Toward Self-Modifying Autonomous Business Process Systems (ABPs)

Achiya Elyasaf[1], **Andreas Metzger[2]**, Sebastian Sardina[3], Arik Senderovich[4], Estefanía Serral Asensio[5] ,Niek Tax[6]

[1] Department of Software and Information Systems Engineering, Ben-Gurion University of the Negev, Israel
[2] paluno – Ruhr Institute for Software Technology, University of Duisburg-Essen, Germany
[3] School of Computing Technologies, RMIT University, Melbourne, Australia
[4] School of Information Technology, York University, Toronto, Canada
[5] Research Centre for Information Systems Engineering (LIRIS), KU Leuven, Belgium
[6] Meta, London, United Kingdom

Presented at PMAI25 - 4th International Workshop on Process Management in the AI era, October 2025, Bologna, Italy, collocated with ECAI 2025

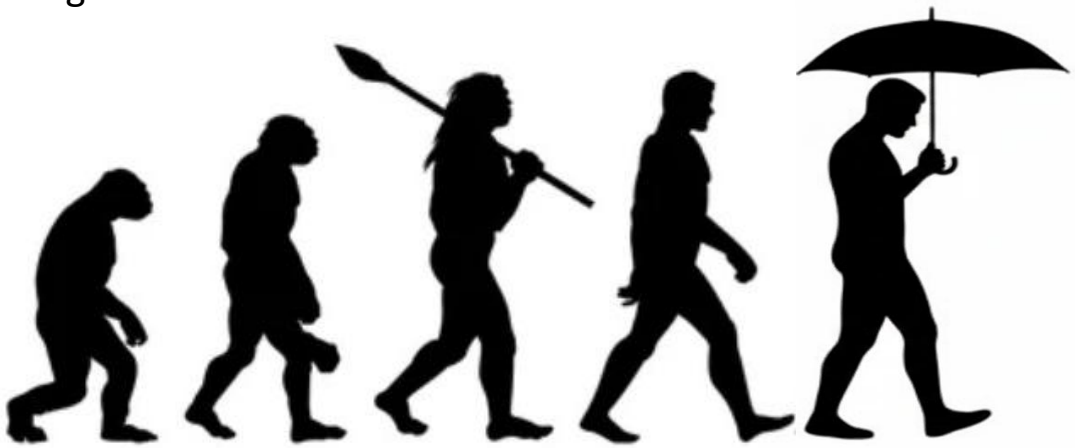

- Background and Motivation
- Taxonomy
- Challenges
- Conclusion

Motivation

Ideation @ Dagstuhl

Dagstuhl Seminar 25192

AUTOBIZ: Pushing the Boundaries of AI-Driven Process Execution and Adaptation

May 04 – May 09, 2025

https://www.dagstuhl.de/25192

Background

AI-augmented Business Processes

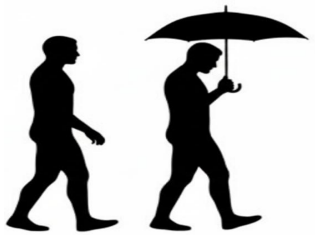

ADAPT

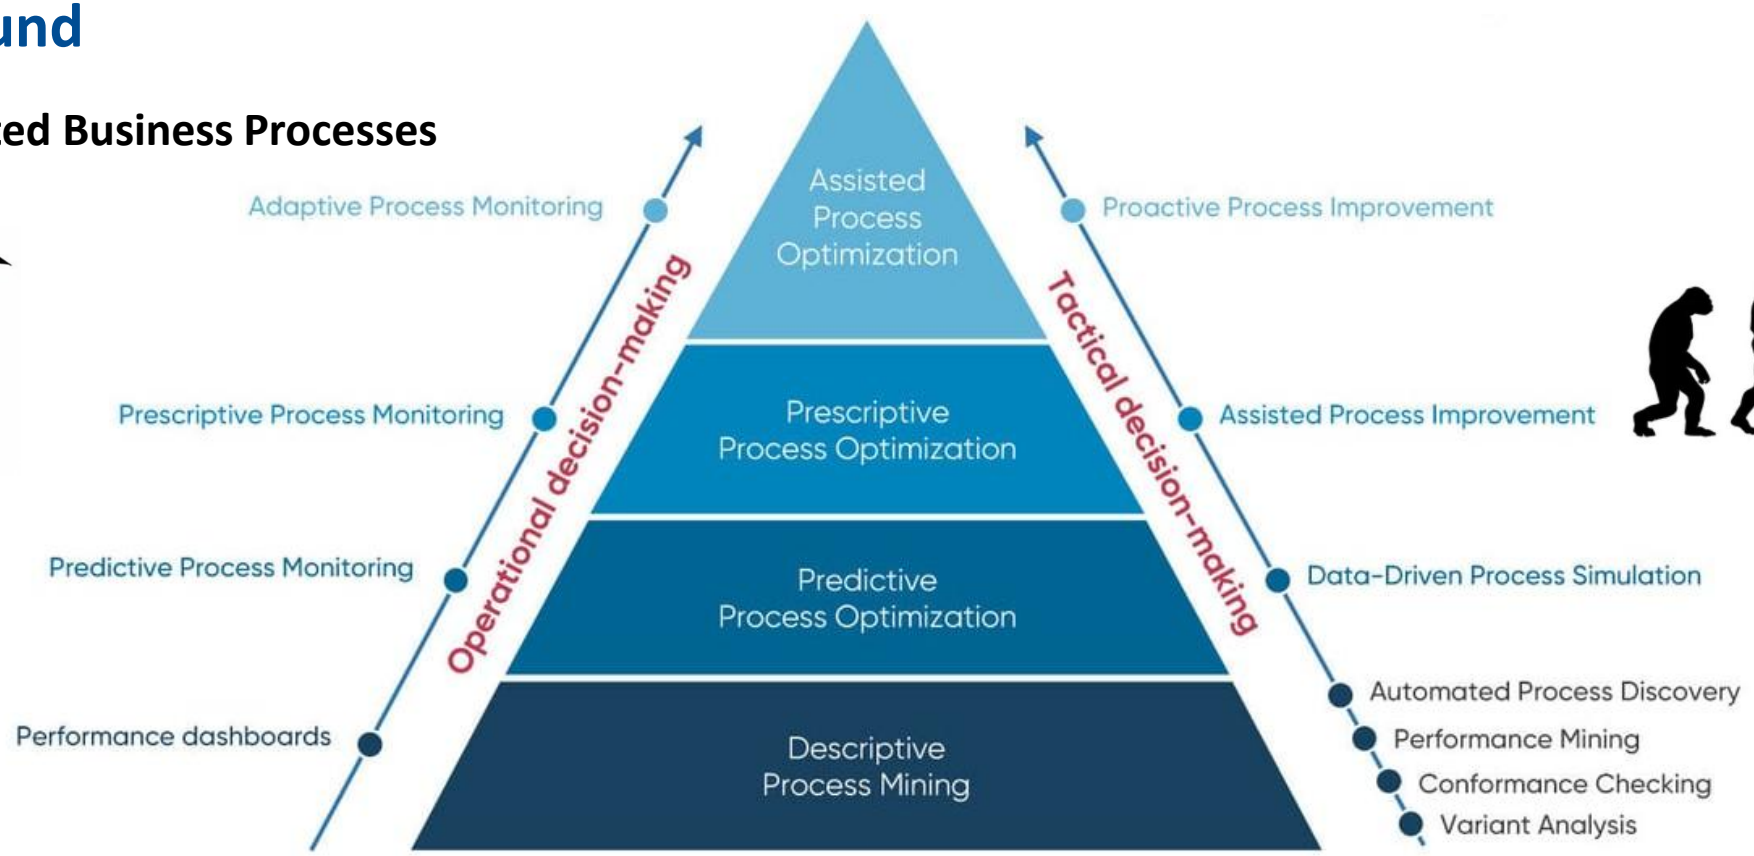

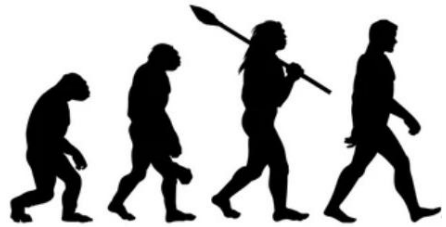

EVOLVE

Autonomous business processes := next generation of AI-Augmented business processes

