PROJECT IDEAS · M.TECH / M.E. · IMAGE PROCESSING

M.Tech projects in *image processing*

These 14 M.Tech projects in image processing and computer vision are each built on a real, open-access IEEE paper from 2023 to 2026. Every card gives the problem, what you build and measure, the dataset and tools, and **one measured extension with the comparison table your thesis would contain**.

Last updated 2026-11-12

We matched each paper below to its Crossref DOI record and read its abstract before writing the card, so a card describes only what the paper itself says. New to base papers? Start with how to select a base paper. For papers outside image processing, see the IEEE base papers list and the IEEE base papers for ECE.

What makes a good M.Tech image processing project?

An M.Tech project in image processing is judged on a comparison, not on a demo. Three things make that comparison possible.

1. **A base paper you can reproduce.** Recent, open access, with a public dataset and enough detail (or code) to rebuild the method. All 14 papers on this page are open access on IEEE Xplore.
2. **One measured extension.** A second dataset, a lighter model on a board, harder test conditions such as night or rain, or a check of where the model looks. Change one thing, so the result can be explained.
3. **A result someone else can repeat.** Fixed train and test splits, the same metric as the base paper, a log of every run and a note of the random seeds. An

examiner who asks how a number was produced should get a clear answer.

Before Stage-I, write the plan as a single sentence that names the base paper, the dataset, your one change and the metric. If your guide reads that sentence and agrees, the topic is ready.

Which image processing projects suit M.Tech in 2026-27?

The 15 below are grouped by area: medical images, agriculture and environment, surveillance and safety, and enhancement and hardware. Each card gives the problem and what you build and measure. Open “The base paper and the M.Tech extension” to see the reference and DOI link, the data and tools, and one extension with the table your thesis would contain.

Topic ideas, not delivered projects

Each card is built on a published, open-access IEEE paper, checked against its DOI record, and describes only what the paper’s abstract states. The project around it is an idea to discuss with your guide. Our 8 delivered case studies are M.Tech and M.E. projects, four of them image classification for health, each shown with its real paper pages, screens and results.

Medical image analysis

Keep every patient in one split only. If images of the same person sit in both the training and the test set, the score looks better than the model really is.

TOPIC IDEA · M.TECH · MEDICAL IMAGING

Skin lesion classification that does not ignore the rare cancers

Problem: In skin lesion datasets benign moles far outnumber melanoma and the other rare classes, so a classifier can score well overall while missing the lesions that matter most.

Output: You rebuild the paper's two-stage Self-Transfer GAN (STGAN) on HAM10000, use it to generate extra images for the minority classes, and report sensitivity, precision and F1 for each class with and without the generated images.

− The base paper and the M.Tech extension

Base paper: A GAN-Based Data Augmentation Method for Imbalanced Multi-Class Skin Lesion Classification, Q. Su et al., IEEE Access, 2024

Data, tools and M.Tech extension: Data and tools: HAM10000, the public seven-class dermoscopy set the paper uses for both generation and classification, with PyTorch on a free cloud GPU and FID to score the generated images. M.Tech extension: replace STGAN with a class-conditional diffusion model, and put per-class recall and FID in one table for four set-ups (no augmentation, flips and crops, STGAN and diffusion), with every lesion kept in a single split.

TOPIC IDEA · M.TECH · MEDICAL IMAGING

Diabetic retinopathy grading with an optimised ensemble, tested on Indian fundus photos

Problem: India has one of the largest diabetic populations in the world, and there are not enough eye specialists to grade every patient's retina photo for diabetic retinopathy.

Output: You extract features from pre-trained ShuffleNet and ResNet-18, feed them to an error-correcting output code (ECOC) ensemble tuned by adaptive differential evolution, and report accuracy for the five-grade and two-grade tasks on APTOS.

− The base paper and the M.Tech extension

Base paper: Diabetic Retinopathy Detection and Grading: A Transfer Learning Approach Using Simultaneous Parameter Optimization and Feature-Weighted ECOC Ensemble, W. K. Wong et al., IEEE Access, 2023

Data, tools and M.Tech extension: Data and tools: APTOS, and the EyePACS plus Messidor-2 combination, both named in the abstract; MATLAB, or Python with scikit-learn, for the ensemble. M.Tech extension: test the trained model without retraining on IDRiD, the Indian Diabetic Retinopathy Image Dataset, add a Vision Transformer as a third feature extractor, and tabulate quadratic weighted kappa and per-grade recall for each extractor on APTOS

and IDRiD.

