

Autonomy in Business Process Execution: Why We Need First-Class Abstractions for Goals and Normative Frames

Diego Calvanese¹, Giuseppe De Giacomo², Timotheus Kampik^{3,4,*}, Yves Lesperance⁵,
Andrea Marrella⁶ and Andrea Matta⁷

¹Free University of Bozen-Bolzano, Bozen-Bolzano, Italy

²University of Oxford, Oxford, United Kingdom

³Umeå University, Umeå, Sweden

⁴SAP, Berlin, Germany

⁵York University, Toronto, Canada

⁶Sapienza Università di Roma, Rome, Italy

⁷Politecnico di Milano, Milan, Italy

Abstract

With the increased deployment of AI-based technologies—recently and most notably large language models—*framing* the autonomy of goal-oriented agents that enact business processes can be expected to be a key challenge. In this paper, we argue that addressing this challenge requires new formal foundations for process specifications. Traditional business process specifications focus on the *how* of business operations and treat neither goals nor norm-based constraints as first-class abstractions. Although goals play a central role in informal notions of business processes, formal definitions tend to treat them as implicit, embedded within procedural specifications that may only partially, and not explicitly, reflect normative boundaries. However, to maximize autonomy within a given *normative frame*, which expands upon the traditional idea of process models as *operational frames*, agents require formally specified goals, from which they can then synthesize their plans and actions, considering the normative frame as a set of deontic constraints. In this paper, we articulate this vision, highlight practical challenges, and propose action items for supporting its implementation.

Keywords

Business process management, AI agents, Autonomy, Autonomous agents and multi-agent systems

1. Introduction

Software systems providing the operational backbone of organizations are becoming increasingly autonomous [1]. This trend is driven in part, but not exclusively, by advances in deep learning-based technologies such as Large Language Models (LLMs). Indeed, the distributed and complex nature of large organizations requires intelligence at the level of autonomous submodules, reflecting how intelligent business decisions are made by humans. In order to deploy autonomous software agents safely and effectively, one must ensure that they comply with normative requirements, while still utilizing their substantial degrees of autonomy to accomplish their goals to the best possible extent [2].

As abstractions for managing guardrails, we propose the notion of (normative) *frames* that—in contrast to the more operational notions of declarative or procedural business processes and rules—focus only on deontic requirements of how organizations should run. Frame representation and reasoning can draw from a wealth of research on deontic logic [3], temporal reasoning [4], planning [5, 6], and normative multi-agent systems [7]. We provide informal definitions of frames, position them in the context of related abstractions, and sketch scenario types describing how frames can be applied to agents enacting business processes. These partially subsymbolic AI agents must then be augmented with symbolic capabilities for synthesizing plans that guarantee frame compliance, as well as for reasoning

PMAT'25: 4th International Workshop on Process Management in the AI era, October 25, 2025, Bologna, Italy

*Corresponding author.

✉ diego.calvanese@unibz.it (D. Calvanese); giuseppe.degiacomo@cs.ox.ac.uk (G. D. Giacomo); tkampik@cs.umu.se (T. Kampik); lesperance@eecs.yorku.ca (Y. Lesperance); marrella@diag.uniroma1.it (A. Marrella); andrea.matta@polimi.it (A. Matta)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

about their own, others', and process-level goals in order to maximize objective satisfaction *within* the frames. Accordingly, on a fundamental level these agents require capabilities for plan and behavior synthesis [8, 9, 10, 11, 12], as well as for goal reasoning [13]. We also highlight a list of challenges that must be addressed to (better) utilize frames in practice. Considering these challenges, we outline action items for laying the formal foundations for framing autonomous business process execution.

2. Framed Autonomy in Business Processes

Framed autonomy requires that an autonomous system operates within its current *frame*. Intuitively, a frame is a set of rules, restrictions, and regulations, which may evolve over time. The frame establishes the boundaries within which the system may operate with maximal flexibility, making autonomous decisions [14]. In Business Process Management (BPM), frames may exist—at least—on *agent type*, *process*, and *organization* levels (as well as potentially across organizations).

More analytically, frames are *normative*: they specify deontic requirements to the process. In contrast, classical process specification languages, such as BPMN and DECLARE are *operational*: they specify behavior required to accomplish a business goal¹. However, in contrast to informal definitions of business processes, e.g., as “sets of activities” performed to “jointly realize a business goal” [15, p. 5], goals are left implicit in these more formal and operationalizable process specification languages.

