Solving All Problems*”

Neural networks are brain-inspired machine learning models using interconnected nodes to learn patterns and make decisions from data

Example: Tensile Strength Model

Adoption: Significant pulp and paper market penetration, many Proof-of-Concept examples

Decline in the 2000s

- **"Black Box" Problem**
- **Ignoring the Downsides**
- **Real-World Problems**
- **Hard to Use**

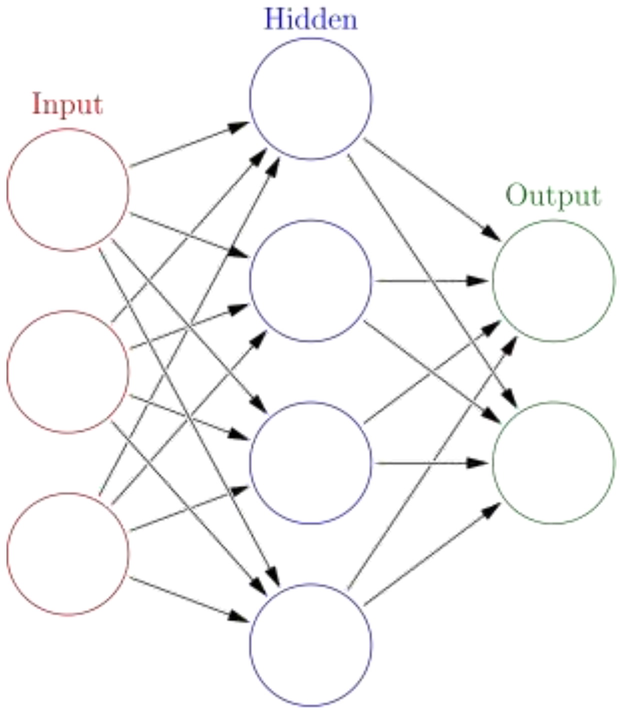

What Did Work?

1980 – Model Predictive Control – Solving Problems

Model Predictive Control (MPC) is an advanced control technique that uses a system model to predict the process response and optimize control actions

Example: Bleach Plant Advanced Control

Adoption: Invented 1982; 5000+ applications in 2000; **1,000,000+ in 2025**

Uses: Process industries, automotive, aerospace, power, robotics, building automation,…

Continuous Growth

- **Better Technology**
- **Energy and Cost Savings**
- **Increasing process complexity**

