X-by-Construction Meets Intersymbolic AI

Maurice H. ter Beek[1(✉)], Clemens Dubslaff[2,3], Juliane Päßler[4], and Ina Schaefer[5]

[1] CNR–ISTI, Pisa, Italy
maurice.terbeek@isti.cnr.it
[2] Eindhoven University of Technology, Eindhoven, The Netherlands
c.dubslaff@tue.nl
[3] TU Braunschweig, Brunswick, Germany
[4] TU Wien, Vienna, Austria
juliane.paessler@tuwien.ac.at
[5] KIT, Karlsruhe, Germany
ina.schaefer@kit.edu

Abstract. X-by-Construction (XbC) encompasses extensions beyond correctness concerns which are common in the traditional Correctness-by-Construction (CbC) paradigm. Like CbC, XbC is a refinement approach to engineer systems that *by construction* satisfy certain (*non-functional*) properties. An ISoLA 2024 track gathered experts to discuss possible synergies between CbC/XbC and symbolic or subsymbolic Artificial Intelligence (AI). A key benefit of symbolic (rule-based) AI is its formal rigor, which comes at the cost of formal modeling effort and computational expensive reasoning. Subsymbolic (data-driven) AI approaches, on the other hand, usually outperform rigorous ones in performance but might lead to unsound results. An AISoLA 2025 track gathered experts to discuss the emerging field of *Intersymbolic AI*, which aims to exploit the benefits of both AI approaches, and the role of formal methods therein. The scope of the ISoLA 2026 track on "X-by-Construction Meets Intersymbolic AI" (XbC-IAI) was to gather researchers and practitioners who work on the topics of these two tracks to discuss possible synergies between CbC/XbC and Intersymbolic AI. These proceedings of the XbC-IAI track contain 13 contributions by participants.

Keywords: Correctness-by-Construction · Formal Methods · Intersymbolic AI

1 Introduction

Correctness-by-Construction (CbC) sees the development of software (systems) as a step-wise refinement process from specification to code, ideally by CbC design tools that automatically generate error-free software (system) implementations from rigorous and unambiguous requirement specifications [7]. X-by-Construction (XbC) encompasses extensions beyond correctness concerns by also considering non-functional properties (X) [4]. XbC is thus concerned with

© The Author(s) 2027
T. Margaria and B. Steffen (Eds.): ISoLA 2026, LNCS 17144, pp. 1–7, 2027.
https://doi.org/10.1007/978-3-032-40121-2_1

a step-wise refinement process from specification to code that automatically generates software (system) implementations that by construction satisfy specific non-functional properties (i.e., concerning security, dependability, reliability, resource or energy consumption, and the like).

On the other hand, there is an increasing use of *artificial intelligence (AI)* methods in software development, which make it difficult to guarantee rigorous properties to hold [20]. In his keynote contribution at the ISoLA 2024 track [3], Platzer called for the study of the field that was coined *intersymbolic AI* [17]. This field targets the combination of *symbolic AI*, whose building blocks have inherent significance/meaning, with *subsymbolic AI*, whose entirety creates significance/effect despite the fact that individual building blocks escape meaning. Symbolic AI, as implemented in rule-based systems, provides formal rigor, but this comes at the cost of increased modeling effort and computationally expensive reasoning. Differently, subsymbolic AI approaches, which typically use data-driven methods from statistical learning, are not as computationally expensive as the rigorous ones but might lead to unsound results. The idea is that intersymbolic AI combines benefits from both symbolic and subsymbolic AI to increase the overall effectiveness, rigor, and explainability of AI compared to either kind of (sub)symbolic AI alone. The combination of Formal Methods and Intersymbolic AI has been the topic of an AISoLA 2025 track [6].

2 Research Questions

All of the above mentioned areas only provide a glimpse of the many important research questions and research opportunities that are emerging from the combination of CbC/XbC with concepts from AI and explainability. On the one hand, AI-based components employed in a software system give rise to new challenges for correctness arguments, both for functional and for non-functional properties, due to their inherent probabilistic and non-deterministic nature. On the other hand, AI-supported software engineering and AI-enhanced verification approaches offer great potential for more efficient development of correct-by-construction software, including non-functional properties.

