Learning as a Diffusion Process: Towards Trustworthy, Multimodal, and Efficient Generative Tutors

Sotirios Chatzis
Department of Electrical Engineering and Computer Science
and Engineering at the Cyprus University of Technology
sotirios.chatzis@cut.ac.cy

Abstract: The current generation of Intelligent Tutoring Systems (ITS) is largely built upon autoregressive large language models, implicitly framing learning as a sequential token prediction task. While powerful, this paradigm remains fundamentally limited in its ability to represent uncertainty, explore alternative reasoning paths, and align with the inherently stochastic, multimodal, and interactive nature of human learning.

In this keynote, I will introduce a forward-looking paradigm shift: modeling learning itself as a diffusion process. Building on recent advances in diffusion-based and latent-variable generative models, I will present a framework in which knowledge acquisition, reasoning, and pedagogical intervention are treated as stochastic trajectories in structured latent spaces. This perspective enables the design of trajectory-space tutors, capable of generating and refining multiple candidate explanations, reasoning paths, and feedback signals under explicit uncertainty modeling.

Within this framework, I will introduce the notion of pedagogical control policies, where tutoring is formulated as the controlled steering of generative processes toward desired learning outcomes. This allows for adaptive curriculum generation, counterfactual exploration of learning strategies, and principled difficulty calibration through controllable generative dynamics. Extending beyond text, I will discuss multimodal diffusion architectures that unify language, vision, and procedural signals, enabling richer and more grounded educational interactions.

A central pillar of this vision is efficiency and deployability. I will highlight recent advances in quantization, compressed generative modeling, and edge-aware architectures that enable diffusion-based systems to operate under realistic computational constraints—supporting the transition from cloud-bound prototypes to scalable, real-world ITS deployments.

Equally critical is the need for trustworthy and aligned educational AI, in line with emerging European priorities on robust and human-centric AI systems. I will show how distribution-aware generative models provide a natural foundation for uncertainty quantification, attribution, bias mitigation, and controllability, addressing key requirements of the EU AI Act and related frameworks for trustworthy AI in high-impact domains such as education.

The talk will conclude with a research agenda at the intersection of generative AI and learning sciences, positioning ITS as a key driver for next-generation AI: systems that do not merely predict the next token, but shape learning as a probabilistic, controllable, and multimodal generative process—efficient, trustworthy, and aligned with human educational values.

Back to the ITS2026 homepage