- Self-executing, (pro)active, agentic, …
- Leveraging advanced AI/ML
- Operate with minimal to no human intervention

https://apromore.com/blog/enhancing-operational-excellence-with-augmented-business-process-management

Background

BPM Life-Cycle

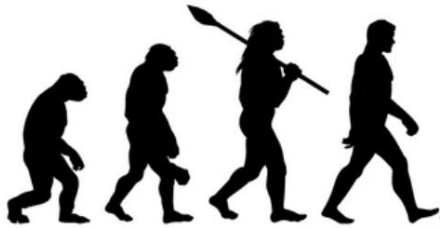

EVOLVE

- Identify tasks and process elements
- Define expected process outcomes
- Determine constraints

DESIGN
Conceptualize

MODEL
Visualize

- Use modelling language to „formalize“ process
- Map task sequences
- Identify bottlenecks

INSTANTIATE
Map

- Map tasks to resources (in an optimal fashion)
- Consider actual process environment
- Plan sequence of tasks

EXECUTE
Enact

- Put into action
- Manual + automated task execution
- Deliver value

MONITOR
Track KPIs

- Collect KPI data
- Track and predict process quality
- Predictive monitoring

MODIFY
Optimize

- Continuous improvement
- Make adjustments to model or running process
- Prescriptive monitoring

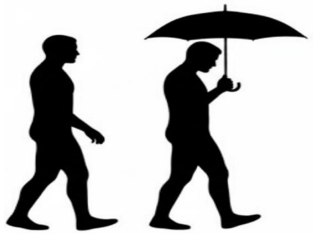

ADAPT

Motivational Example

The Automated Warehouse

Scenario: A large company operates an automated warehouse

Process: Fleets of robots retrieve and transport shelves (pods) to human pickers

Problem: During the busy Christmas season, a robot malfunctions, crashes, and thus blocks an aisle

→ This single event may trigger two distinct types of modification...

