Research Statement

Anindya Sarkar

Ph.D. Candidate, Department of Computer Science & Engineering, Washington University in St. Louis

anindyasarkar.ece@gmail.com | anindyasarkariith.github.io

1 Overview

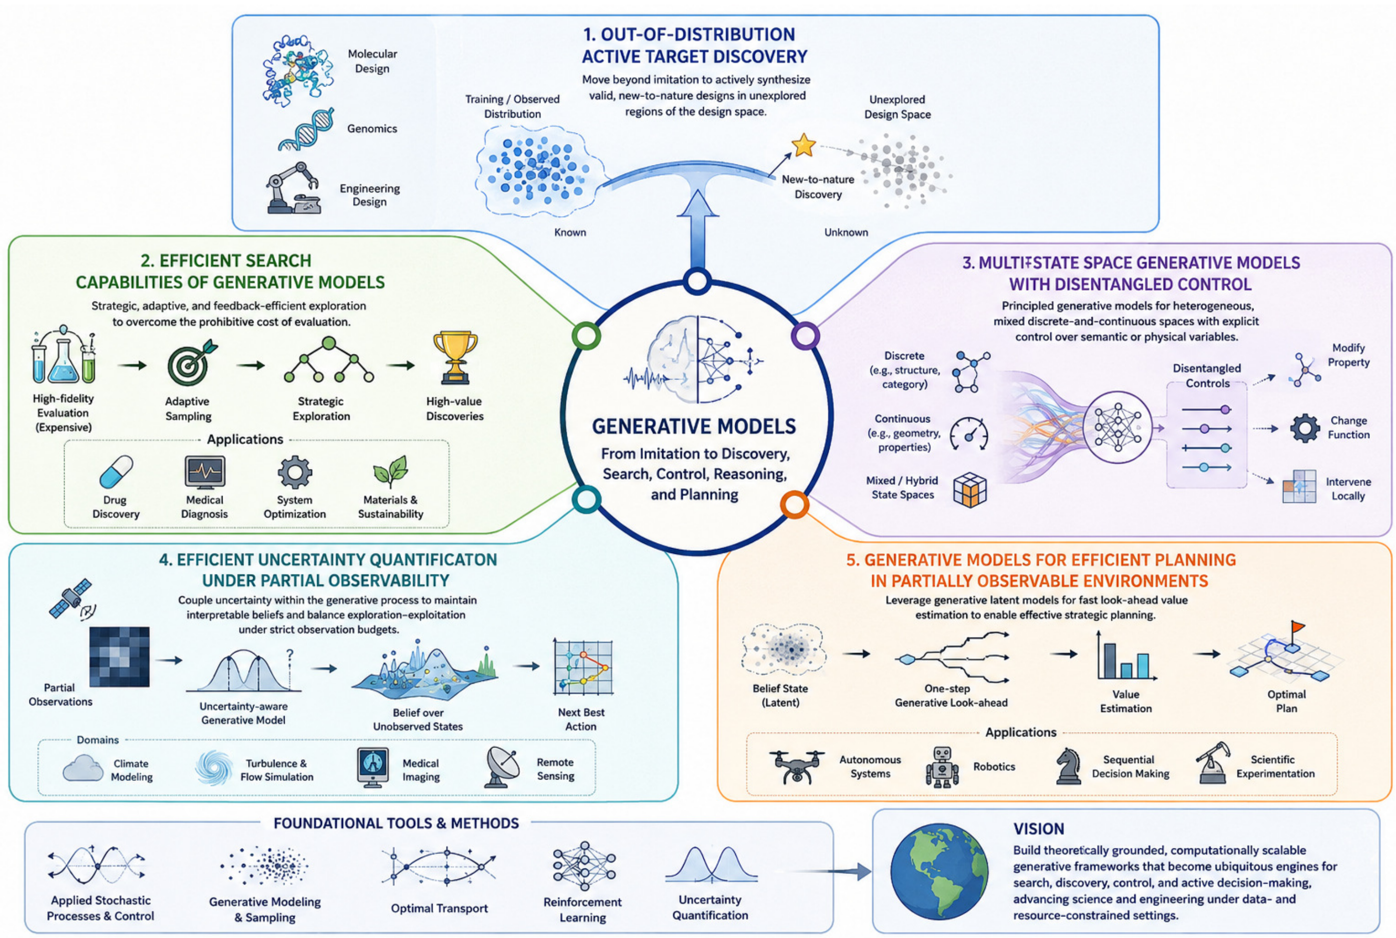

Figure 1: An overview of research themes. Generative Models: From imitation to discovery (Image Credit: ChatGPT)

