PROJECT IDEAS · IEEE BASE PAPERS · 2026

IEEE base papers for projects, *by domain*

These 28 IEEE base papers for B.Tech final year and M.Tech projects were published from 2023 to 2025, and all are open access, so you can read the full text free. Each card links the paper by its DOI, names its dataset and says **what to reproduce for B.Tech and what to change for M.Tech**.

Last updated 2026-10-22

Every paper below was checked against its DOI record on Crossref on 22 October 2026, so the title, journal and year are exactly as IEEE registered them. This page lists papers; how to judge a paper of your own is covered in our guide on how to select a base paper.

Which IEEE base papers are on this page?

The table lists all 28 with the domain, the dataset, the journal and the year. All are open access: 27 are in IEEE Access and one is in the IEEE Open Journal of Engineering in Medicine and Biology. Tap a paper to open its card.

28 IEEE base papers for projects, with the domain, dataset, journal, year and access

PAPER	DOMAIN	DATASET	JOURNAL	YEAR	OPEN ACCESS
Skin lesion GAN augmentation	Medical imaging	HAM10000	IEEE Access	2024	Yes

PAPER	DOMAIN	DATASET	JOURNAL	YEAR	OPEN ACCESS
Diabetic retinopathy ensemble	Medical imaging	APTOS, EyePACS + Messidor-2	IEEE Access	2023	Yes
Brain tumour ResNet50 + EfficientNet-B7	Medical imaging	Crystal Clean: Brain Tumor MRI	IEEE Access	2024	Yes
Pneumonia domain-adapted ViT	Medical imaging	5,856 chest X-rays	IEEE Access	2025	Yes
Lung nodule LoRA	Medical imaging	LIDC, NLSTx	IEEE Open Journal of Engineering in Medicine and Biology	2025	Yes
Plant disease on a phone	Computer vision	PlantVillage, TomatoVillage	IEEE Access	2025	Yes
Driver drowsiness ViT	Computer vision	NTHU-DDD, UTA-RLDD	IEEE Access	2025	Yes
Deepfake video ViViT	Computer vision	Celeb-DF v2	IEEE Access	2024	Yes
Smoke and wildfire YOLO	Computer vision	Foggia dataset	IEEE Access	2023	Yes
Fall detection YOLOv8	Computer vision	CUCAFall, DiverseFALL10500	IEEE Access	2025	Yes
Fake news BERT hybrid	NLP and LLMs	WELFake	IEEE Access	2024	Yes
RAG vs fine-tuning	NLP and LLMs	SQuAD, MS MARCO	IEEE Access	2025	Yes
Dravidian code-mixed hate speech	NLP and LLMs	DravidianLangTech, HASOC	IEEE Access	2024	Yes
SMS spam evasion	NLP and LLMs	The authors' 68,000+ SMS set	IEEE Access	2024	Yes
LLM vulnerability detection	NLP and LLMs	Java code	IEEE Access	2025	Yes, and on arXiv
Phishing URL voting hybrid	Cybersecurity	11,000+ website URLs	IEEE Access	2023	Yes
Explainable IIoT intrusion detection	Cybersecurity	ToN_IoT	IEEE Access	2023	Yes
IoT botnet stacked model	Cybersecurity	UNSW-NB15	IEEE Access	2024	Yes

PAPER	DOMAIN	DATASET	JOURNAL	YEAR	OPEN ACCESS
Android malware co-existing features	Cybersecurity	Drebin, Malgenome, MalDroid2020	IEEE Access	2023	Yes
SecurityBERT	IoT and edge	Edge-IIoTset	IEEE Access	2024	Yes, and on arXiv
Federated IoT attack detection	IoT and edge	CIC IoT 2023	IEEE Access	2023	Yes
ESP32 smart irrigation	IoT and edge	Your own sensor readings	IEEE Access	2025	Yes
TinyML solar forecasting	IoT and edge	Solar farm yield (see paper)	IEEE Access	2024	Yes
Battery state of health, federated	Forecasting and energy	NASA battery dataset	IEEE Access	2025	Yes
Engine remaining useful life	Forecasting and energy	C-MAPSS	IEEE Access	2023	Yes
Bearing fault diagnosis	Forecasting and energy	CWRU bearing data	IEEE Access	2023	Yes
Electricity theft KNN-LSTM	Forecasting and energy	SGCC smart-meter data	IEEE Access	2025	Yes
Solar PV 1D CNN-GRU	Forecasting and energy	Hydro-Floating Solar Plant	IEEE Access	2024	Yes

“Open access” means the full text is free to read on IEEE Xplore. Datasets are as the paper’s abstract or title names them; read the paper’s data section and the dataset’s licence before you commit.