TOPIC IDEA · M.TECH · MEDICAL IMAGING

Brain tumour type from MRI with a two-network ensemble, checked for duplicate images

Problem: A brain MRI report needs the tumour type, not only whether a tumour is present, and glioma, meningioma and pituitary tumours can look alike on a single slice.

Output: You train ResNet50 and EfficientNet-B7 on the four-class Crystal Clean MRI set, combine them into the paper's ensemble, and report per-class recall and a confusion matrix for each network and for the ensemble.

− The base paper and the M.Tech extension

Base paper: Ensemble Deep Learning Models for Enhanced Brain Tumor Classification by Leveraging ResNet50 and EfficientNet-B7 on High-Resolution MRI Images, R. Singh et al., IEEE Access, 2024
Data, tools and M.Tech extension: Data and tools: the Crystal Clean: Brain Tumor MRI Dataset, which the abstract names, with over 22,000 images in four classes; PyTorch or Keras on a free cloud GPU. M.Tech extension: find and remove near-duplicate images with perceptual hashing before you split, retrain, and show in one table how accuracy and per-class recall change for each model before and after the clean-up.

TOPIC IDEA · M.TECH · MEDICAL IMAGING

Pneumonia detection on chest X-rays with a domain-adapted transformer

Problem: Pneumonia is a leading cause of death among children under five, and spotting it on a chest X-ray depends on subtle patterns that need a trained eye.

Output: You rebuild DAViT, a Vision Transformer joined to a shallow CNN and adapted on chest X-rays, compare it with an ImageNet-only model on the 5,856-image test set, and check its Grad-CAM heatmaps.

− The base paper and the M.Tech extension

Base paper: DAViT: A Domain-Adapted Vision Transformer for Automated Pneumonia Detection and Explanation Using Chest X-Ray Images, M. Fu et al., IEEE Access, 2025
Data, tools and M.Tech extension: Data and tools: the 5,856 chest X-rays the paper evaluates on, and the training code and pre-trained models the authors put on GitHub, as the abstract states. M.Tech extension: test the trained model on X-rays from a second public source without retraining, and tabulate F1 and AUC on both sources beside how often the heatmap's peak falls inside the lung fields.

Agriculture and environment

Field photos differ from benchmark photos in light, background and angle. A test on images from a real field gives your thesis a result the base paper did not have.

TOPIC IDEA · M.TECH · AGRICULTURE AND ENVIRONMENT

Tomato leaf disease detection on a phone with knowledge distillation

Problem: Tomato is grown across India, and a disease checker has to run on an ordinary phone in the field, where lab-trained models lose accuracy and large networks do not fit.

Output: You train a four-network soft-voting ensemble on the paper's 15-class PlantVillage and TomatoVillage benchmark, distil it into ShuffleNetV2, quantise it to INT8, and report accuracy, model size and inference time at each step.

− The base paper and the M.Tech extension

Base paper: Deployable Deep Learning for Cross-Domain Plant Leaf Disease Detection via Ensemble Learning, Knowledge Distillation, and Quantization, M. J. Hasan et al., IEEE Access, 2025

Data, tools and M.Tech extension: Data and tools: the open cross-domain benchmark and the code the authors released on GitHub; PyTorch for training, then TensorFlow Lite or ONNX for the phone. M.Tech extension: photograph tomato leaves in a local field, label a small test set, and compare the ensemble, the distilled model and the quantised model on it beside the benchmark test set in one table.

TOPIC IDEA · M.TECH · AGRICULTURE AND ENVIRONMENT

Spotting new construction in satellite image pairs with a Siamese network

Problem: Indian cities are growing fast, and town planners need to know where new buildings have come up between two dates without comparing satellite images by eye.

Output: You build SIAM-CDNET, a Siamese network on a modified U-Net++ with attention, an edge-detection branch and a composite loss for change detection, and report F1 and accuracy on the LEVIR-CD, CDD and GZ-CD benchmarks.

