Perspective

Stefan Rolf Benz*, Jule Reitz, Anne Benz, Paul Benz, Philipp Reitz and André Mihaljevic

The neuro-symbolic anatomy engine: a novel architecture for universal anatomical modeling

https://doi.org/10.1515/iss-2025-0045
Received September 21, 2025; accepted April 3, 2026;
published online April 28, 2026

Abstract: We present the neuro-symbolic anatomy engine, a novel artificial intelligence architecture that integrates structured anatomical knowledge with neural pattern recognition for comprehensive human anatomy modeling. Rather than focusing on individual procedures, the anatomy engine represents a Copernican shift: it encodes a universal model of human anatomy from which any surgical procedure can be simulated. This hybrid approach combines symbolic reasoning over anatomical relationships with neural network processing of text, images, and medical data. It enables transparent, traceable decision-making suited to surgical and educational applications. The system addresses critical limitations in current surgical training and has the potential to support lifelong learning, expand access to complex procedures, and drive surgical innovation.

Keywords: anatomy engine; neurosymbolic; surgical simulation; artificial intelligence

Introduction

Surgical education faces a crisis characterized by diverging challenges [1]. Healthcare systems demand higher surgical quality despite economic constraints, while the field struggles with recruitment and retention. Surgical techniques are becoming increasingly complex, requiring specialized training that traditional approaches are often unable to adequately provide. The rapid pace of surgical innovation further outpaces current training paradigms, resulting in a significant narrowing of the surgical competence spectrum.

Workforce pressures like shorter working hours, demographic shifts, and career interruptions widen the gap between training capacity and clinical demands. Most trainees can learn to perform laparoscopic cholecystectomy or ileostomy takedown, but advancing to robotic CME colectomy, rectal resection, or pancreatectomy is far more challenging. We describe this as the “surgical training trap”: early training yields satisfaction and certification, yet surgical competence stagnates well below the level needed for high-complexity procedures. Many surgeons continue to expand their repertoire by reading, watching videos, and then performing complex operations without dedicated practical training. This results in a substantial proportion of complex surgery taking place on a learning curve.

These ongoing learning curves pose risks to patient safety through increased complications and mortality [2, 3]. At the same time, rising ethical standards render unstructured on-the-job learning increasingly unacceptable.

The financial impact is also significant. Complications linked to inexperience lead to longer hospital stays, additional interventions, and greater resource utilization ultimately inflating healthcare costs substantially [4].

This crisis underscores the urgent need for innovative solutions, particularly advanced surgical simulation systems. Such systems should not only bridge the gap between training and practice but also enable the simulation of entirely new procedures. This would foster surgical innovation in a risk-free environment, allowing surgeons to experiment, refine techniques, and expand their horizons without jeopardizing patient safety. Additionally, high-fidelity simulation has the potential to produce substantial cost savings by reducing complications, shortening learning curves, and optimizing resource use within healthcare.

To effectively address these challenges, we must adopt a Copernican shift in our approach to surgical simulation. Instead of focusing solely on individual procedures, we should develop a comprehensive anatomical workspace that enables the simulation of any procedure. This paradigm

***Corresponding author: Stefan Rolf Benz**, Universitätsklinik für Allgemeine-, Viszeral- und Transplantationschirurgie, Hoppe-Seyler Str. 3, 72076 Tübingen, Germany; and Simultare GmbH, Nagold, Germany, E-mail: stefanrolfbenz@web.de
Jule Reitz and Anne Benz, Simultare GmbH, Nagold, Germany
Paul Benz, Fakultät Maschinenbau, KIT – Karlsruher Institut für Technologie, Karlsruhe, Germany; and Simultare GmbH, Nagold, Germany
Philipp Reitz, Fakultät Maschinenbau, KIT – Karlsruher Institut für Technologie, Karlsruhe, Germany
André Mihaljevic, Universitätsklink für Allgemeine-, Viszeral- und Transplantationschirurgie, Tübingen, Germany

change could revolutionize surgical training and innovation, leading to improvements in surgical competence and efficiency. Over the long term, such advancements could significantly reduce healthcare costs while enhancing patient outcomes and resource allocation.

