

# AI&DS INSIGHTS



*"Design is a journey of  
discovery."*

*– Alvin Huang*



## & Student Achievements

Empowering Minds, Inspiring Success: Students Who  
Make Us Proud.

To read refer page 3

# Student Achievements

Ms. Ishwari Shrirao, a first-year AIDS student, has earned selection for the Indian Team Trials in Rifle Shooting after competing at the 68th National Shooting Championship in Bhopal.



# Team RoboRebels Clinches National Victory at Smart India Hackathon 2025 Hardware Grand Finale

Team RoboRebels, featuring Diksha Sadashiv Shingne and Chetana Ingale, secured a prestigious national-level 🏆 First Place at the Smart India Hackathon (SIH) 2025 – Hardware Edition Grand Finale held at BHU Varanasi Nodal Centre, for their groundbreaking portable Embedded Edge-AI Microscope developed in response to a challenge from the Ministry of Earth Sciences.

The innovative system integrates embedded hardware, computer vision, edge AI, and agentic AI workflows to enable real-time, autonomous identification and counting of microscopic marine organisms—dramatically reducing manual effort in marine biodiversity research—while overcoming intense hackathon pressures through redesigns, debugging marathons, and expert mentorship.



Their win, accompanied by a ₹75,000 prize, highlights exceptional engineering skill, resilience, and teamwork, bringing pride to their institution and setting a powerful example for aspiring innovators nationwide.

## Manav Dantkale Wins Gold at Maharashtra State Shooting Championship, Qualifies for Nationals

Manav Dantkale, a First-Year B.Tech student in Artificial Intelligence and Data Science, delivered an outstanding sporting performance by winning the 🏆 Gold Medal in the Junior Men category at the 49th Maharashtra State Shooting Championship 2025 in the 10m Air Pistol event.

Along with earning a cash prize of ₹10,000, he has secured qualification for the prestigious 68th National Rifle and Pistol Championship scheduled for December 2025.

His display of exceptional focus, precision, and composure under pressure has brought immense pride to his institute and serves as an inspiring example for fellow students, as he prepares to compete at the national level and continue his promising shooting career.



# Faculty Achievement Spotlight

Dr. Parikshit Mahalle earns international recognition with a second Post-Doctoral Fellowship and Honorary Professorship at Peru's UMA.




## Dr. Parikshit Mahalle Delivers Global PBL/PCL Initiative at UMA Focusing on SDGs and Technology Readiness

Dr Parikshit N Mahalle has conducted an innovative course on PBL/PCL impacting SDGs and TRL at UMA, Lima from 24<sup>th</sup> November to 5<sup>th</sup> December, 2025 for the leadership at UMA, Dr. Parikshit is Professor in Department of Artificial Intelligence and Data Science and Dean Academics at Vishwakarma Institute of Technology (VIT), Pune. UMA , proudly announces the launch of a forward-looking and outcome-driven course titled “Project-Based and Project-Centric Learning Impacting Sustainable Development Goals (SDGs) and Technology Readiness Level (TRL).” The course is designed to strengthen experiential learning practices and align student competencies with global development priorities and emerging technological benchmarks.

This course aims to empower students and faculty with the ability to conceptualize, design, implement, and evaluate projects that address real-world challenges derived from the United Nations Sustainable Development Goals (SDGs). It also integrates the Technology Readiness Level (TRL) framework to ensure that innovations progress systematically from ideation to deployment.

The program emphasizes hands-on learning, interdisciplinary collaboration, critical thinking, innovation, and societal impact. By engaging learners in authentic problem-solving, the course supports the institution’s vision of nurturing technologically skilled, socially responsible, and globally relevant professionals.

## Faculty Article



AI-Driven Predictive Analytics in Healthcare refers to the use of artificial intelligence (AI) and machine learning (ML) techniques to analyse historical and real-time healthcare data in order to predict future health outcomes, risks, and trends. It helps clinicians, hospitals, and policymakers move from reactive care to proactive and preventive care.

# AI-Driven Predictive Analytics in Healthcare

AI-Driven Predictive Analytics in Healthcare is one of the most impactful application areas of Artificial Intelligence and Data Science in recent years. The healthcare sector generates enormous volumes of data through electronic health records, medical imaging systems, laboratory reports, wearable devices, and patient monitoring systems.