Notice that sometimes the operational specifications have been called *frames* as well [1]. Indeed, they can be considered a sort of operational frame. Here, however, our focus of “frames” is on the normative specification. When we need to distinguish, we call the two frames *normative frame* and *operational frame*, respectively.

Observe that if there are no choices to be made (no autonomous decision-makers), then the normative frame is just an additional condition over the operational frame; but if decision-making is possible then the operational frame requires finding a strategy to satisfy the objective, whereas the normative frame requires choosing a strategy that remains within what is allowed (with respect to the frame).

Strategies for achieving goals under framed autonomy are associated with decision-makers, including software agents, giving rise to several problem setups, for centralized as well as distributed intelligence. **Centralized intelligence.** We consider the “AI agents” as a single entity orchestrating the *process* that is executed in a mutually fully observable and coordinated manner. The *environment* may be stochastic and not fully observable. The frame is over the process. The single entity may have active or passive responsibility for the frame. If we have multiple agents we may break down the problem into several of the above scenarios.

Distributed intelligence. We consider AI agents as distributed entities that enact the process as *resources*. This has wide-ranging implications: a resource may have only partial observability of what other resources are doing; coordination may be effortful, and resource-level goals may be mutually inconsistent, or inconsistent with process-level goals. In such scenarios, we can frame individual resources, groups of resources, or the entire process. Accordingly, we need to assign responsibility to individual agents or groups thereof, and there may be strategic interactions affecting responsibility.

From these problem setups, we can derive three different blueprint scenarios for framed autonomy in business processes (see Figure 1): (i) we have a single decision-maker and place a frame on process behavior; (ii) we have multiple decision-makers and place frames on individual decision-makers; (iii) we have multiple decision-makers and place frame(s) on process behavior or parts thereof.

In practice, there may be additional variance to the scenarios. For example, normative frames may be partially represented within operational process specifications, restricting overall agent autonomy. An example is a purchasing process where purchase orders can only be created and paid through a central IT system that enforces normative rules, e.g. regarding four-eyes approval policies. Other parts of the global normative frame can potentially be projected to local agent-level norms. For example, overall spending limits may apply on the global level, but could be operationalized locally.

¹BPMN is *imperative*, specifying—at least supposedly—exactly what needs to be done, while DECLARE is *declarative*, specifying constraints that need to be satisfied by otherwise flexible behaviors; still, both are *operational*.




Figure 1: Different scenarios of framed autonomy in business processes.

3. Practical Challenges

Achieving framed autonomy in business processes comes with practical challenges. Below, we list (and briefly discuss) three such challenges that we consider of particular importance.

What is a pragmatic notion of an agent in the context of business process execution? Before the broad adoption of LLMs, the notion of an agent did not play a major role in the engineering of business information systems and the processes that run them. Consequently, practitioners cannot be expected to be familiar with the depth and sophistication of agent-related abstractions. To the contrary, a practitioner may consider as an agent a software tool that makes use of an LLM, without much thought about further properties. Defining a more precise and robust notion of an agent that is still intuitively understandable by business process practitioners can thus be considered a key prerequisite.

How to elicit and specify frames? The elicitation and specification of frames requires a *frame meta-model*, and one or several specification languages. To this end, existing specification languages can be reused; potentially, several languages and their underlying concepts can be combined. For example, declarative approaches to process specification—such as DECLARE [16] and in more practical contexts business rule and query languages with temporal reasoning capabilities [17]—can be augmented with deontic notions in order to promote normativity to a first-class abstraction. For elicitation, both symbolic and subsymbolic approaches can be used and fused. LLMs can generate frames or parts thereof from natural language text, whereas rule mining approaches can be applied to infer normative constraints from the traces of well-behaved agents and multi-agent systems.

How to operationalize frames on real-world symbolic data? Once specified, frames need to be integrated with business information systems, to ensure systems' frame-compliance during runtime. A short- to mid-term prerequisite is the operationalization of frames using technologies that do in fact run in large organizations. Here, explainability is a necessity, considering the practical intricacy of normative requirements, as well as the scale of real-world symbolic queries and data.

