ECE 6083 Graduate

Deep Generative Model

Probabilistic foundations and modern neural generative model families.

A course on how generative models represent, sample, transform, and evaluate probability distributions, covering the main objective families behind autoregressive models, VAEs, flows, GANs, energy-based models, score models, and diffusion models. The course emphasizes objective derivations, sampling procedures, and evaluation trade-offs across model families.

TOPIC MAP

Probabilistic Foundations

- Likelihood, latent variables, entropy, and KL divergence.
- JS, Wasserstein, MMD, and related distances.
- Explicit, implicit, latent-variable, and energy-based modeling views.
- Numerical stability in generative objectives.

Inference and Sampling

- ELBO and variational inference.
- Reparameterization and gradient estimators.
- Monte Carlo estimation, MCMC, and importance sampling.
- Variance, bias, and tractability in approximate inference.

Explicit-Likelihood Models

- Autoregressive factorization and tractable likelihoods.
- Variational autoencoders and amortized inference.
- Normalizing flows, invertibility, and change of variables.

Implicit and Energy-Based Models

- Generative adversarial networks and min–max training.
- Energy-based models and unnormalized densities.
- Contrastive, score, and sampling-based learning ideas.
- Training stability and sample quality trade-offs.

Score and Diffusion Models

- Score functions and denoising score matching.
- SDE and ODE views of diffusion.
- DDPMs, noise schedules, and reverse processes.
- Samplers, distillation, and fast generation.

Evaluation

- Likelihood and held-out density estimates.
- FID, KID, precision-recall, and diversity.
- Memorization, calibration, and coverage.
- Ablation checks and qualitative failure modes.

Autoregressive Models / VAEs / Flows / GANs / EBMs / Diffusion