What does each base paper solve, and what do you build from it?

Each card gives the problem the paper solves and the output you build and show. Open “The paper and what to change” for the full reference with its DOI link, the dataset, what to reproduce for B.Tech and one change to measure for M.Tech.

Topic ideas, not delivered projects

Each card is built on a published IEEE paper, checked against its DOI record. The project around it is an idea to discuss with your guide. Our 8 delivered case studies are M.Tech and M.E. projects, each shown with its real paper pages, screens and results.

Medical imaging and health AI

Where one patient or lesion gives several images, split by patient or lesion, not by image, or the same person ends up in both training and test data.

TOPIC IDEA · DEEP LEARNING · MEDICAL IMAGING

Skin lesion classification with GAN-made images for rare classes

Problem: Skin lesion datasets hold many images of harmless moles and few of rare cancers, so a classifier tends to miss the rare classes.

Output: Reproduce the paper's two-stage GAN on HAM10000, add its synthetic images to the rare classes, and show recall for each class with and without them.

− The paper and what to change

Base paper: A GAN-Based Data Augmentation Method for Imbalanced Multi-Class Skin Lesion Classification, Q. Su et al., IEEE Access, 2024
For B.Tech / for M.Tech: Dataset: HAM10000, used in the paper for both generating and classifying images. B.Tech: reproduce a CNN classifier on the real images first, then add the paper's synthetic images and compare recall for each class. M.Tech: replace the GAN with a diffusion model, or test on an ISIC challenge set the model never saw, keeping every lesion in only one split. See our delivered HAM10000 Vision Transformer case study.

TOPIC IDEA · DEEP LEARNING · MEDICAL IMAGING

Diabetic retinopathy grading with a transfer learning ensemble

Problem: Grading retina photos for diabetic retinopathy takes an eye specialist's time, and there are too few specialists to screen every diabetic patient.

Output: Features from ShuffleNet and ResNet-18 fed to an error-correcting output code ensemble that grades a fundus photo, with the 5-grade and 2-grade results shown on APTOS.

− The paper and what to change

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
For B.Tech / for M.Tech: Datasets: APTOS, and EyePACS combined with Messidor-2, as the abstract states. B.Tech: reproduce the two pre-trained networks and the ensemble on APTOS. M.Tech: swap the feature extractors for a Vision Transformer, train on APTOS and test on Messidor-2, and add heatmaps a doctor can check. See our delivered diabetic retinopathy ViT ensemble.

TOPIC IDEA · DEEP LEARNING · MEDICAL IMAGING

Brain tumour classification with a ResNet50 and EfficientNet-B7 ensemble

Problem: Sorting brain MRI scans by tumour type takes a radiologist's time, and a single network can confuse one tumour type with another.

Output: An ensemble of ResNet50 and EfficientNet-B7 that labels an MRI as no tumour, glioma, meningioma or pituitary tumour, with recall for each class and a confusion matrix.

− The paper and what to change

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
For B.Tech / for M.Tech: Dataset: the Crystal Clean: Brain Tumor MRI Dataset named in the abstract, over 22,000 images in four classes. B.Tech: reproduce one network first, then the ensemble. M.Tech: check the set for near-duplicate images before you split, test on a second public MRI set, and compare with a lighter model that runs on a laptop CPU.

TOPIC IDEA · DEEP LEARNING · MEDICAL IMAGING

Pneumonia detection with a domain-adapted Vision Transformer

Problem: Pneumonia is diagnosed from chest X-rays, and models pre-trained only on everyday photos can miss the patterns that matter in medical images.

Output: A hybrid of a Vision Transformer and a shallow CNN, adapted on chest X-rays, that labels an X-ray as normal or pneumonia and shows where it looked, beside a plain CNN on the same split.