AI and Data Science techniques enable the effective analysis of this data to identify hidden patterns, predict health risks, and support timely medical interventions. Predictive analytics helps healthcare professionals move from reactive treatment to proactive and preventive care, thereby improving patient outcomes and operational efficiency.

- **Predictive Analytics Principle:** Utilizes historical and real-time healthcare data to forecast disease risks and patient outcomes using machine learning and statistical models.
- **Machine Learning & Deep Learning Principle:** Algorithms such as classification, regression, and neural networks are applied for early disease detection and medical image analysis.
- **Data-Driven Decision-Making:** AI systems support doctors and healthcare administrators by providing accurate insights, trends, and predictions for effective clinical decisions.

- **Personalization Principle:** Enables personalized medicine by tailoring treatment plans based on patient history, lifestyle, and genetic data.
- **Scalability & Big Data Principle:** Technologies like Hadoop, Spark, and cloud platforms allow efficient handling of large-scale healthcare datasets.
- **Ethical AI & Data Privacy:** Emphasizes data security, patient privacy, fairness, and transparency while deploying AI models in healthcare systems.
- **Industry & Academic Relevance:** Provides real-world applications for AI & DS students, supporting project-based learning, research, and industry readiness.



## Major Applications in Healthcare

### Disease Prediction & Early Diagnosis

- Predicting diabetes, heart disease, cancer, or sepsis at an early stage
- Identifying high-risk patients before symptoms worsen

### Patient Monitoring & Risk Scoring

- ICU patient deterioration prediction
- Fall risk prediction in elderly patients
- Mortality and complication risk assessment

### Hospital Resource Management


- Predicting patient admissions and bed occupancy
- Optimizing staff scheduling and equipment usage

### Personalized Treatment & Medicine

- Predicting how a patient will respond to a specific treatment
- Dosage optimization based on patient data

### Readmission & Length of Stay Prediction

- Identifying patients likely to be readmitted within 30 days
- Reducing healthcare costs and penalties



Prof. Badasaheb Jadhav

Assistant Professor  
AI & DS Department



## Diffusion Models – From Random Noise to Stunning Creations

Okay, so let's be real—haven't we all spent way too much time this year playing with AI image generators? I know I have. Those dreamy Studio Ghibli-style scenes with soft watercolors, whimsical forests, and glowing spirits are impossible to resist. The sharp anime/manga characters, the realistic video frames from Sora that look like actual movie clips, photorealistic landscapes, neon cyberpunk cities, elegant digital paintings, clean vector art, detailed 3D renders, and gentle pastel illustrations—there's a style for everything. It's made creating beautiful visuals accessible to anyone with a simple text prompt, no professional skills needed.



This technology has truly democratized creativity. People who never thought they could draw or design are now making stunning art just for fun or for projects. Beyond personal use, it's helping teachers create custom visuals to explain complex topics in class, making lessons more engaging and easier to understand for students. In professional fields, designers and filmmakers use it to brainstorm concepts quickly, saving time and money. Even in science, researchers are using diffusion models to generate possible molecular structures for new medicines, speeding up drug discovery in ways that could save lives one day.



The generative AI market has grown massively because of this—expected to hit around \$60–70 billion globally by the end of 2025. Companies are investing heavily, and it's creating new jobs while pushing innovation forward.

Diffusion models trace their roots to earlier work like GANs in 2014, but the real breakthrough came in 2020 with the Denoising Diffusion Probabilistic Models (DDPM) paper by Jonathan Ho and his team. They took inspiration from physics—how particles naturally spread out over time—and reversed the process to create images.


## How Diffusion Models Work

The core idea is surprisingly elegant. During training (forward diffusion), a clean image from a huge dataset is gradually corrupted by adding small amounts of Gaussian noise over hundreds or thousands of steps until it becomes pure random static. The model learns to predict exactly what noise was added at each step.

When generating a new image (reverse diffusion), we start from pure random noise and run the process backward: the model repeatedly predicts and removes noise, step by step, while being guided by the user's text prompt (converted into embeddings, often using something like CLIP). Modern versions work in a compressed "latent" space to make the whole process much faster and less resource-heavy. The result? Coherent, high-quality images that match the description.

