Generative Models and GANs in Vision and Drug Discovery: Application Notes and Protocols

Author: BenchChem Technical Support Team. **Date**: December 2025

Compound of Interest		
Compound Name:	BMVC	
Cat. No.:	B3029348	Get Quote

For Researchers, Scientists, and Drug Development Professionals

This document provides detailed application notes and protocols on the use of generative models, with a focus on Generative Adversarial Networks (GANs), as presented in the British Machine Vision Conference (**BMVC**) proceedings and their broader applications in drug discovery.

Application in Medical Image Analysis (Based on BMVC Proceedings)

Generative models, particularly diffusion models and adaptations of large pre-trained models, have been a key focus in recent **BMVC** proceedings for tasks in medical image analysis. These models offer powerful solutions for image segmentation, synthesis, and 3D reconstruction from limited data.

Morphology-Driven Learning with Diffusion Transformer for Medical Image Segmentation

A novel approach presented at **BMVC** 2024, the Diffusion Transformer Segmentation (DTS) model, demonstrates robust performance in segmenting noisy medical images. This method replaces the commonly used U-Net encoder in diffusion models with a transformer architecture to better capture global dependencies in the image.
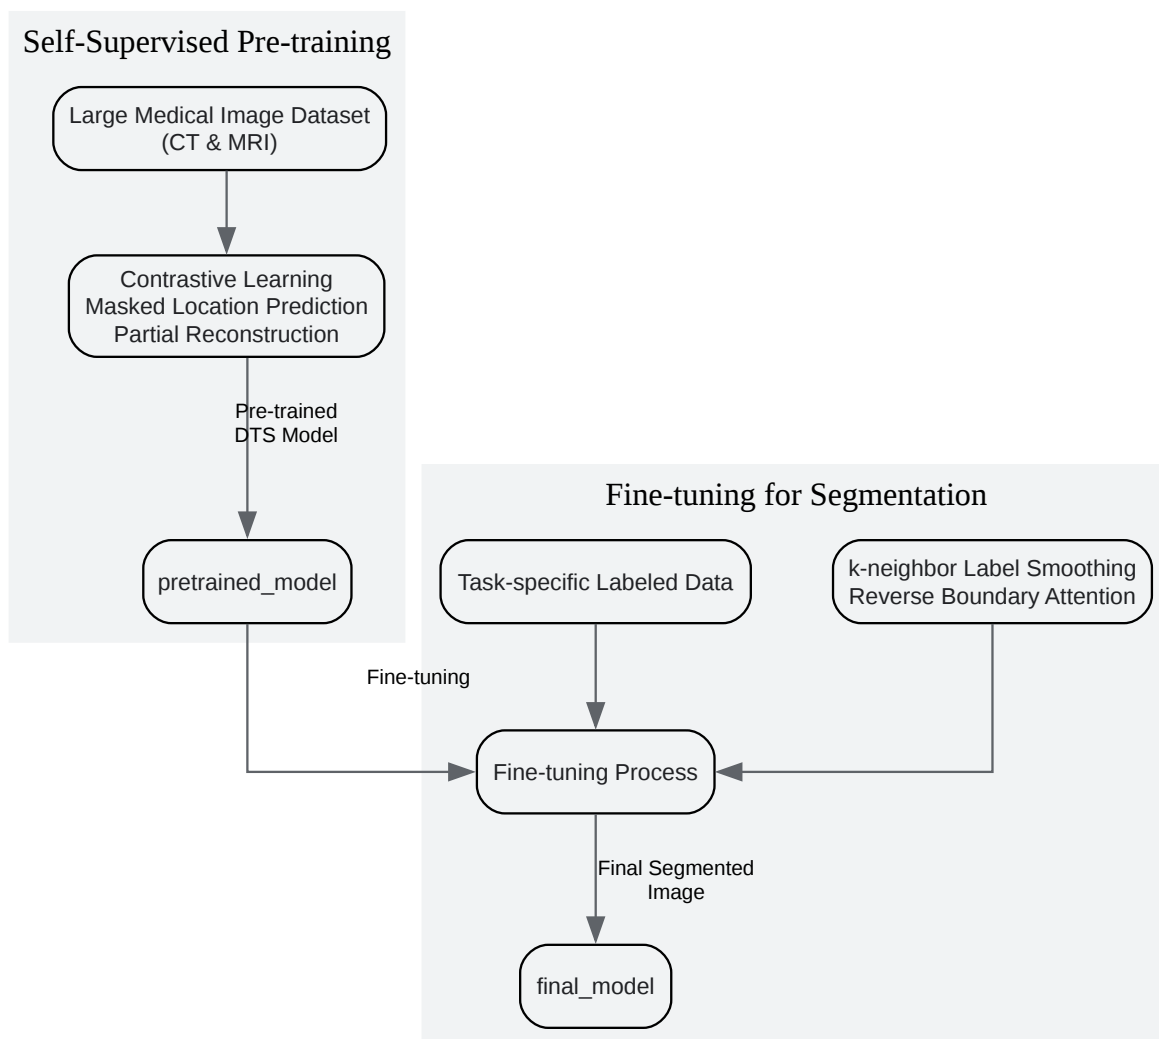
Experimental Protocol:

- Pre-training: The model is pre-trained on a large dataset of medical images, including CT and MRI scans (3,358 and 6,970 subjects, respectively), without using corresponding labels. This step leverages self-supervised learning.

- Self-Supervised Learning: The pre-training employs a combination of:

  - Contrastive Learning (SimCLR): To learn discriminative feature representations from augmented samples.

  - Masked Location Prediction: To predict the location of masked regions in the input.

  - Partial Reconstruct Prediction (SimMIM): To reconstruct masked patches of the input volume.

- Fine-tuning for Segmentation: The pre-trained model is then fine-tuned on specific medical image segmentation tasks.

- Morphology-Driven Techniques: To enhance performance on complex structures, the following are incorporated:

  - k-neighbor Label Smoothing: Leverages the relative positions of organs for smoother label transitions.

  - Reverse Boundary Attention: Focuses the model's attention on the boundaries of regions of interest.

- Loss Function: A combination of losses is used to train the model, including reconstruction loss, masked location prediction loss, and contrastive learning loss, with empirically determined weights (e.g., $\lambda_1=0.1$, $\lambda_2=0.01$ for specific tasks).[1]

Quantitative Data Summary:

Model	Dataset	Dice Score
DTS (Ours)	Multi-organ CT	91.2
nnU-Net	Multi-organ CT	89.5
Swin-UNETR	Multi-organ CT	90.1

Note: The table presents a simplified summary of reported results. For detailed metrics across different organs and datasets, please refer to the original publication.

Experimental Workflow Diagram:

Self-Supervised Pre-training

Large Medical Image Dataset
(CT & MRI)

↓

Contrastive Learning
Masked Location Prediction
Partial Reconstruction

Pre-trained
DTS Model

↓

pretrained_model

Fine-tuning for Segmentation

Task-specific Labeled Data

k-neighbor Label Smoothing
Reverse Boundary Attention

Fine-tuning

Fine-tuning Process

Final Segmented
Image

↓

final_model

Click to download full resolution via product page

DTS Model Training Workflow

Adapting Segment Anything Model (SAM) to Medical Images

The "Segment Anything Model" (SAM) has shown remarkable zero-shot segmentation capabilities. However, its performance on out-of-distribution medical images can be suboptimal.

The AutoSAM paper from **BMVC** 2023 proposes a method to adapt SAM for medical imaging without fine-tuning the entire model.[2]

Experimental Protocol:

- Architecture Modification: The prompt encoder of the original SAM architecture is replaced with a new encoder (e.g., a Harmonic Dense Net).[2][3] This new encoder takes the input image itself as the prompt.

- Training: The original SAM model is frozen. The new prompt encoder is trained using the gradients provided by the frozen SAM.[2]

- Mask Generation: A shallow deconvolutional network is also trained to decode the output of the new encoder into a segmentation mask. This provides a lightweight solution for direct segmentation.