Generative models' transformative impact spans a vast spectrum of disciplines, driving innovation in foundational engineering applications like vision and text, while simultaneously unlocking unprecedented capabilities in complex scientific domains such as molecular design and genomics. However, while modern generative models—such as diffusion models, flow matching, and LLMs—have mastered the art of resembling the training distribution, scientific breakthroughs demand more than mere imitation. In science and engineering, the most valuable discoveries lie far beyond the observed data, hidden in vast, unexplored regions of the design space. The true frontier is *generative discovery*: moving beyond replicating the known to actively synthesize valid, new-to-nature designs that standard models currently ignore as mathematically improbable. I refer to this line of research as *Out-of-Distribution Active Target Discovery*. Moreover, across many resource-constrained domains—progress is fundamentally bottlenecked by the staggering cost of evaluation. Whether synthesizing novel therapeutics (such as drug discovery) or optimizing interactive systems (such as medical diagnosis), acquiring high-fidelity feedback is severely resource-constrained. Overcoming this barrier demands sampler capable of strategic, adaptive, and feedback-efficient exploration. By equipping generative models with enhanced exploration capabilities, I aim to minimize the prohibitive physical and human feedback cost, unlocking targeted, global search in complex, high-dimensional environments. I refer to this research direction as *Efficient Search Capabilities of Generative Models*. Beyond search efficiency, a critical limitation of current generative models is their lack of fine-grained controllability; altering a single attribute typically forces a complete resampling of the entire output. Moreover, real-world applications—ranging from scientific discovery to clinical diagnostics rely on isolated interventions across heterogeneous state spaces with complementary effects on the target sample. For example, rather than redesigning molecules from scratch to adjust a single property, chemists require precise interventions. They must hold a core scaffold fixed while tuning solubility, or preserve binding while optimizing accessibility—balancing competing objectives across heterogeneous state spaces. To address this, I aim to develop principled generative models that natively support heterogeneous, mixed discrete-and-continuous state spaces with explicitly addressable control variables, enabling precise interventions on individual semantic or physical variables. I refer to this research direction as *Multi-state space Generative Models with Disentangled Control*. Beyond discovery, search, and controllability, I investigate their capacity to drive reasoning in complex, partially observable environments, such as turbulence-flow simulation, climate modeling, medical imaging, and remote sensing. Under strict observation budgets and uninformative priors, effective reasoning hinges entirely on reliable uncertainty quantification to dynamically balance exploration and exploitation. To this end, this line of research aims to couple uncertainty estimation directly within the generative stochastic process. By repurposing generative models to maintain interpretable belief distributions over unobserved states, this approach aims to provide the robust uncertainty metrics necessary for strategic planning and reasoning in partially observable domains. I refer to this research direction as *Efficient Uncertainty Quantification under Partial Observability*. Additionally, during my Ph.D., I explored how generative latent representations can guide sequential decision-making. To extend this, I aim to develop principled generative models tailored specifically for robust value estimation. By advancing recent one-step generators like Wasserstein gradient flows and flow maps, this research will enable rapid look-ahead evaluations to drive highly effective strategic planning. I refer to this research direction as *Generative Models for Efficient Planning Under Uncertainty*. In a nutshell, as a postdoctoral researcher, my goal is to develop theoretically grounded and computationally scalable machine learning frameworks that push the boundaries of generative models in data- and resource-constrained domains. By drawing on foundational tools from applied stochastic processes & control, Generative Modeling & Sampling, Optimal Transport, and Reinforcement Learning, I aim to contribute to the development of a fundamentally new class of generative models and to advance them into ubiquitous engines for search, discovery, control, and active decision-making across science and engineering. In the following sections, I detail my recent contributions and outline my future vision for each of these core research directions.

2 Search Capabilities of Generative Models

Inference-time search can be framed as sampling from a reward-tilted distribution. While current methods—such as optimal control frameworks (e.g., MFM [1]), SMC-based samplers (e.g., DAS [2]), Feynman-Kac correctors (e.g., FKC [3])—attempt to address this, they are fundamentally bottlenecked by mode collapse, weight degeneracy, and reward over-optimization, preventing efficient coverage of the true target distribution. To overcome these limitations, we introduce **IMPFM** [4] (Sequentially-Controlled Interactive Multi-Particle Flow-Maps). IMPFM replaces the standard single-particle dynamics of training-free guidance with an *interacting* particle system. By correcting each particle's drift with an ensemble-aggregated, kernel-weighted value signal and an explicit repulsive term to discourage redundant exploration, IMPFM induces an interaction-aware Feynman-Kac corrector. We show this formulation effectively steers the ensemble toward a KL-tilted target distribution while strictly preventing the mode collapse and over-optimization that plague independent particle-based guidance. As shown in Fig. 2 (left), IMPFM promotes substantially greater sample diversity- and sustains a persistently high ESS across all correction steps— failure modes endemic to the FKC samplers. Similarly, Fig. 2 (right) demonstrates that incorporating multi-particle interactions into an optimal control framework (referred to as MPC) mitigates mode collapse and reward over-optimization, as evidenced by richer sample diversity and consistently higher VQA scores measured via InstructBLIP.