Such questions were posed and addressed in the ISoLA 2026 track on "X-by-Construction Meets Intersymbolic AI" (XbC-IAI) we will report on in the next section. However, it is clear that a single track cannot deliver all answers to the interplay between CbC/XbC (and formal methods at large) with concepts from AI and explainability (in particular from the emerging field of intersymbolic AI). Indeed, this ISoLA 2026 track fits perfectly in the tradition of the following related and highly successful (A)ISoLA tracks:

ISoLA 2016 Correctness-by-Construction and Post-hoc Verification: Friends or Foes?
ISoLA 2018 X-by-Construction
ISoLA 2020/2021 X-by-Construction Meets Probability
ISoLA 2022 X-by-Construction Meets Runtime Verification
ISoLA 2024 X-by-Construction Meets AI
AISoLA 2025 Formal Methods for Intersymbolic AI
ISoLA 2026 X-by-Construction Meets Intersymbolic AI

3 Contributions

The XbC-IAI track at ISoLA 2026 [5] was a two-day event. Researchers and practitioners working on Formal Methods and AI presented advances spanning trustworthy and explainable AI (XAI), neuro-symbolic methods, reinforcement learning, autonomous systems, and practical tools and frameworks grounded in formal methods. These proceedings contain 13 contributions by participants in the track. These contributions are briefly discussed below, grouped thematically.

Keynote: Towards Correct-by-Construction Legal Explanations for Intersymbolic Conversational Interfaces to Governmental Agencies

The track featured a keynote address by the invited speaker, Thomas Hildebrandt (University of Copenhagen, Denmark). In his invited talk, based on [12], Hildebrandt presented a Conversational Hybrid AI (CHAI) architecture for building trustworthy natural-language interfaces to governmental agencies by combining LLMs with symbolic DCR graph models of legal processes. The approach uses LLMs to translate legal texts into executable symbolic models and to mediate user interactions, while the symbolic execution engine guarantees accurate answers and explanations. By grounding explanations in the formal semantics of DCR graphs and explicit links to legal sources, the framework aims to provide correct-by-construction explanations for governmental cases.

Session 1: Neuro-Symbolic AI and Decision-Making

In [2], Bartocci et al. advocate a model-centric neuro-symbolic framework that combines formal methods with learning-based decision making. They argue that symbolic models, whether designed or learned from data, are essential for enabling trust by design in autonomous systems operating in complex and uncertain environments. Their vision emphasizes a continuous feedback loop between model learning and formal analysis, allowing models and decision policies to be refined, verified, and monitored throughout deployment.

In [15], Liu et al. develop SpecRepair, a neuro-symbolic framework that combines LLMs with formal reactive synthesis to iteratively construct and repair realizable natural-language requirements. The approach uses LLMs to translate requirements into structured intent models, applies temporal-logic-based realisability analysis to detect inconsistencies, and leverages counterstrategy-guided explanations to generate interpretable repair suggestions before translating the repaired models back into natural language.

In [11], Ciancia and Tuosto introduce a lightweight calculus of located task graphs (LOTS) for specifying compliant-by-construction hybrid AI workflows in regulated domains. The approach represents workflows as located task graphs that explicitly capture data dependencies, storage and computation locations, and institutional boundaries. Through a medical imaging case study, the authors illustrate how compliance-relevant requirements can be incorporated into executable workflow specifications already at the design stage rather than addressed only after deployment.

Session 2: Explainable Formal Methods and AI

In [13], Izza et al. design XChecker, a methodology for validating the correctness of formal XAI tools by cross-checking their explanations against independent formal explainers, proof traces, and the underlying machine learning model. The framework distinguishes provable from likely explanation errors and is used to uncover bugs in an existing formal explainer. The results highlight the importance of systematic validation techniques to ensure the reliability of practical implementations of formal XAI.

