PROJECT IDEAS · AI AND ML

Machine learning projects for final *year*

The best AI and machine learning projects for final year pair **a public dataset, one honest metric and one change you can measure**.

Last updated 2026-10-18

Below are 49 ideas for 2026–27 across classic machine learning, deep learning, computer vision, NLP, generative AI and LLMs, explainable AI and MLOps. Each lists a level (B.Tech mini, B.Tech major or M.Tech), a named dataset, the metric to report and a research extension, and where we have built one, a link to the delivered version.

KEY POINTS

- **49 ideas for 2026–27** across classic machine learning, deep learning, computer vision, NLP, generative AI and LLMs, explainable AI and MLOps.
- **Each idea lists** a level (B.Tech mini, B.Tech major or M.Tech), a named dataset, the metric to report and a research extension.
- **Your level sets the scope.** A mini project trains and evaluates one model, a major project ships a working system, and an M.Tech project must improve on published work.
- **Three checks save most projects:** can you download the data today, is the split fair, and is there a simple baseline to beat?
- **Let the data decide the model.** On tabular data, gradient boosting is usually strong and easier to explain; for images, audio and text, fine-tune a pretrained deep learning model.

How do you choose an AI project for final year?

Choosing between AI projects for final year starts with your level. A mini project proves you can train and evaluate one model; a major project ships a working system; an M.Tech project must improve on published work. Every machine learning and deep learning project idea on this page is tagged with the level it fits best.

What an AI project needs at each level

LEVEL	SCOPE	COMPUTE	EXAMINERS LOOK FOR
B.Tech mini	One model on one public dataset, in a notebook or a simple page	Laptop or free Colab	A clean split, one metric, an honest confusion matrix
B.Tech major	Data to model to app, with two or three models compared	Free cloud GPU	A live demo, a comparison table, report and PPT
M.Tech / M.E.	A base paper reproduced, plus one measured extension	A GPU for hours, not weeks	Same-data baseline, ablation, a thesis and often a paper

Rules differ between universities. Your guide's requirements come first.

Three checks save most projects: can you download the data today, is the split fair (by patient, speaker or user when one person has many rows), and is there a simple baseline to beat? For CSE projects outside AI, see our final year projects for CSE.

In an AI&DS or AIML branch? Every idea here fits B.Tech programmes in Artificial Intelligence and Data Science or AI and Machine Learning. Because your course already covers the basics, start from the B.Tech major rows and add an explainable AI or MLOps angle, which shows engineering beyond a notebook.

Where do you find datasets for AI projects?

- UCI Machine Learning Repository: small, well-documented tabular and sensor sets.

- Kaggle Datasets and past competitions, which come with a fixed metric.
- Hugging Face Datasets: most NLP, speech and LLM benchmarks.
- PhysioNet: ECG, EEG and clinical signals; some sets need credentialed access.
- AI4Bharat: Indian-language text and speech resources.
- Open Government Data Platform India: public data for local problems.

Topic ideas, not delivered projects

The 49 ideas below are suggestions to discuss with your guide. Where a row says "see a delivered version", that links to a separate project we built; our delivered work is in the case studies.

Which machine learning project ideas work on tabular data?

Customer, hospital and sensor records are tables, and gradient-boosted trees are usually the model to beat on them. These seven ideas reward careful features and fair evaluation over big networks.

Topic ideas: classic machine learning

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Churn prediction with a cost-based threshold Full guide	B.Tech mini	Telco Customer Churn (IBM sample, on Kaggle)	PR-AUC, recall on churners	Set the threshold from retention cost, not 0.5; explain each customer with SHAP
Term-deposit prediction without leakage	B.Tech mini	Bank Marketing (UCI)	ROC-AUC, lift in the top 10%	Show how the call-duration column inflates scores; the dataset notes say to drop it for a realistic model
30-day hospital readmission	B.Tech major	Diabetes 130-US Hospitals (UCI)	ROC-AUC, Brier score	Split by patient, not by visit, and calibrate the probabilities
Customer segmentation	B.Tech mini	Online Retail II (UCI)	Silhouette score, stability across months	RFM features; k-means against Gaussian mixtures and HDBSCAN