− The paper and what to change

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
For B.Tech / for M.Tech: Data: the abstract describes a set of 5,856 chest X-rays for testing and a separate chest X-ray set for the domain adaptation; check the paper's data section for both sources. B.Tech: reproduce the normal-versus-pneumonia result with a pre-trained ViT and a CNN baseline. M.Tech: test on X-rays from a second source, and check whether the explanations fall on the lungs or on text and markers at the edge.

TOPIC IDEA · DEEP LEARNING · MEDICAL IMAGING

Lung nodule malignancy classification with LoRA on a large vision model

Problem: Telling a cancerous lung nodule from a harmless one needs biopsy-confirmed data, and fully fine-tuning a large vision model on a small medical set is slow and costly.

Output: A large pre-trained vision model adapted with low-rank adaptation (LoRA) that classifies nodules as benign or malignant, compared with full fine-tuning on results, trainable parameters and training time.

− The paper and what to change

Base paper: Low-Rank Adaptation of Pre-Trained Large Vision Models for Improved Lung Nodule Malignancy Classification, B. P. Veasey and A. A. Amini, IEEE Open Journal of Engineering in Medicine and Biology, 2025
For B.Tech / for M.Tech: Datasets: LIDC and NLSTx; the authors released the NLSTx annotations on GitHub with the paper. B.Tech: suits a strong team with GPU time; reproduce LoRA against full fine-tuning on LIDC. M.Tech: compare LoRA with other parameter-efficient methods, or train on one dataset and test on the other.

Computer vision

TOPIC IDEA · DEEP LEARNING · COMPUTER VISION

Plant leaf disease detection that runs on a phone

Problem: Leaf disease models trained on lab photos lose accuracy in the field, and the best models are too large for a farmer's phone.

Output: A four-model ensemble distilled into a small ShuffleNetV2 and then quantized, with accuracy, model size and speed shown at each step and a test across datasets.

− The paper and what to change

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
For B.Tech / for M.Tech: Datasets: PlantVillage and TomatoVillage, which the paper joins into one tomato benchmark. B.Tech: reproduce one CNN and the distilled ShuffleNetV2, and run it in a phone app. M.Tech: test on field photos you collect, or use a Vision Transformer as the teacher, and report what accuracy the small model gives up.

TOPIC IDEA · DEEP LEARNING · COMPUTER VISION

Driver drowsiness detection with a Vision Transformer on a Raspberry Pi

Problem: Drowsy driving causes fatal crashes, and systems that track only the eyes struggle when drivers wear glasses or the light changes.

Output: A face detector and a Vision Transformer that label each video frame drowsy or alert, running on a Raspberry Pi with an infrared camera and a buzzer.

− The paper and what to change

Base paper: A Real-Time Vision Transformers-Based System for Enhanced Driver Drowsiness Detection and Vehicle Safety, A. Jarndal et al., IEEE Access, 2025
For B.Tech / for M.Tech: Datasets: NTHU-DDD and UTA-RLDD, two public drowsiness datasets. B.Tech: reproduce the classifier on one dataset and the Raspberry Pi demo with a camera. M.Tech: train on one dataset and test on the other, report frames per second on the Pi, and compare with a smaller CNN.

TOPIC IDEA · DEEP LEARNING · COMPUTER VISION

Deepfake video detection with a Video Vision Transformer

Problem: Deepfake videos are hard to spot by eye, and detectors that do well on easy datasets often fail on harder ones.

Output: A Video Vision Transformer fed with facial landmark regions from each video that labels it real or fake, with an ablation that shows what each part adds.

− The paper and what to change

Base paper: Improving Video Vision Transformer for Deepfake Video Detection Using Facial Landmark, Depthwise Separable Convolution and Self Attention, K. N. Ramadhani et al., IEEE Access, 2024
For B.Tech / for M.Tech: Dataset: Celeb-DF version 2. B.Tech: reproduce the pipeline on a subset of Celeb-DF v2, with a frame-level CNN as the baseline. M.Tech: train on Celeb-DF and test on a second deepfake dataset, and measure how video compression changes the result.

TOPIC IDEA · DEEP LEARNING · COMPUTER VISION

Smoke and wildfire detection: comparing YOLO versions

Problem: Wildfires spread fast, and a camera system must spot smoke early without raising false alarms on cloud or fog.

