BPM-Tutor: A Conversational, AI-based Learning Platform for BPMN Modeling Education

Seyyid A. Ciftci[1,*], Roman Geist[1] and Christian Janiesch[1]

[1]TU Dortmund University, Dortmund, Germany

Abstract
Learning to model business processes demands active practice, adaptive feedback, and pedagogical scaffolding calibrated to each learner's progress. The BPM-Tutor introduces a conversational, AI-based learning platform for BPMN process modeling education built on scientific research on human-AI collaboration in process modeling. The platform integrates five pedagogically distinct AI agent archetypes, from Socratic mentors to autonomous modeling partners, alongside a context-aware element-level visual feedback system. Dedicated, configurable modules support instructors in managing tasks and AI-assisted assessment, and enable researchers to design controlled empirical studies with integrated behavioral data collection.

Keywords
Business Process Management, Learning Platform, Agentic Business Process Modeling, Human-AI Collaboration

1. Introduction

The Business Process Model and Notation (BPMN) is the established standard for modeling organizational workflows across industry and academia [1]. Developing proficiency in BPMN is a cognitively demanding undertaking that requires the simultaneous application of formal notation rules, logical process structure, and domain-specific knowledge [2, 3]. Students frequently encounter difficulties when translating informal business descriptions into structurally valid and semantically correct process models, particularly in environments where instructor feedback is limited or delayed.

Recent years have produced a growing body of AI-supported process modeling tools. Natural language to BPMN generation systems, conversational modeling assistants, and agentic process modeling frameworks have demonstrated strong capabilities in assisting professional modelers in creating and refining process models from textual descriptions [4, 5, 6]. BPMN education, however, presents a set of requirements beyond those addressed by tools designed for professional use or single-purpose feedback systems [7]: adaptive pedagogical support calibrated to learner level, transparent and controllable agent behavior that preserves human agency, and real-time feedback anchored to model elements.

We present the BPM-Tutor[1], a conversational, AI-based learning platform for BPMN modeling education. The platform integrates an empirically validated agentic architecture [8, 9], a context-aware element-level visual feedback system, instructor-facing assessment tools, and a configurable research study module in a single, self-hostable application. The agent architecture builds on an exploratory study comparing AI collaboration archetypes in BPMN modeling [8] and a design science project that refined the findings into a self-reflective agentic system validated through benchmarks and user studies [9]. Rather than proposing a new conversational modeling agent, the BPM-Tutor packages this agent architecture as a deployable educational platform, contributing pedagogical configurability, assessment support, and research instrumentation absent from prior tools, thereby lowering the barrier to AI-augmented BPMN courses and to studies on human-AI collaboration in process modeling.

Proceedings of the BPM 2026 Best Dissertation Award, Doctoral Consortium, and Demonstration Resources Forum (BPM 2026), Toronto, Canada, September 27th to October 2nd, 2026.

*Corresponding author.
seyyid.ciftci@tu-dortmund.de (S. A. Ciftci); roman.geist@tu-dortmund.de (R. Geist); christian.janiesch@tu-dortmund.de (C. Janiesch)
0009-0000-7840-4644 (S. A. Ciftci); 0009-0008-3480-883X (R. Geist); 0000-0002-8050-123X (C. Janiesch)

[1]Video: https://doi.org/10.23728/b2share.x04rf-6jf70, Tool: https://bpm-demo.cs.tu-dortmund.de/,
Repository: https://doi.org/10.5281/zenodo.20844389

2. The BPM-Tutor Platform

2.1. AI Agent Support and Modeling Environment

The modeling environment presents students with a three-panel interface combining a task description area, a bpmn.io modeling canvas [10], and a conversational AI chat window (see Figure 1). At each interaction step, the agent receives the current BPMN canvas state alongside the task description and full conversation history, grounding all responses in the student's modeling progress. This context awareness enables the agent to provide targeted guidance, generate model elements, or review the current model state.