- Evaluation: The adapted model is evaluated on various medical image segmentation benchmarks.

Quantitative Data Summary:

Model	Dataset	Mean IoU
AutoSAM (Ours)	GlaS	0.892
SAM (point prompt)	GlaS	0.756
nnU-Net	GlaS	0.881

Note: This table provides a high-level comparison. Refer to the original paper for a comprehensive evaluation across multiple datasets and metrics.

Logical Relationship Diagram:
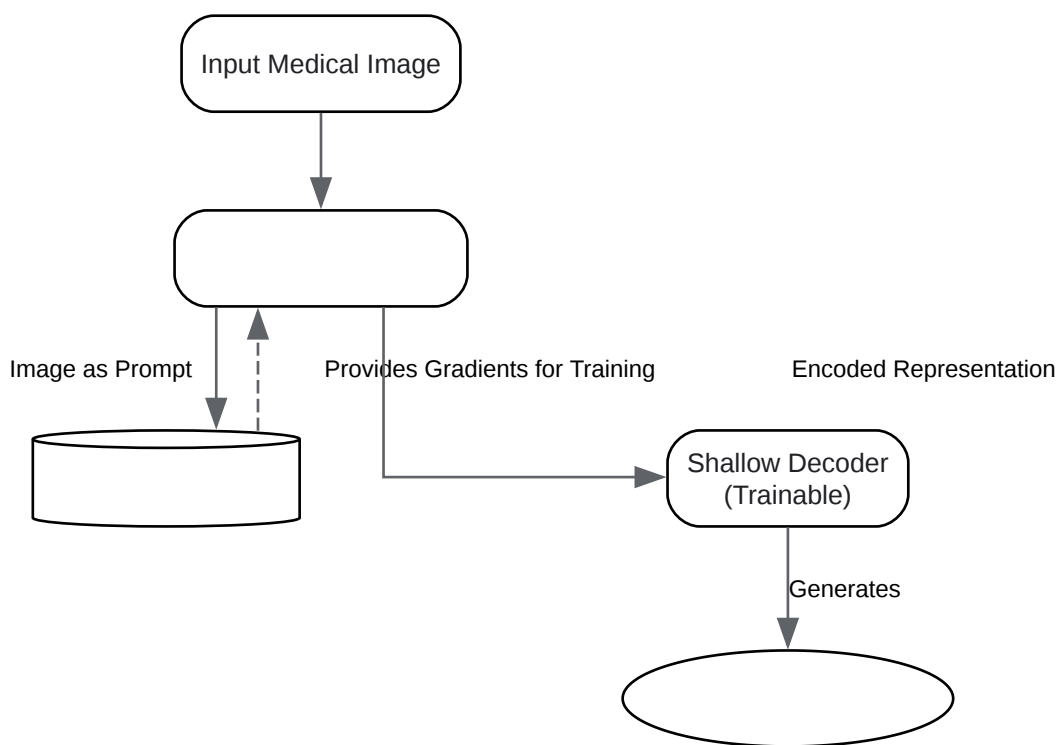
AutoSAM Adaptation Logic

Application in Drug Discovery: De Novo Molecule Generation

Generative Adversarial Networks are being increasingly utilized for de novo drug design, which involves generating novel molecular structures with desired chemical and pharmacological properties. Models like MolGAN and ORGAN have been instrumental in this domain.

MolGAN: An Implicit Generative Model for Small Molecular Graphs

MolGAN is designed to generate molecular graphs directly, avoiding the need for sequential string-based representations like SMILES. It employs a generator, a discriminator, and a reward network to produce molecules with optimized properties.[4]

Experimental Protocol:

- Molecular Representation: Molecules are represented as graphs with a node feature matrix (atom types) and an adjacency tensor (bond types).[4][5]

- Generator: A multi-layer perceptron (MLP) takes a sample from a standard normal distribution and generates a dense adjacency tensor and an atom feature matrix. Categorical sampling is then used to obtain a discrete molecular graph.[6]

- Discriminator: A relational graph convolutional network (GCN) is used to distinguish between real molecules from a dataset and generated molecules.