1980 – Model Predictive Control – Solving Problems

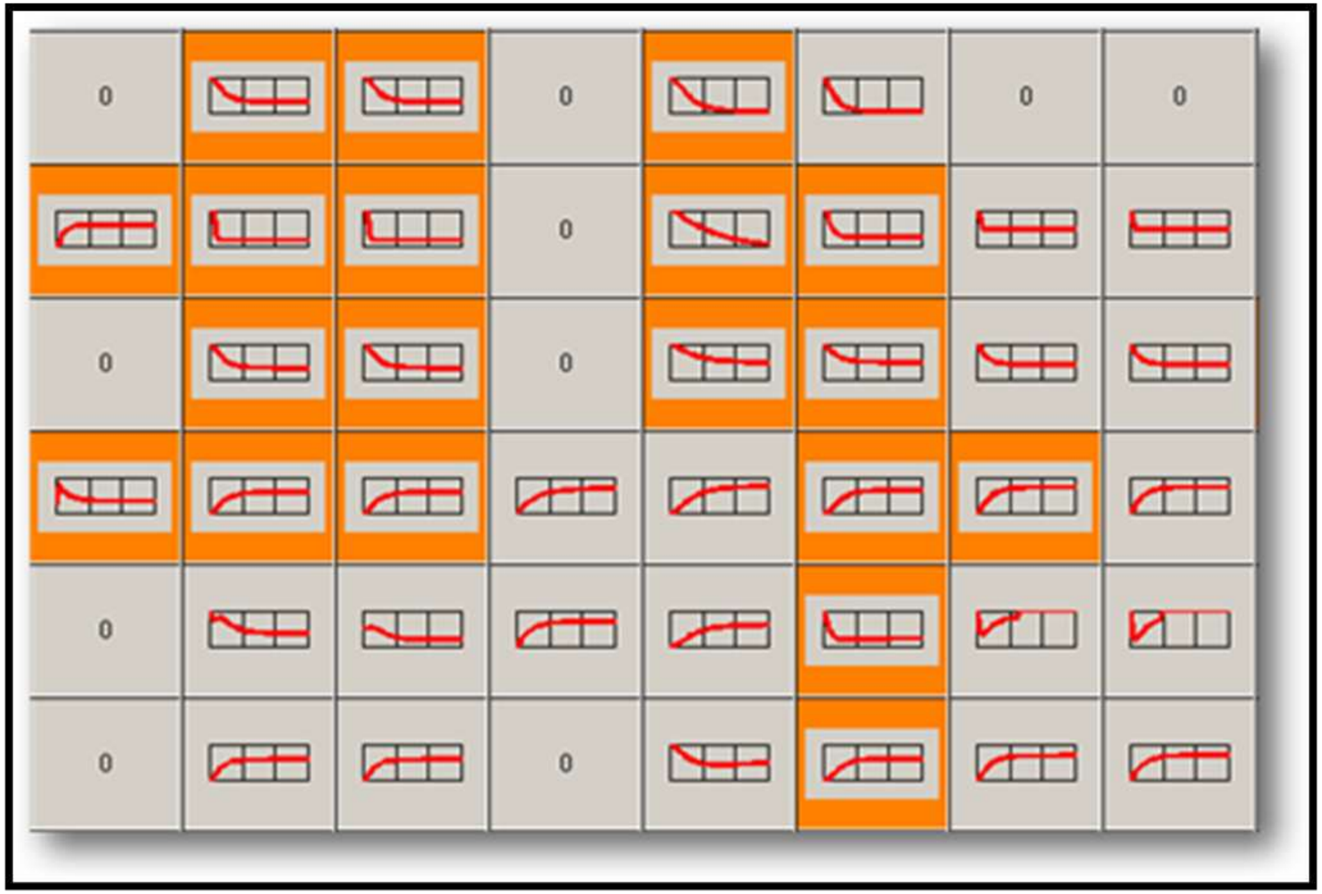

- The process is represented by a collection of dynamic response models
- The model matrix can have hundreds or more sub-models
- Well established implementation methodology