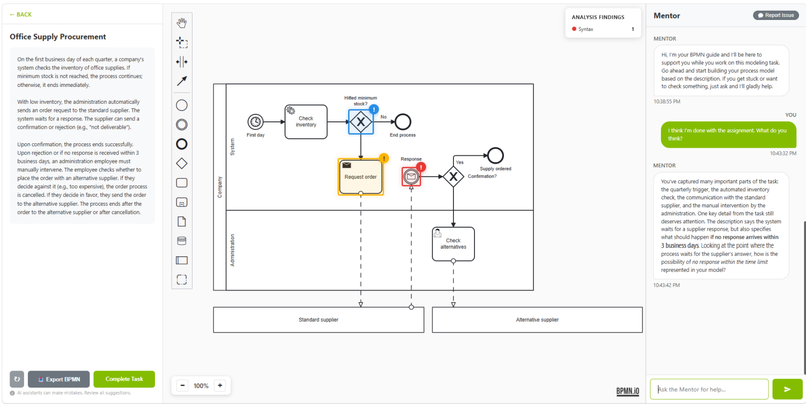

Figure 1: The BPM-Tutor modeling environment during an active task session. The AI agent provides element-level visual feedback directly on the bpmn.io canvas, highlighting a syntax violation (red), a semantic deviation from the task description (yellow), and a non-mandatory structural suggestion (blue).

The BPM-Tutor provides five configurable agent archetypes that distribute initiative and modeling effort differently between student and agent on a spectrum from dependent to autonomous, developed and empirically validated across two prior studies [8, 9]. On the dependent end, the *mentor* archetype applies a reactive Socratic approach, responding to student queries with targeted questions that activate reasoning without providing direct answers, while the *assistant* archetype delivers concise, context-aware factual information only upon explicit request. Toward the center of the spectrum, the proactive *supervisor* archetype monitors the evolving model and delivers feedback both textually via chat and visually via element-level markers, adopting a reactive posture by default to avoid disrupting the student's concentration. On the autonomous end, the co-modeling *colleague* archetype models alongside the student on the shared canvas through a structured task-division protocol, while the *delegate* archetype generates a complete process model when instructed and presents the validated draft for critical review. All archetypes adhere to the Seven Process Modeling Guidelines [11]. To mitigate the risk of erroneous agent responses, all model-modifying actions are validated against BPMN syntax rules before rendering, and the underlying self-reflective agent architecture was benchmarked for output accuracy in prior work [9]. Instructors may customize each archetype's system prompt, knowledge base, persona, and intervention style, or define additional custom agents tailored to specific course objectives. Whether feedback is provided proactively or only upon request thus depends on the archetype: the supervisor monitors the model continuously, whereas reactive archetypes assess it only when queried.

The modeling canvas is augmented by an element-level visual feedback system that annotates individual BPMN elements directly in the diagram with color-coded markers in three severity categories, as illustrated in Figure 1. *Red markers* indicate structural violations of BPMN notation rules, such as missing terminating events or gateways with invalid flow configurations. *Yellow markers* flag semantic mismatches between the model and the task description, such as activities placed in incorrect pools or lanes, or unresolved control-flow race conditions. *Blue markers* present non-mandatory structural suggestions and naming convention recommendations. Hovering over a marked element reveals a concise contextual explanation, while clicking opens a detailed diagnostic panel with repair guidance and, where applicable, an interactive simulation of the faulty process behavior. This localized feedback design minimizes extraneous cognitive load [12] by anchoring diagnostic information to the relevant model elements rather than requiring students to parse extended textual summaries [9].

2.2. Instructor Mode

The BPM-Tutor provides instructors with a dedicated dashboard for managing the complete lifecycle of BPMN courses. Task creation involves specifying a textual business process description, optional supplementary materials such as domain background or process constraints, and an optional reference model for automated and manual assessment. Tasks are assigned to individual students or cohorts with configurable availability windows and deadlines. Each task is associated with an agent configuration, allowing instructors to vary the level and type of AI support across assignments or course stages.

Assessment is configurable at the task level through three grading modes. A *practice* mode provides formative support without a formal grade. A *point-based* mode computes scores from the presence, type, and structural correctness of BPMN elements relative to the reference model. A *pass-fail* mode applies threshold-based criteria. Instructors may delegate an initial assessment pass to the AI, which grades the submission according to the selected mode and generates structured feedback anchored to individual model elements. Instructors retain full authority to review, modify, or reject AI-generated assessments before releasing grades.

The instructor dashboard further provides comprehensive configuration controls for the platform environment. Instructors define the agent library available in a given course, assign default agents to specific tasks, configure the interface language and visual theme, and manage course-level access permissions. Agent prompts, knowledge bases, and behavioral parameters are adjustable, enabling pedagogical control over the type and degree of AI support offered at each stage of a course.