− The base paper and the M.Tech extension

Base paper: SIAM-CDNET: A Remote Sensing Image Change Detection Network for Optimized Edge Detection and Mitigation of Pseudo Changes, Y. Zhang and W. Xue, IEEE Access, 2024

Data, tools and M.Tech extension: Data and tools: LEVIR-CD, CDD and GZ-CD, the three public change detection sets named in the abstract; PyTorch on a free cloud GPU, with the large images cut into tiles. M.Tech extension: train on LEVIR-CD only and test on GZ-CD without retraining, then fine-tune on a small share of GZ-CD, and tabulate F1 and IoU for each setting against a plain Siamese U-Net.

TOPIC IDEA · M.TECH · AGRICULTURE AND ENVIRONMENT

Early smoke and wildfire detection: choosing which YOLO to deploy

Problem: Forest fires break out in the dry season in hill states such as Uttarakhand, and a watch camera must flag the first smoke without raising alarms on mist or cloud.

Output: You train YOLOv5, YOLOv6, YOLOv7, YOLOv8 and YOLO-NAS under the same settings on the 8,974–image Foggia dataset, and compare precision, recall, F1, mAP and inference time.

− The base paper and the M.Tech extension

Base paper: Assessing the Effectiveness of YOLO Architectures for Smoke and Wildfire Detection, E. Casas et al., IEEE Access, 2023
Data, tools and M.Tech extension: Data and tools: the Foggia smoke and fire dataset used in the paper; the Ultralytics and YOLO-NAS libraries on a free cloud GPU. M.Tech extension: build a test subset of distant and poorly lit scenes, which the abstract names as the models' weak point, add a newer YOLO version, and tabulate recall, false positives and frames per second on a small board for every model.

Surveillance, traffic and safety

Camera projects raise privacy questions. Train on public datasets, blur faces in any footage you record yourself, and say so in the thesis.

TOPIC IDEA · M.TECH · SURVEILLANCE AND SAFETY

Driver drowsiness alarm on a Raspberry Pi, tested across two datasets

Problem: Truck and bus drivers on long highway routes often drive through the night, and an alarm that tracks only the eyes fails when the driver wears glasses or the light changes.

Output: You detect the face in each video frame, classify the whole face as drowsy or alert with a Vision Transformer, and deploy it to a Raspberry Pi fitted with an infrared camera and a buzzer, reporting accuracy and frame rate.

The base paper and the M.Tech extension

Base paper: A Real-Time Vision Transformers-Based System for Enhanced Driver Drowsiness Detection and Vehicle Safety, A. Jarndal et al., IEEE Access, 2025
Data, tools and M.Tech extension: Data and tools: NTHU-DDD and UTA-RLDD, the two public datasets in the paper; PyTorch for training, and an infrared camera on a Raspberry Pi for the prototype. M.Tech extension: train on NTHU-DDD and test on UTA-RLDD, then the reverse, and compare the Vision Transformer with a MobileNet in one table of cross-dataset accuracy, frames per second and power drawn on the Pi.

TOPIC IDEA · M.TECH · SURVEILLANCE AND SAFETY

Fall detection at home from a room camera, without storing faces

Problem: A fall at home is most dangerous when nobody sees it, and wearable sensors are uncomfortable, need upkeep and can give false readings.

Output: You chain YOLOv8 person detection with a time-space transformer that reads the motion across frames, and report mAP and F1 for fall and non-fall events on CUCAFall and DiverseFALL10500.

The base paper and the M.Tech extension

Base paper: Hybrid Deep Learning for Human Fall Detection: A Synergistic Approach Using YOLOv8 and Time-Space Transformers, Y. Sanjalawe et al., IEEE Access, 2025
Data, tools and M.Tech extension: Data and tools: CUCAFall and DiverseFALL10500, the two benchmarks in the paper; Ultralytics YOLOv8 and PyTorch. M.Tech extension: feed pose keypoints instead of image crops to the temporal model, so no face is ever stored, and tabulate F1, false alarms and frames per second for both input types on a camera angle held out from training.

TOPIC IDEA · M.TECH · SURVEILLANCE AND SAFETY

Pothole detection on a Raspberry Pi 5 with phase congruency and YOLOv8n

Problem: Potholes appear on Indian roads every monsoon, and manual inspection is slow, needs staff and cannot watch a whole city in real time.