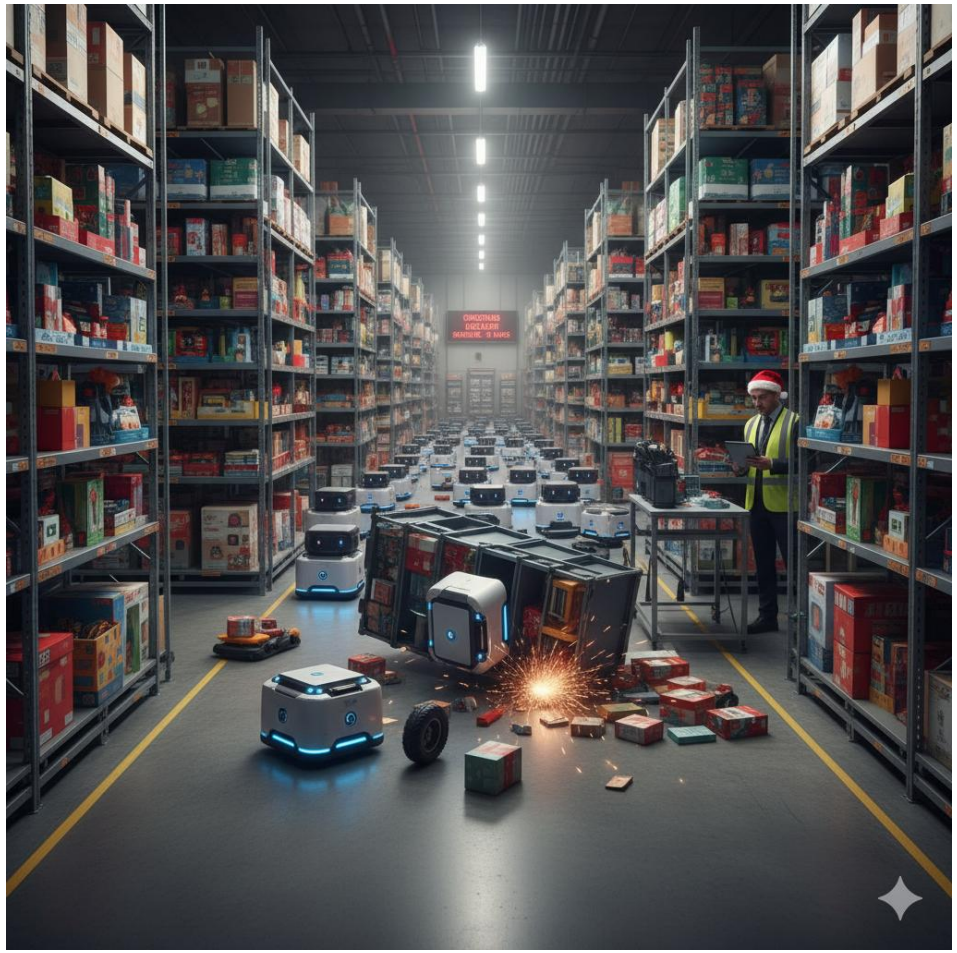

Motivational Example

UNIVERSITÄT DUISBURG ESSEN
Offen im Denken

The Automated Warehouse

Immediate Response:

- Nearby robots instantly **reroute** their paths to avoid the blocked aisle
- The pending order is **reassigned** to a functioning robot
- Human workers receive updated instructions

→ This is an ADAPTATION:

- A **short-term, instance-specific** modification
- It addresses an immediate, unforeseen issue
- It does **not** alter the underlying process model

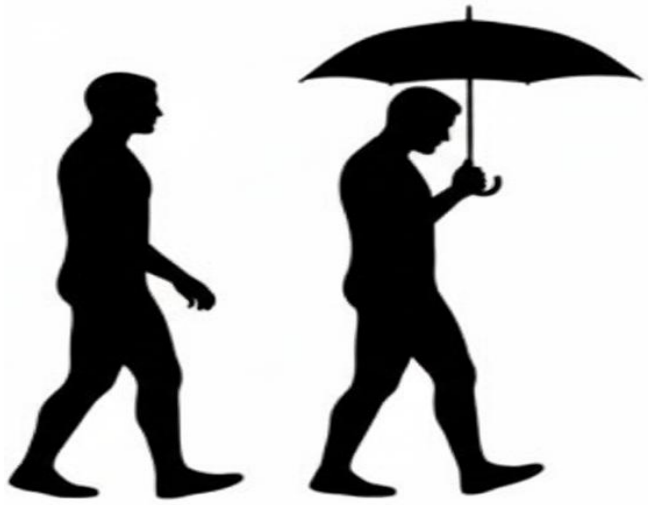

ADAPT

Motivational Example

The Automated Warehouse

Longer-term Analysis:

- The warehouse system logs the failure event for review
- Engineers (or the system) discover a **pattern**: similar failures occur after ~1,000 picks

- A **new maintenance rule** is introduced: "Mandate preemptive inspection after every 900 operations"

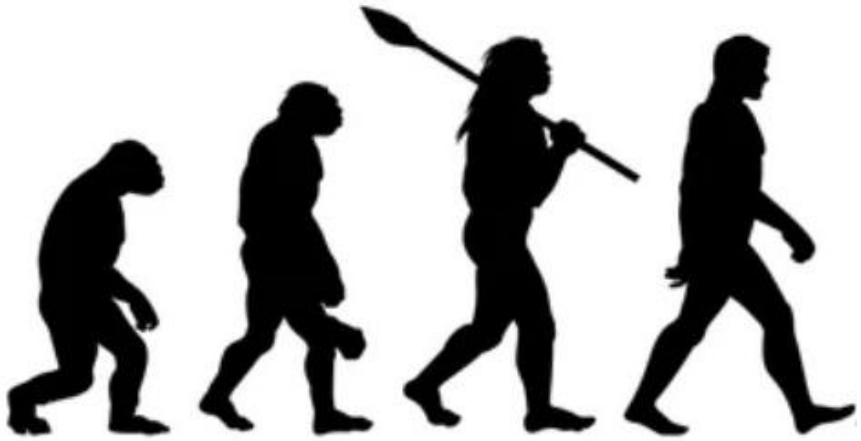

EVOLVE

→ This is an EVOLUTION:

- A **long-term modification** of the process logic/model
- It's informed by aggregated insights and patterns over time
- It affects **all future instances** of the process

- Background and Motivation
- Taxonomy
- Challenges
- Conclusion

Taxonomy

Overview

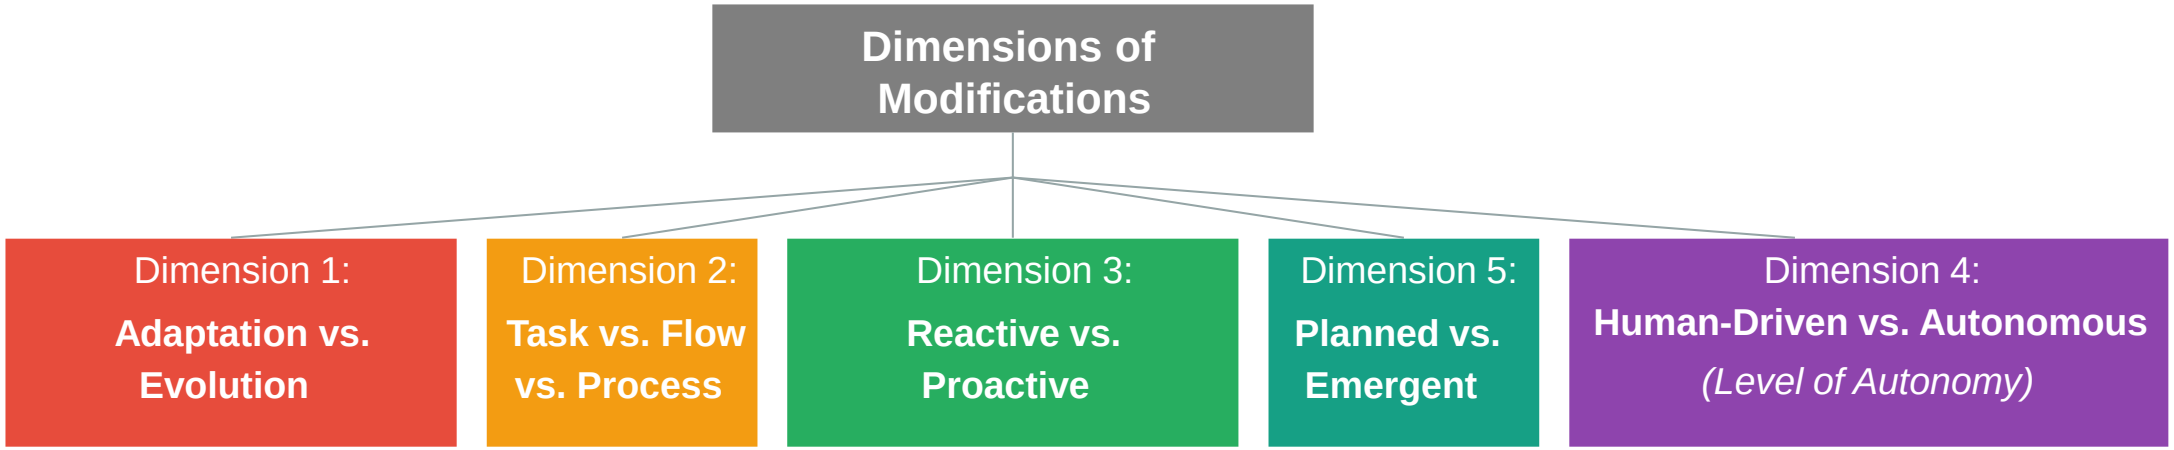

Taxonomy

Dimension 1: **Adaptation vs. Evolution**	Dimension 2: **Task vs. Flow vs. Process**	Dimension 3: **Reactive vs. Proactive**	Dimension 5: **Planned vs. Emergent**	Dimension 4: **Human-Driven vs. Autonomous** *(Level of Autonomy)*