Proposed solution: the neurosymbolic anatomy engine

As a solution aligned with this Copernican turn, we propose the Neuro-Symbolic Anatomy Engine, a hybrid AI architecture that generates a symbolic knowledge-based representation of the entire human anatomy in high resolution. Current surgical simulators work like rigid, pre-built 3D maps: fixed, procedure-specific, and unable to reason about anatomy. The anatomy engine instead constructs a living model of the human body capable of understanding topographic relations and reasoning. It is built from rules derived from text and images, not just geometric meshes.

These existing simulation systems are primarily based on limited anatomical representations within fixed coordinate matrices such as VTK (Visualization Toolkit) or Gmsh. Manipulation within these systems is typically achieved using the Finite Element Method (FEM) and related approaches, exemplified by the SOFA framework [5]. Data input for VTK is facilitated by medical image processing pipelines involving segmentation and surface extraction methods to construct polygonal meshes. Gmsh input, conversely, relies more on mesh generation tools that convert geometric/anatomical shapes into volumetric FEM-ready meshes.

While these frameworks have enabled the creation of highly useful procedure-specific simulation systems, they are subject to important restrictions:

VTK and Gmsh datasets:
- Cannot be directly retrieved from anatomical textbooks and scientific literature, thereby precluding the leveraging of existing anatomical knowledge.
- Lack the capability to logically reason about the relationships and properties of anatomical structures.
- Are not scalable to the high resolution required for realistic simulation of truly sophisticated procedures, neither in terms of anatomical representation nor for the computational power needed for FEM simulation.

In contrast, the Neuro-Symbolic Anatomy Engine is envisioned as a world model of human anatomy. It comprises an integrated framework designed to process anatomical and physiological information, as well as to support surgical reasoning tasks. The knowledge representation layer maintains a comprehensive model of human anatomy, where anatomical entities are encoded with their relative topographical relationships and are logically linked, including physiological and physical properties, as well as known cause-effect chains.

The symbolic reasoning engine will perform logical inference over anatomical relationships, maintain spatial consistency, and enable transparent decision-making through rule-based operations. This component is responsible for ensuring that anatomical simulation remains consistent with established medical knowledge and physiological constraints and can be linked to surgical manipulation.

The neural perception modules handle multimodal input processing. This primarily involves natural language processing, through which topographical and causal relationships are derived from text; image processing of anatomical atlases (utilizing CLIP [6]); medical imaging data from CT and MRI; surgical video streams; and sensor data from tracking systems employed in modern surgical environments. These neural networks are designed to extract relevant inputs and translate them into symbolic representations amenable to processing by the reasoning engine.

An integration interface provides bidirectional communication between the symbolic and neural components. This fusion mechanism enables the system to ground neural perceptions in anatomical knowledge while allowing symbolic reasoning to guide neural attention and processing.

Simulation using the Anatomy Engine can be understood as video generation from a neuro-symbolic world model where instrument inputs represent prompts. In this manner, simulation could be achieved as described by Atharva Sehgal [7], who established COSMOS, an advanced AI framework for video generation and simulation based on a neuro-symbolic architecture. Learning to simulate physical (physiological) phenomena is an inherent property of such a system.

The architecture is exemplified in a limited proof-of-concept model:

We converted anatomical prose into TSV triples of the form <head>tab<relation>tab<tail> (e.g., “Aorta → passes_-through → aortic hiatus”) strictly from explicit statements and independently validated the result. Rules were extracted by GPT-5 in a clean, single-turn chat, given only the passage and a “no prior knowledge/assumptions” instruction, with no project-specific priming. Repeating the extraction in fresh, context-free sessions yielded materially identical rule sets, supporting reproducibility.

Using the vetted rules, a standalone Python program produced a two-dimensional knowledge graph that visualizes spatial/topographic relations (left/right; superior/inferior) alongside functional relationships as seen in Figure 1.