Output: You compute phase congruency maps, which stay stable when the lighting changes, feed them into a channel of YOLOv8n, and run the detector on a Raspberry Pi 5 with a Pi Camera, a Telegram alert and a web dashboard, reporting mAP@0.5, precision and F1.

The base paper and the M.Tech extension

Base paper: Real-Time Pothole Detection System Using Phase Congruency Enhanced YOLO Model on Edge Device, H. Padala et al., IEEE Access, 2026
Data, tools and M.Tech extension: Data and tools: the abstract does not name the training images, so read the paper's data section and add road photos of your own; Ultralytics YOLOv8n, an open phase congruency implementation, a Raspberry Pi 5 and a Pi Camera v2. M.Tech extension: collect test images of local roads in daylight, shadow, rain and at night, and tabulate mAP@0.5 and frames per second on the Pi for plain YOLOv8n against the phase congruency version in each condition.

TOPIC IDEA · M.TECH · SURVEILLANCE AND SAFETY

Helmet detection for every rider on a two-wheeler, from traffic cameras

Problem: Traffic police in Indian cities issue e-challans for riding without a helmet, and a camera system has to check every person on the bike, including the pillion rider.

Output: You build the paper's detector, a Swin Transformer backbone with a Feature Pyramid Network and Cascade R-CNN, that finds each rider and whether they wear a helmet, and compare its mAP with other detectors on the same split.

The base paper and the M.Tech extension

Base paper: A Swin Transformer-Based Approach for Motorcycle Helmet Detection, A. Bouhayane et al., IEEE Access, 2023
Data, tools and M.Tech extension: Data and tools: the abstract does not name the dataset, so read the paper's experiments section and use a public set with helmet labels for each rider; MMDetection has Swin backbones and Cascade R-CNN ready to configure. M.Tech extension: label a few hundred frames of Indian street traffic with pillion and triple riding, and tabulate mAP and frames per second for the Swin model and a single-stage YOLO detector on the public test set and on your Indian frames.

Image enhancement and hardware

These three suit M.Tech students in ECE, signal processing and VLSI. For more hardware topics, see our M.Tech projects for ECE.

TOPIC IDEA · M.TECH · ENHANCEMENT AND HARDWARE

Low-light image enhancement with an intuitionistic fuzzy Retinex method

Problem: Photos and CCTV frames taken at night or in dim rooms lose detail and colour, and simply brightening them makes noise worse and washes out the colour.

Output: You implement the paper's intuitionistic fuzzy generator with histogram equalisation Retinex, and score the output with full-reference metrics such as PSNR, AMBE and mutual information, and no-reference metrics such as entropy and BRISQUE.

The base paper and the M.Tech extension

Base paper: Low-Light Image Enhancement via New Intuitionistic Fuzzy Generator-Based Retinex Approach, M. S. Ragavendirane and S. Dhanasekar, IEEE Access, 2025
Data, tools and M.Tech extension: Data and tools: MATLAB, or Python with OpenCV, and public low-light sets such as LOL, which pairs each dark image with a normal-light one for the full-reference metrics; no GPU is needed. M.Tech extension: compare the method with CLAHE and a small deep low-light network on the same images, and tabulate PSNR, SSIM, BRISQUE and run time per image on a laptop CPU.

TOPIC IDEA · M.TECH · ENHANCEMENT AND HARDWARE

Real-time image denoising on an FPGA with an approximate bilateral filter

Problem: Machine-vision cameras on factory lines need noise removed at full frame rate, and an exact bilateral filter costs too much logic and power for a low-cost FPGA.

Output: You design a reconfigurable approximate bilateral filter that adapts to the noise level at run time, and compare its PSNR and SSIM with an exact software filter, alongside LUTs, flip-flops, DSP blocks, frame rate and power on a Zynq-7020.

The base paper and the M.Tech extension

Base paper: Design of Approximate Bilateral Filters for Image Denoising on FPGAs, F. Spagnolo et al., IEEE Access, 2023
Data, tools and M.Tech extension: Data and tools: standard grayscale test images with added Gaussian noise; Verilog or high-level synthesis in Vivado, a PYNQ-Z2 or another Zynq-7020 board, and Python or MATLAB as the software reference. M.Tech extension: add a noise estimator that picks the filter setting by itself, and tabulate PSNR, SSIM, resources and power for the exact, fixed approximate and adaptive versions at three noise levels.