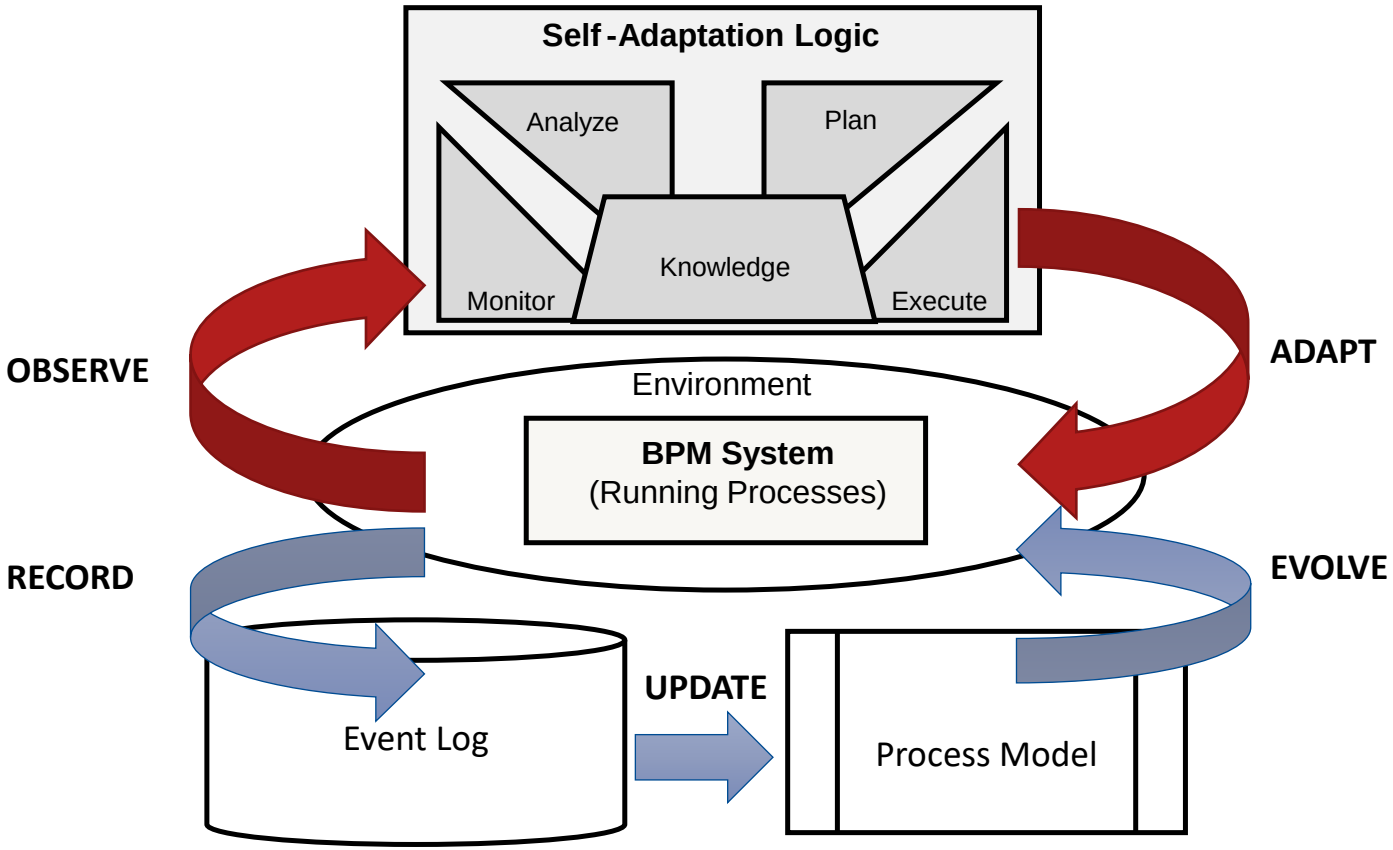

Taxonomy

Dimension 1: **Adaptation vs. Evolution**	**Dimension 2: Task vs. Flow vs. Process**	Dimension 3: **Reactive vs. Proactive**	Dimension 5: **Planned vs. Emergent**	Dimension 4: **Human-Driven vs. Autonomous** *(Level of Autonomy)*

Defines *scope* and *impact* of modification

Task-level: Modifies how *individual tasks* are performed

- *Example:* Changing a task's duration, logic, or resource assignment

Control Flow-level: Adjusts the *routing or sequencing* of tasks

- *Example:* Rerouting orders, skipping tasks (like the warehouse robots)

Process-level: Alters the *entire process structure* or its central resources

- *Example:* Introducing new roles, changing coordination logic, or revising the maintenance schedule

Taxonomy

Dimension 1: **Adaptation vs. Evolution**	Dimension 2: **Task vs. Flow vs. Process**	Dimension 3: **Reactive vs. Proactive**	Dimension 5: **Planned vs. Emergent**	Dimension 4: **Human-Driven vs. Autonomous** *(Level of Autonomy)*

Defines *trigger* of modification

Reactive: Modification in response to event(s) or failure(s)

- *Examples:* robot malfunction, process bottleneck, …

Proactive: Modification based on forecasts or predictions

- *Examples:* scheduling maintenance *before* a predicted failure, updating process model before release of AI Act, …

Taxonomy

Dimension 1: **Adaptation vs. Evolution**	Dimension 2: **Task vs. Flow vs. Process**	Dimension 3: **Reactive vs. Proactive**	Dimension 5: **Planned vs. Emergent**	Dimension 4: **Human-Driven vs. Autonomous** *(Level of Autonomy)*

Defines *reason* for modification

Planned: Deliberate, top-down redesign

- **Enhansive**
  - *Example:* capturing new customer requirements in the business process
- **Adaptive** (≠ Adaptation)
  - *Example:* rolling out a new process model to capture AI Act

Substructure based on different types of maintenance in software engineering

