Some particle physics research is predicated on comparing recorded particle collisions with computer simulations to validate theoretical models. While traditional simulations provide accurate results, they are computationally slow and expensive. To address simulation costs, researchers at Fermilab investigated an AI-based diffusion model to generate particle showers. Building upon this work, the University of Chicago's Data Science Clinic has been working on a latent diffusion model to reduce computing power requirements while maintaining a high degree of accuracy.

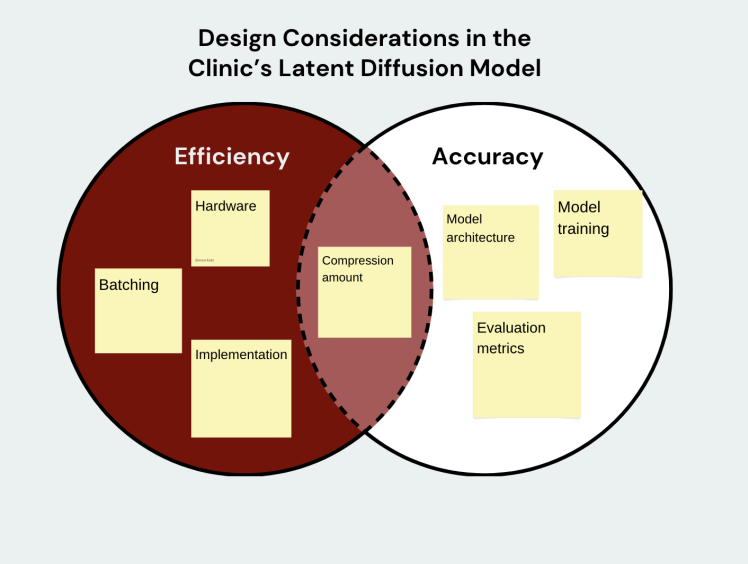

Previous quarters' teams focused on measuring and optimizing efficiency through compression amount and batching; in this quarter, the team focused on improving accuracy through model architecture and evaluation metrics. They implemented non-integer compression factors through interpolation for the autoencoder, allowing for more "gentle" compression that better handles smaller particle simulations. The team also experimented with Fréchet physics distance as a new metric for evaluating autoencoder performance throughout model training. Lastly, they explored using a different architecture, a variational autoencoder, intended to encode semantic meaning in the latent diffusion space.