Natural Resources Canada | Ressources naturelles Canada

AI and Enterprise Business Processes + NL-Driven Enterprise

Alexei Lapouchnian, PhD, Manager, Digital Accelerator, Natural Resources Canada

Canada

AI and the Enterprise

- Based on research conducted at the University of Toronto in partnership with a large technology company
- AI and the Enterprise
  - How do we introduce AI into the IT landscape of an enterprise?
  - *Specifically*: How do we inject AI into enterprise business processes (BPs)?
  - *More specifically*: How do we augment decisions made as part of BPs with AI?

AI's Adoption in Enterprises: Some Issues

- Adoption and operation requires specialized skills (e.g., in ML), user training, etc.
- Where will AI create value in organizations and where should it not be deployed?
- How do we integrate AI-driven systems into existing processes?
- How do we ensure that black-box AI systems comply with various regulations? Audit trails may be needed!
- How do we guard against unintended consequences of using AI systems that may have financial/reputational damage?
- How do we make sure that AI systems are trusted and accepted within the organization, by both management and individual users?

Improving Adoption of AI in Enterprises

- Understand and manage complexity
- Build and maintain trust in AI
  - Transparency (when needed/desired)
  - Ability to provide explanations/justifications for AI's recommendation to users
  - Clear evolution paths for AI-driven systems and user engagement with them
- Customized/personalized, context-aware interactions

AI-Enhanced Business Processes

- Integration of AI into business processes
  - BPs – typical way organizations define work to be done
- BPs contain decisions (and activities)
  - E.g., whether to approve a loan or an insurance claim
- Currently – most decisions are done by humans
- With the help of AI-based systems
  - Introduce AI Advisors to support human decision making
  - Many options for such systems: from gathering/filtering relevant information for humans to supporting fully automated decisions

AI and Business Processes – Complexity

- Introduction of AI affects:
  - The process containing the decisions aided by AI
  - Other processes are introduced or are impacted as well, e.g.:
    - Design of the AI-based system, data collection, model training, etc.
    - Monitoring of decision/recommendation outcomes and their effects on *business goals*
    - Context monitoring for situation awareness
    - **User engagement** with AI-driven systems
    - Etc.
- AI and organizational BPs both evolve
  - BPs: due to changes in requirements, business context, etc.
  - Systems: increased sophistication, changes in recommendation quality

Business Process Architectures

- Need to model and analyze multiple related BPs – i.e., need to focus on **BP Architectures** (BPAs)
  - Business Process Architecture - model showing BPs and their relationships
  - We use a previously proposed notation called *hiBPM*