In [16], Perez-Palacin et al. propose an approach for making machine learning-based decision-making uncertainty-aware by construction by representing confidence scores using fuzzy theory. Rather than reducing uncertainty to a single calibrated value, the approach combines multiple calibration and reliability indicators into fuzzy confidence representations that preserve richer uncertainty information. This enables an interpretable fuzzy decision-maker to adapt its interpretation of uncertainty to changing system goals.

In [1], Baier et al. report on recent advances of PMC-VIS, an interactive visualization framework for the development, analysis, and debugging of probabilistic systems. The framework enhances the explainability of probabilistic model checking by integrating state responsibility measures and witnessing subsystems through a general-purpose data import interface. A case study demonstrates how combining formal analysis techniques with interactive visualization helps users understand, diagnose, and debug Markov decision process models.

Session 3: Reinforcement Learning and Hybrid Systems

In [18], Pranger and Könighofer aim to make formal shielding readily accessible within standard reinforcement learning workflows by integrating shield synthesis into an end-to-end infrastructure. They introduce TEMPESTPY, a Python library that connects TEMPEST-based shield synthesis directly with the Gymnasium API, and extend TEMPEST with sound algorithms for synthesizing shields for stochastic multiplayer games. Together, these contributions enable practical shielded reinforcement learning while retaining formal safety guarantees.

In [9], Blohm et al. extend an existing formal approach for quantitatively analyzing stochastic Simulink models by incorporating reinforcement learning (RL) components. They introduce a formal representation of learned RL policies into a transformation from Simulink models to Stochastic Hybrid Automata, enabling their integration into the analysis model. The approach supports the quantitative evaluation of safety and reliability properties of intelligent cyber-physical systems combining stochastic uncertainty and learning-based behavior.

In [19], Stübbe et al. describe an approach for analyzing Stochastic Hybrid Models (SHMs) by combining symbolic state-space analysis with reinforcement learning to address the challenge of resolving nondeterminism in complex hybrid systems. Their approach integrates flowpipe-based reachability analysis with Q-learning to guide system exploration while pruning unreachable regions in a correct-by-construction manner. This enables more efficient quantitative verification without sacrificing the fidelity of the underlying continuous hybrid model.

Session 4: Agents and LLM-based Engineering

In [10], Bollig et al. devise a domain-specific language based on message sequence charts to formally specify the coordination of LLM-based multi-agent systems independently of the nondeterministic behavior of LLMs. The authors equip the language with a formal semantics and a syntax-directed projection that automatically generates deadlock-free local agent programs while preserving coordination guarantees. They also extend the approach to runtime-generated workflows, ensuring structural guarantees for LLM-planned coordination.

In [8], Blando et al. examine how integrating LLM-based decision making into agent-based models (ABMs) affects simulation reliability, computational cost, and emergent behavior. The authors introduce a controlled methodology for studying the interplay between symbolic agent rules and LLM-driven behavior by replacing a local decision rule with an LLM-based one in an otherwise standard ABM and comparing the resulting simulation dynamics using statistical model checking with MultiVeStA.

In [14], Kodetzki et al. investigate the use of LLMs to automatically generate program code within the CbC programming methodology, reducing the manual effort required for formally verified software development. Four LLMs are evaluated and compared against a classical synthesis approach with respect to correctness and reproducibility. The best-performing model is integrated into the WebCorC development environment, where it provides automated code synthesis during the CbC refinement process.

References

1. Baier, C., et al.: On the integration of explainable formal methods tools into PMC-VIS. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 136–157. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_8
2. Bartocci, E., Češka, M., Hudák, D., Novotný, P., Tappler, M.: Towards robust and safe decision-making under uncertainty: a model-centric view. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 31–54. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_3
3. ter Beek, M.H., Cleophas, L., Dubslaff, C., Schaefer, I.: X-by-Construction meets AI. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 12th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Software Engineering Methodologies (ISoLA'24). LNCS, vol. 15222, pp. 155–161. Springer (2024). https://doi.org/10.1007/978-3-031-75387-9_10
4. ter Beek, M.H., Cleophas, L., Schaefer, I., Watson, B.W.: X-by-Construction. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 8th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Modeling (ISoLA'18). LNCS, vol. 11244, pp. 359–364. Springer (2018). https://doi.org/10.1007/978-3-030-03418-4_21