Output: YOLOv5, YOLOv6, YOLOv7, YOLOv8 and YOLO-NAS trained the same way and compared on precision, recall, mAP and speed, with examples of what each misses.

− The paper and what to change

Base paper: Assessing the Effectiveness of YOLO Architectures for Smoke and Wildfire Detection, E. Casas et al., IEEE Access, 2023
For B.Tech / for M.Tech: Dataset: the Foggia dataset of 8,974 images. B.Tech: reproduce two of the YOLO versions and a live demo on video. M.Tech: add a newer YOLO version, test on distant and low-light scenes, which the paper names as the weak point, and report speed on a small board.

TOPIC IDEA · DEEP LEARNING · COMPUTER VISION

Human fall detection with YOLOv8 and a time-space transformer

Problem: Older people who live alone can lie for hours after a fall, and wearable alarms are often not worn.

Output: A camera pipeline in which YOLOv8 finds the person and a time-space transformer decides whether a fall happened, with results on two fall datasets.

The paper and what to change

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
For B.Tech / for M.Tech: Datasets: CUCAFall and DiverseFALL10500. B.Tech: reproduce YOLOv8 detection with a simple LSTM on top as the baseline. M.Tech: test on videos from a room and camera angle the model never saw, and use pose keypoints instead of raw frames to protect privacy.

NLP and large language models

TOPIC IDEA · NLP · LLMS

Fake news detection with BERT and a bidirectional layer

Problem: False news spreads faster than it can be checked, and simple models often learn the style of a source instead of whether a story is true.

Output: BERT with a bidirectional deep learning layer on top, compared with classical machine learning, deep learning and plain BERT on the same data, and a page that scores a pasted article.

− The paper and what to change

Base paper: A Hybrid Transformer-Based Model for Optimizing Fake News Detection, F. Al-Quayed et al., IEEE Access, 2024
For B.Tech / for M.Tech: Dataset: WELFake, which combines four news datasets. B.Tech: reproduce plain BERT and the hybrid model. M.Tech: hold out one of the four source datasets for testing, and test on Indian news articles in English to see whether the model still works.

TOPIC IDEA · NLP · LLMS

RAG, fine-tuning or both: comparing ways to adapt an LLM

Problem: A college or company that wants an LLM to answer from its own documents must choose between retrieval, fine-tuning or both, and each costs differently.

Output: Retrieval-augmented generation, fine-tuning and the two combined, compared on question answering with the same model and the same metrics.

− The paper and what to change

Base paper: Unveiling the Power of Large Language Models: A Comparative Study of Retrieval-Augmented Generation, Fine-Tuning, and Their Synergistic Fusion for Enhanced Performance, G. Budakoglu and H. Emekci, IEEE Access, 2025
For B.Tech / for M.Tech: Datasets: SQuAD, MS MARCO and a set of SQL CREATE TABLE statements, as the abstract lists. B.Tech: build a RAG chatbot over SQuAD passages with an open model and compare it with the model alone. M.Tech: repeat the comparison on your college's own documents, such as the syllabus and rules, and measure how often each method gives an unsupported answer.

TOPIC IDEA · NLP · LLMS

Hate speech detection in code-mixed Tamil, Malayalam and Kannada

Problem: Indian social media posts mix English with regional languages, and English-only filters miss offensive comments written that way.

Output: MuRIL embeddings with a support vector machine and cost-sensitive learning that flag offensive code-mixed comments, compared across six datasets in three languages.

− The paper and what to change

Base paper: Detection of Hate Speech and Offensive Language CodeMix Text in Dravidian Languages Using Cost-Sensitive Learning Approach, K. Sreelakshmi et al., IEEE Access, 2024
For B.Tech / for M.Tech: Datasets: six openly available sets, including DravidianLangTech 2021 and HASOC 2021 Malayalam-English. B.Tech: reproduce MuRIL with an SVM on one language. M.Tech: fine-tune MuRIL or IndicBERT end to end, test on a language left out of training, and show the words behind each decision.

TOPIC IDEA · NLP · LLMS

SMS spam detection that holds up against evasion tricks

Problem: Spammers change spellings and wording to slip past filters, and the old SMS datasets most students use are small and dated.