- Reward Network: This network has the same architecture as the discriminator and is trained to predict a specific chemical property of a molecule (e.g., drug-likeness).

- Training: The generator and discriminator are trained adversarially using the Wasserstein GAN (WGAN) objective. The generator is also optimized using reinforcement learning, with the reward provided by the reward network.

- Dataset: The model is typically trained on datasets of small molecules like QM9 or ZINC.

Quantitative Data Summary (QM9 Dataset):

Metric	MolGAN	ORGAN	VAE
Validity (%)	98.1	95.7	93.5
Novelty (%)	94.2	99.9	95.1
Uniqueness (%)	99.9	100	99.9
Drug-likeness (QED)	0.83	0.78	0.81

Note: This table summarizes typical performance metrics. Actual values may vary based on specific experimental setups.

MolGAN Architecture Diagram:

Tech Support
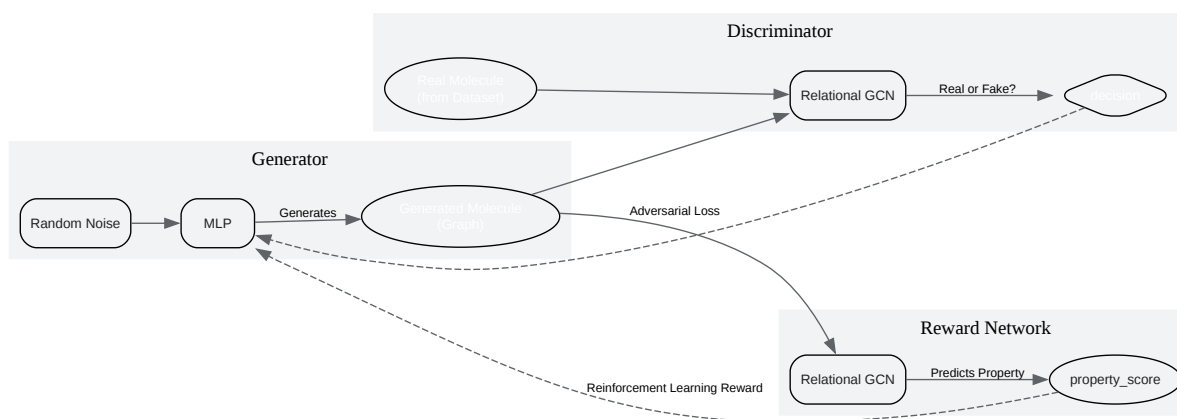
MolGAN Architecture

ORGAN: Objective-Reinforced Generative Adversarial Networks for Sequence Generation

ORGAN is designed to generate novel molecular sequences (SMILES strings) that are optimized for desired properties. It combines a GAN with reinforcement learning to guide the generation process.[7][8]

Experimental Protocol:

- Representation: Molecules are represented as SMILES strings.

- Generator: A recurrent neural network (RNN), typically with LSTM or GRU cells, is used as the generator to produce SMILES strings character by character.

Tech Support

- Discriminator: A convolutional neural network (CNN) is often used as the discriminator to classify SMILES strings as real or fake.

- Reinforcement Learning: The generator is treated as a reinforcement learning agent. The "reward" for generating a particular molecule is a combination of the discriminator's output (how "real" it looks) and a score from an external objective function (e.g., predicted solubility, synthesizability, or drug-likeness).[8]

- Training: The discriminator is trained to distinguish real from generated SMILES. The generator is then trained using policy gradients (e.g., REINFORCE algorithm) to maximize the expected reward.

- Pre-training: The generator is often pre-trained on a large dataset of known molecules (e.g., from ChEMBL) using maximum likelihood estimation before the adversarial training begins.