Although tree-based samplers (e.g., DTS [5]) avoid intermediate reward evaluations by backpropagating ter-

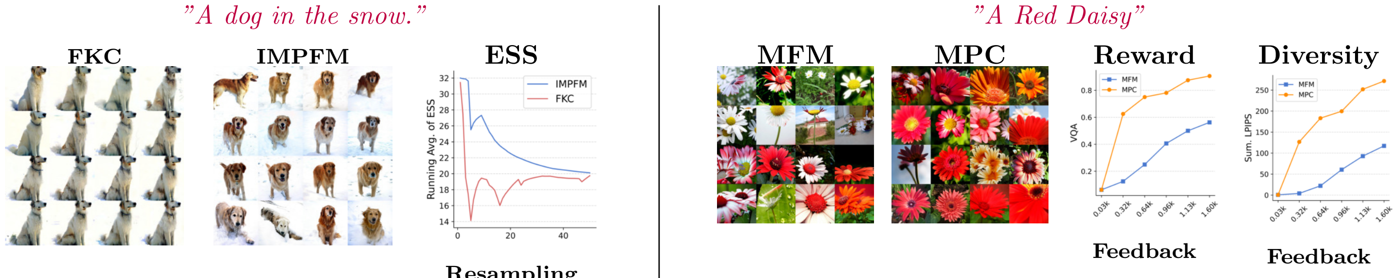

(a) Impact of Particle Interaction on FKC.

(b) Impact of Particle Interaction on Optimal Control.

Figure 2: Impact of IMPFM.

minal rewards— mitigating key issues in SMC, FKC-style methods—they remain severely bottlenecked by the massive number of function evaluations (NFEs) required for node valuation. To overcome this, we introduce Bootstrap Flow-Map-Tree (**BFMT**) [6], a computationally efficient framework for budget-constrained inference-time search and alignment. Using a novel bootstrap sufficient statistic mechanism, BFMT constructs full tree paths from any depth with a single NFE, drastically reducing overhead while preserving sequential foresight. Furthermore, dynamic time-step scheduling—powered by the bootstrap sufficient-statistic—equips BFMT with adaptive search capabilities. This ensures optimal budget allocation, seamlessly transitioning from broad global exploration to fine-grained local refinement. See Fig. 3 for an example visualization of BFMT's exploration strategy.

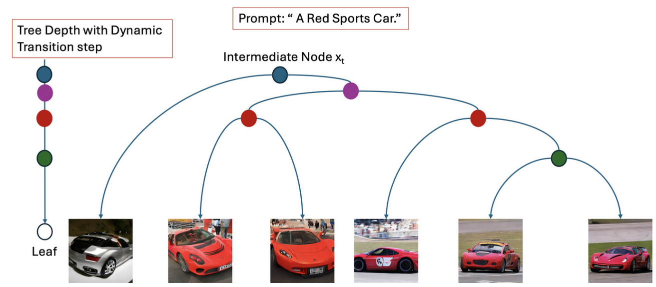

Figure 3: BFMT's Hierarchical Search Strategy.

During my Ph.D., I also investigated the adaptive search capabilities of generative models in online, budget-constrained settings. Here, the target is initially unknown and only revealed through costly sequential feedback. The core challenge is to generate target-aligned samples while minimizing these expensive interactions. To tackle this, it requires a sampler that strategically navigates the underlying distribution based on online feedback. To this end, we develop **PAPA** [7], a principled framework that efficiently fine-tunes the generative model using online feedback (exploitation). By introducing an objective that restricts the updated model from drifting too far from its pre-trained reference (exploration), PAPA seamlessly balances rapid adaptation with the preservation of high-quality, diverse samples.

Looking forward, I am driven to translate these foundations into high-impact scientific discoveries. Through active collaboration with the AstraZeneca Molecule Discovery team, I am expanding this framework into the complex realms of molecule and protein design. However, a critical barrier remains: current generative approaches falter when reward signals are exceptionally sparse, leaving valuable samples trapped within rare modes of the prior distribution. I aim to develop principled frameworks to dismantle these bottlenecks, which is the key to unlocking the real-world, practical deployment of generative models.

3 Out-of-Distribution Active Target Discovery

Every method in Section 2 shares an implicit assumption: the target (e.g., a molecule with a set of desired properties or an image with certain features) being searched for is reachable by reweighting or steering the support of an already-trained generative model. This assumption is reasonable when the target is rare-but-typical — an unusual combination of otherwise familiar features. To validate this hypothesis, we recently proposed a bootstrap stochastic interpolant tree sampler (**BSIT**)[Under Review] that decouples child node generation from explicit parent dependencies. By expanding the degrees of freedom during child node sampling, BSIT encourages diverse trajectory rollouts and unlocks broader path-space diversity—the exact catalyst needed for efficient exploration. Crucially, we show that BSIT's improved exploration enables broader distributional coverage than samplers that rely on DDPM-like transitions for child-node generation, all without incurring additional function evaluations or reward feedback, thereby providing a highly efficient basis for rare target search. Fig. 4 visualizes target samples with rare compositional features that are scarce or absent in the training data (e.g., images of "three cats" (left) or "a white fish" (right) are rare or absent in the training set of ImageNet). While existing tree samplers struggle to synthesize these rare structural compositions, BSIT generates them effectively. Motivated by the same objective, we introduce the Lévy Adaptive Tree Sampler