Emergent: Arises bottom-up from past process observations

- **Corrective**
  - *Example:* process that always adapts due to failed task → change task/process model
- **Perfective**
  - Example: improving process performance

Taxonomy

Dimension 1: **Adaptation vs. Evolution**

Dimension 2: **Task vs. Flow vs. Process**

Dimension 3: **Reactive vs. Proactive**

Dimension 5: **Planned vs. Emergent**

Dimension 4: **Human-Driven vs. Autonomous** *(Level of Autonomy)*

Defines *autonomy level* of modification, such as

Level 0: No Automation
- All tasks are fully manual
- Users detect, decide, and implement the modification

Level 1: Process Assistance
- ABPs provide recommendations or highlight anomalies (e.g., predictive & prescriptive monitoring)

Level 2: Partial Autonomy
- ABPs independently execute isolated tasks within predefined boundaries

Level 3: Contextual Autonomy
- ABPs autonomously perform most tasks and orchestrate flows to achieve overall foals
- Human intervention only in exceptional cases

- Background and Motivation
- Taxonomy
- Challenges
- Conclusion

Challenges

1: Governance, Oversight, and Human Interaction

When to handover control?

- When to shift control between ABPs and humans?
  - → *learning to defer*

How to validate?

- How can user validation be incorporated into real-time adaptation without bottle-necking autonomy of ABPs?
  - → *explainable ABPs*

How to align goals?

- How can ABPSs optimize multiple objectives while remaining within formal and ethical constraints?

How to ensure quality?

- What are the data quality and coverage thresholds for safe, autonomous decision-making in ABPs?
- How to evaluate the success or failure of modifications when human validation is unavailable?

Challenges

2: Continuous Learning and Adaptation Management

Learning from Experience:

- How can ABPSs continuously record adaptations and assess their effectiveness over time?