Quantitative Data Summary (Targeting High Drug-likeness):

Model	Average QED Score	% Novel	% Valid
ORGAN	0.91	98.5	96.2
Pre-trained Generator	0.75	95.1	97.8
Baseline GAN	0.82	97.3	94.5

Note: This table illustrates the improvement in a target property (QED) with ORGAN compared to baseline models.

ORGAN Training Workflow:
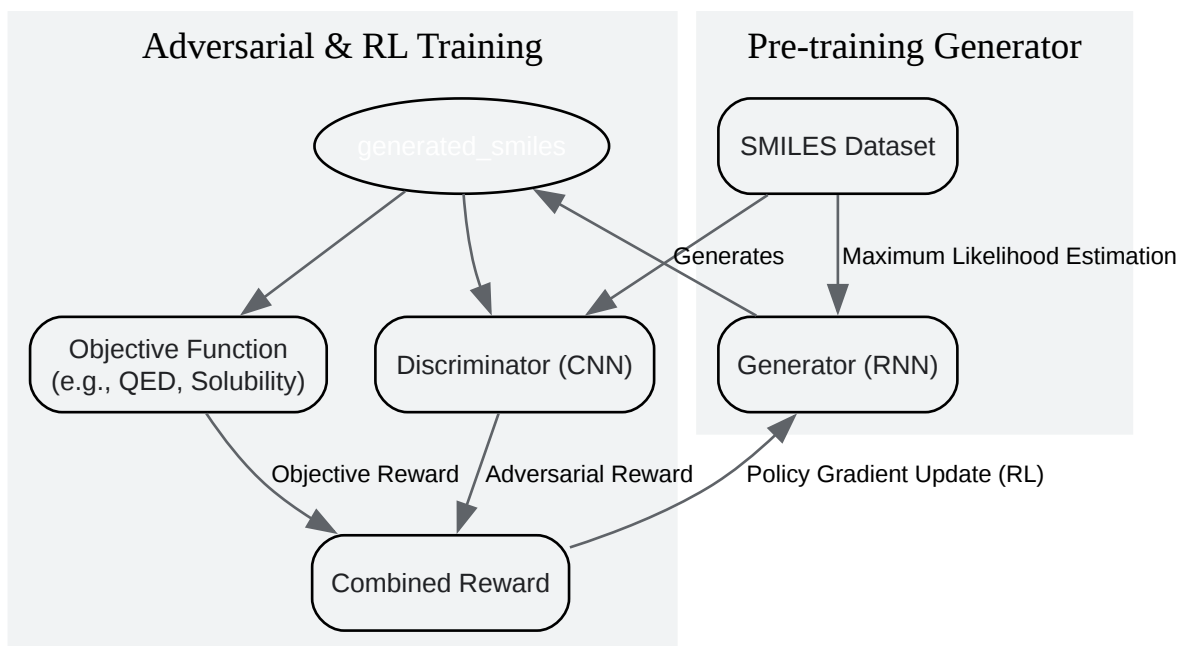
Click to download full resolution via product page

ORGAN Training Workflow

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopiclabeling.

Email: info@benchchem.com or Request Quote Online.

References

- 1. Exploring the Advantages of Quantum Generative Adversarial Networks in Generative Chemistry - PMC [pmc.ncbi.nlm.nih.gov]

- 2. researchgate.net [researchgate.net]

- 3. papers.bmvc2023.org [papers.bmvc2023.org]

- 4. arxiv.org [arxiv.org]

- 5. chemrxiv.org [chemrxiv.org]

- 6. DeepChem [deepchem.io]

- 7. chemrxiv.org [chemrxiv.org]

- 8. shortscience.org [shortscience.org]

- To cite this document: BenchChem. [Generative Models and GANs in Vision and Drug Discovery: Application Notes and Protocols]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b3029348#generative-models-and-gans-in-bmvc-proceedings]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [Contact our Ph.D. Support Team for a compatibility check]

Need Industrial/Bulk Grade?   Request Custom Synthesis Quote

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd

Ontario, CA 91761, United States

Phone: (601) 213-4426

Email: info@benchchem.com