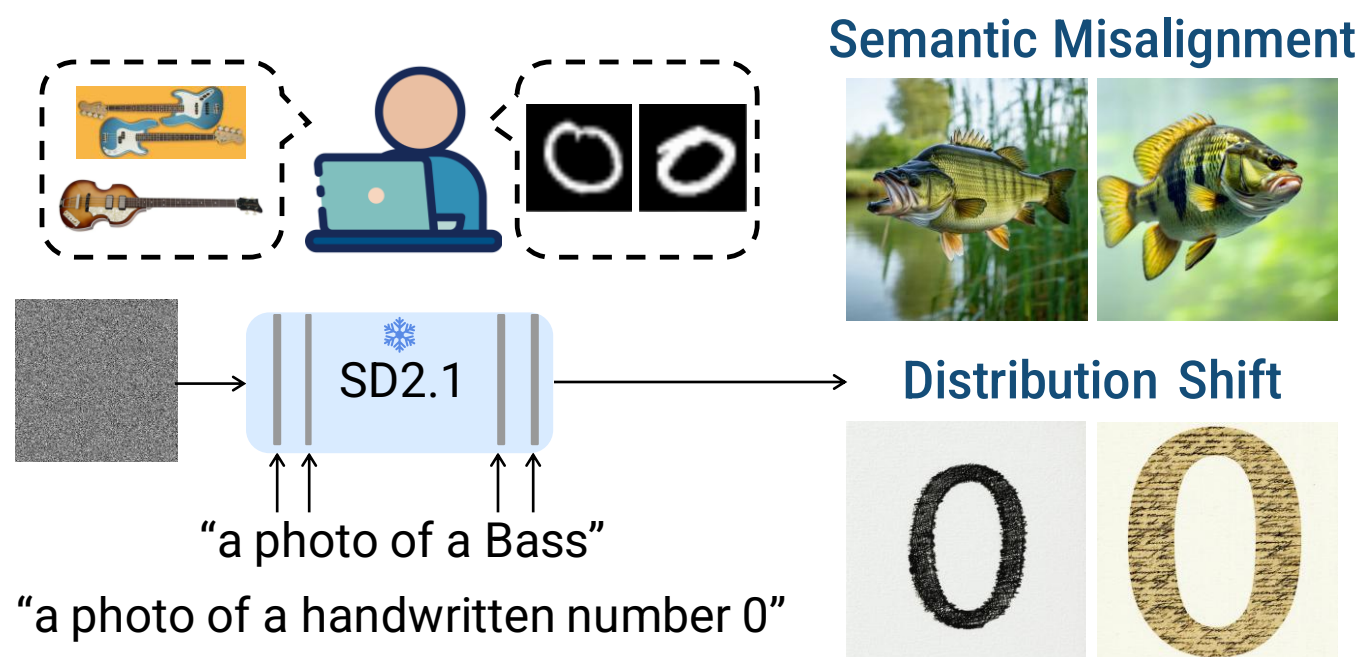


ChimeraLoRA: Multi-Head LoRA-Guided Synthetic Datasets

LoRA A = Class-wise = Diversity
LoRA B = Image-wise = Fidelity

Limitations of Zero-shot Diffusion Models



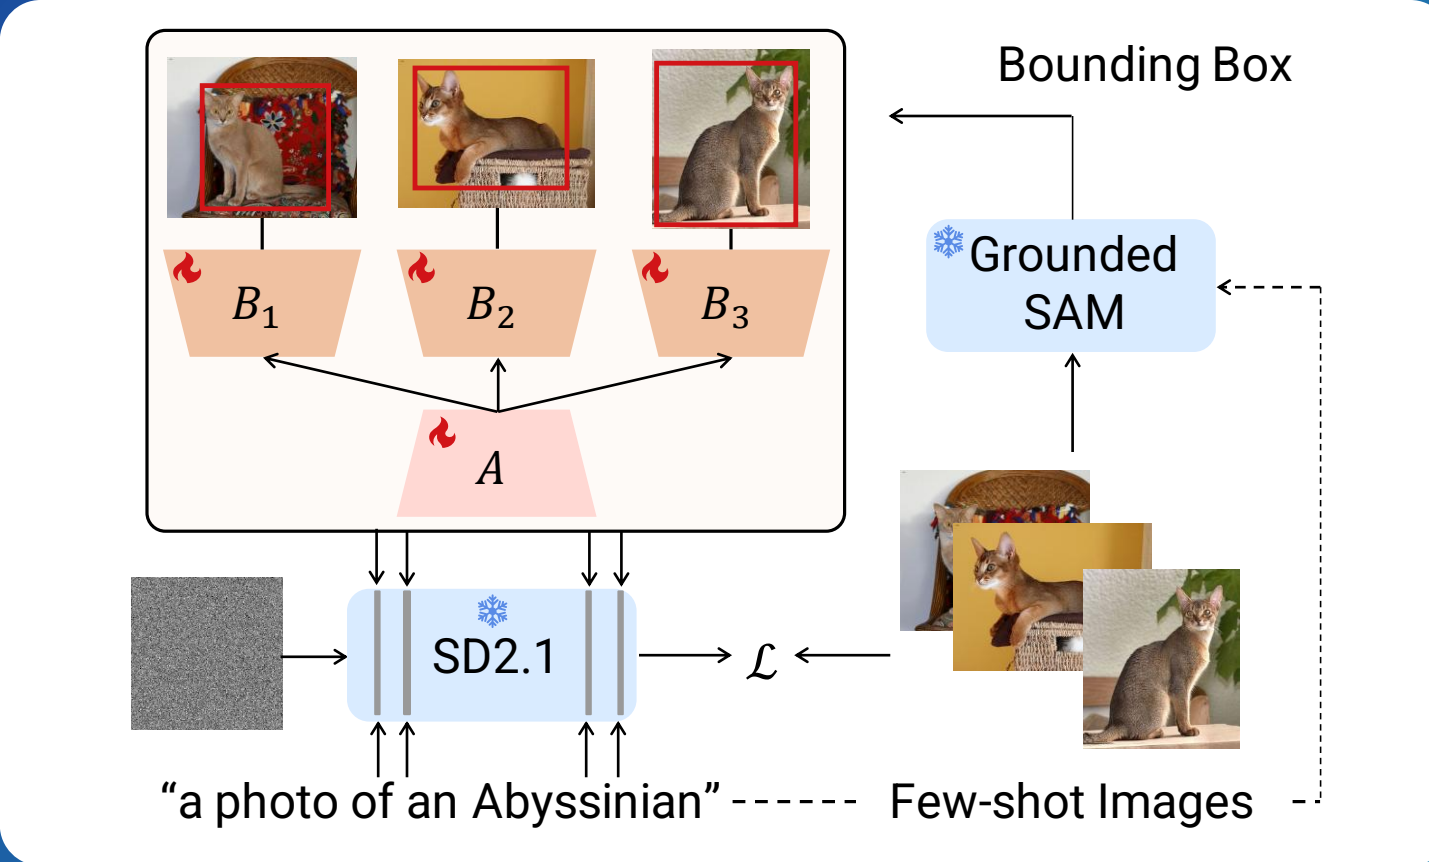
From few-shot images to *realistic* synthetic images