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Air-quality forecasting for Indian cities	B.Tech major	Air Quality Data in India (Kaggle, compiled from CPCB)	MAE, RMSE per city	Gradient boosting against an LSTM; test on a city the model never saw
Machine failure prediction	B.Tech major	AI4I 2020 Predictive Maintenance (UCI)	Recall and F1 on the failure class	SMOTE against class weights, and name the failure mode. See a delivered M.E. / M.Tech version
Boosted trees versus deep tabular models	M.Tech	OpenML-CC18 benchmark suite	Mean rank across datasets, Friedman test	XGBoost, LightGBM and CatBoost against FT-Transformer or TabPFN

What are good deep learning project ideas beyond images?

Signals, audio, time series and graphs make strong deep learning projects for final year and M.Tech, because simple baselines exist and the data is public.

Topic ideas: deep learning for signals, time series and graphs

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
12-lead ECG classification	M.Tech	PTB-XL (PhysioNet)	Macro ROC-AUC on the recommended test fold	A single-lead model for wearables, and what accuracy it costs
Environmental sound classification	B.Tech major	ESC-50	Accuracy over its 5 official folds	Pretrained audio embeddings against a CNN on log-mel spectrograms
Speech emotion recognition	B.Tech major	RAVDESS and CREMA-D	Unweighted accuracy, speaker-independent split	Train on one corpus, test on the other
Human activity recognition	B.Tech mini	Human Activity Recognition Using Smartphones (UCI)	Macro-F1 on unseen volunteers	A 1D CNN on raw signals against the provided handcrafted features
Electricity load forecasting	M.Tech	ElectricityLoadDiagrams 2011-2014 (UCI)	MAE and MASE per horizon	A trained LSTM against a zero-shot pretrained forecaster such as Chronos

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Traffic speed forecasting on road graphs	M.Tech	METR-LA and PEMS-BAY	MAE, RMSE, MAPE at 15, 30 and 60 min	A graph neural network against an LSTM, with sensors removed at test time
Molecule property prediction	M.Tech	ogbg-molhiv (Open Graph Benchmark)	ROC-AUC on the scaffold split	GIN against GCN, then test whether a pretrained molecular encoder helps

Which computer vision projects suit final year?

Vision projects demo well, and pretrained CNNs or Vision Transformers train on a free GPU. The strongest ones test on data the model has not seen, not just a random split.

Topic ideas: computer vision and medical imaging

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Plant disease detection that works in the field	B.Tech major	PlantVillage to train, PlantDoc to test	Macro-F1	Measure the lab-to-field drop, then reduce it with augmentation or field photos
Road damage and pothole detection	B.Tech major	RDD2022 (includes Indian roads)	mAP@0.5	Run a YOLO model on a phone or Jetson and report frames per second
Segmentation of Indian road scenes	M.Tech	India Driving Dataset (IDD)	Mean IoU	Train on Cityscapes, test on IDD, then adapt with a few labels
Multi-label chest X-ray screening	M.Tech	NIH ChestX-ray14	Per-class AUROC, patient-wise split	Check whether Grad-CAM points at the lungs or at shortcuts such as text markers
Diabetic retinopathy grading	B.Tech major	APTOS 2019 Blindness Detection (Kaggle)	Quadratic weighted kappa	Ensemble pretrained models and add attention heatmaps. See a delivered M.Tech version
Parkinson's screening from drawings	B.Tech major	Parkinson's Drawings (Kaggle: spirals and waves)	ROC-AUC, sensitivity	One model per drawing type, with Grad-CAM. See a delivered M.Tech version
Handwritten Devanagari recognition	B.Tech mini	Devanagari Handwritten Character Dataset (UCI)	Accuracy, per-class confusion	Test on your own handwriting, then move from characters to words

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Skin lesion classification	M.Tech	HAM10000	Balanced accuracy, macro-F1	A Vision Transformer against CNNs, with class-imbalance handling. See a delivered M.Tech version

Picked an idea?

Send it with your level and deadline. The consultation is free.

Message before 7 pm for a reply today.

B.Tech project → WhatsApp

Which NLP project ideas use Indian-language data?

Indian languages and code-mixed text give an NLP project a local angle that generic sentiment projects lack. Speech counts here too.