1980 – Model Predictive Control – Solving Problems

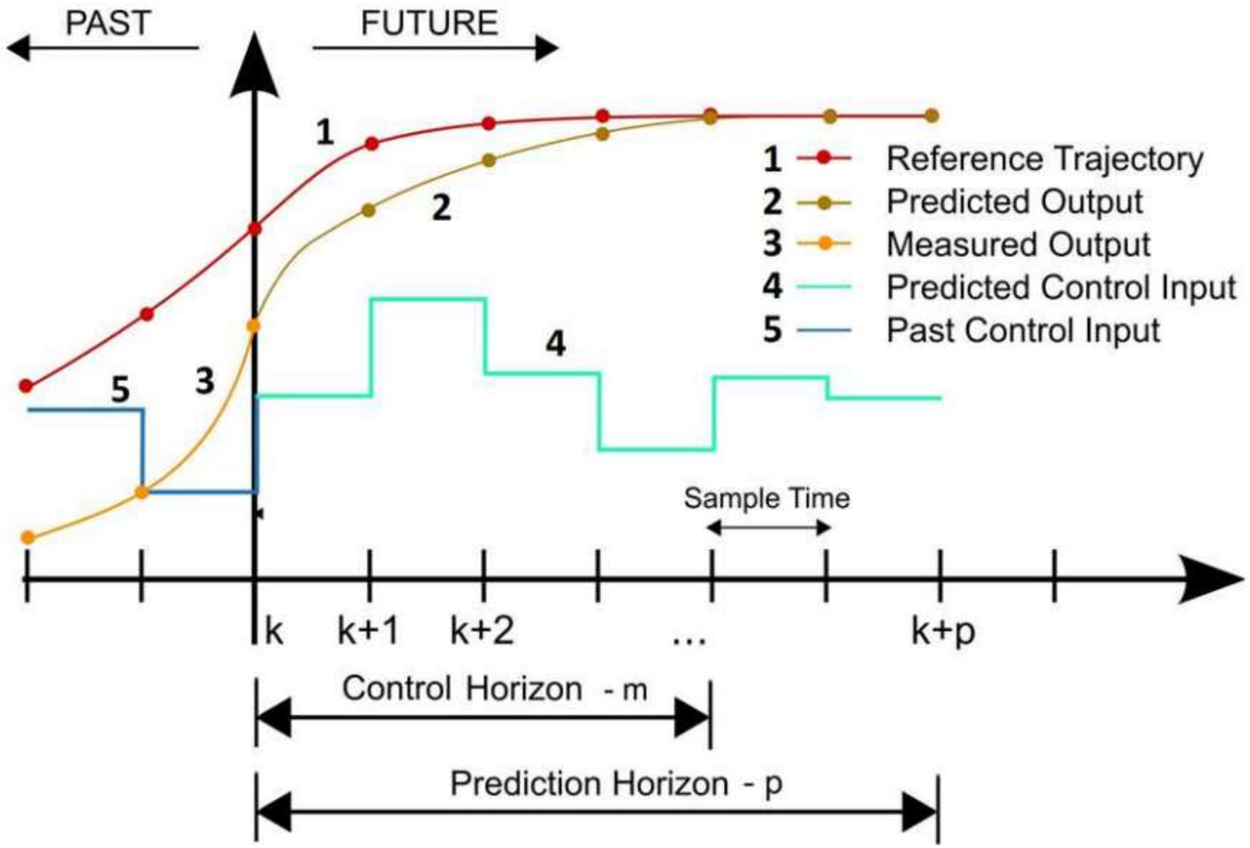

S N Andriyashin et al 2021 J. Phys.: Conf. Ser. 2032 012065

Bleach Plant Savings with MPC

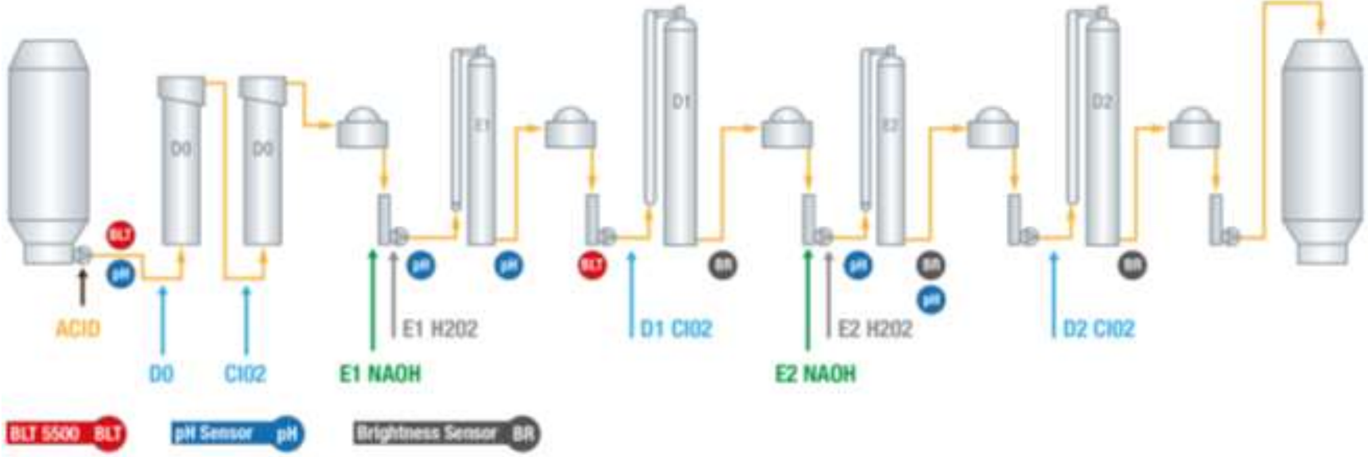

CVs	MVs									FFs						
	D0 CLO2	E1 NAOH	E1 H2O2	E1 O2	D1 NAOH	D1 CLO2	E2 NAOH	D2 NAOH	D2 CLO2	FF D0 BLT	FF D0 CLO2	FF E1 BLT K	FF E1 BLT BR	FF D1 CLO2	FF D2 BR IN	FF D2 ClO2
E1 BLT K#	-		-	-						+						
E1 BLT BR	+		+	+						-						
E1 PH		+									-					
D1 PH					+									-		
E2 PH							+							-		
D2 BR IN						+						-	+			
D2 PH								+								-
D2 BR OUT									+						+	