TOPIC IDEA · M.TECH · ENHANCEMENT AND HARDWARE

YOLO object detection on a low-cost Kria FPGA board

Problem: Running detection inside a roadside camera saves sending video to a server, but a low-cost FPGA holds little memory on chip, and every trip to off-chip RAM costs time and power.

Output: You prune YOLOv3–Tiny a few filters at a time, apply quantisation-aware training with a scale for each layer, keep the data in on-chip BRAM and URAM, and report detection accuracy, frames per second, latency and power on a Kria KV260.

− The base paper and the M.Tech extension

Base paper: Low-Cost FPGA-Enhanced CNN Accelerator for Real-Time YOLO Object Detection and Classification, J. S. Fata and W. M. Elmannai, IEEE Access, 2026
Data, tools and M.Tech extension: Data and tools: PASCAL VOC or another public detection set, PyTorch for pruning and quantisation, and Vitis or Vivado with a Kria KV260 starter kit. M.Tech extension: retrain on the India Driving Dataset (IDD), which includes autorickshaws and two-wheelers, and tabulate mAP, frames per second and watts for the vendor's standard flow against your on-chip memory design.

Want topics in other areas? See our M.Tech project ideas across CSE, AI and ECE, and the AI-generated image detection guide for one more vision topic worked out in full.

KEY POINTS

- **14 M.Tech projects in image processing and computer vision,** each on an open-access IEEE paper from 2023 to 2026.
- **Four areas:** medical image analysis (4), agriculture and environment (4), surveillance and safety (4), and image enhancement and hardware (3).
- **Every card names one extension** and the comparison table it produces, because an M.Tech result is a measured difference.
- **Four cards run on hardware:** a Raspberry Pi, a Zynq-7020 or a Kria KV260. For microcontroller vision, see our M.Tech embedded systems projects.
- **No accuracy figures are quoted here.** Take the numbers from each paper, and reproduce them in your own runs before you rely on them.

Shortlisted a vision project?

Send the card you like or your own base paper, your specialisation and your stage dates. The consultation is free.

Message in the next 7 h 19 min for a reply today.

Free consultation →

WhatsApp

How does an M.Tech vision project run, stage by stage?

Most universities split the M.Tech project into two stages, each with a report and a review. For an image processing project the work divides like this.

How an M.Tech image processing project splits across Stage-I and Stage-II

STAGE	IMAGE PROCESSING WORK	WHAT YOU HAND IN
Stage-I	Literature review around the base paper; download the dataset and fix the splits; reproduce the baseline with the paper's metric; write the problem statement and the extension plan	The Stage-I report, your numbers beside the paper's in one table, and a review presentation
Stage-II	Build the extension; run the base method and your version on one shared test set, several times; add confusion matrices, heatmaps or sample outputs; port to the board if the card has one	The results chapter, the full thesis, a paper where your university asks for one, and viva preparation

Stage names, dates and report rules differ between universities; follow your own regulations.

The hardest part of Stage-I is usually matching the paper's number. The base paper implementation guide covers what to do when your result falls short, and the base paper presentation guide covers the first review. For Stage-II, the M.Tech thesis format guide sets out the chapters and both stage reports, and the IEEE paper format guide covers the paper.

Which datasets, GPUs and boards will you need?

- **Public datasets.** Every card trains on public data; some extensions add photos you take yourself. Read each dataset's licence before you start, because many medical sets allow research use only.
- **Free GPUs.** Google Colab and Kaggle notebooks give a free GPU with limits on session length and weekly hours, and those limits change. Save checkpoints often, test with small images first, and keep the heaviest models (EfficientNet-B7, Swin, change detection at full size) for a college GPU or paid cloud time.
- **Boards.** Five cards need hardware. Order the board early, because the port to the board always takes longer than the training.