Output: Shallow and deep models compared on a large new SMS dataset, with a test of how each holds up when the spam text is deliberately altered.

− The paper and what to change

Base paper: Investigating Evasive Techniques in SMS Spam Filtering: A Comparative Analysis of Machine Learning Models, M. Salman et al., IEEE Access, 2024
For B.Tech / for M.Tech: Dataset: the paper's own set of more than 68,000 SMS messages, which the authors released for research. B.Tech: reproduce TF-IDF with logistic regression and one deep model. M.Tech: add recent Indian scam texts about KYC, UPI and parcel delivery, and test a multilingual transformer on Hinglish messages.

TOPIC IDEA · NLP · LLMS

Finding security bugs in code with LLM prompting

Problem: Many software vulnerabilities are not caught before release, and traditional static analysis tools miss some of them.

Output: Several LLMs prompted with chain of thought, tree of thought and self-consistency to find vulnerabilities in Java code, compared with traditional static analysis tools on recall, precision and F1.

The paper and what to change

Base paper: Harnessing Large Language Models for Software Vulnerability Detection: A Comprehensive Benchmarking Study, K. Tamberg and H. Bahsi, IEEE Access, 2025

For B.Tech / for M.Tech: Data: Java source code; an author version of the paper is on arXiv. B.Tech: reproduce two prompting strategies with an open model and one static analysis tool. M.Tech: move to Python or C code, or add a checking step that cuts the false positives the paper names as the LLMs' weakness.

Cybersecurity

TOPIC IDEA · MACHINE LEARNING · SECURITY

Phishing website detection from URL features

Problem: Phishing sites copy banks and shops, and users cannot tell a fake address from a real one at a glance.

Output: Eight classifiers compared on URL features, then the paper's hybrid of logistic regression, a support vector classifier and a decision tree with soft and hard voting, and a page that checks a typed URL.

− The paper and what to change

Base paper: Phishing Detection System Through Hybrid Machine Learning Based on URL, A. Karim et al., IEEE Access, 2023
For B.Tech / for M.Tech: Dataset: phishing and legitimate URL attributes from more than 11,000 websites, taken from a public repository the paper names. B.Tech: reproduce the single models and the voting hybrid. M.Tech: test on phishing URLs collected after the dataset was made, and add features from the page content, not only the address.

TOPIC IDEA · MACHINE LEARNING · SECURITY

Explainable intrusion detection for industrial IoT

Problem: Attacks on factory networks can stop production, and operators do not trust an alarm when nobody can say why it fired.

Output: An ensemble deep learning intrusion detector with SHAP and LIME explanations for its decisions, compared with an extreme learning machine baseline.

− The paper and what to change

Base paper: An Explainable Ensemble Deep Learning Approach for Intrusion Detection in Industrial Internet of Things, M. M. Shtayat et al., IEEE Access, 2023
For B.Tech / for M.Tech: Dataset: ToN_IoT. B.Tech: reproduce the baseline, one deep model and the ensemble, then SHAP explanations for a few alerts. M.Tech: train on ToN_IoT and test on a second IoT attack dataset such as Edge-IIoTset, and check whether the explanations stay the same.

TOPIC IDEA · MACHINE LEARNING · SECURITY

IoT botnet attack detection with a stacked deep model

Problem: Hacked IoT devices are joined into botnets, and their attack traffic hides among normal traffic.

Output: A stack of ANN, CNN, LSTM and RNN models (the paper's ACLR) compared with each model on its own, using k-fold cross-validation and precision-recall curves.

− The paper and what to change

Base paper: Hybrid Machine Learning Model for Efficient Botnet Attack Detection in IoT Environment, M. Ali et al., IEEE Access, 2024
For B.Tech / for M.Tech: Dataset: UNSW-NB15, with normal traffic and nine attack types. B.Tech: reproduce the four single models and the stack. M.Tech: report recall for each rare attack type, not only overall accuracy, and test on a newer IoT dataset such as CIC IoT 2023.

TOPIC IDEA · MACHINE LEARNING · SECURITY

Android malware detection from permissions and API calls that occur together

Problem: Malicious apps ask for many of the same permissions as safe apps, so single features are weak signals.