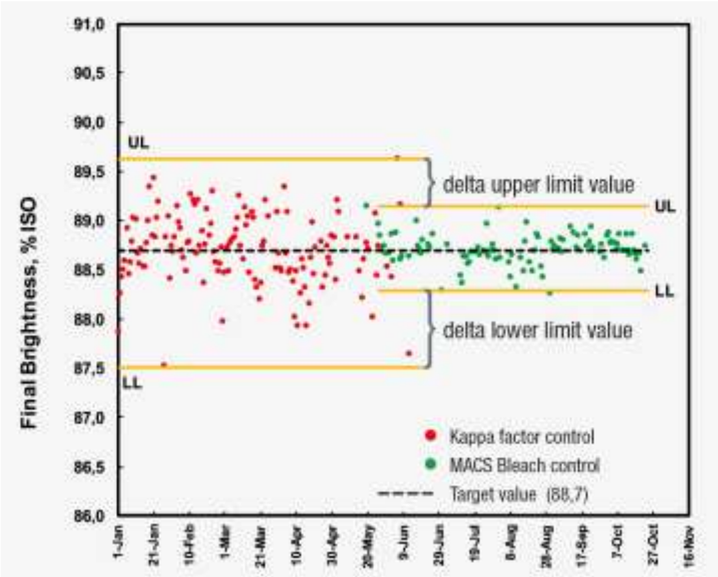

- Chemical use optimized and final brightness variability significantly reduced
- Lower variability enables target shift and additional chemical savings
- << 12-month ROI from bleach chemical savings