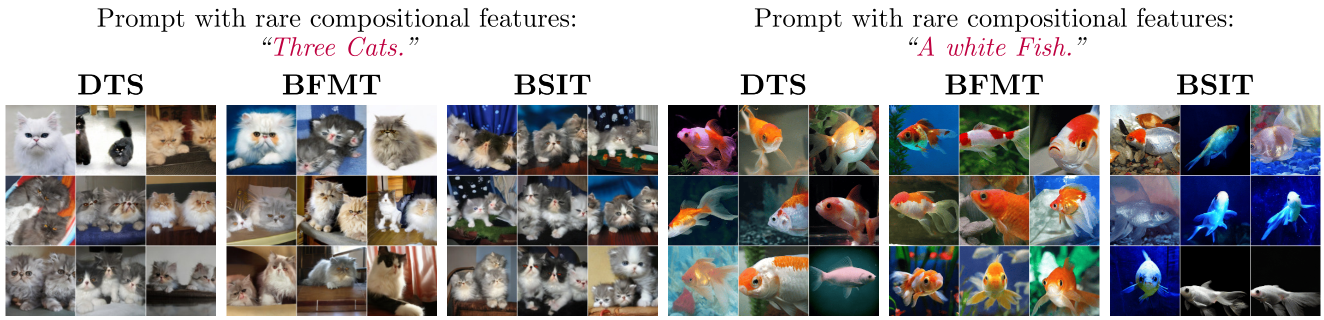

Figure 4: BSIT's capability of generating samples with rare compositional prompts.

(**LATS**)[Under Review]. By leveraging a Lévy process for tree construction, LATS achieves broad exploration to uncover rare modes. Paired with value-guided node selection for targeted exploitation, LATS successfully generates target-aligned samples from low-likelihood tail regions of the underlying distribution.

However, these approaches break down for a large and scientifically important class of problems in which the target lies genuinely *outside* the training distribution: a material with a property never seen in the training corpus, a molecular scaffold structurally distant from anything the model was fit on, an emergent pathogen variant. Active search machinery that only reweights an existing density has no mechanism for reaching such targets; what is needed is principled *expansion* of the region a generative model can be trusted to describe, not just efficient search within it.

This is the focus of an ongoing project where we actively expand the Wasserstein gradient flow by factorizing it into transport and explicitly controllable expansion maps. Crucially, we train this expansion map via a meta-learning objective to robustly learn the dynamics of controllable support expansion. By inspecting the principle components of the resulting expansion velocity, we extract a precise novelty signal that tracks movement outside the training manifold. This meta-learned framework systematically extends a generative model's support to uncover target-aligned, out-of-distribution (OOD) examples.

Beyond extending these methods to the sciences, my future research aims to maximize the feedback efficiency of this out-of-distribution (OOD) expansion framework. Because ground-truth evaluations are prohibitively expensive in real-world applications, we must systematically quantify how far our model's extrapolations can be trusted at every interaction step. I view this trust-calibrated expansion as the essential bridge between active search and genuine scientific discovery—where the most transformative answers are, by definition, those a static training set could never reveal.

4 Efficient Uncertainty Quantification Under Partial Observability

Beyond establishing generative models as powerful engines for search and discovery, I plan to pursue a complementary research direction focused on a critical challenge: how can these models efficiently quantify their uncertainty over unobserved target states to enable informed decision-making in partially observable environments? Concretely: *can we design an algorithm to uncover a target while minimizing expensive ground-truth feedback and avoiding task-specific supervised training?* Consider MRI diagnosis: scanning the entire brain is prohibitively costly and risky. Instead, a doctor must adaptively select subsequent scan regions based on prior observations to localize a tumor using as few scans as possible. To this end, we propose **DiffATD** [8], which repurposes a pretrained diffusion model as an online belief-tracker over the unobserved search space: At each step, it runs a batch of reverse-diffusion particles conditioned on the measurements collected so far, and *disagreement/agreement* among these particles' denoised predictions directly quantifies *uncertainty/likelihood* in each unmeasured region — no separate uncertainty model is needed. It then samples a candidate location by balancing a diffusion-derived exploration score with an exploitation score. The latter integrates a diffusion-based likelihood over the unobserved measurement locations with a lightweight reward network updated after each measurement, ultimately enabling label-free, budget-aware active search driven by the generative model's internal uncertainty. Fig. 5 illustrates how DiffATD dynamically adapts its exploration strategy based on the sequential feedback across two different stages of the search process. Despite its advantages, DiffATD faces a fundamental limitation: its strict reliance on an inherently informative diffusion prior. In scenarios where data is exceptionally scarce or the target is entirely novel, this prior collapses, severely limiting the model's effectiveness. Diagnosing a rare tumor perfectly exemplifies this challenge—without sufficient data to train a reliable prior, the search strategy is forced to rely on an uninformative prior. Addressing this critical vulnerability is the primary motivation for our subsequent framework, **EM-PTDM** [9]. Concretely, EM-PTDM addresses active discovery when no informative prior exists to begin with: it pairs a frozen, pre-trained diffusion model as permanent memory (generalizable structure learned once, never updated) with a lightweight transient memory module that learns only a task-specific correction to the score via Doob's h-transform — since the true conditional score given observations is analytically intractable, the h-transform learns exactly the residual term needed to steer the frozen diffusion dynamics toward consistency with what's been measured so far, without ever backpropagating through the frozen model. Fitting this correction is cast as Expectation-Maximization: the E-step draws posterior samples over the unobserved search space using the current permanent+transient score, and the M-step updates the transient (h-transform) parameters to make those samples more likely under the next round of observations — each EM cycle provably raises the expected log-evidence of the prior (a monotonic-improvement guarantee), so the belief the agent searches with only ever gets sharper as measurements accumulate, even starting from an uninformative prior. Fig. 6b (right) demonstrates how Doob's h-transform rapidly corrects the posterior estimate using sparse observations under an uninformative prior. Furthermore, Fig. 6a (left) contrasts the sampling strategies across different search phases, demonstrating EM-PTDM's superior ability to uncover significantly more distinct target regions than DiffATD under an uninformative prior for a fixed sampling budget.