Output: Features built from permissions and API calls that appear together in an app, tested with classical machine learning against single-feature baselines.

− The paper and what to change

Base paper: A Novel Machine Learning Approach for Android Malware Detection Based on the Co-Existence of Features, E. Odat and Q. M. Yaseen, IEEE Access, 2023
For B.Tech / for M.Tech: Datasets: APK samples from Drebin, Malgenome and MalDroid2020. B.Tech: reproduce the co-existence features and a random forest on Malgenome. M.Tech: test on apps released after the training data, and on obfuscated versions of the same apps.

IoT and edge AI

TOPIC IDEA · IOT · EDGE AI

SecurityBERT: a compact language model that detects IoT attacks

Problem: IoT and industrial devices face many kinds of attack, and large models are too heavy to run close to the devices.

Output: A compact BERT-based model that reads network traffic as text and sorts it into fourteen attack types, with its model size and CPU inference time reported.

− The paper and what to change

Base paper: Revolutionizing Cyber Threat Detection With Large Language Models: A Privacy-Preserving BERT-Based Lightweight Model for IoT/IIoT Devices, M. A. Ferrag et al., IEEE Access, 2024
For B.Tech / for M.Tech: Dataset: Edge-IIoTset. An author version of the paper is on arXiv. B.Tech: reproduce the traffic-to-text encoding and a small BERT. M.Tech: run the model on a Raspberry Pi and report the delay, or test it on CIC IoT 2023.

TOPIC IDEA · IOT · EDGE AI

Federated learning for IoT attack detection

Problem: IoT networks in different places cannot share raw traffic, yet one site alone sees too few examples of each attack.

Output: A federated model trained across simulated edge clients without sharing raw data, compared with the same model trained on all the data in one place.

− The paper and what to change

Base paper: A Novel Federated Edge Learning Approach for Detecting Cyberattacks in IoT Infrastructures, S. Abbas et al., IEEE Access, 2023
For B.Tech / for M.Tech: Dataset: CIC IoT 2023. B.Tech: simulate a few clients with the Flower framework or a simple federated averaging loop. M.Tech: give each client different attack types (non-IID data), and measure the drop in results and the data sent per round.

TOPIC IDEA · IOT · EDGE AI

Smart irrigation with an ESP32 and regression

Problem: Farmers often water by habit, wasting water on some days and under-watering on others.

Output: An ESP32 node with temperature, humidity and water-level sensors that predicts how much water the crop needs with a regression model and switches the pump.

The paper and what to change

Base paper: An Innovative Smart Irrigation Using Embedded and Regression-Based Machine Learning Technologies for Improving Water Security and Sustainability, A. Morchid et al., IEEE Access, 2025
For B.Tech / for M.Tech: Data: readings the system collects from its own sensors, so you log your own. B.Tech: build the node, log a few weeks of readings and fit linear regression. M.Tech: add a soil moisture sensor and weather forecasts, compare regression with a small neural network on the board, and measure the water saved.

TOPIC IDEA · IOT · EDGE AI

Solar energy forecasting on a microcontroller (TinyML)

Problem: Homes and small plants need a solar forecast on a cheap controller, without a cloud server.

Output: LSTM, bidirectional GRU, bidirectional LSTM and bidirectional RNN models compared for solar energy yield, with the trade-offs of running them on a microcontroller.

The paper and what to change

Base paper: Intelligent Solar Forecasts: Modern Machine Learning Models and TinyML Role for Improved Solar Energy Yield Predictions, A. M. Hayajneh et al., IEEE Access, 2024
For B.Tech / for M.Tech: Data: solar farm energy yield records; the abstract does not name the source, so read the paper's data section before you commit. B.Tech: compare two of the four models and run the smaller one with TensorFlow Lite for Microcontrollers. M.Tech: add weather inputs, quantize the model, and report the error, memory and time on the board.

Forecasting, energy and maintenance

For finished examples of time-series work, see our delivered network traffic prediction with Bi-LSTM and predictive maintenance with random forest and XGBoost.

TOPIC IDEA · TIME SERIES · ENERGY

Battery state of health prediction with federated learning

Problem: Battery packs lose capacity as they age, and owners and makers may not want to share raw battery data.