### How does diffusion model operate while predicting system behavior

This slide provides information regarding diffusion model that assesses and predicts system behavior & generated similar data on which it was trained. The diffusion models are integrated with AI to ensure noise removal and offer high-quality images from random seeds.




The journey from research to everyday tools has been fast. After the 2020 foundation, improvements in 2021 made outputs sharper. 2022 brought major releases: DALL-E 2 (April), Stable Diffusion going open-source (August), and Midjourney becoming popular on Discord. 2023 delivered higher-resolution images, 2024 introduced Sora for realistic video generation, and throughout 2025 we've seen Flux fix long-standing issues like inaccurate hands or text, while Midjourney v7 and newer video models handle longer, more coherent sequences.

### AI Evolution Timeline: Generative & Agentic AI (2020-2025+)

| 2020                                                                                                                                                            | 2021                                                                                                                                                            | 2023                                                                                                                                                                         |                                                                                                                                                                       |                                                                                                                                                                  | 2024                                                                                                                                                                    | 2025+                                                                                                                                                                               |                                                                                                                                                                |                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <br><b>GPT-3</b><br>June 2020<br>Advanced text generation by OpenAI            | <br><b>DALL-E</b><br>January 2021<br>AI image generation from text prompts     | <br><b>MidJourney</b><br>July 2023<br>Creative AI for visual art generation                 | <br><b>Claude</b><br>March 2023<br>Safety focused target language model              | <br><b>Bard</b><br>(now Gemini)<br>May 2022<br>Conversational and multimodal AI | <br><b>GPT-4o</b><br>May 2024<br>Multimodal AI model handling text, image, and audio | <br><b>Claude 3.7 Sonnet</b><br>by Anthropic<br>February 2023<br>Enhanced Reasoning Capabilities |                                                                                                                                                                |                                                                                                                                                                  |
| <br><b>ChatGPT</b><br>November 2022<br>Conversational AI Interface by OpenAI | <br><b>ChatGPT</b><br>November 2022<br>Conversational AI Interface by OpenAI | <br><b>Bard</b> (now Gemini)<br>by Google<br>May 2023<br>Conversational and multimodal AI | <br><b>Bing AI / Copilot</b><br>November 2024<br>AI-integrated search engine tools | <br><b>Runway Gen-2</b><br>April 2024<br>AI-powered image generation model    | <br><b>Grok</b><br>by xAI<br>November 2024<br>Conversational AI assistant           | <br><b>LLaMA</b><br>by Meta<br>September 2024<br>Multilingual generative AI model              | <br><b>StackAI</b><br>May 2025<br>No-code platform for building AI agents | <br><b>Figma AI Tools</b><br>Integrated AI tools for design and prototyping |


Despite the progress, challenges remain. Training these models requires enormous datasets and powerful GPU clusters, consuming significant energy. Biases present in training data can show up in outputs, and achieving perfect consistency—especially with fine details or complex compositions—still needs work. Researchers and companies are continually addressing these through better data curation, efficiency tricks, and ethical guidelines.

On the positive side, the real-world impact keeps growing. In creative industries, diffusion tools help independent artists, game developers, and filmmakers produce professional-quality work without big budgets. In education, teachers generate tailored diagrams and illustrations, making abstract subjects more concrete. In healthcare and chemistry, generating plausible protein or molecular structures helps scientists explore new possibilities much faster than traditional methods.



The technology is also sparking important conversations about copyright, artist livelihoods, and fair use—leading to new policies, compensation models, and reskilling programs. Overall, the benefits far outweigh the concerns for most users.

As we are in AI&DS, getting hands-on with diffusion models feels exciting. It's amazing to see how a clever reversal of a physical process has become a tool that empowers creativity, accelerates science, and touches so many parts of life. From that 2020 paper to the tools we use daily in 2025, it's a perfect example of how research can quickly change the world.



What kind of image will you create next?

# OUR TEAM



Dr Minal Barhate



Vidhi Rahangdale  
SY AI&DS-E



Jatin Shende  
SY AI&DS -E

For Feedback  
and Query,  
Scan this:

