

SEMilar: Zero-Shot SEM Image Generation via Fixed-Point Equilibrium and CLS Embedding Transfer

Nir Billfeld, Applied Materials

Scanning Electron Microscopy (SEM) is critical in domains such as medical diagnostics, microbiology, and semiconductor manufacturing. However, SEM datasets are often limited due to high acquisition costs and restricted accessibility, posing challenges for training data-hungry vision models. We introduce SEMilar, a zero-shot generative framework for synthesizing high-fidelity SEM images from a single, unseen reference sample. Guided by user-defined structural constraints, SEMilar enables controllable and layout-aware image generation without requiring paired supervision. At its core is a fixed-point equilibrium principle, which enforces semantic consistency between generated outputs and reference embeddings. SEMilar disentangles style from structure using hierarchical FiLM modulation, allowing flexible editing operations such as style injection, feature elimination, and structural shuffling. By fine-tuning only lightweight modulation parameters, the framework adapts to new samples without retraining the generator, making it highly scalable. Our approach supports geometry-aware image synthesis and efficient SEM data augmentation, offering a practical solution for vision tasks in low-data regimes. It also opens new directions for unsupervised meta-learning and cross-domain generalization in scientific imaging. SEMilar bridges the gap between generative modeling and domain-specific constraints, contributing a novel tool for advancing computer vision in specialized applications.