Topic ideas: NLP and speech

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Hindi-English code-mixed sentiment	B.Tech major	SemEval-2020 Task 9 (SentiMix, Hinglish)	Macro-F1	MuRIL or IndicBERT against XLM-R, with and without transliteration
Fact-checking claims against evidence	M.Tech	FEVER	Label accuracy, FEVER score	Swap the retriever for dense embeddings and measure evidence recall
Hindi news summarisation	B.Tech major	XL-Sum (Hindi split)	ROUGE-1, ROUGE-2, ROUGE-L	Fine-tune a small mT5, then flag summary facts not in the article
English to Indian-language translation	M.Tech	FLORES-200 devtest (now maintained as FLORES+)	chrF++	Fine-tune on one domain, such as health notices, and measure the gain
Named-entity recognition for Indian languages	M.Tech	Naamapadam (AI4Bharat, 11 languages)	Entity-level F1	Train on Hindi, test zero-shot on Marathi or Bengali

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Chatbot intent detection	B.Tech mini	CLINC150	In-scope accuracy, out-of-scope recall	Reject questions the bot cannot handle instead of guessing
Toxic comment classification Full guide	B.Tech major	Jigsaw Toxic Comment Classification (Kaggle)	Mean column-wise ROC-AUC	Measure bias against identity words with the Jigsaw Unintended Bias data
Hindi speech recognition	M.Tech	Mozilla Common Voice (Hindi)	Word error rate (WER)	Fine-tune Whisper small and compare with the zero-shot model

What generative AI and LLM projects can students build in 2026?

RAG, agents and small-model fine-tuning are the live topics. A chat window is a demo, not a result, so each idea names a benchmark and a number to report. Most run on open models without a paid API.

Topic ideas: generative AI, RAG, agents and fine-tuning

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
RAG for multi-hop questions	M.Tech	HotpotQA	Exact match, F1, retrieval recall@k	BM25, dense and hybrid retrieval compared, then a re-ranker
Hallucination detector for RAG answers	M.Tech	RAGTruth	Span-level F1	A small NLI model against an LLM judge, on accuracy and cost
Natural-language to SQL agent	B.Tech major	Spider	Execution accuracy	Let the agent run its query, read the error and retry; count the fixes
Tool-calling agent on a small open model	M.Tech	Berkeley Function Calling Leaderboard (BFCL)	Call accuracy by category	Constrained JSON decoding against plain prompting
QLoRA fine-tuning for maths word problems	B.Tech major	GSM8K	Final-answer accuracy	4-bit QLoRA against LoRA and a prompted larger model, on accuracy and GPU memory
LLM-labelled data for a small classifier	M.Tech	AG News	Accuracy against human labels, cost per 1,000 labels	Active learning to choose which examples the LLM labels

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Consistent characters across generated scenes	M.Tech	DreamBooth dataset (30 subjects)	DINO and CLIP-I for the subject, CLIP-T for the prompt	LoRA against anchor-image conditioning for multi-scene stories. See a delivered M.Tech version
Interactive story generator with memory	B.Tech major	WritingPrompts	Human coherence ratings, consistency-check pass rate	Track characters and plot threads between turns. See a delivered M.Tech version
Visual question answering for blind users	M.Tech	VizWiz-VQA	VQA accuracy, answerability F1	Fine-tune a small vision-language model that says "unanswerable" when unsure

Which explainable AI projects stand out?

Explainable AI (XAI) turns a plain classifier into research: you test whether an explanation is faithful, stable or fair, not just draw a heatmap.

Topic ideas: explainable and responsible AI

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Do saliency maps show what the model uses?	M.Tech	CUB-200-2011 (with part locations)	Deletion and insertion AUC, pointing game	Grad-CAM, Integrated Gradients and attention rollout compared
SHAP versus LIME stability	B.Tech major	Adult / Census Income (UCI)	Rank correlation across runs	Add random features and check that both methods rank them last
Fairness audit of credit scoring	B.Tech major	Default of Credit Card Clients (UCI)	ROC-AUC, equal-opportunity difference	Mitigate with Fairlearn and report the accuracy cost
Counterfactual explanations for loans	M.Tech	Statlog German Credit (UCI)	Validity, proximity, sparsity	DiCE limited to features a person can actually change
Concept-based skin lesion explanations	M.Tech	Derm7pt (seven-point checklist)	Concept accuracy, diagnosis accuracy	A concept bottleneck model that a clinician can correct

What MLOps project ideas show real engineering?