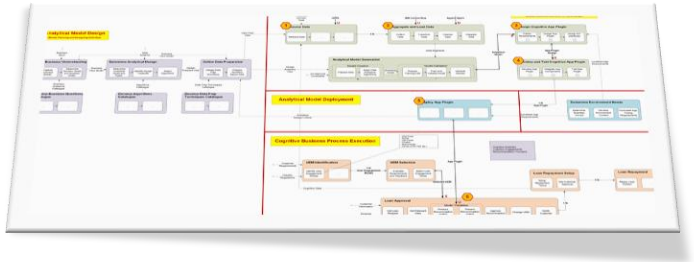

Business Process Architectures

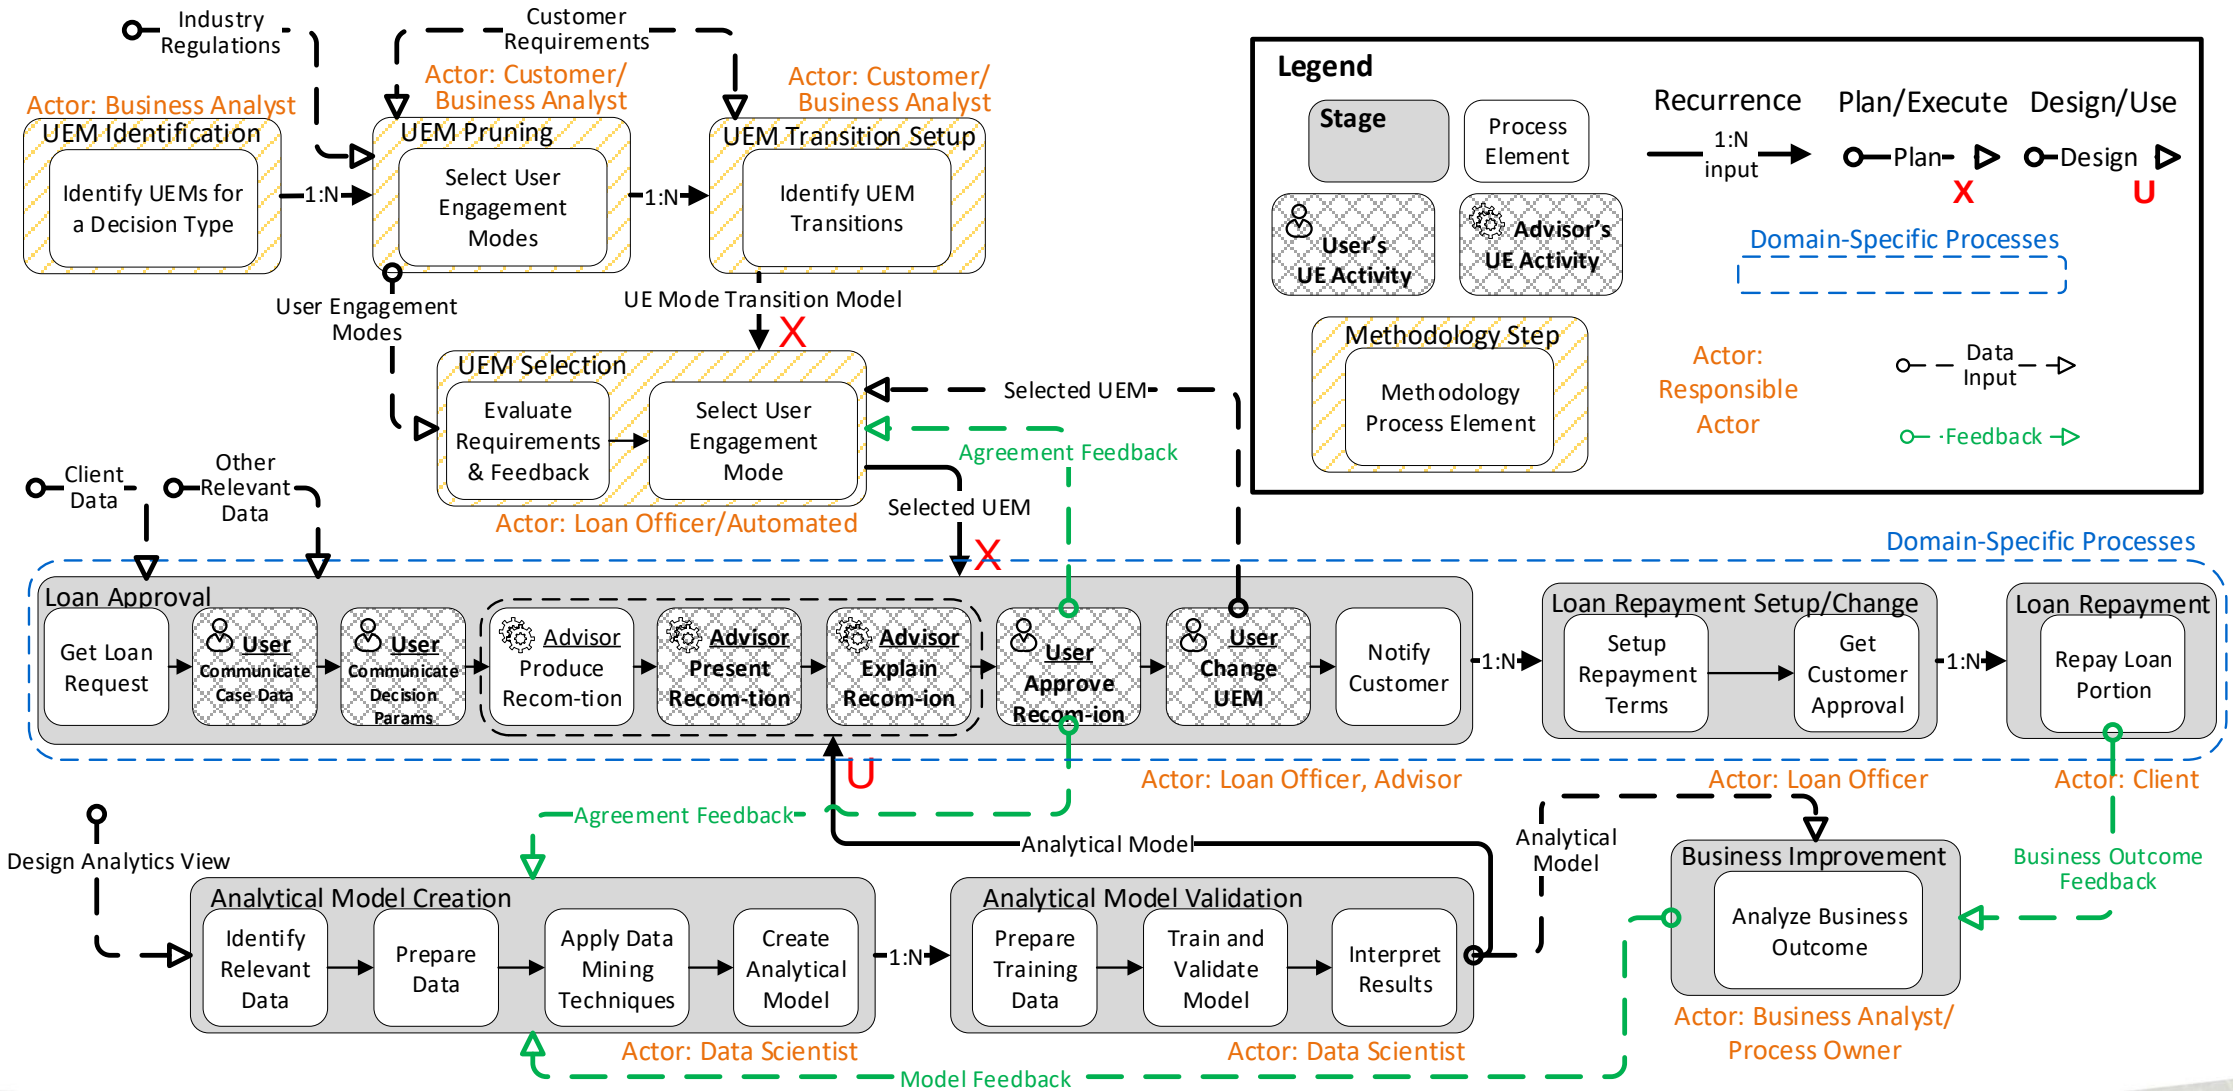

Business Process Architectures

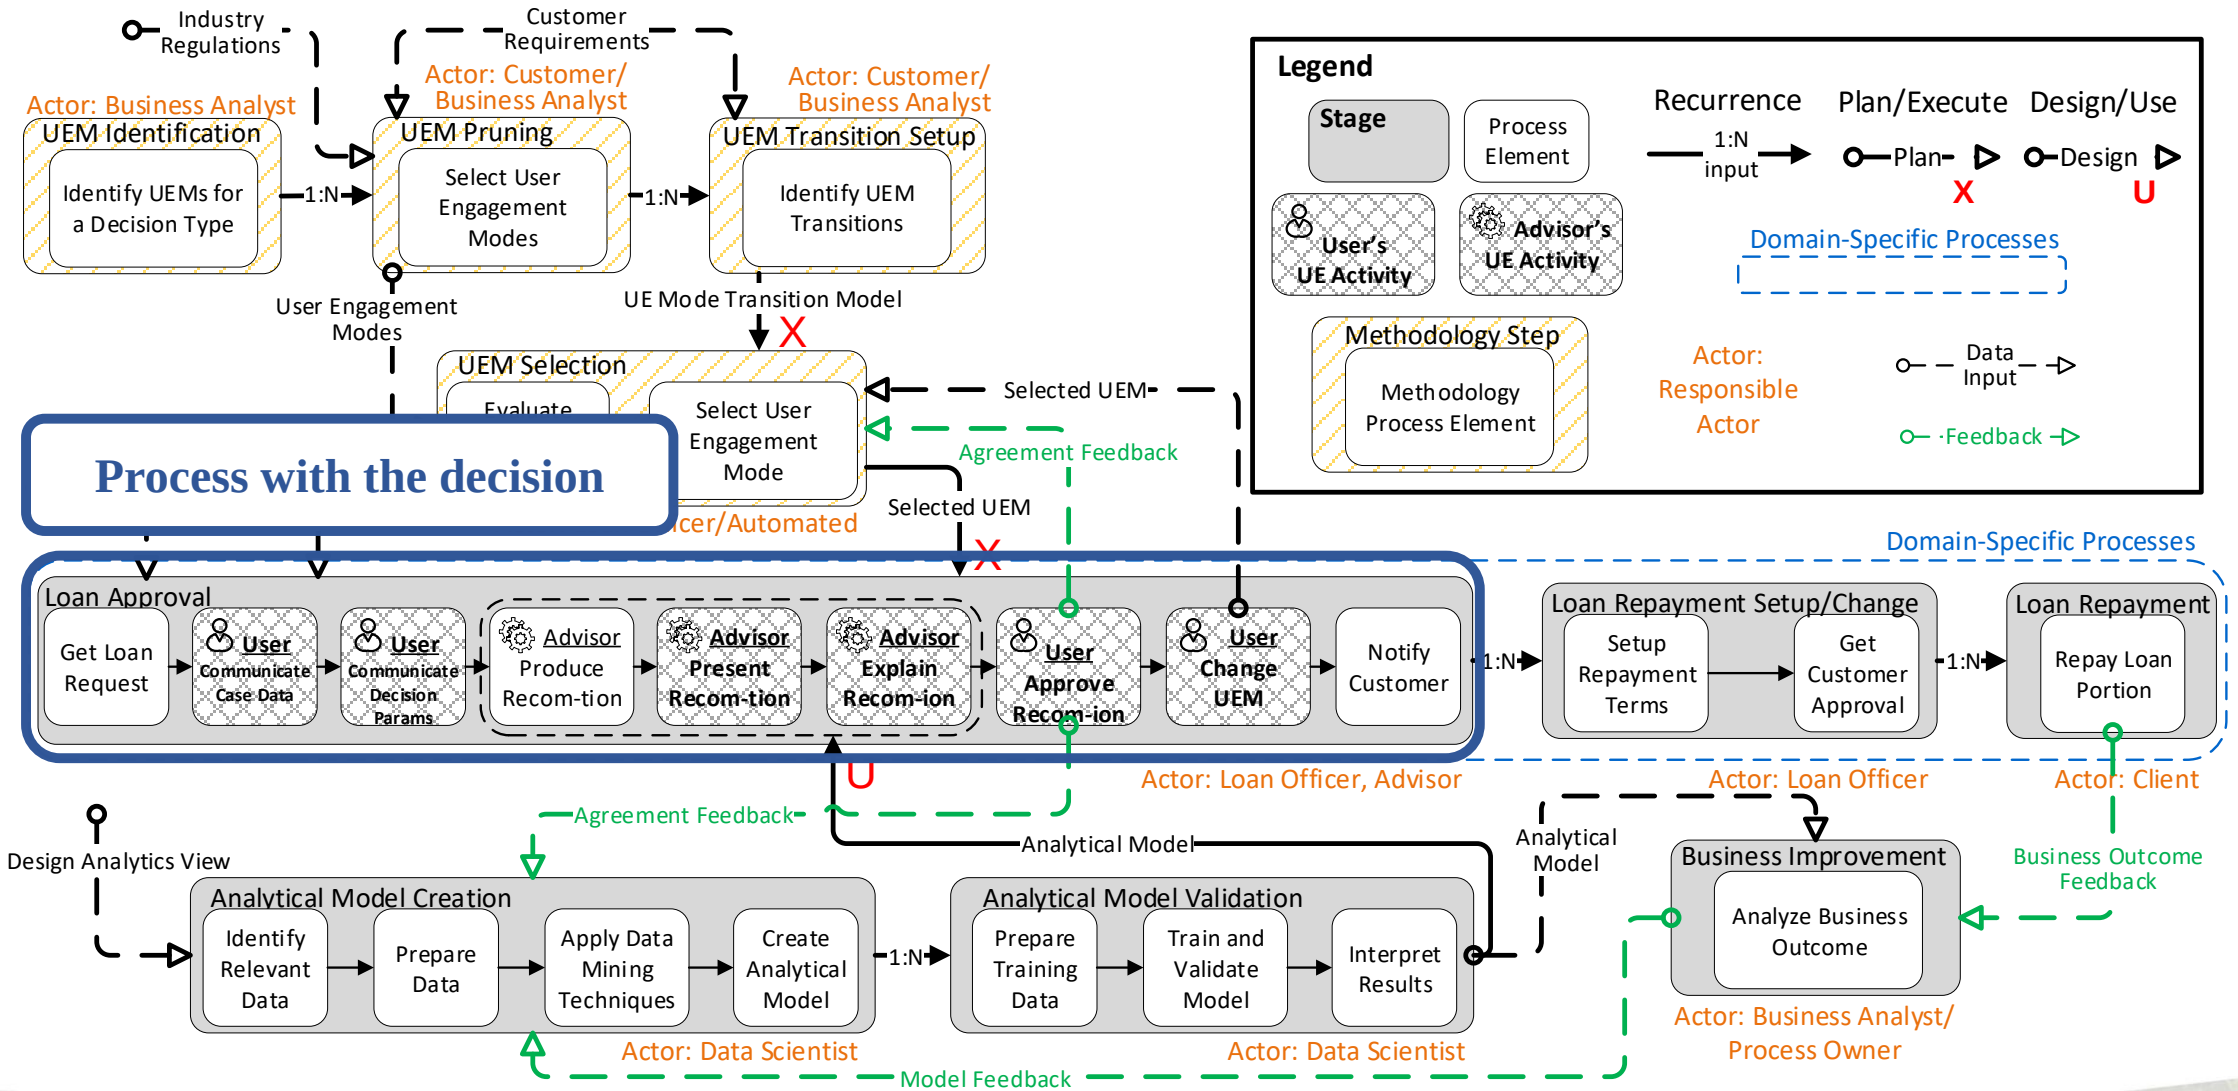

Business Process Architectures

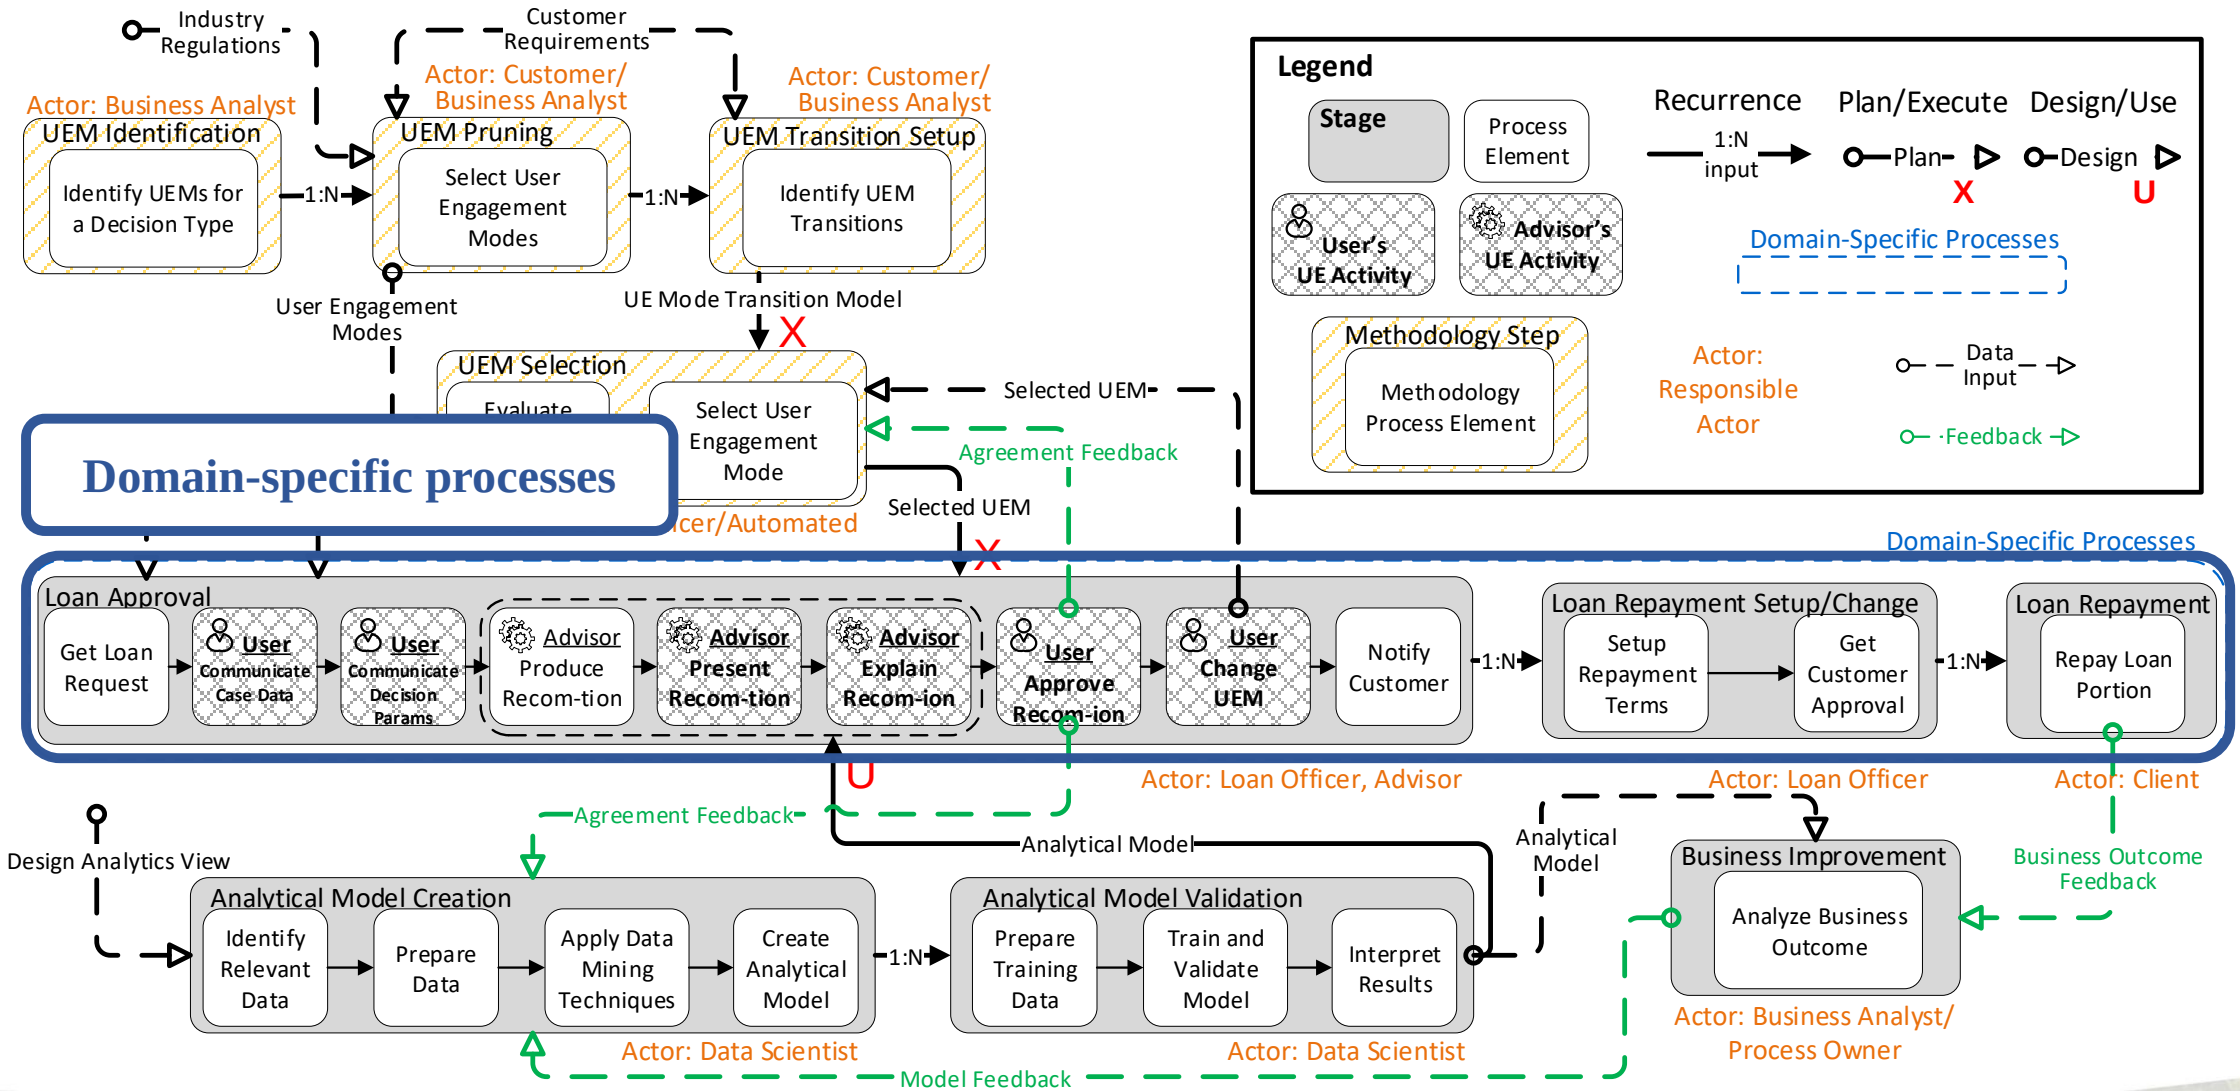

Business Process Architectures

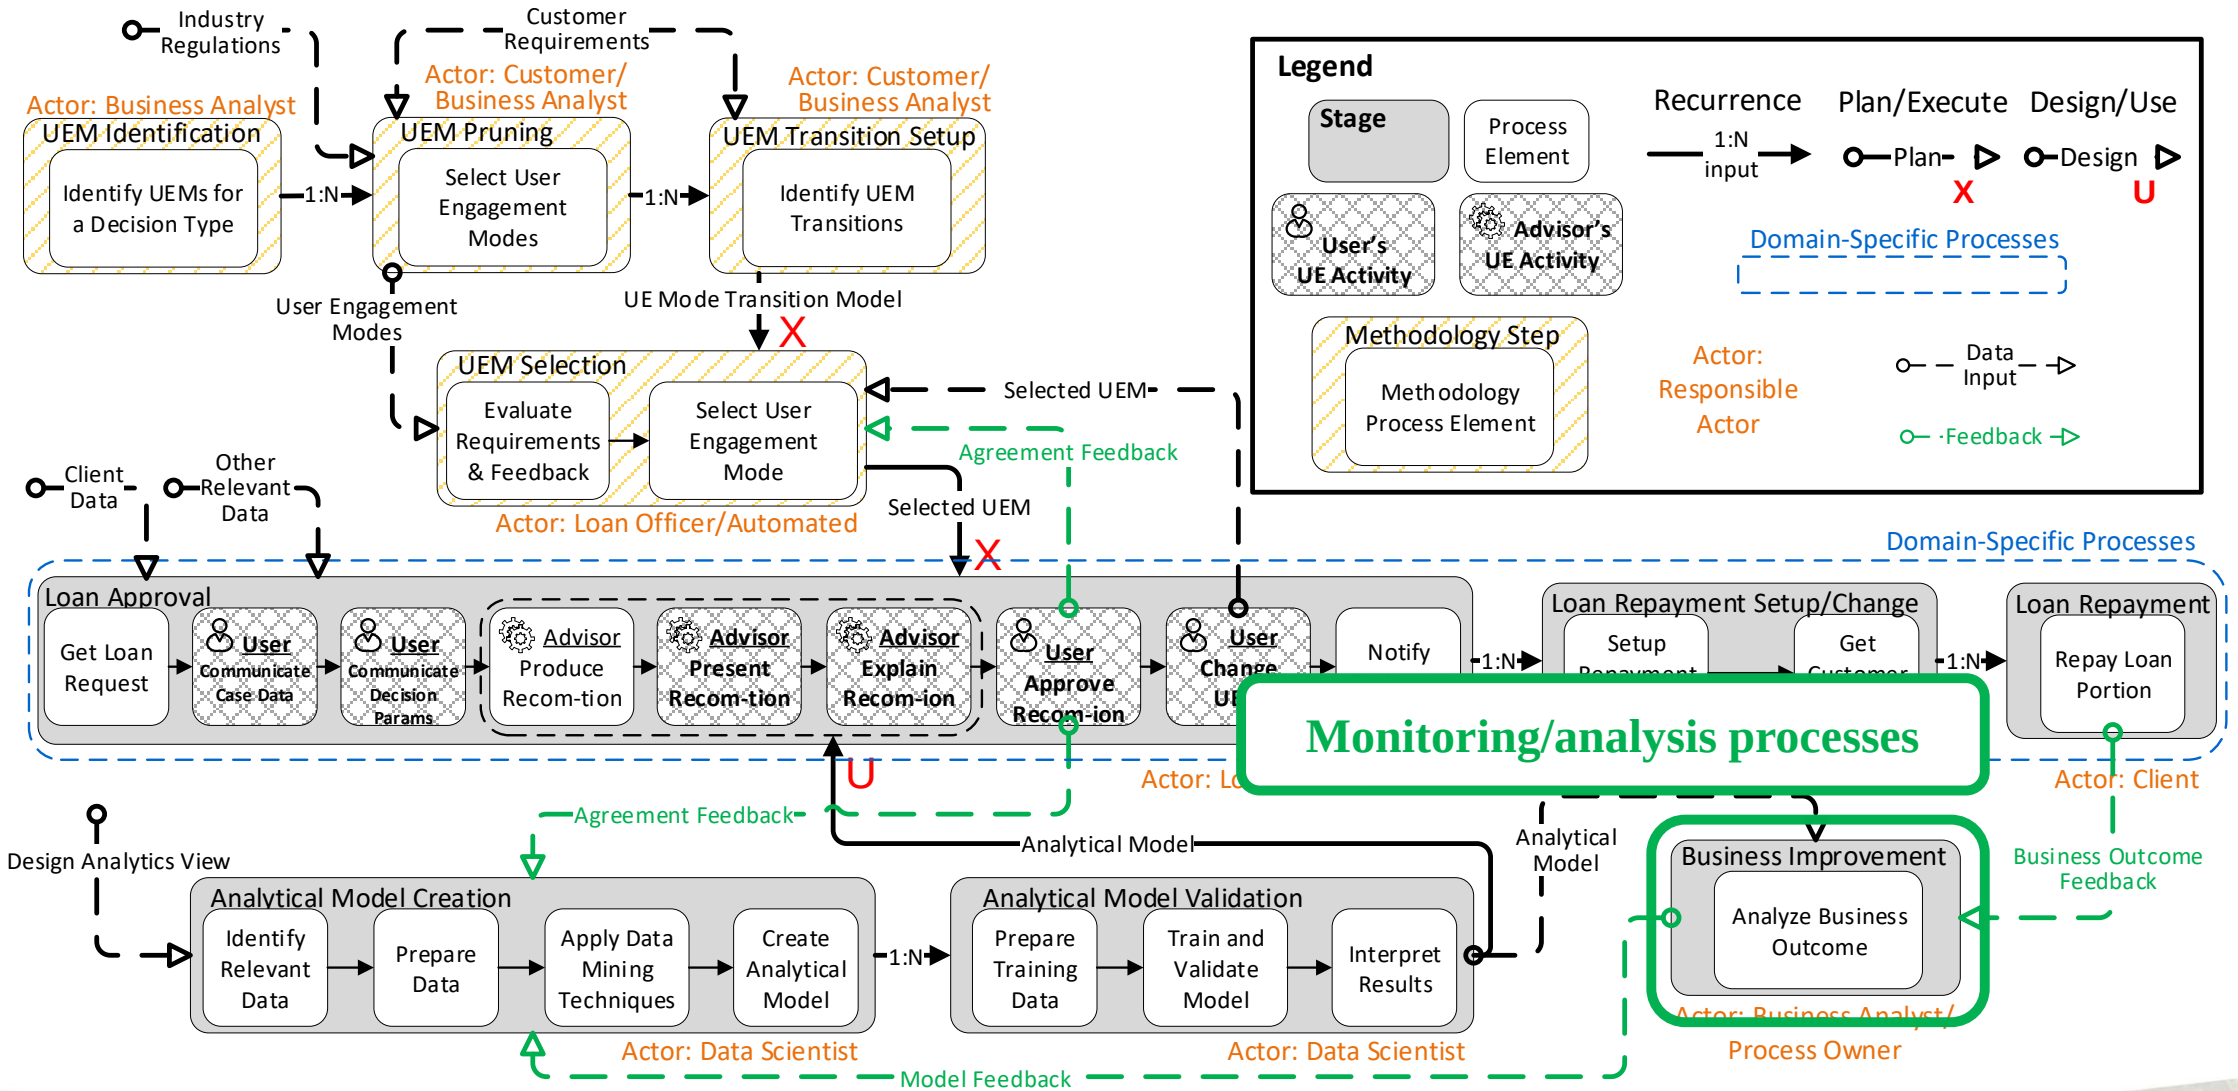

Business Process Architectures

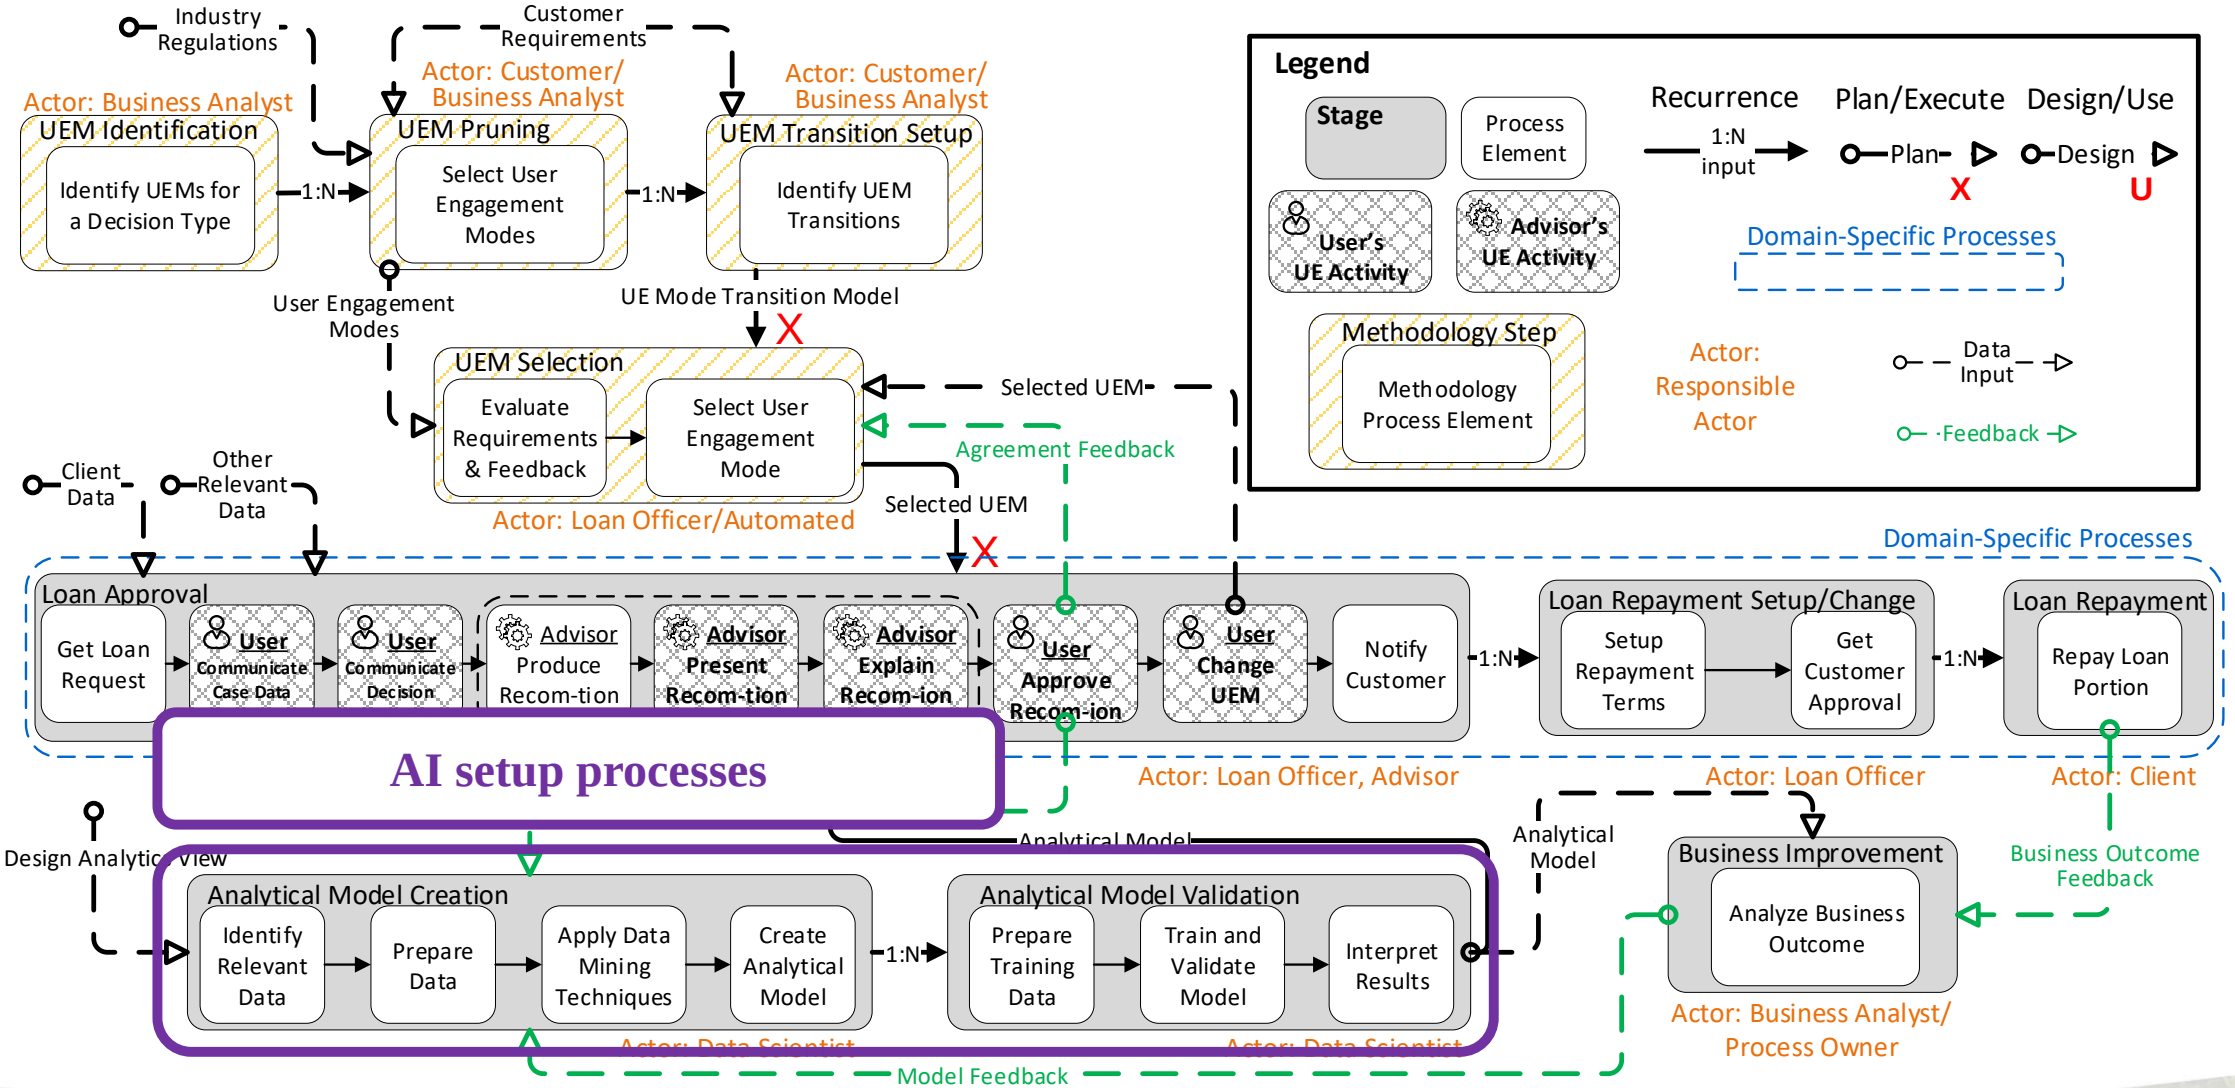

Business Process Architectures

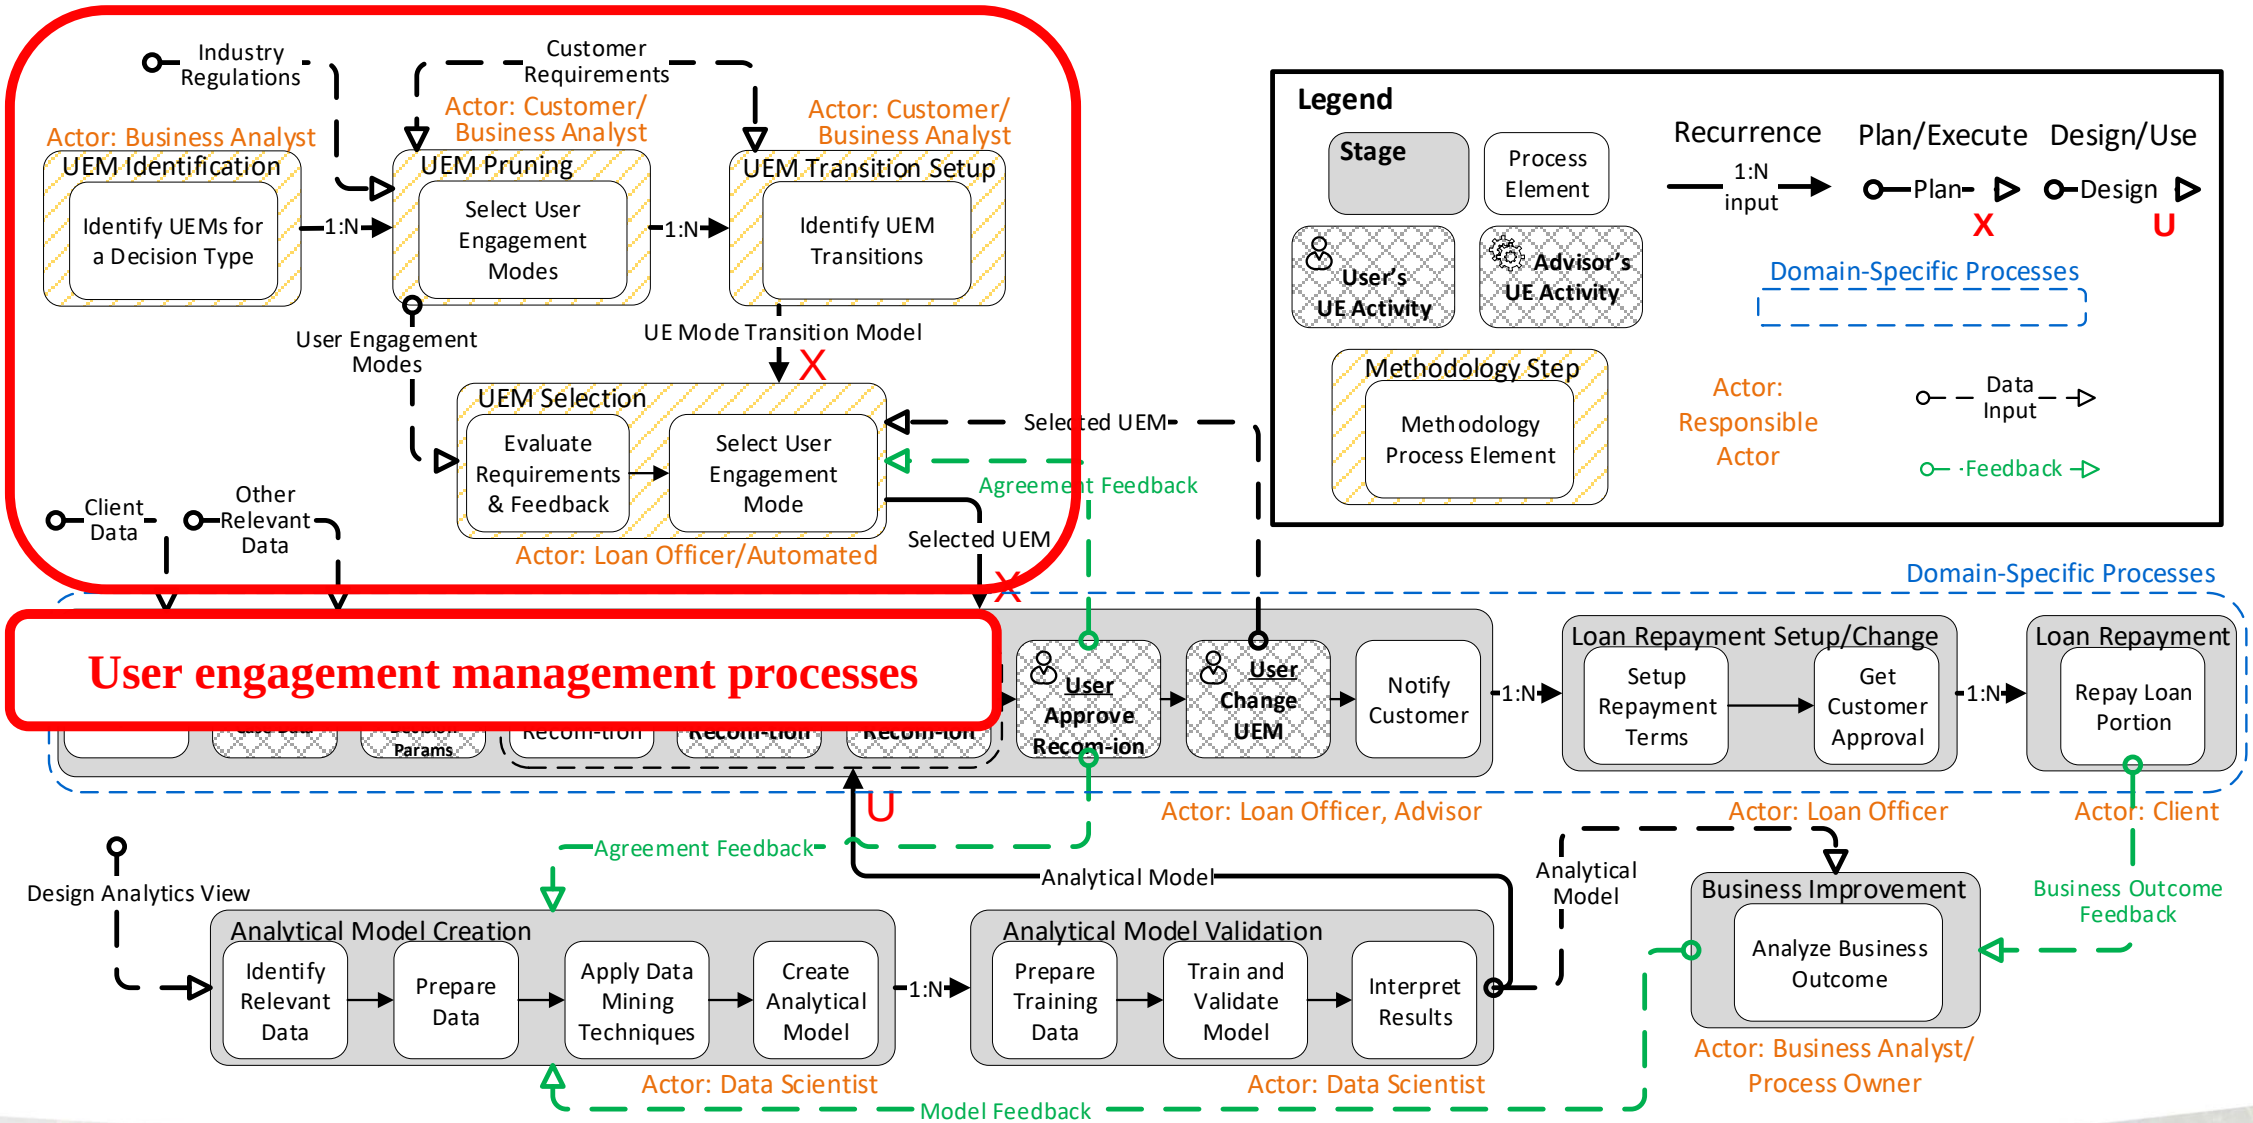

User Engagement

- Our view of user engagement:
  - How users are interacting with AI and how AI is interacting with users **in the context of decision making**
- Objectives for designing user engagement with AI:
  - Minimize disruptions to users and to existing BPs
  - Comply with organizations' internal requirements and constraints as well as those from their respective business domains
  - Change based on different situations – contexts
  - Support desired level of transparency of systems' operations
  - Build and maintain the desired level of trust in AI-driven systems

Building User Engagement

- What are the building blocks of UE?
- User Engagement Actions (UEAs)
  - Elementary activities/interactions that comprise UE with AI
  - Different combinations of UEAs define particular user engagement configurations
  - **NOTE**: The set of UEAs depends on the nature of the task – e.g., on the type of decision to be made
    - Can also be influenced by the business context, regulations, etc.

User Engagement Actions

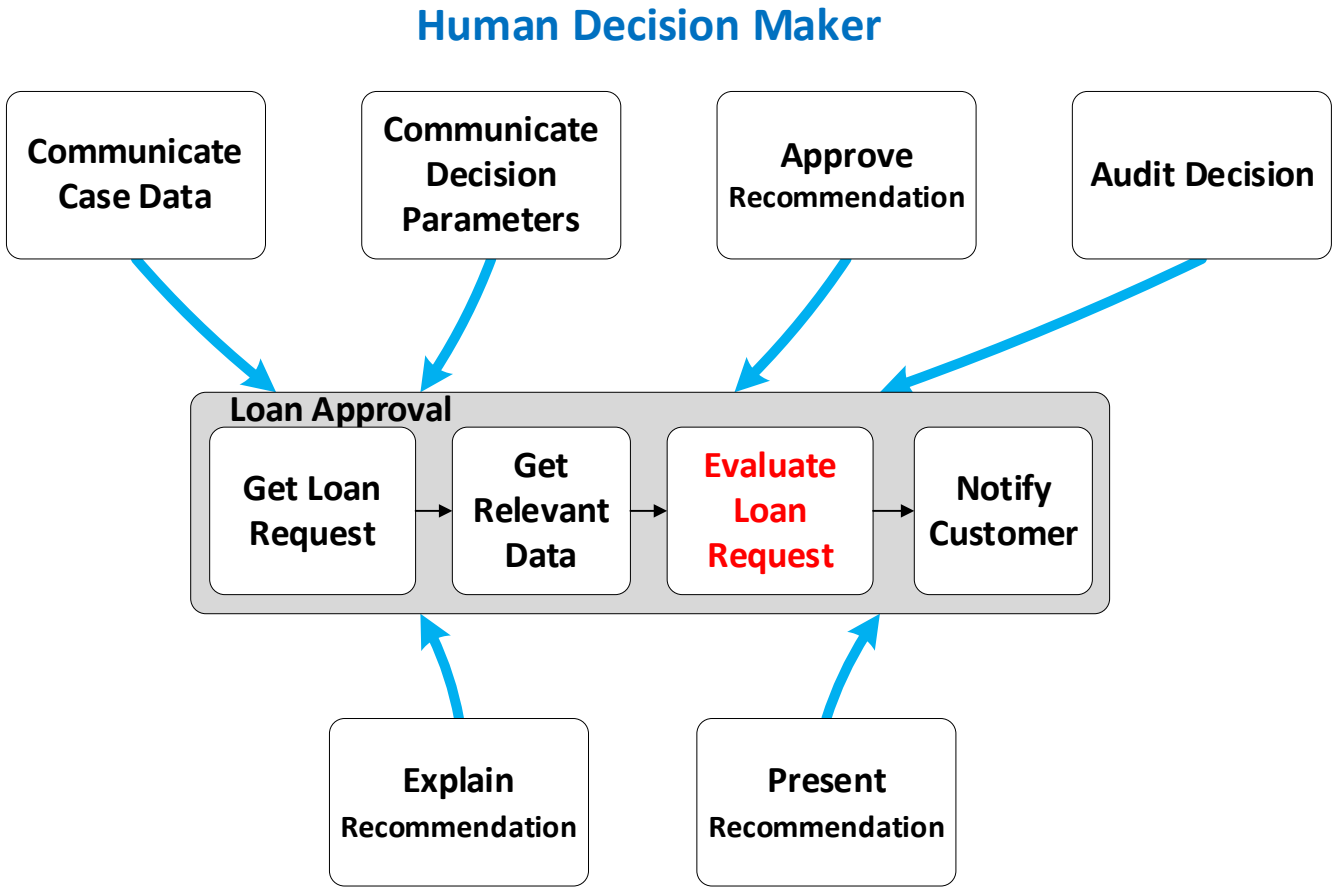

User Engagement Actions (UEAs) [1]

- Sample UEAs for decision making
  - Actors: human decision maker (**H**) and AI assistant (**AIA**)
- **Communicate Case Data**
  - H communicates the details of the particular decision instance to AIA. Alternatively, AIA can obtain the relevant data itself.
- **Communicate Decision Parameters**
  - H communicates decision parameters to AIA, including the criteria for making it, the desired confidence levels, etc.
- **Present Recommendation**
  - AIA presents the recommended decision alternative to H. For decisions with more than two potential alternatives (those that are non-boolean), a ranked list of alternatives may be given.

User Engagement Actions (UEAs) [2]

- **Approve Recommendation**
  - H either approves or rejects the previously presented recommendation. Variants for decisions with more than two options may include the ability for the user to pick an alternative option.
- **Explain Recommendation**
  - Explanation and justification of the recommendation are presented by AIA to H.
- **Present Decision**
  - AIA presents previously made decision to H (or to a specially designated auditor). Variants include the presentation of batches of decisions for subsequent audits.
- **Audit Decision**
  - H (or a designated auditor) audits the decision previously made by AIA. Variants include the audit of batches of decisions.