Output: 1D CNN, CNN-LSTM and CNN-GRU models compared on one machine, then the 1D CNN trained by federated averaging across five clients, with predicted and true state of health plotted.

The paper and what to change

Base paper: Lithium-Ion Battery State of Health Degradation Prediction Using Deep Learning Approaches, T. Alharbi et al., IEEE Access, 2025
For B.Tech / for M.Tech: Dataset: the NASA battery dataset of charge and discharge cycles. B.Tech: reproduce the 1D CNN and the plot of predicted against true state of health. M.Tech: test on a second battery dataset, or give each client batteries that aged differently to make the federated setting harder.

TOPIC IDEA · TIME SERIES · ENERGY

Remaining useful life of aircraft engines with LSTM-Transformer fusion

Problem: Engines serviced on a fixed schedule either lose good parts early or run failing ones too long.

Output: LSTM-Transformer fusion models that estimate the cycles left for each engine, scored with RMSE and the C-MAPSS score, beside an LSTM baseline.

The paper and what to change

Base paper: Novel Transformer-based Fusion Models for Aero-engine Remaining Useful Life Estimation, Q. Hu et al., IEEE Access, 2023
For B.Tech / for M.Tech: Dataset: the C-MAPSS turbofan engine benchmark. B.Tech: reproduce an LSTM baseline and one fusion model on one C-MAPSS subset. M.Tech: train on one subset and test on another with different operating conditions, and add a prediction interval to each estimate.

TOPIC IDEA · TIME SERIES · ENERGY

Bearing fault diagnosis from vibration signals

Problem: Bearings fail inside motors and pumps, and a fault caught early avoids a breakdown.

Output: Envelope analysis features from vibration signals fed to several machine learning models that name the bearing fault, compared to find the best one.

The paper and what to change

Base paper: Bearing Fault Diagnosis With Envelope Analysis and Machine Learning Approaches Using CWRU Dataset, M. Alonso-González et al., IEEE Access, 2023
For B.Tech / for M.Tech: Dataset: the public CWRU bearing dataset. B.Tech: reproduce the envelope features and two classifiers. M.Tech: train at one motor load and test at another, or compare the features with a 1D CNN on the raw signal.

TOPIC IDEA · TIME SERIES · ENERGY

Electricity theft detection from smart-meter data

Problem: Meter tampering and illegal connections cost utilities heavily, and rule-based checks miss new methods.

Output: LSTM combined in turn with KNN, random forest, XGBoost and logistic regression on daily consumption records, with the best combination chosen.

The paper and what to change

Base paper: Hybrid KNN-LSTM Framework for Electricity Theft Detection in Smart Grids Using SGCC Smart-Meter Data, S. Ness, IEEE Access, 2025
For B.Tech / for M.Tech: Dataset: SGCC smart-meter consumption from January 2014 to October 2016, as the title and abstract state. B.Tech: reproduce the KNN with LSTM combination and one other. M.Tech: theft cases are rare, so report precision, recall and AUC rather than accuracy alone, and find out why one combination scored far below the others.

TOPIC IDEA · TIME SERIES · ENERGY

Short-term solar power forecasting with a 1D CNN-GRU

Problem: Solar output changes with cloud and season, and grid operators need to know it hours ahead.

Output: A 1D CNN-GRU model compared with CNN, GRU, LSTM and CNN-GRU models, tested separately in winter, summer and the rainy season.

− The paper and what to change

Base paper: Enhancing Short-Term Solar Photovoltaic Power Forecasting Using a Hybrid Deep Learning Approach, N. Thipwangmek et al., IEEE Access, 2024
For B.Tech / for M.Tech: Data: the Hydro-Floating Solar Plant dataset from a 45 MW floating plant at Sirindhorn Dam, Thailand. Check that you can get it; if not, use a public solar plant dataset and say so. B.Tech: reproduce the five-model comparison. M.Tech: add prediction intervals, or train on one plant and test on another.

Want ideas before you pick a paper? See our lists of B.Tech project ideas and M.Tech project ideas, and the 25 projects with Kaggle datasets.

