



IISER Bhopal

INDIAN INSTITUTE OF SCIENCE EDUCATION AND RESEARCH BHOPAL

Last updated in July 2026

J Harishwar Rao

📍 Bhopal / Hyderabad | ✉ jamindar23@iiserb.ac.in | 🌐 harish-jhr.github.io | 🐙 GitHub | in LinkedIn |

Personal Statement

I am a Computer Science undergraduate at IISER Bhopal working in **Computer Vision Research**, with a focus on **Generative Image Synthesis**, **Image Editing**, **Representation Learning**, and **Mechanistic Interpretability of Generative Models**. I have been working extensively with **Diffusion Models**, particularly **score based** and **SDE-based** formulations, **Conditional Flow Matching**, and **Latent Diffusion**, building end to end systems from scratch to understand training dynamics, conditioning, controllability, and failure modes. A growing focus of my work is **interpreting what generative models learn**: applying **Sparse Autoencoders** and **probing** techniques to understand internal representations in diffusion U-Nets. My most recent Research Experience at **CVIT, IIIT-H** involves Conditional Diffusion Models with object and part-level awareness for semantic image generation.

Publications

PLATO++: Pose-Conditioned Part-Aware Object Generation via Residual Structure Learning

CVPRW 2026 🌐

Varghese P Kuruvilla, **Harishwar Rao J.**, Ravi Kiran Sarvadevabhatla
IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops —
CVEU Workshop, CVPR 2026

Crosscoders over SAEs for Concept Steering in T2I Diffusion Models
Harishwar Rao J.

Manuscript in Preparation

Research Experience

Research Intern, Spectrum Lab, IISc Bengaluru

May 2026 to Present

Advisor: Prof. Chandra Sekhar Seelamantula

- Working on **Flow Matching** — based generative approaches to solve scarcity of high quality annotated Eye Fundus Imagery.
- Core Focus is on developing a model that can produce high quality Fundus Imagery and the associated segmentation maps.

Research Intern, CVIT, IIIT Hyderabad

Apr 2025 to Mar 2026

Advisor: Prof. Ravi Kiran

- Working on **Part-Aware Diffusion Models** — conditioning generation on structured lists of object parts to achieve fine-grained **controllability** and **compositional diversity**.
- Studying how **structured conditioning** via object and part-level layouts affects the internal representations of diffusion U-Nets; probing attention patterns and cross-attention maps for interpretability.
- Implemented novel **layout generation methods** and **layout-conditioned diffusion pipelines** where structured spatial layouts guide the training and sampling process.
- Analysing limitations of conditional generative models under **distribution shift** between training layouts and inference prompts.

- Applied **Convolutional Neural Networks** to detector image data from the **OPERA particle physics experiment** at CERN.
- Implemented **CNN-based classification pipelines** and developed **visualization tools** for analyzing experimental neutrino and muon track data.
- Gained experience working with real-world **scientific datasets** containing noise and domain-specific constraints.

Education and Achievements

Indian Institute of Science Education and Research, Bhopal
BS in Computer Science

Aug 2023 to Present

- **GPA: 8.86 / 10.0**
- Relevant coursework:
 - **Computer Science:** Computer Vision, Deep Learning, Machine Learning, Artificial Intelligence, Data Structures and Algorithms, Algorithms, Theory of Computation, Data Science in Practice.
 - **Mathematics:** Single Variable Calculus, Multivariable Calculus, Probability and Statistics, Linear Algebra, Discrete Mathematics, Complex Analysis.
- **Recipient of the Reliance Undergraduate Scholarship.**
- Winner of institute-wide competitive coding contest.
- Member of the IISER Bhopal Coding Community.
- I previously led the IISER Bhopal Quiz Club. ([IBQC](#)).

Technologies and Skills

Languages: Python, C++, C, Java, Lua

Frameworks and Libraries: PyTorch, HuggingFace Diffusers, torch-geometric, FAISS, NumPy, Matplotlib, Scikit-learn, TensorFlow

Tools: Git, GitHub, L^AT_EX, MATLAB, VS Code, Weights & Biases

Research Areas: Generative Image Synthesis, Diffusion Models, Conditional Flow Matching, Representation Learning, Mechanistic Interpretability, Scientific ML

Projects

DiffSteer: Concept Steering in Stable Diffusion via SAE and CrossCoder

[GitHub](#)

- Suppressed unsafe visual concepts in **Stable Diffusion 1.4** at inference time using **sparse dictionary learning**, with no retraining of the Diffusion model being required.
- Trained a **K-Sparse Autoencoder (SAE)** and a **CrossCoder** on CLIP text encoder activations; identified concept-aligned latent directions and subtracted them via **forward hooks** amplified through CFG.
- CrossCoder steering reduced nudity rate from **55% to 5%** (50 pp reduction) on a large eval prompt set; released checkpoints on HuggingFace and reproducible e2e scripts.

DiffuseSeg: Synthetic Data and Segmentation from a Single Diffusion Model

[GitHub](#)

- Reused **frozen DDPM U-Net multi-scale features** for downstream **dense prediction** — no task-specific backbone training.
- Lightweight **pixel-wise segmentation heads** on diffusion decoder features achieve **mIoU > 0.35 on**

CelebA-HQ with only **100** labeled images.

- Released trained models, synthetic datasets with segmentation masks, and reproducible scripts.

Deepfake Detection using Diffusion Reconstruction Error

[GitHub](#)

- Explored **DDIM inversion + reconstruction error** as a feature representation for **deepfake detection** — real vs. AI-generated image discrimination.
- **PCA** of per-pixel reconstruction error maps shows clear separability; lightweight **CNN classifier** on error maps achieves **96.65% test accuracy**.

FlowMatching: Conditional Flow Matching for Face Generation

[GitHub](#)

- Implemented **OT-CFM (Conditional Flow Matching)** from scratch on **CelebA-HQ 64×64** — learns straight-path velocity fields between noise and data, enabling fast ODE-based sampling.
- Built two interchangeable backbones: a **UNet** (HuggingFace Diffusers) and a **DiT-S/4** (patch-size 4, adaLN time conditioning).
- Implemented both **Euler** and second-order **Heun** ODE solvers; Heun achieves comparable visual quality with fewer function evaluations.
- Training infrastructure: **AdamW + cosine schedule**, **EMA** (decay 0.999), **mixed precision (AMP)**, **W&B** logging.

DiffNowcast: Large scale Diffusion based Cloud Nowcasting on INSAT-3DR Data

[GitHub](#)

- Worked on developing large scale **Autoregressive Cloud Nowcasting** using INSAT-3DR Sattelite data with Diffusion based Generative Approaches.
- Developed an physics inspired Optical Flow conditioned Diffusion Model which beats existing baselines.
- Demonstrated viable **3-hour autoregressive forecasting** on sparse 30-min resolution Indian weather data, with correct modeling of cloud dissipation across extended rollouts.