Safe Generalization:

- What metrics and techniques enable an ABPS to safely generalize learned behavior across varying contexts?
- How to avoid overgeneralization from too few instances?

Bounded Knowledge:

- Retaining or analyzing complete history of long-running or high-frequency processes might not be feasible
- How can an ABPs maintain bounded knowledge representations (e.g., summaries, sliding windows)?

Challenges

3: Modeling and Measuring Uncertainty

Differentiating Uncertainty:

- How can ABPs differentiate / handle epistemic and aleatoric uncertainty during execution?
  - Aleatoric: Inherent randomness
    (e.g., "how long will this task take?").
  - Epistemic: Lack of knowledge
    (e.g., "what will happen if I try this new route?").

Quantifying Uncertainty:

- How can ABPs combine qualitative and quantitative uncertainty metrics for robust decision-making?
  - Knowing the level and type of uncertainty is key to deciding when to act vs. when to defer to a human.

Communicating Uncertainty:

- What are effective representations of uncertainty (e.g., using probabilistic or fuzzy paradigms)?
- How should ABPs communicate uncertainty and associated risk to users in a transparent and actionable way?

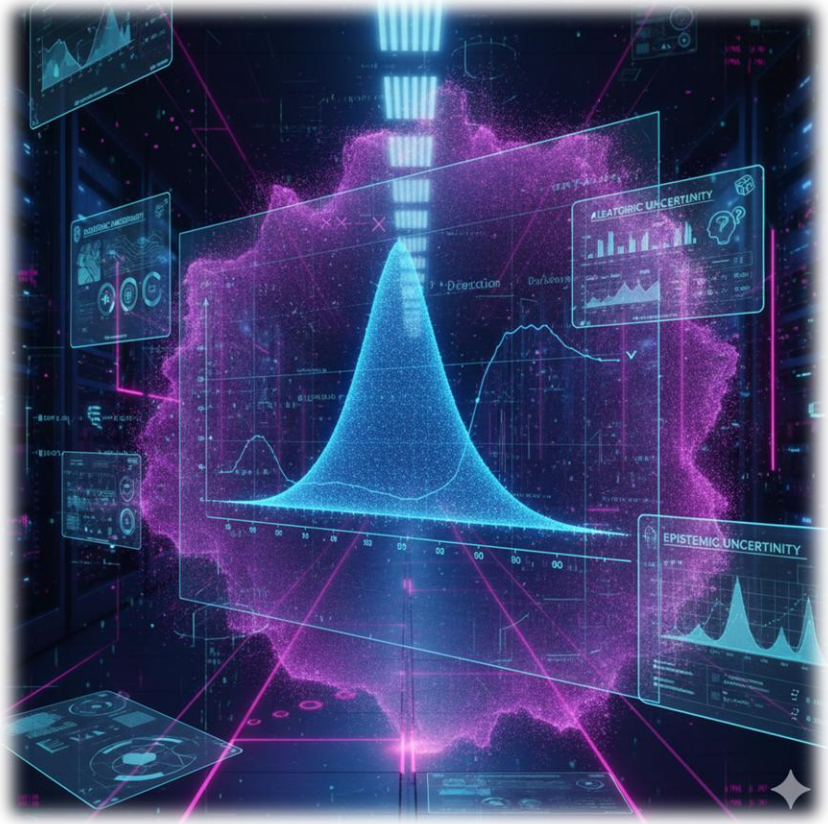

- Background and Motivation
- Taxonomy
- Challenges
- Conclusion

Conclusion

We defined: Self-modification in ABPs, distinguishing **adaptation** (short-term) from **evolution** (long-term)

We proposed: Taxonomy for modifications in ABPs

We identified 3 groups of core research challenges:

1. Governance & Human Oversight
2. Continuous Learning & Adaptation Management
3. Uncertainty Modeling & Communication

Future Vision: ABPs as

- Collaborative, transparent agents
- Balancing autonomy with accountability,
- Integrating AI, ML, XAI, and process mining