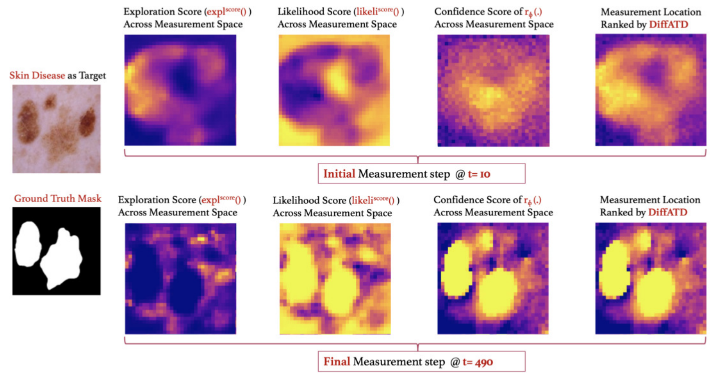

Figure 5: Explanation of *DiffATD*. Brighter indicates a higher value and higher rank.

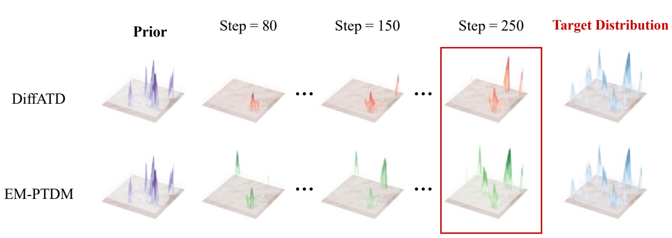

(a) Exploration Behavior Across Discovery Phases. Using a Black-capped Chickadee species distribution as the prior, EM-PTDM most accurately discovers the regions containing the target species (Cedar Waxwing).

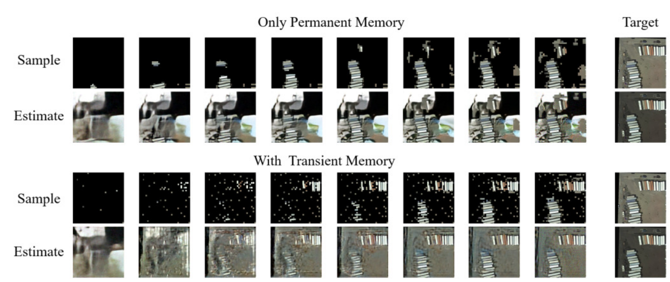

(b) Adaptability of Doob's h-model. Sampled regions and the corresponding Posterior estimates across active discovery phases for overhead objects (i.e., trucks), utilizing an ImageNet-trained prior.

While these frameworks lay the foundation for efficient, uncertainty-aware exploration under partial observability, my future research aims to extend them into complex, high-dimensional scientific domains. To achieve this, I will focus on three critical bottlenecks. First, current uncertainty quantification relies on computationally heavy iterative denoising; I plan to drastically accelerate this process by integrating recent advances in distilled one-step generators [10]. Second, these frameworks currently assume noise-less observations. To enable practical, real-world deployment, my future work will generalize these models to operate robustly in realistic, noisy environments. Finally, in physical domains such as fluid turbulence, posterior samples must strictly adhere to governing partial differential equations (PDEs). Therefore, I aim to develop principled frameworks for uncertainty estimation that natively enforce these physical constraints, ensuring reliable and physically valid exploration across constrained scientific domains.