KEY POINTS

- **28 IEEE base papers from 2023 to 2025,** each checked against its DOI record for the exact title, journal and year.
- **Six domains:** medical imaging (5), computer vision (5), NLP and LLMs (5), cybersecurity (4), IoT and edge (4), and forecasting and energy (5).
- **All 28 are open access:** 27 in IEEE Access and one in the IEEE Open Journal of Engineering in Medicine and Biology.
- **Every card names the dataset,** what to reproduce for B.Tech and one change to measure for M.Tech.
- **We quote no accuracy figures.** Read the paper's own results, then match them yourself before you claim anything.

Picked a base paper?

Send its DOI with your level (B.Tech or M.Tech) and your review dates. The consultation is free.

Message in the next 6 h 27 min for a reply today.

Free consultation →

WhatsApp

How do you use a base paper in your project?

In four steps, in this order:

1. **Reproduce.** Rebuild the paper's method on its dataset, with the same split and the same metric.
2. **Match the baseline.** Put your numbers beside the paper's in one table and explain any gap.
3. **Make one change.** The M.Tech line on each card suggests one: a newer model, a second dataset, explanations or a lighter model.

4. **Compare.** Run the original and your version on the same test set, and show what your change added.

A B.Tech project can stop after step 2 with a working demo. An M.Tech project needs all four, written up as a thesis and often a paper. For how to find, check and test a paper of your own, read our guide on how to select a base paper; for the papers you cite around it, see our literature review guide.

Where can you read these papers for free?

- **IEEE Access and the IEEE Open Journals.** These are open access, so the full text and PDF of all 28 papers on this page are free on IEEE Xplore. The DOI link on each card opens the paper's page.
- **Other IEEE papers.** IEEE Transactions and most conference papers need a subscription. Use your college library's IEEE Xplore access, on campus or through its remote login.
- **Author versions.** Two papers here also have an author version on arXiv: SecurityBERT and the LLM vulnerability detection study. For other papers, Google Scholar's "All versions" link sometimes leads to a copy on the author's university page.

Cite the published IEEE version in your report even when you read an author version. Our IEEE paper format guide shows how to write the reference.

How can we help with an IEEE base paper project?

For a B.Tech team, we rebuild the paper's method on its dataset, put your results beside the paper's and add a working demo, with the black book, the PPT and viva preparation. For an M.Tech student, we check the paper first, reproduce it, then

plan, run and write up one measured extension as your thesis and, if you need one, an IEEE-format paper.

See how it runs on our IEEE projects page, with scope by level on final year projects for CSE and M.Tech projects.

Frequently asked questions

Where can I find IEEE base papers for 2026 projects?

Start with the 28 on this page, all from 2023 to 2025 and all free to read. For more, search IEEE Xplore with your task and a dataset name, limit the years to the last three, and open the IEEE Access results first because they are open access. Our guide to selecting a base paper covers the full search and the checks.

Are IEEE Access papers good enough as base papers?

Yes, for most B.Tech and M.Tech projects. IEEE Access is a peer-reviewed IEEE journal listed on IEEE Xplore, and its papers are open access, so you, your guide and your examiner can all read the full text. Check that the paper's data is public before you choose it. If you would rather start from a dataset than a paper, see our Kaggle dataset projects.

Which base paper suits a B.Tech final year project?

One you can reproduce in a few weeks on a laptop or a free cloud GPU, with a public dataset and a result you can demo. The plant disease, phishing URL, bearing fault and SMS spam papers above fit well. For what a complete project includes, see our final year projects for CSE.

What should an M.Tech student change in a base paper?

One thing you can measure: a newer model on the same split, a test on a second dataset, explanations you check, or a lighter model for a phone or a small board.

Every card above names one such change. Our M.Tech projects service shows how the extension and the thesis are planned.

Do I have to write my project paper in IEEE format?

Only if your college or guide asks for it. The layout of your own paper is separate from where your base paper was published, and some colleges ask for the IEEE two-column layout whatever the base paper. Our IEEE paper format guide lists the rules.

Can The Ultimate Project World help me implement one of these papers?

Yes. Send the DOI, your level and your review dates in the free consultation. We will tell you exactly which parts we can take on, such as reproducing the baseline, the extension, the report, an IEEE-format paper and viva preparation. Our IEEE projects service explains the process.

RELATED SERVICES

From a topic to a project you can present

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

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

Final-year projects
Code, a working demo, the black book, a PPT and viva prep. →