Proven Applications in Pulp & Paper

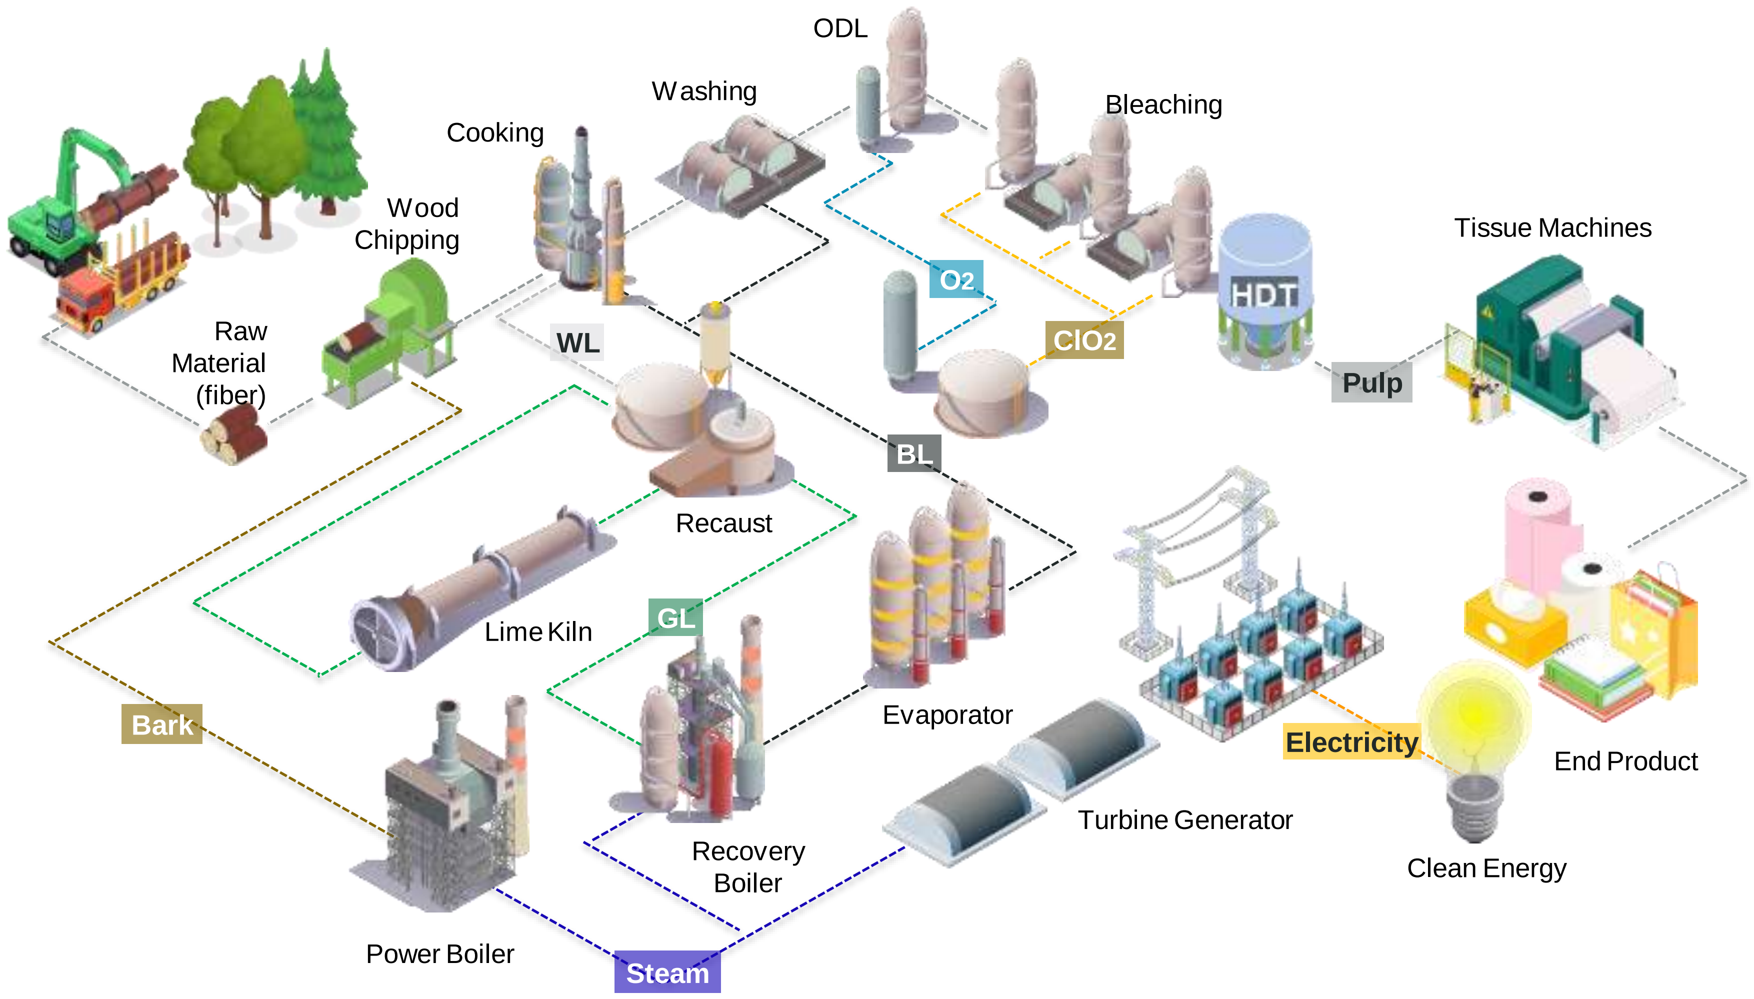

Path to Success – Why MPC works

- Built on base of instrumentation, data platforms, and technical capability
- Well established methodology
- Realistic understanding of required engineering effort
- Projects focused on ROI