5 Multi-state space Generative Models with Disentangled Control

Another fundamental limitation of current generative models, distinct from the questions above, is *controllability across heterogeneous state space*: most conditional generative models treat a change in a high-level semantic factor as a request to resample the entire generative trajectory. Ask a generative model to modify a single functional group to improve a molecule's solubility, and in general the entire structure is redrawn—3D conformation, synthetic accessibility, and binding affinity all shift uncontrollably along with the local edit. This is more than a cosmetic annoyance. In scientific and clinical applications, generation is rarely about a single free-floating sample; it is about *editing one factor while explicitly controlling the rest.* The factors that matter are causally or physically meaningful and often span heterogeneous state spaces—some continuous (a dosage, a binding affinity, a physical parameter), some discrete (a functional group, a scaffold topology). Crucially, these factors are deeply interrelated: altering one control variable might cleanly adjust its corresponding attribute and leave the rest untouched, or it might inevitably and adversely impact another, such as when optimizing for solubility inadvertently compromises binding affinity.

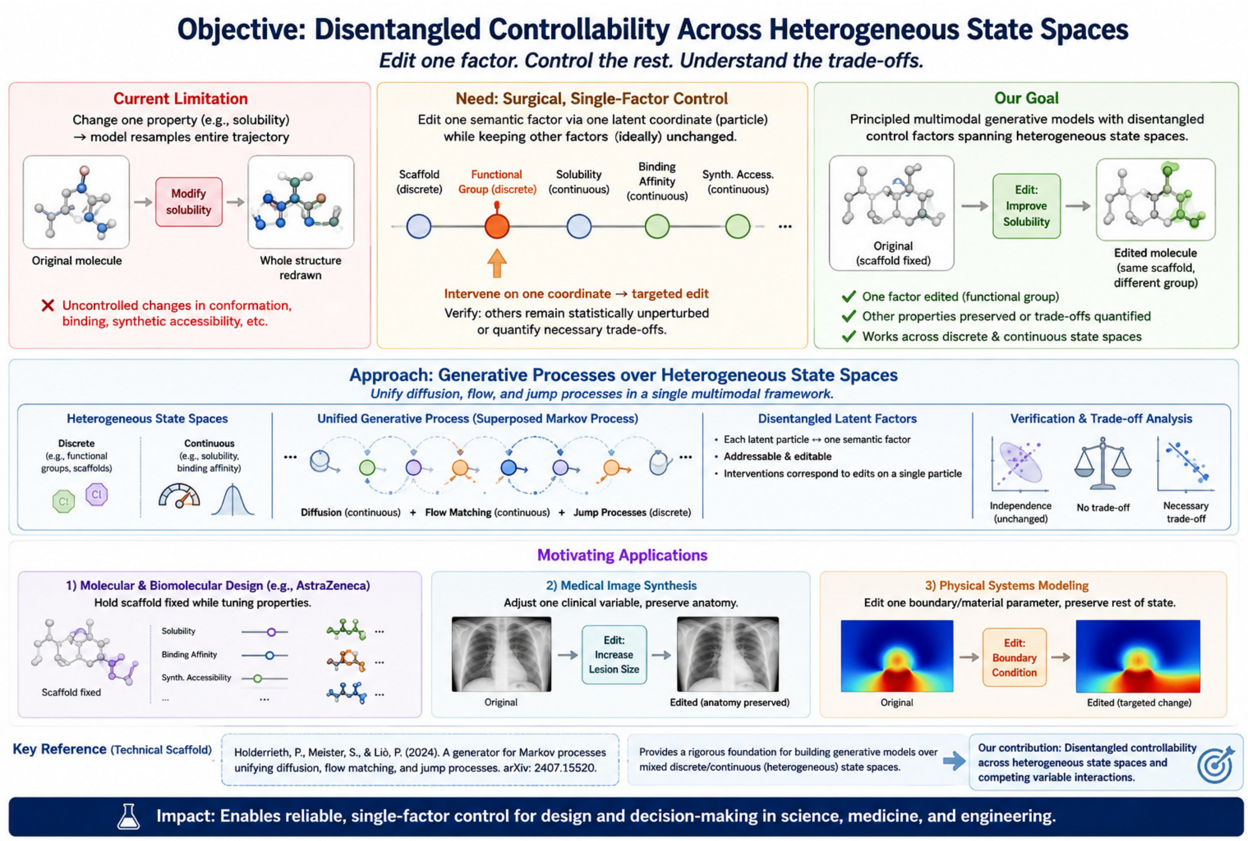

Figure 7: Disentangled Control Across Heterogeneous State-spaces (Image Credit: ChatGPT)

I aim to develop principled multimodal generative models with *disentangled control factors spanning heterogeneous state spaces*, in which a change to one semantic attribute corresponds to an intervention on a single, identifiable latent particle or coordinate. The goal is to isolate these interventions—so that substituting a functional group requires editing exactly one control variable, rather than resampling the entire molecule—while formally accounting for when and how variables compete. Realizing this requires generative models that natively support heterogeneous, mixed discrete-and-continuous state spaces with an explicit, addressable factorization of the control space, together with a mechanism for verifying whether an edit to one factor genuinely leaves the others statistically unperturbed, or if it strictly necessitates a trade-off. Recent work on unifying generative modeling across arbitrary state spaces—for instance [11, 12], which shows how different process spanning different state spaces can be superimposed into generalized multimodal process—provides useful technical scaffolding for building generative processes over mixed discrete/continuous modalities; however, equipping these multi-state spaces generative models with disentangled controllability across different state spaces and enabling competing variable interactions is key for deployment of these models in critical applications across science and engineering. See Fig. 7 for an overview of the objective. In a recent collaboration with the AstraZeneca Molecule Discovery team, we are developing a principled multi-state space generative model with disentangled controllability for multi-objective molecule design. This need for surgical, single-factor control extends beyond chemistry into medical image synthesis, remote sensing, and physical-systems modeling. In my view, disentangled multimodal control across heterogeneous state spaces is the critical missing capability for deploying generative models as reliable design tools in these domains.