5. ter Beek, M.H., Dubslaff, C., Päßler, J., Schaefer, I.: X-by-Construction Meets Intersymbolic AI. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: X-by-Construction Meets Intersymbolic AI (ISoLA'26). LNCS, vol. 17144, pp. 1–7. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_1
6. ter Beek, M.H., Dubslaff, C., Schaefer, I.: Formal Methods for Intersymbolic AI. In: Steffen, B. (ed.) Proceedings of the 3rd International Conference on Bridging the Gap between AI and Reality (AISoLA'25). LNCS, vol. 16220, pp. xvii – xx. Springer (2025). https://doi.org/10.1007/978-3-032-07132-3
7. ter Beek, M.H., Hähnle, R., Schaefer, I.: Correctness-by-Construction and post-hoc verification: friends or foes? In: Margaria, T., Steffen, B. (eds.) Proceedings of the 7th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Foundational Techniques (ISoLA'16). LNCS, vol. 9952, pp. 723–729. Springer (2016). https://doi.org/10.1007/978-3-319-47166-2_51
8. Blando, S., Guerrazzi, E., Porcedda, R., Squillace, G., Tschaikowski, M., Vandin, A.: Towards agentic agent-based models: feasibility, performance, and statistical model checking. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 256–276. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_13
9. Blohm, P., Remke1, A., Herber, P.: Towards quantitative analysis of intelligent Simulink models using stochastic hybrid automata. In: Margaria, T., Beckert, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 173–195. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_10
10. Bollig, B., Függer, M., Nowak, T.: Provable coordination for LLM agents via message sequence charts. In: Margaria, T., Beckert, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 230–255. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_12
11. Ciancia, V., Tuosto, E.: Compliant-by-construction hybrid medical AI workflows via located task graphs. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 76–91. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_5
12. Hildebrandt, T.: Towards correct-by-construction intersymbolic conversational interfaces to governmental agencies. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 8–30. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_2
13. Izza, Y., Huang, X., Ignatiev, A., Marques-Silva, J.: Towards the validation of formal explainers. In: Margaria, T., Beckert, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 92–114. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_6
14. Kodetzki, M., Grimm, H., Demmler, F., Schaefer, I.: LLM-based program synthesis in correctness-by-construction engineering. In: Margaria, T., Beckert, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of

Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 277–299. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_14
15. Liu, T., Bremer, A., Panov, I., Weigl, A., Tyszberowicz, S., Vandin, A.: Realisability-driven repair of natural-language requirements via neuro-symbolic reasoning. In: Margaria, T., Beckert, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 55–75. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_4
16. Perez-Palacin, D., Scotti, V., Maleki, A., Grassi, V., Mirandola, R.: Interpretable and uncertainty-aware by construction decisions with fuzzy theory. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 115–135. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_7
17. Platzer, A.: Intersymbolic AI: Interlinking symbolic AI and subsymbolic AI. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 12th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation: Software Engineering Methodologies (ISoLA'24). LNCS, vol. 15222, pp. 162–180. Springer (2024). https://doi.org/10.1007/978-3-031-75387-9_11
18. Pranger, S., Könighofer, B.: TEMPESTPY: easy-to-use shielded reinforcement learning. In: Margaria, T., Steffen, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 158–172. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_9
19. Stübbe, J., Lauenroth, F., da Silva, C., Remke1, A.: Guided discrete-event simulation for SHM via symbolic precomputation and Q-learning. In: Margaria, T., Beckert, B. (eds.) Proceedings of the 13th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'26). LNCS, vol. 17144, pp. 196–229. Springer (2026). https://doi.org/10.1007/978-3-032-40121-2_11
20. Wing, J.M.: Trustworthy AI. Commun. ACM **64**(10), 64–71 (2021). https://doi.org/10.1145/3448248

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.