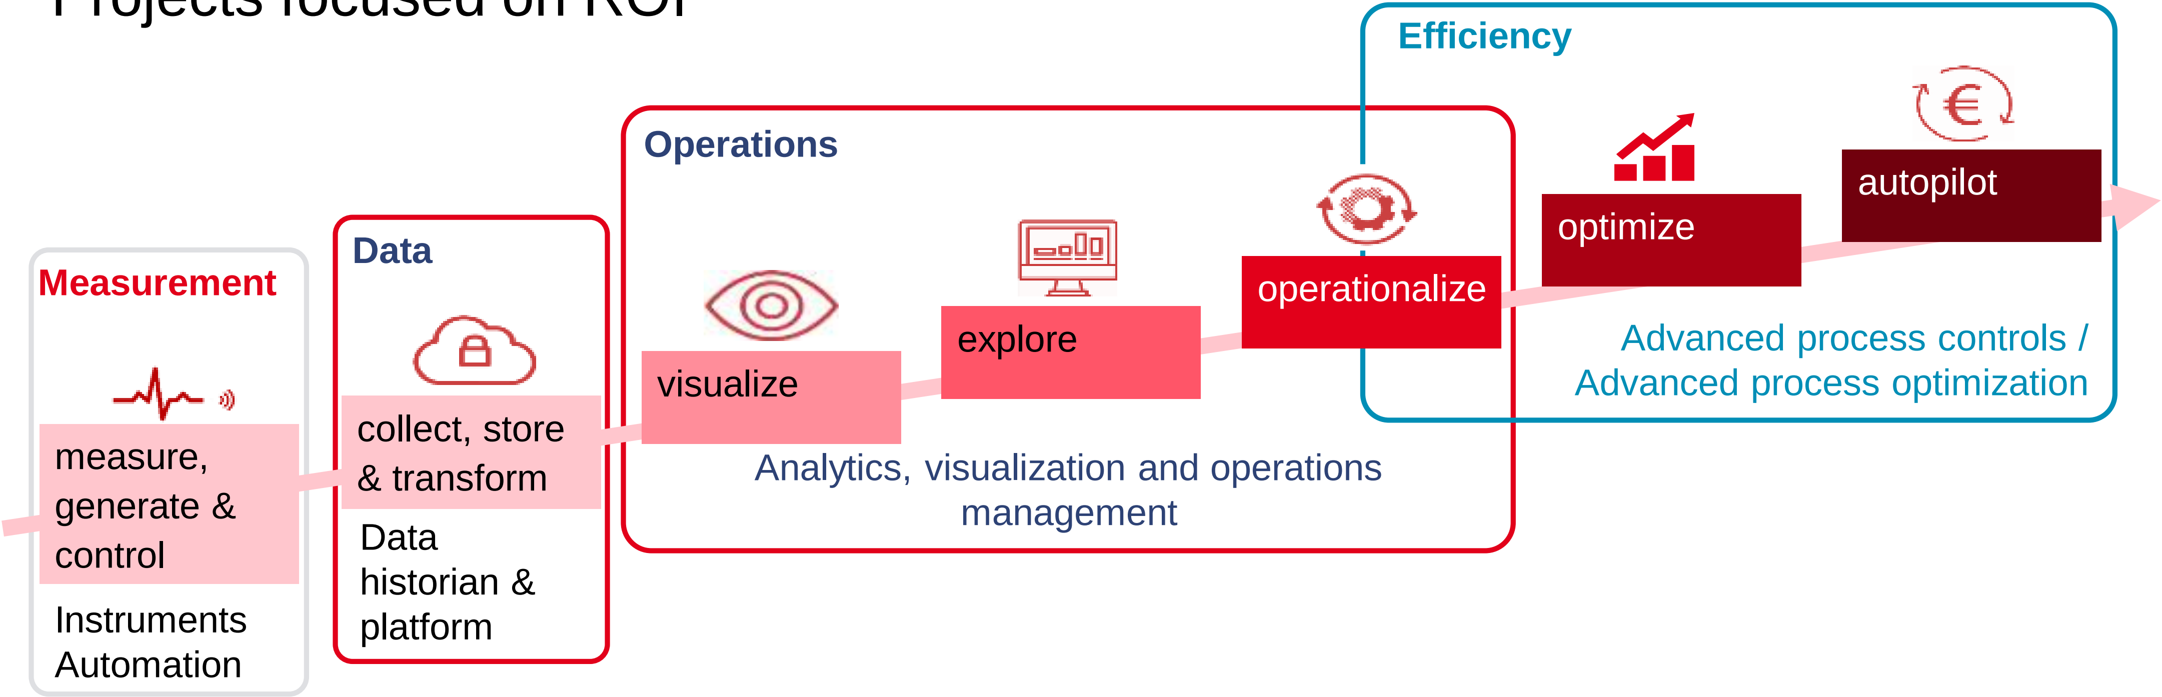

Making MPC Better with AI

- MPC algorithm “understands” how the process responds based on an embedded model

but….

- Algorithm does not include a methodology to “learn”
- Modern AI techniques hold promise to add dynamic “learning”

What is Reinforcement Learning?

Learning by Trial and Error

- A type of Machine Learning (ML) where an agent learns to make decisions in an environment
- The goal is to maximize cumulative reward over time
- Learns through interaction and feedback
- Think of it like training a dog: reward good behavior, discourage bad behavior