6 Generative Models for Efficient Planning Under Uncertainty

I also investigate how the representation learned from generative models aids in planning in partially observable environments. We first explored this in **GOMAA-Geo** [13]. Specifically, GOMAA-Geo tackles active geo-localization: an aerial agent (e.g., a drone-like camera) starts at an unknown position over a large area and must find a specific target location (specified either as natural language text, or a ground-level image or overhead image) within a limited number of movement steps, using only sequential, partial aerial views it acquires along the way. To tackle this, in the first phase, GOMAA-Geo trains an autoregressive model (e.g., an LLM) to predict, at each step, the action that brings the agent closer to a goal conditioned on the trajectory history together with the goal specification. In the next phase, the learned LLM's representation is frozen and handed to a downstream PPO [14] agent as the state — both the actor and the critic condition directly on this learned embedding rather than on raw observations. The subtle but crucial finding is that a representation trained only to imitate "what action gets me closer to the goal" ends up implicitly encoding history-aware, goal-conditioned information about proximity to the goal — i.e., something functionally close to a value signal — even though it was never trained with a value objective. That implicit value-like structure is what makes the representation so effective for the PPO agent's planning under partial observability. In Fig. 8, we present visualizations of an example exploration strategy for the GOMAA-Geo agent across different goal-modality targets.

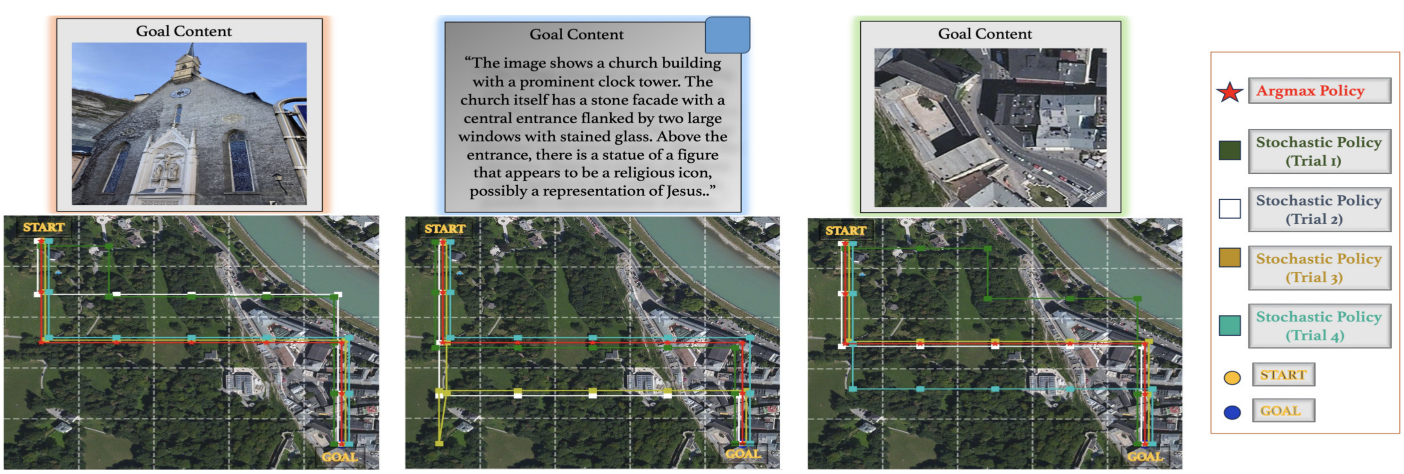

Figure 8: **Example exploration behavior of *GOMAA-Geo* across different goal modalities.** The stochastic policy selects actions probabilistically, whereas the *argmax* policy selects the action with the highest probability.

I extend this line of inquiry to *visual active search* [15, 16], a setting where targets are distributed across multiple regions and agents possess global mobility. Here, the primary objective is to uncover the maximum number of target-containing regions within a strict observation budget. This structure mirrors critical real-world applications, from search-and-rescue operations to anti-poaching efforts. To address this, we developed **DiffVAS** [17], a framework that couples a diffusion-based generative reconstructor with an RL planner. By continuously building a plausible estimate of the entire search space from narrow partial observations, DiffVAS equips the planner with an evolving world model. As this reconstruction sharpens, the agent makes increasingly informed decisions to localize multiple targets within a tight query budget.
While current approaches often use generative models to learn representations for a separate RL planning module, the generative model can theoretically be repurposed to compute values directly, removing the need for a distinct RL component. Although recent work explores using generative models as planners [18, 19]—typically guided by reward signals—efficient value estimation remains a critical bottleneck. Single-step generators [10] offer a promising route for cheap look-ahead value computation; however, jointly learning the environment dynamics and seamlessly coupling them with efficient value estimation remains an unsolved challenge for deploying generative models as the sole planning engine in partially observable environments.