Synthetic-to-Real Gap Analyses

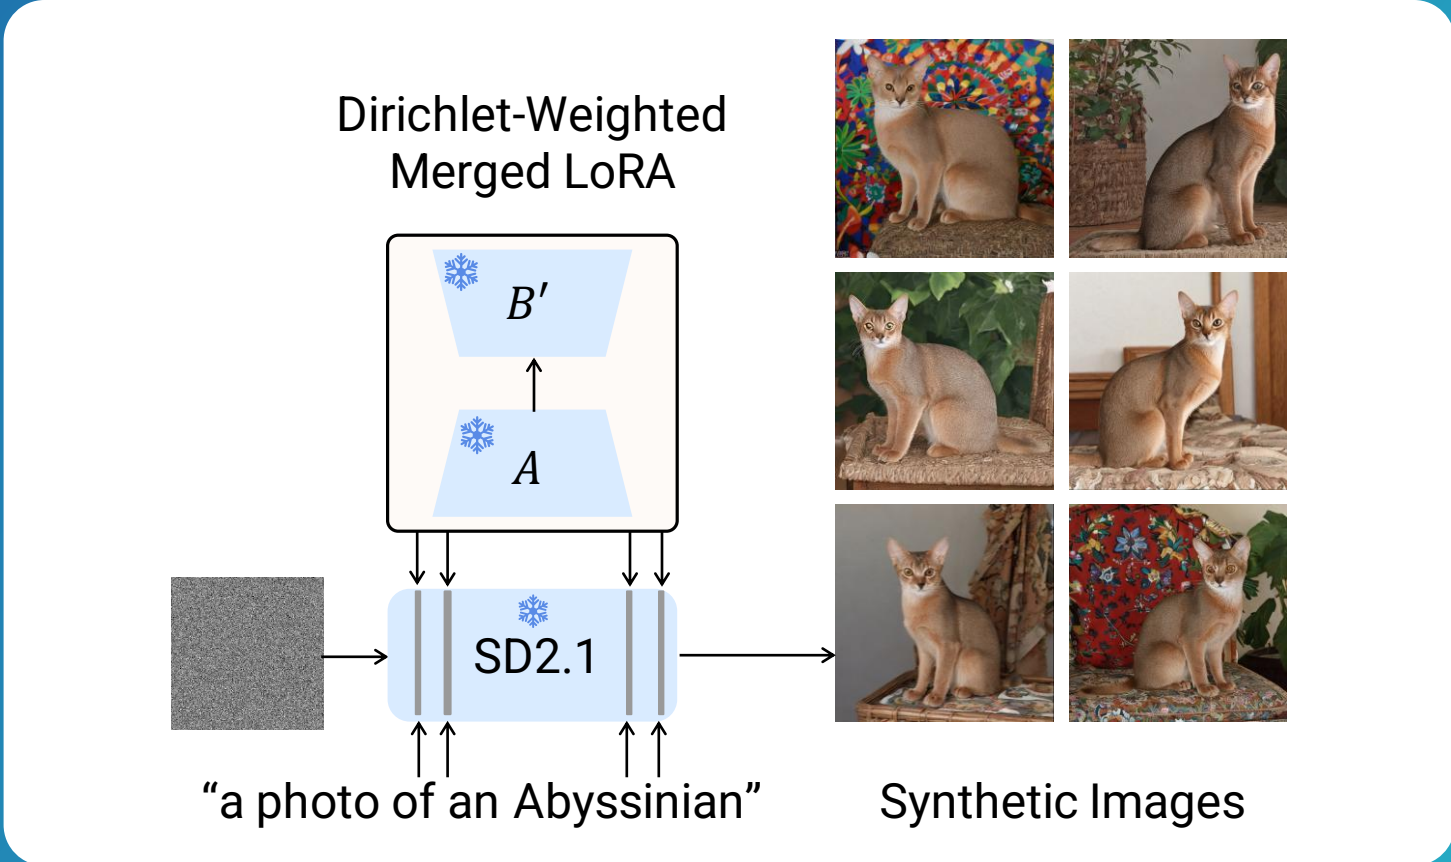


Methods	FID@4 ↓	CLIP sim. ↑
Real (4-shots)	0.00	29.48
Image-wise	0.22	30.04
Class-wise	0.23	29.67
ChimeraLoRA	0.20	30.31

Multi-Head LoRA Fine-tuning



Multi-Head LoRA-Guided Image Generation



Qualitative Results of Synthetic Images

