

Chance-Constrained Stochastic Systems and Generative AI: Shared Mathematical Foundations and Pathways to Safe AI Models

Ashkan Jasour, 2025

In “stochastic systems”—such as Partially Observable Markov Decision Processes (POMDPs) and Stochastic Differential Equations (SDEs)—randomness in behavior leads to system states being represented by probability distributions. Achieving the desired behavior in such systems (e.g., meeting safety constraints and attaining specific goals) requires designing system parameters, such as control inputs, to maximize the probability of success or limit the risk of failure under uncertainties arising from state probability distributions [1]. This process involves reshaping the probability distributions of system states through parameter adjustments to satisfy probabilistic constraints, commonly formulated as chance constraints [1, 2].

Coming from a background in mathematics, stochastic systems, and AI, my insight is that generative AI models and chance-constrained stochastic systems share the same underlying mathematical principles. In both cases, the objective is to obtain a desired probability distribution that captures the target behavior—such as achieving a probability distribution that maximizes the probability of success in stochastic systems or achieving a probability distribution to enable high-quality, diverse data generation in generative AI models.

More specifically, in “generative AI”, the goal is to learn the underlying distribution of a dataset and generate new data by sampling from it. One example is autoregressive models, such as large language models (LLMs), which are trained to learn a probability distribution over the next token given the preceding context. Another example includes diffusion models and flow-matching models, where the process starts with pure noise and learns a transformation—denoising process or a continuous flow—that maps the noise distribution to the data distribution, enabling the generation of new data samples [3]. This process of transforming an initial noise distribution into a desired one closely parallels what occurs in stochastic systems, where the goal is to steer the system’s initial probability distribution toward a target distribution that satisfies risk or safety constraints [4]. In generative AI, this transformation is accomplished through learned mechanisms—such as denoising processes, continuous flows, or autoregressive training—whereas in stochastic systems, it

is guided by (uncertainty-aware) planning and control algorithms. In both cases, the fundamental challenge is the same: controlling the evolution of probability distributions—whether to ensure safety in stochastic systems or to generate high-quality, diverse data in generative.




Figure: Stochastic Systems Vs Generative AI - [Right] Generative AI systems are at the forefront of the recent AI revolution, driving advancements across multiple domains. The primary objective of GenAI is to model complex data distributions, either explicitly—through methods such as variational autoencoders (VAEs), autoregressive models, normalizing flows, and diffusion models—or implicitly, as seen in generative adversarial networks (GANs). For instance, denoising diffusion and flow matching approaches have emerged as state-of-the-art techniques in various fields, including image generation (e.g., Stable Diffusion), video generation (e.g., Sora), and scientific applications (e.g., AlphaFold3); These models generate new data by transforming simple initial distributions (e.g., Gaussian noise) into complex data distributions, which are then used in the sampling process to produce realistic outputs. To achieve this, they construct/learn a mapping—represented by neural ordinary differential equations (ODEs) or stochastic differential equations (SDEs)—that iteratively refines noise into meaningful data. [Left] Risk-aware autonomy involves transforming the initial distribution of an stochastic system's states—represented by ODEs or SDEs— into desired probability distributions that represent the system's safe and optimal behavior. Under these probability distributions, the system must satisfy safety constraints and achieve optimal behavior with high probability, ensuring risk-bounded performance.

Given this observation, we can leverage similar reasoning and algorithmic tools from stochastic systems to develop risk-aware AI models. For example, this involves: 1) Uncertainty characterization, where input uncertainty arises from adversarial distributions, noise, or perturbations, 2) Uncertainty propagation through the AI model to quantify the

resulting output distribution for safety and robustness analysis, 3) Risk-aware training, where the propagated uncertainties are incorporated during training to enforce safety and improve robustness. Such an uncertainty-aware framework for AI models enables the estimation of the full range of possible outputs corresponding to an input distribution that captures a variety of potential input scenarios. For example, the input distribution may represent a set of perturbed images or semantically similar text samples. By propagating this distribution through the model, we can analyze the full spectrum of resulting outputs or behaviors. Unlike traditional safety approaches that focus on specific adversarial or stress-test scenarios, this framework supports a more comprehensive safety and robustness analysis by considering a family of input-output behaviors. It also allows us to quantify robustness under input perturbations by estimating the risk of producing undesirable outputs—for instance, the risk of generating harmful or disallowed tokens, or the risk of misclassifying images or text under adversarial attacks. Additionally, the framework enables evaluation of robustness to semantically similar inputs by assessing the model’s consistency in its responses to inputs that differ slightly in form but are equivalent in meaning.

[1] Ashkan Jasour, “Risk Aware and Robust Nonlinear Planning”, MIT 16.S498 Graduate Course, <https://ocw.mit.edu/courses/16-s498-risk-aware-and-robust-nonlinear-planning-fall-2019/>

[2] Ashkan Jasour, Necdet S Aybat, and Constantino M Lagoa. “Semidefinite programming for chance constrained optimization over semialgebraic sets”. In: SIAM Journal on Optimization 25.3 (2015), pp. 1411–1440.

[3] Ashkan Jasour, "Generative AI Foundations: Algorithms and Architectures", Course <https://jasour.github.io/generative-ai-course> , 2024-2025

[4] <https://www.ajasour.com/risk-aware-ai>