From a 207-word passage we obtained 46 rules; consequently, several nodes are constrained by a single statement, which limits spatial determinacy. This limitation is expected for sparse inputs and diminishes as the corpus expands and the density of explicit rules increases. However, for a true topographical model considerable improvements are required.

Discussion

With the neurosymbolic anatomy engine we describe for the first time a plausible architecture for encoding the entire human anatomy and physiology in an AI-system in a way that makes it manipulable, extendable and especially amendable for anatomical and surgical reasoning. In contrast to current efforts using annotated surgical videos to train AI-systems to recognize surgical anatomy, our approach aims to leverage the existing anatomical knowledge from textbooks and atlases as well as digital sources. In analogy to a human surgeon, we aim at creating a deep understanding of anatomy that allows for predicting the topography of anatomical structures even if they are not directly visible by inference in terms of relative positions

We believe that our architecture for the anatomy engine is an especially suitable use case for a neuro-symbolic AI due to the following aspects.

Construction of a topographical world model

Just like a surgeon, who navigates the body through learned spatial relationships rather than memorized coordinates, the anatomy engine encodes anatomy as a system of logical rules, not a fixed 3D map.

From a scientific point of view, topographical anatomy is a rule and logic-based system of relative positions in a confined space. This means that every structure can be described as a set of steps starting from a reference point, e.g., the portal vein lies dorsal to the gastroduodenal artery which is the first large branch of the common hepatic artery which is the right brand of the celiac trunk which is the first anterior branch of the aorta below the diaphragm. This set of rules can be thought of as knowledge graphs that can be arranged topographically according to the rules described in the anatomy textbooks, like outlined by the group of Wermter[3].