2.3. Research Study Module

The BPM-Tutor originated from experimental research on human-AI collaboration [8, 9]. The platform retains this dual purpose by design: the same environment that supports classroom learning also serves as instrumented infrastructure for empirical studies. The research study module enables researchers to configure and conduct controlled studies on BPMN modeling and human-AI collaboration directly within the platform. Researchers define the study type as within-subject or between-subject, design the experimental sequence through a visual step editor, and assign tasks, agent conditions, and survey instruments to each step. Agent assignments can be randomized for counterbalancing, anonymized vis-à-vis participants, or explicitly labeled for transparency. Multi-week scheduling supports longitudinal and repeated-measures designs. Access to the study flow is managed through configurable participant codes and session tokens, and study progress is tracked in the instructor view.

The module collects multi-modal behavioral trace data throughout modeling sessions. Recorded streams include BPMN element creation, editing, and deletion events with timestamps, full agent interaction logs with message transcripts, and optionally mouse movement and timing data. These behavioral streams can be combined with self-report questionnaires constructed through the built-in survey builder. The survey builder supports multiple item types, including Likert scales for measuring constructs such as perceived cognitive load via standardized instruments like the NASA Task Load Index (NASA-TLX) [12, 13], multiple-choice items, free-text fields, and slider responses. Surveys are embedded directly into the study flow and can be presented before, during, or after modeling tasks, enabling fine-grained measurement of cognitive and motivational states at any point in the experimental sequence. All collected data are exportable for analysis in standard statistical environments.

2.4. Architecture

The BPM-Tutor is organized around two user roles: students and instructors. The platform integrates the bpmn.io [10] graphical editor as the primary modeling canvas, together with an AI agent layer powered by configurable large language models (LLMs). The default configuration connects to the

OpenAI GPT model family. Instructors may specify a custom API endpoint to route requests to a locally hosted model, ensuring full data sovereignty and compliance with data protection requirements.

All platform data are stored locally in an encrypted database, covering user sessions, BPMN submissions, agent configurations, study responses, and interaction logs. The self-hosted deployment model removes dependency on external cloud services and enables institutions to retain complete ownership of collected data. At each interaction step, the AI layer receives a dynamically assembled context window containing the task description, the serialized BPMN canvas state, and the full conversation history. This design ensures that agent responses remain continuously grounded in the student's actual modeling progress throughout a session.

Agent communication with the modeling canvas relies on a lightweight, action-based data format developed and benchmarked in prior work [9]. By encoding modeling operations as parameterized command lists rather than complete document trees, this format reduces token consumption and response latency compared to XML-based approaches, while eliminating structural generation errors that arise when verbose serialization formats approach the underlying model's effective context window limit. The canvas renders agent-issued actions step-by-step with intentional delays, enabling students to follow the agent's modeling logic incrementally as the process structure takes form.

3. Maturity of the Tool

The BPM-Tutor is built on the agent architecture and modeling support system developed through a three-cycle design science research process described in detail in [9]. The first cycle involved a within-subject laboratory study with 15 participants from computer science, information systems, and data science, in which participants completed BPMN modeling tasks using three initial agent archetypes: a reactive assistant, a co-modeling colleague, and a proactive supervisor [8]. Data collection combined concurrent think-aloud protocols, NASA-TLX workload ratings, and post-study questionnaires assessing autonomy and competence. Results revealed substantial differences in cognitive load and motivation across archetypes. Participants preferred the reactive assistant configuration, which preserved full human agency and produced the best balance between task support and perceived human autonomy. The supervisor archetype was consistently associated with the highest reported workload, as participants described unsolicited text-heavy feedback as disruptive to the modeler's concentration. The co-modeling colleague configuration induced high cognitive coordination overhead and led some participants to accept AI-generated model contributions without critical review.

The resulting platform is currently deployed in a *Business Process Management* course at TU Dortmund University with 116 students, covering 29 voluntary modeling tasks ranging from simple sequential processes to multi-pool collaboration models with exception handling. Feedback collected through this ongoing educational use continuously informs the platform's further development and refinement, allowing the system to evolve iteratively based on real-world classroom experience. The platform is a web-based application, requiring no local installation, lowering the barrier to entry for instructors and students. Moreover, the platform is openly accessible in a demo version, allowing interested users, researchers, and practitioners to explore its core capabilities without prior setup or commitment. Institutions seeking earlier adoption may already deploy the platform in an experimental capacity within their own configurations.