The 14 projects with their data and the compute or board each needs

PROJECT	DATA	COMPUTE OR BOARD
Skin lesion rare classes	HAM10000	GPU for the GAN or diffusion model

PROJECT	DATA	COMPUTE OR BOARD
Diabetic retinopathy grading	APTOS, EyePACS + Messidor-2, IDRiD	Free GPU
Brain tumour MRI ensemble	Crystal Clean: Brain Tumor MRI	GPU (EfficientNet-B7 is large)
Pneumonia chest X-ray	5,856 chest X-rays (see paper)	Free GPU; authors' code on GitHub
Tomato leaf on a phone	PlantVillage + TomatoVillage	Free GPU, then an Android phone
Satellite change detection	LEVIR-CD, CDD, GZ-CD	GPU
Smoke and wildfire YOLO	Foggia dataset	Free GPU, then a small board
Driver drowsiness	NTHU-DDD, UTA-RLDD	Free GPU, then a Raspberry Pi with an IR camera
Fall detection	CUCAFall, DiverseFALL10500	Free GPU
Pothole detection	Public pothole images + your own	Free GPU, then a Raspberry Pi 5
Two-wheeler helmet detection	See paper + your labelled frames	GPU (Swin with Cascade R-CNN)
Low-light enhancement	LOL and other low-light sets	Laptop CPU
Bilateral filter on FPGA	Standard test images	Zynq-7020 board such as PYNQ-Z2
YOLO on a Kria FPGA	PASCAL VOC, IDD	GPU for training, then a Kria KV260

Datasets are as each paper's abstract names them, or a public set suggested for the extension. Check that you can download the data before you write the synopsis.

How can we help with your M.Tech image processing project?

We start by checking your base paper: that the data is public, that the method can be rebuilt in your time frame, and that the extension can be measured. Then we reproduce the baseline, build and run the extension, write the results chapter and the thesis in your university's format, and prepare an IEEE-format paper and your viva if you need them.

See how it runs, with timelines and what you receive, on our M.Tech projects for CSE and AI page. For FPGA, microcontroller and other hardware image processing, see M.Tech projects for ECE.

Our delivered vision work includes a HAM10000 skin lesion classifier with a Vision Transformer and a diabetic retinopathy Vision Transformer ensemble with heatmaps, each shown with its paper pages and screens. The full set is on our work page.

More M.Tech domains, each built on verified IEEE base papers: M.Tech projects in machine learning, M.Tech projects in embedded systems and IoT, M.Tech projects in cyber security. For the full service, see our M.Tech projects for CSE and AI.

Frequently asked questions

Which image processing topic is best for an M.Tech project?

One with a recent open-access base paper, a public dataset you can download today and an extension you can measure. The medical imaging cards suit CSE and AI students with GPU access, and the FPGA and microcontroller cards suit ECE and VLSI students. Our guide to selecting a base paper lists the checks to run before you commit.

Should I use MATLAB or Python for an M.Tech image processing project?

Python with PyTorch suits the deep learning cards, because most base paper code is written in Python. MATLAB still suits classical enhancement and filtering, such as the fuzzy Retinex card, and many ECE guides prefer it. The hardware cards add Verilog or high-level synthesis. Pick the language your guide can review.

Can I do an M.Tech computer vision project on a free GPU?

Most of these, yes, on Colab or Kaggle notebooks, if you save checkpoints often and test with smaller images first. The free tiers limit session length and weekly hours, and the limits change, so large models such as EfficientNet-B7 or a Swin detector may need a college GPU or paid cloud time for the final runs.

What should the M.Tech extension be in an image processing project?

One change you can measure against the base paper on the same test set: a second dataset, a lighter model on a board, harder conditions such as night or rain, or a check of where the model looks. Every card above names one, with the table it would produce. See how to implement a base paper for the reproduction that comes first.

Can an M.Tech image processing project lead to a published paper?

Yes, if the extension gives a clear comparison. A results table against the base paper on the same data is the core of a short paper. Check which venues your university accepts before you submit. Our research paper writing service covers the IEEE format and a journal shortlist.

Can The Ultimate Project World build my M.Tech image processing project?

Yes. Send the card you like or your own base paper, your specialisation and your Stage-I and Stage-II dates in the free consultation. We tell you which parts we can take on, from reproducing the paper to the extension, the thesis, a paper and viva preparation. See our M.Tech projects service for how it runs.

RELATED SERVICES

From a topic to a project you can present

M.Tech projects
A base paper, one measured extension, a paper and the thesis. →

IEEE projects
Start from a published base paper and extend it. →


Research papers
A paper in IEEE, Springer or journal format, and a journal shortlist.