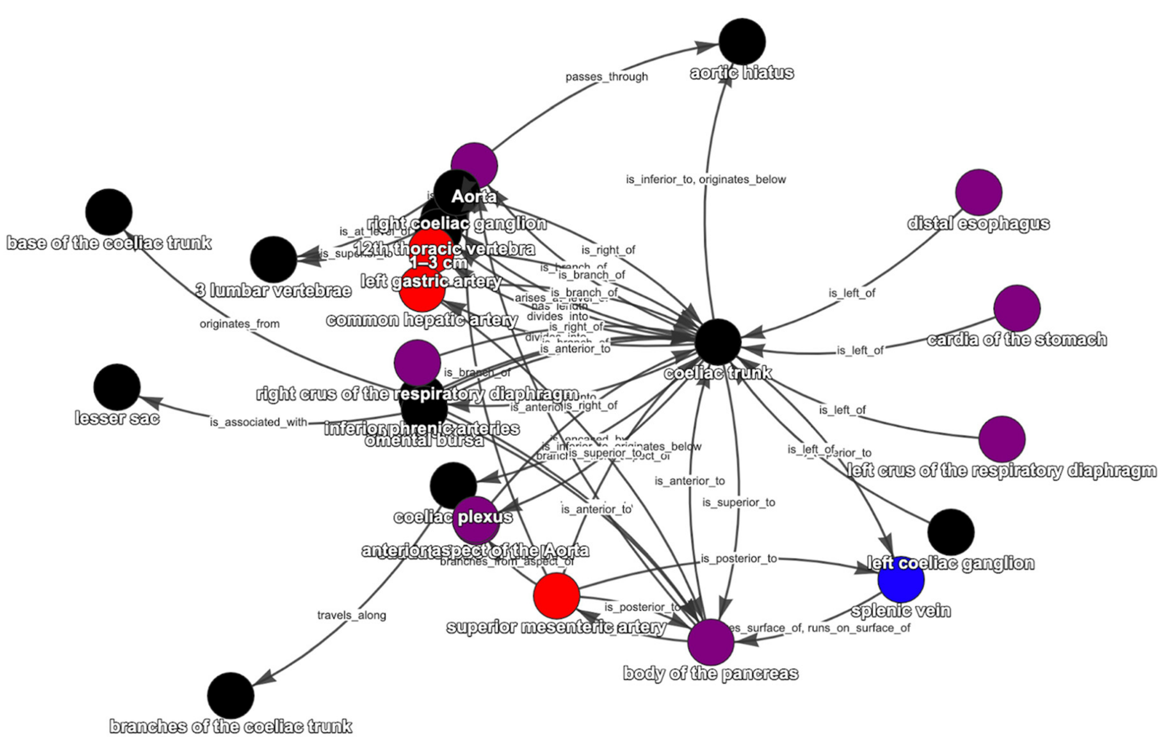

Figure 1: Proof-of-concept. 2D graph from LLM extracted rules; poisons set by left/right and suprior/inferior, author relations as edge labels. Colours label categories as identified by the LLM.

This rule and logic-based description, however, does not fit into VTK or GMSH frameworks because it lacks exact spatial positions. However, this is the way surgeons navigate through anatomy because it allows them to compensate for the variabilities in size and even structure if alternative rules can be applied. In addition to topographical features there are also physiological aspects like elasticity or adherence to neighboring structures (e.g., fascias) that can be coded with every structure which is a prerequisite for simulation.

Interpretability considerations

Unlike black-box AI systems, the anatomy engine can show its reasoning and when errors occur, they can be traced and corrected manually.

The interpretability advantages of the neuro-symbolic anatomy engine require careful qualification based on current evidence from neuro-symbolic AI research. A comprehensive survey of 191 neuro-symbolic studies found that although these approaches hope to enhance explainability through symbolic reasoning, the actual improvements are often less evident than theoretically expected [8]. The effectiveness of interpretability enhancements depends heavily on specific architectural choices and implementation details.

The neuro-symbolic anatomy engine addresses these limitations through explicit symbolic representation of anatomical relationships, constraint logic programming for transparent decision validation, and traceable inference paths through the anatomical knowledge graph. Consequently, errors or misaligned representations within the symbolic layer can be systematically identified and corrected through manual intervention, which constitutes a distinct advantage for applications in the medical domain. However, achieving meaningful interpretability and enabling efficient debuggability in practice will require careful design of the symbolic reasoning components and thorough validation of the transparency benefits in clinical settings.

Data efficiency potential

Because the system reasons from symbolic knowledge rather than relying solely on examples, it can make accurate predictions even where training data is scarce. This is critical for rare procedures or unusual anatomical variants.

Evidence for improved data efficiency in neuro-symbolic approaches shows promise in structured domains where symbolic knowledge provides strong inductive biases [9]. Studies demonstrate that coupled neuro-symbolic systems can achieve more data-efficient learning by using symbolic planning and constraint guidance. In medical domains specifically, neuro-symbolic approaches can effectively learn and reason under conditions of limited data while leveraging transfer learning capabilities.

The data efficiency mechanism operates through symbolic reasoning that allows models to make accurate predictions even when training data is limited. This is particularly relevant for rare surgical procedures or anatomical variations where collecting large training datasets is impractical or impossible. However, the construction of symbolic knowledge often requires substantial manual effort from domain experts, potentially offsetting some of the data efficiency gains.

Scalability and limitations

Scaling the anatomy engine to clinical deployment poses substantial challenges.

Real-time surgical applications demand inference latencies below 100 ms. Current neuro-symbolic architectures have not yet reliably met this requirement at the scale of a full anatomical knowledge graph. This will necessitate dedicated hardware acceleration and algorithm optimization.

Building and maintaining the symbolic knowledge base requires sustained, high-level expert input: anatomical relationships must be manually curated, validated, and updated as medical understanding evolves, representing an ongoing organizational and scientific commitment rather than a one-time engineering effort.

Finally, integrating the anatomy engine into existing clinical and surgical systems introduces significant deployment complexity, including interoperability with established medical imaging pipelines, compliance with regulatory standards, and adaptation to varied institutional workflows. Neither can be addressed through architectural design alone.

Conclusions

This paper presents the neuro-symbolic anatomy engine, a comprehensive architecture that addresses fundamental limitations in current surgical simulation systems through the integration of structured anatomical knowledge with neural pattern recognition. The proposed system makes three key technical contributions: a universal anatomical knowledge representation that enables procedure-agnostic simulation, a neural-symbolic fusion mechanism that grounds perceptual inputs in structured anatomical knowledge, and a real-time reasoning engine capable of maintaining anatomical consistency during dynamic surgical procedures.

The neuro-symbolic anatomy engine's primary innovation lies in its anatomy-centered rather than procedure-centered design philosophy, potentially enabling simulation of novel surgical techniques through compositional reasoning over anatomical relationships.

However, several significant challenges must be addressed for practical deployment, including the substantial manual knowledge engineering, the computational complexity of and the need for extensive clinical validation to demonstrate safety and efficacy in medical applications. However, taking into account the rapid evolution of AI technology we are convinced that it is due time to start the endeavour of creating a full scale manipulable and dynamic AI-model of human anatomy and physiology to bring excellence and safety in surgery to a completely different level, and to speed up surgical education to make it compatible with today's requirements.

Research ethics: N/A.
Informed consent: N/A.
Author contributions: SB: Conceptualization of the model together with PB, AB and PR, writing the manuscript, PB: Proof-of-concept-model, Python programming, Rule xtrction from LLM. AB Manuscript writing, literature research, PR reviewing manuscript. AM: Reviewing Manuscript. All authors have accepted responsibility for the entire content of this manuscript and approved its submission.
Use of Large Language Models, AI and Machine Learning Tools: The manuscript was based on a synthesis of a presentation deck the Claude® LLM. The deck, however, was established by the authors. The text was then extensively rewritten and modified by the first author. The model itself is based on extraction of anatomical rules by LLMs as described in the man text (GPT5). The first version of the code was generated by GPT5 and was modified by PB.
Conflict of interest: SB, AB, PB and PR are founders and shareholders of SBM/Simultare GmbH that intends to commercialize the anatomy model. AM has no conflict of interest. Currently, no revenue is generated.
Research funding: None declared.
Data availability: The complete code of the proof-of-principle model is shared as supplemental material.

References

1. Yuen J. What is the single greatest challenge to training the future generation of surgeons? How can we overcome it? Surgery 2015;33: 290–3.
2. Claassen L, van Workum F, Rosman C. Learning curve and postoperative outcomes of minimally invasive esophagectomy. J Thorac Dis 2019;11: S777–s85.
3. Hopper AN, Jamison MH, Lewis WG. Learning curves in surgical practice. Postgrad Med J 2007;83:777–9.
4. Yaacobi E, Shachar T, Segal D, Agabaria A, Halima G, Marom O, et al. Unveiling the economic toll of surgical learning curve in elderly hip fractures. J Clin Med 2023;12. https://doi.org/10.3390/jcm12154880.
5. Consortium S. Modular FEM mechanics 2026. https://www.sofa-framework.org/applications/plugins/modular-fem-mechanics/ [accessed 03 December 2026].
6. Radford A, Kim JW, Hallacy C, Ramesh A, Goh G, Agarwal S, et al., editors. Learning transferable visual models from natural language supervision. International Conference on Machine Learning; 2021.
7. Sehgal A, Grayeli A, Sun JJ, Chaudhuri S. Neurosymbolic grounding for compositional world models. arXiv preprint; arXiv:231012690. 2023.
8. Zhang X, Sheng VS. Neuro-symbolic ai: explainability, challenges, and future trends. arXiv preprint; arXiv:241104383. 2024.
9. Velasquez A, Bhatt N, Topcu U, Wang Z, Sycara K, Stepputtis S, et al. Neurosymbolic AI as an antithesis to scaling laws. PNAS Nexus 2025;4: pgaf117.

Supplementary Material: This article contains supplementary material (https://doi.org/10.1515/iss-2025-0045).