4. Conclusion and Future Work

We presented the BPM-Tutor, an intelligent and configurable platform for AI-supported BPMN modeling education. The platform addresses a gap in BPM educational tooling by combining an empirically grounded multi-agent interaction layer, context-aware element-level visual feedback, flexible instructor-facing assessment support, and a research-ready empirical study module in a self-hostable application. The agent architecture and modeling support at the platform's core were developed and validated in prior research on human-AI collaboration and cognitive efficiency in process modeling [8, 9].

Future development will pursue three directions: The first direction involves extending LLM provider support beyond the current OpenAI integration to include API-compatible open-source models and fully self-hosted deployments to reduce institutional dependency. The second direction encompasses three related extensions: broadening the platform's configuration capabilities to accommodate a wider range of institutional and course-specific requirements; developing adaptive pedagogical support in which the agent dynamically adjusts its archetype and intervention level to the student's observed state, for example mitigating cognitive overload or counteracting uncritical delegation to the AI; and expanding the scope of supported tasks beyond BPMN modeling to cover further activities in the BPM lifecycle, such as process analysis, simulation, or process mining. The third direction involves training domain-specific, fine-tuned models optimized for BPMN modeling, enabling real-time agent assistance with reduced latency and higher accuracy on complex process structures, and overcoming the technical constraints that currently limit seamless human-AI co-modeling in educational contexts.

Declaration on Generative AI

The authors have not employed any generative AI tools.

References

[1] M. Dumas, M. La Rosa, J. Mendling, H. A. Reijers, Fundamentals of business process management, Springer, Heidelberg, 2018.

[2] M. Neurauter, J. Pinggera, M. Martini, A. Burattin, M. Furtner, P. Sachse, B. Weber, The influence of cognitive abilities and cognitive load on business process models and their creation, in: F. D. Davis, R. Riedl, J. vom Brocke, P.-M. Léger, A. B. Randolph (Eds.), NeuroIS 2015, Springer, Cham, 2015, pp. 107–115.

[3] K. Figl, R. Laue, Cognitive complexity in business process modeling, in: H. Mouratidis, C. Rolland (Eds.), CAiSE 2011, Springer, Heidelberg, 2011, pp. 452–466.

[4] F. Friedrich, J. Mendling, F. Puhlmann, Process model generation from natural language text, in: H. Mouratidis, C. Rolland (Eds.), CAiSE 2011, Springer, Heidelberg, 2011, pp. 482–496.

[5] L. F. Hörner, M. Möller, M. Reichert, Automatically generating BPMN 2.0 process models from natural language process descriptions: Challenges, framework, quality assessment, Business & Information Systems Engineering 68 (2026) 59–83.

[6] N. Klievtsova, J.-V. Benzin, T. Kampik, J. Mangler, S. Rinderle-Ma, Conversational process modelling: State of the art, applications, and implications in practice, in: C. Di Francescomarino, A. Burattin, C. Janiesch, S. Sadiq (Eds.), BPM Forum 2023, Springer, Cham, 2023, pp. 319–336.

[7] J. Sànchez-Ferreres, L. Delicado, A. Abbad Andaloussi, A. Burattin, G. Calderón-Ruiz, B. Weber, J. Carmona, L. Padró, Supporting the process of learning and teaching process models, IEEE Transactions on Learning Technologies 13 (2020) 552–566.

[8] S. A. Ciftci, Designing synthetic work relationships: An exploratory study of agentic AI roles, cognitive load, and psychological need satisfaction, in: HHAI 2026, IOS Press, 2026, pp. 118–131.

[9] S. A. Ciftci, C. Janiesch, Agentic business process modeling: Co-creating business processes with human-centric AI agents, in: BPM 2026, Springer, Cham, 2026. To appear.

[10] Camunda Services GmbH, bpmn.io: Web-based tooling for BPMN, DMN, CMMN, and forms, https://bpmn.io/, 2026. Accessed: 2026/06/12.

[11] J. Mendling, H. Reijers, W. van der Aalst, Seven process modeling guidelines (7PMG), Information and Software Technology 52 (2010) 127–136.

[12] J. Sweller, Cognitive load theory, in: J. P. Mestre, B. H. Ross (Eds.), Cognition in Education, volume 55 of *Psychology of Learning and Motivation*, Academic, San Diego, CA, 2011, pp. 37–76.

[13] S. G. Hart, L. E. Staveland, Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research, in: P. A. Hancock, N. Meshkati (Eds.), Human Mental Workload, volume 52 of *Advances in Psychology*, Elsevier, Amsterdam, 1988, pp. 139–183.