References

[1] Peter Potaptchik, Adhi Saravanan, Abbas Mammadov, Alvaro Prat, Michael S Albergo, and Yee Whye Teh. Meta flow maps enable scalable reward alignment. *arXiv preprint arXiv:2601.14430*, 2026.

[2] Sunwoo Kim, Minkyu Kim, and Dongmin Park. Test-time alignment of diffusion models without reward over-optimization. *arXiv preprint arXiv:2501.05803*, 2025.

[3] Marta Skreta, Tara Akhound-Sadegh, Viktor Ohanesian, Roberto Bondesan, Alán Aspuru-Guzik, Arnaud Doucet, Rob Brekelmans, Alexander Tong, and Kirill Neklyudov. Feynman-kac correctors in diffusion: Annealing, guidance, and product of experts. *arXiv preprint arXiv:2503.02819*, 2025.

[4] Binglin Ji, Anindya Sarkar, Hengchang Lu, Jens Sjölund, and Yevgeniy Vorobeychik. Sequentially-controlled interactive multi-particle flow-maps for online feedback-driven search. *arXiv preprint arXiv:2607.01144*, 2026.

[5] Vineet Jain, Kusha Sareen, Mohammad Pedramfar, and Siamak Ravanbakhsh. Diffusion tree sampling: Scalable inference-time alignment of diffusion models. *arXiv preprint arXiv:2506.20701*, 2025.

[6] Binglin Ji, Anindya Sarkar, Hengchang Lu, Jens Sjölund, and Yevgeniy Vorobeychik. Bootstrap flow-map tree sampling enables online feedback driven search. *arXiv preprint arXiv:2607.02915*, 2026.

[7] Anindya Sarkar, Nasik Muhammad Nafi, Isaac Lyngaas, Muralikrishnan Gopalakrishnan Meena, and Yevgeniy Vorobeychik. Papa: Online personalized active preference alignment. *arXiv preprint arXiv:2607.00486*, 2026.

[8] Anindya Sarkar, Binglin Ji, and Yevgeniy Vorobeychik. Online feedback efficient active target discovery in partially observable environments. *Advances in Neural Information Processing Systems*, 38:160003–160040, 2026.

[9] Anindya Sarkar, Binglin Ji, and Yevgeniy Vorobeychik. Active target discovery under uninformative priors: The power of permanent and transient memory. *Advances in Neural Information Processing Systems*, 38:64560–64598, 2026.

[10] Nicholas M Boffi, Michael S Albergo, and Eric Vanden-Eijnden. How to build a consistency model: Learning flow maps via self-distillation, 2025. *URL https://arxiv. org/abs/2505.18825.*

[11] Peter Holderrieth, Marton Havasi, Jason Yim, Neta Shaul, Itai Gat, Tommi Jaakkola, Brian Karrer, Ricky TQ Chen, and Yaron Lipman. Generator matching: Generative modeling with arbitrary markov

processes. In *International Conference on Learning Representations*, volume 2025, pages 52153–52219, 2025.

[12] John Nguyen, Marton Havasi, Tariq Berrada, Luke Zettlemoyer, and Ricky TQ Chen. Oneflow: Concurrent mixed-modal and interleaved generation with edit flows. *arXiv preprint arXiv:2510.03506*, 2025.

[13] Anindya Sarkar, Srikumar Sastry, Aleksis Pirinen, Chongjie Zhang, Nathan Jacobs, and Yevgeniy Vorobeychik. Gomaa-geo: Goal modality agnostic active geo-localization. *Advances in Neural Information Processing Systems*, 37:104934–104964, 2024.

[14] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

[15] Anindya Sarkar, Michael Lanier, Scott Alfeld, Jiarui Feng, Roman Garnett, Nathan Jacobs, and Yevgeniy Vorobeychik. A visual active search framework for geospatial exploration. In *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 8301–8310. IEEE, 2024.

[16] Anindya Sarkar, Nathan Jacobs, and Yevgeniy Vorobeychik. A partially-supervised reinforcement learning framework for visual active search. *Advances in Neural Information Processing Systems*, 36:12245–12270, 2023.

[17] Anindya Sarkar, Srikumar Sastry, Aleksis Pirinen, Nathan Jacobs, and Yevgeniy Vorobeychik. Diffvas: Diffusion-guided visual active search in partially observable environments. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems*, pages 791–799, 2026.

[18] Anurag Ajay, Yilun Du, Abhi Gupta, Joshua Tenenbaum, Tommi Jaakkola, and Pulkit Agrawal. Is conditional generative modeling all you need for decision-making? *arXiv preprint arXiv:2211.15657*, 2022.

[19] Michael Janner, Yilun Du, Joshua B Tenenbaum, and Sergey Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.