The Reinforcement Learning Loop
Interaction and Learning

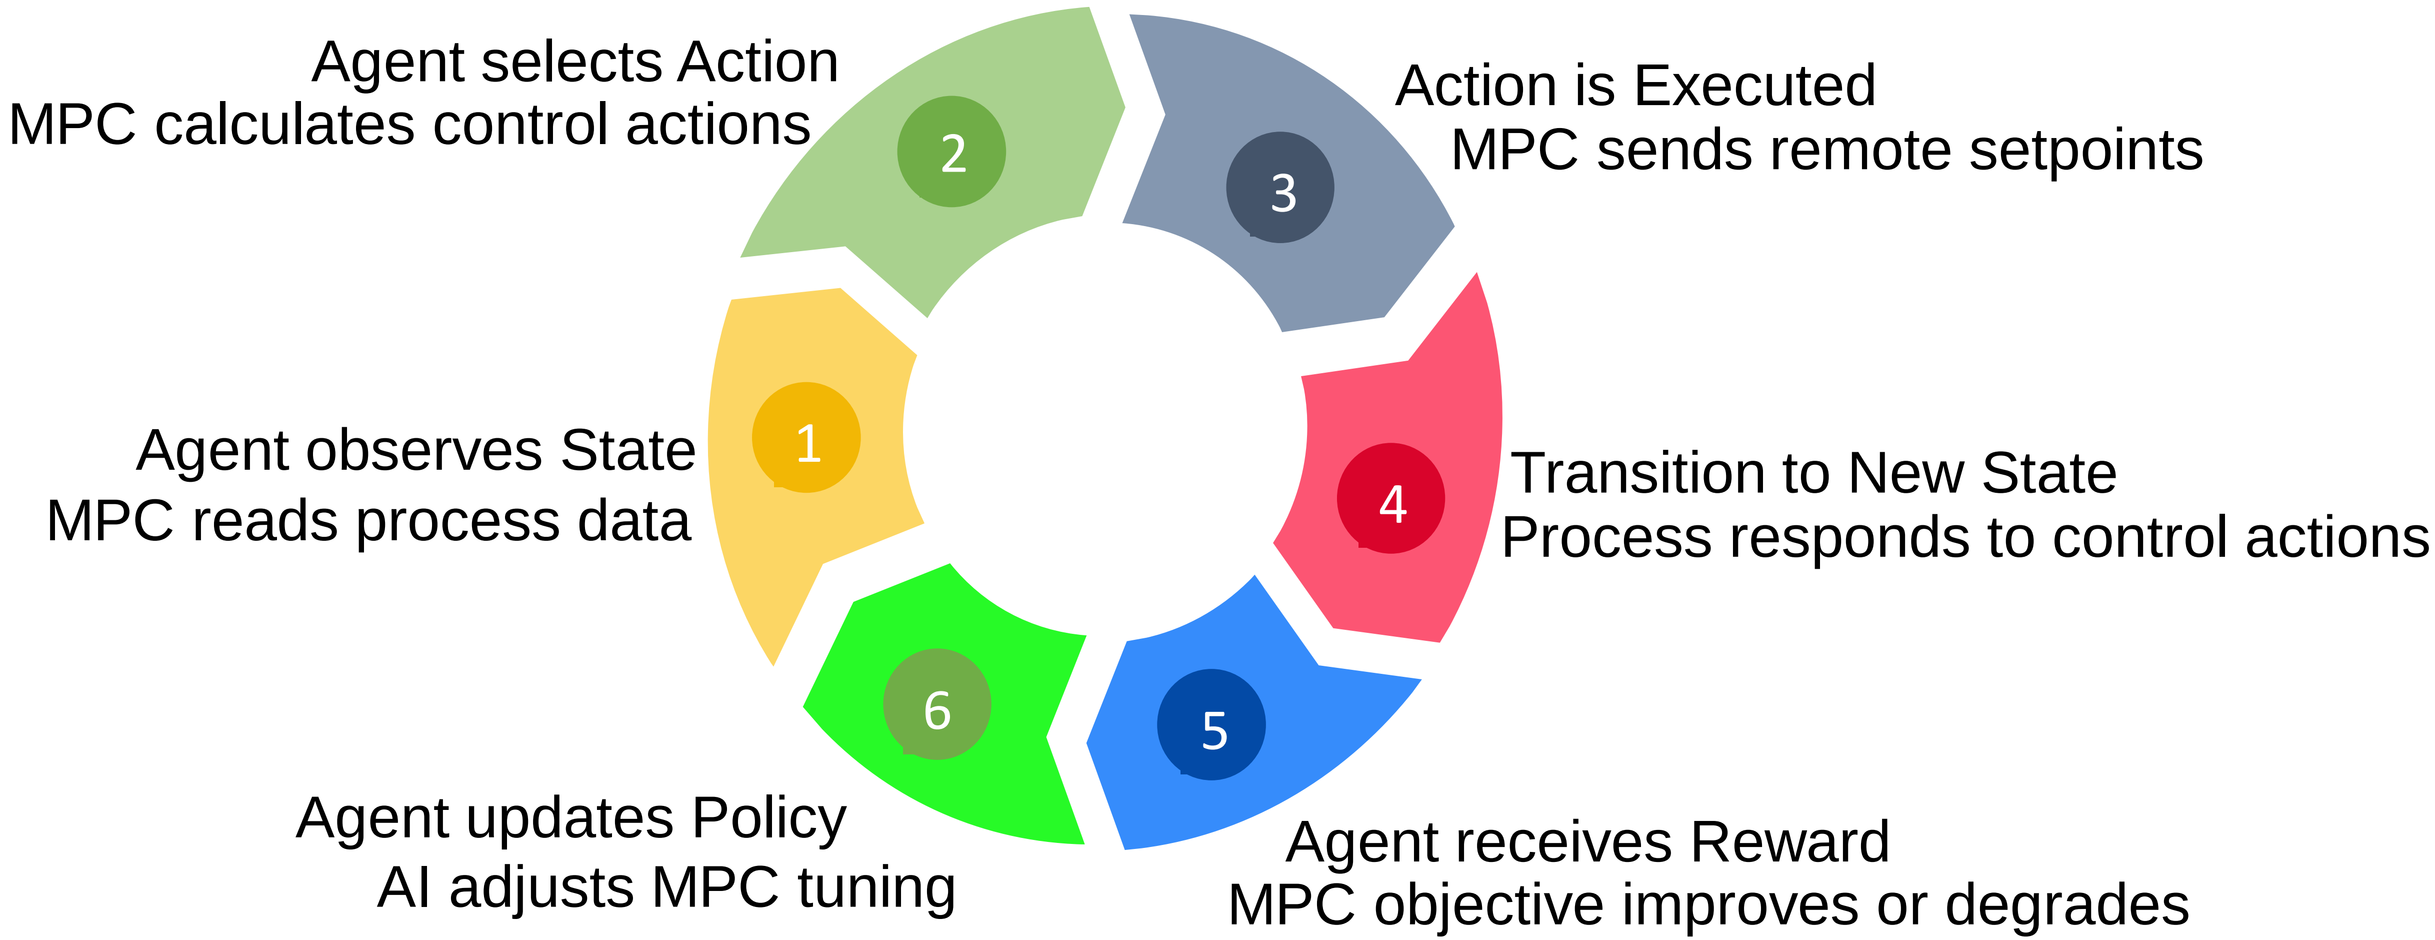

MPC - Reinforcement Learning (RL) Agent

Reward: RL can be used to optimize the MPC. The reward signal reflects how well the MPC is performing according to a higher-level objective. This could include:

- Tracking a desired setpoint
- Minimizing energy consumption
- Maintaining stability
- Optimizing throughput

Challenges

It’s Not Always Easy

- **Exploration vs. Exploitation**: Finding the right balance between trying new things and sticking with what works.