User Engagement Options

- Many ways for users to engage with AI systems
  - Different ways to communicate
  - Different granularity/precision of communication – e.g., per decision instance, per set of instances, for all instances, per user role, etc.
    - Results in different positions of UEAs in the BP architecture
  - The type of decision/problem – e.g., approval vs. resource allocation

Representing and Analyzing User Engagement Options

- Need a systematic way of analyzing user engagement
- Need to support the analysis of:
  - Organizational requirements, industry and other regulations
  - Individual decision makers' objectives
- Need to be able to represent:
  - The space of options for UE
  - The criteria for selecting the best options and triggers for change
  - Trade-offs

Possible Solution

- A pattern-based UE representation framework
  - Identify a limited number UE patterns for particular decision types
    - UE pattern – a typical, tried-and-tested collection UEAs arranged in a particular fashion in a BP architecture, i.e., it is a particular way to engage with an AI-based system
- Customized for particular decisions and organizations
  - Reduces the space of options for UE design
  - Provides transparency and predictability

UE Patterns for Decisions

- **P1: Supervised learning** – AI (AIA) monitors the work of human decision maker (H)
  - Decisions made by a human expert. AIA is using case data and context + the decision outcome as the input to a supervised learning algorithm.
- **P2: AIA as an Advisor** – Decisions made by a human expert. AIA's recommendations are presented as advice.
- **P3: AIA-made decisions approved by a human** – Human approves/rejects CA's recommendations (per instance).
- **P4: Human informed of AIA-made decisions** – AIA makes decisions, while a human is informed (per instance).
- **P5: AIA's decisions with (batch) human audits** – Humans audit automatically produced decisions once per N number of decision instances, once in a time interval, etc.
- **P6: AIA's decisions with human audits on request** – Humans have the opportunity to audit AIA's recommendations whenever they wish.
- **P7: AIA's decisions with automated self-audits** – No humans are involved by default. Humans can review the self-audits.

Pattern-Based Evolution of User Engagement

- Define upfront the conditions (triggers) that would cause changes in user engagement (transitions between UE Patterns)
  - Predefined user engagement evolution paths
  - Changes – not only towards more automation!
    - Identify situations when to change to a more automated UE Pattern – e.g., when the decision quality is at an acceptable level, when user trust in the AIA is adequate, etc.
    - But also specify under which conditions to **revert back to a more manual configuration** – e.g., when decision quality goes below a certain threshold and/or when the business context changes significantly
      - May need to bring humans back into the loop

Triggers

- What triggers a change from one UE Pattern to another?
  - Quality of decisions produced by an AIA
    - Can use known metrics for algorithm-produced results (confidence levels, AUC ROC, etc.)
  - Level of trust – whole organization and/or particular user – in an AIA
    - Difficult to measure automatically
    - May need to have a manual switch that would allow users to change UE patterns based on their level of trust
  - Capability/ability/skills – not triggers per se, but necessary conditions
    - The ability of AIAs and users to perform certain activities would be the necessary conditions for some transitions
      - E.g., ability to audit decisions, having enough training data, etc.
  - Context – detect changes that make the model obsolete (model drift)
    - Revert to less automated UE patterns – likely all the way to P1: Supervised Learning

Specifying Transitions

- State transition models seem to be a good notation to represent changes in user engagement

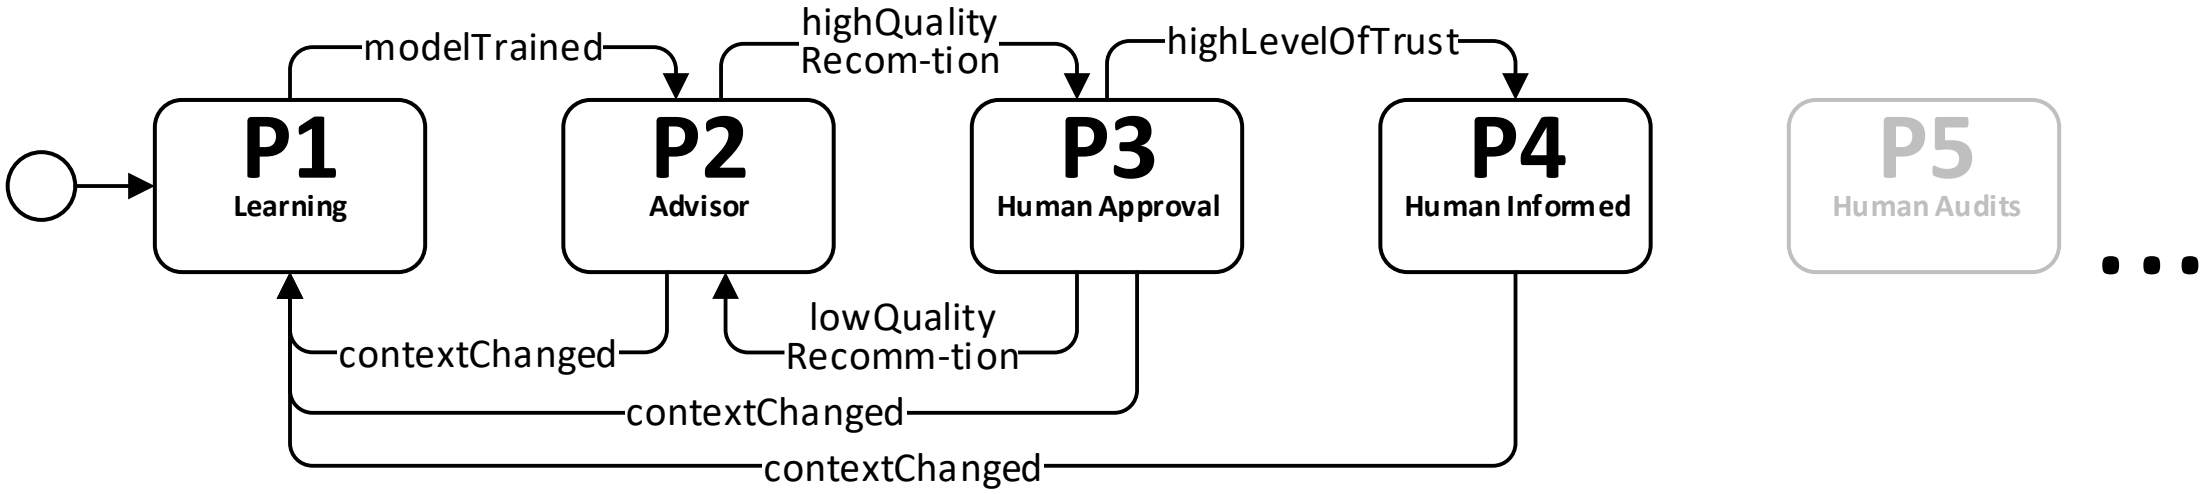

Findings

- Automated decision making helps go from human scale to machine scale decision making
  - Machine scale – decisions are much faster than humans can process them
  - If humans are involved (e.g., reviewing each decision) – still human scale
- Care is needed when introducing automation
  - Many additional processes needed, existing BPs are affected (e.g., responsibility of human users)
  - Human employees will be affected
  - Many new requirements/complexities
  - Context is very important
  - Trust is needed for the solution acceptance
  - Complex, dynamic user engagement with AI systems

Natural Language-Driven Enterprise

- Chatbots (and other NL-driven interfaces) are now primarily used to:
  - Get domain-specific or general information using a human-like interface
  - Execute some simple actions (e.g., set up alarms/reminders, play music/video)
- Enterprises are interested in:
  - Creating enterprise-level assistants that
    - Meet enterprise requirements – e.g., know how to handle security, privacy, etc. issues
      - For instance, prevent users from sending certain information outside of the organization
    - Know the context
      - Business domain-level context
      - Company-level context
      - Department-level context
      - Team-level context
      - Personal context of the user
      - Domain vocabulary

Natural Language-Driven Enterprise

- Help employees instantiate and even create processes/workflows using NL communication
- Identify user intent in all sorts of NL communication (emails, instant messaging, etc.) and map it into executable actions/processes in any of the organizations' systems
  - E.g., producing sales reports from Salesforce, setting up a recurrent meeting with the team on Wednesdays at 1pm
- Create long-running (possibly custom) workflows/processes based on the intents in NL communications
  - E.g., initiate a procurement process for a new laptop or a new hiring process
  - These will need to be executed (involving potentially many people), tracked/monitored in the background, etc.
  - Successful custom processes can be adopted by the organization