4. Call to Action: Goals and Frames for Processes

When autonomy is included in a business process execution system, the notion of *normative frame* becomes essential to guardrail autonomous decision-making. Normative frames have a deontic nature and are concerned with the sets of strategies that an agent can choose from while satisfying the frame. Accordingly, when goal-oriented agents synthesize their operational strategies, these strategies are implicitly mapped to those at the normative level and checked against the frame. AI agents—whether

based on symbolic or subsymbolic methods—that enact business processes must be able to synthesize such strategies so that frame compliance can be guaranteed and exceptional violations can be justified. Currently BPM lacks formal foundations for framed autonomy. Accordingly, we suggest to (i) introduce first-class abstractions for goals and normative frames to BPM; (ii) develop and evaluate algorithms for synthesizing provably frame-compliant and performant operational specifications from frames, goals, and environmental information; (iii) demonstrate the applicability of the abstractions and algorithms in the context of real-world business information systems.

Acknowledgments

The paper is derived from working group notes as published in the report of Dagstuhl Seminar 25192, “AUTOBIZ: Pushing the Boundaries of AI-Driven Process Execution and Adaptation”. T. Kampik was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. A. Marrella was supported by the Sapienza project FOND-AIBPM. A. Marrella and D. Calvanese were supported by the PNRR MUR project PE0000013-FAIR.

Declaration on Generative AI

The authors have not employed any Generative AI tools when writing this paper.

References

- [1] M. Dumas, F. Fournier, L. Limonad, A. Marrella, M. Montali, J. Rehse, R. Accorsi, D. Calvanese, G. D. Giacomo, D. Fahland, A. Gal, M. L. Rosa, H. Völzer, I. Weber, AI-augmented Business Process Management systems: A research manifesto, *ACM Trans. Manag. Inf. Syst.* 14 (2023) 11:1–11:19.
- [2] T. Kampik, A. Mansour, O. Boissier, S. Kirrane, J. Padget, T. R. Payne, M. P. Singh, V. Tamma, A. Zimmermann, Governance of autonomous agents on the web: Challenges and opportunities, *ACM Trans. Internet Technol.* 22 (2022).
- [3] D. Gabbay, J. Horty, X. Parent, R. Van der Meyden, L. van der Torre, et al., *Handbook of Deontic Logic and Normative Systems*, Volume 2, College Publications, 2021.
- [4] G. De Giacomo, M. Y. Vardi, Linear Temporal Logic and Linear Dynamic Logic on finite traces, in: *IJCAI*, 2013, pp. 854–860.
- [5] G. De Giacomo, S. Rubin, Automata-theoretic foundations of FOND planning for LTL_f and LDL_f goals, in: *IJCAI*, 2018.
- [6] M. Ghallab, D. S. Nau, P. Traverso, *Acting, Planning and Learning*, Cambridge Univ. Press, 2025.
- [7] A. Chopra, L. van der Torre, H. Verhagen, *Handbook of Normative Multiagent Systems*, College Publications, 2018.
- [8] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: *POPL*, ACM, 1989, pp. 179–190.
- [9] B. Finkbeiner, F. Klein, N. Metzger, Live synthesis, *Innov. Syst. Softw. Eng.* 18 (2022) 443–454.
- [10] G. De Giacomo, M. Y. Vardi, Synthesis for LTL and LDL on finite traces, in: *IJCAI*, 2015.
- [11] S. Zhu, G. De Giacomo, Act for your duties but maintain your rights, in: *KR*, 2022.
- [12] S. Zhu, G. De Giacomo, Synthesis of maximally permissive strategies for LTL_f specifications, in: *IJCAI*, 2022.
- [13] D. W. Aha, Goal reasoning: Foundations, emerging applications, and prospects, *AI Mag.* 39 (2018).
- [14] G. Acitelli, A. Alman, F. M. Maggi, A. Marrella, Achieving framed autonomy in AI-augmented business process management systems through automated planning, *Inf. Syst.* 133 (2025).
- [15] M. Weske, *Business Process Management: Concepts, Languages, Architectures*, Springer, 2024.
- [16] C. Di Ciccio, M. Montali, Declarative process specifications: Reasoning, discovery, monitoring, in: W. M. P. van der Aalst, J. Carmona (Eds.), *Process Mining Handbook*, volume 448 of *LNBIP*, 2022.
- [17] T. Kampik, C. Okulmus, Expressive power and complexity results for SIGNAL, an industry-scale process query language, in: *BPM 2024 Forum*, volume 526 of *LNBIP*, Springer, 2024, pp. 3–19.