- **Reward Shaping**: Designing a good reward function is crucial but difficult. A bad reward function can lead to unintended consequences.
- **Generalization**: Learning to perform well in unseen situations.
- **Sample Efficiency**: RL algorithms can require a lot of data to learn effectively.

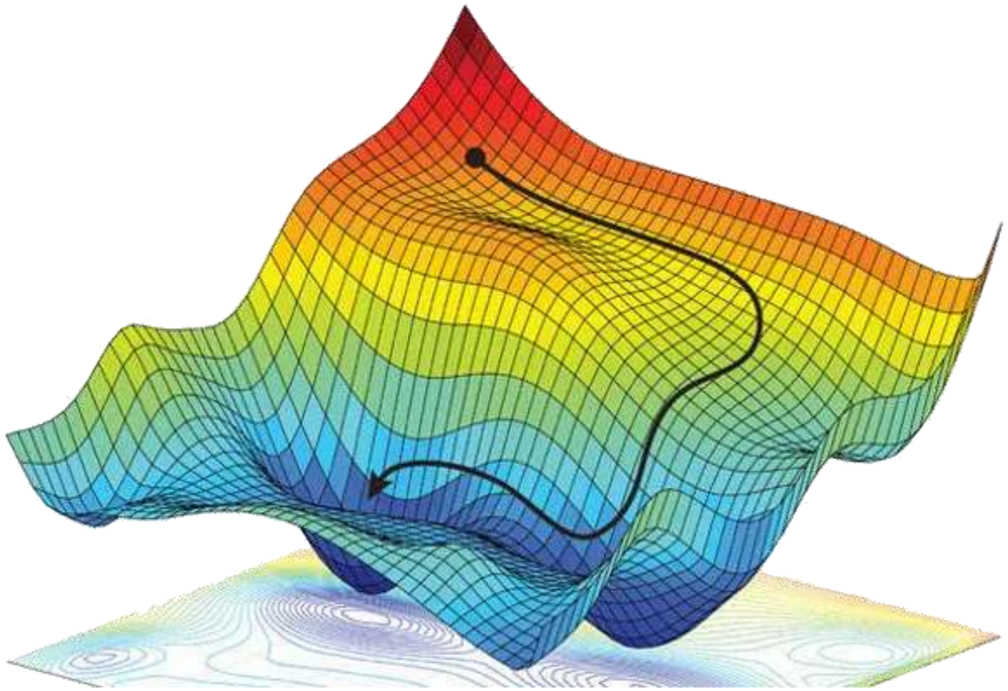

Conclusion – Take the "Next" Step on the Digital Journey

- Many "next steps" on the Digital Journey
- Digital technology available today can generate excellent ROI
- R&D underway to significantly extend the journey