MLOps projects suit students who like systems more than models: the result is a pipeline that keeps a model fast, current and safe to deploy.

Topic ideas: MLOps, edge and federated learning

IDEA	LEVEL	DATASET	METRIC	RESEARCH EXTENSION
Drift monitoring with automatic retraining	B.Tech major	NYC TLC Trip Record Data (monthly files)	PSI or KS drift score, MAE over time	Retrain on a drift trigger against a fixed schedule, tracked in MLflow
Faster model serving on a CPU	M.Tech	Imagenette (ImageNet subset)	p95 latency, throughput, accuracy drop	ONNX Runtime and INT8 quantisation, then in-browser inference. See a delivered M.E. version
CI/CD pipeline for a model	B.Tech major	Bike Sharing (UCI)	RMSE, share of injected data faults caught	Data-validation tests that block a bad model from deploying
Federated learning on skewed data	M.Tech	FEMNIST (LEAF benchmark)	Accuracy against communication rounds	FedAvg against FedProx as client data grows more skewed
Cost and latency monitoring for an LLM app	M.Tech	MS MARCO queries	p95 latency, tokens per answer, cache hit rate	Semantic caching, and how often a cached answer is wrong

Which machine learning projects have a full step-by-step guide?

The tables above give each idea one row. Nine projects that students search for by name have a guide of their own, with the problem, the output, the dataset, the build steps, the metrics, the mistakes to avoid and the viva questions.

- Credit card fraud detection: handling a dataset where only 0.172% of transactions are frauds.
- Customer churn prediction: from the Telco data to a ranked call list.
- Loan default prediction: calibrated risk scores, reasons and a fairness check.

- Demand forecasting: beating a naive baseline on the M5 retail data.
- Toxic comment classification: six labels on the Jigsaw comments, from TF-IDF to a transformer.
- Waste classification using deep learning: transfer learning on TrashNet, with a smart-bin option.
- Defect detection using deep learning: training on good parts only, with heatmaps, on MVTec AD.
- AI-generated image detection: real against fake images, tested on generators the model never saw.
- Unknown object and class detection: a classifier that can answer "unknown".

How do you turn an idea into a project your guide approves?

1. **Download the data first**
   Check that it downloads, its licence allows your use and it fits your compute.
2. **Fix the split and the metric**
   Use the official or a subject-wise split, and the metric in the table, before any tuning.
3. **Build the baseline**
   A simple model, or the base paper's method for M.Tech.
4. **Add one change at a time**
   The research extension is your contribution; an ablation shows what helped.
5. **Write it up and rehearse the viva**
   Report or thesis, PPT, and a paper if your guide wants one.

Starting from a published paper? See how to select a base paper, IEEE base paper implementation and M.Tech projects. Need a paper at the end? See research paper writing and publication. Planning a PhD after that? See our PhD research topics in AI and machine learning.

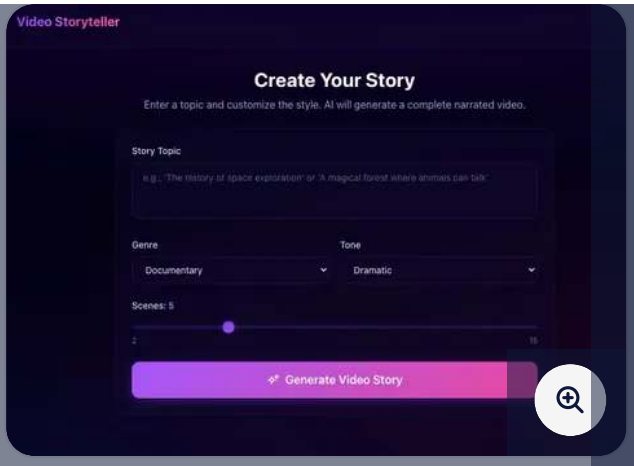
Follow your university's academic rules

Our work is building, writing support, guidance and explanation. Use it to learn the method, cite every dataset and paper, and check what your institution allows.

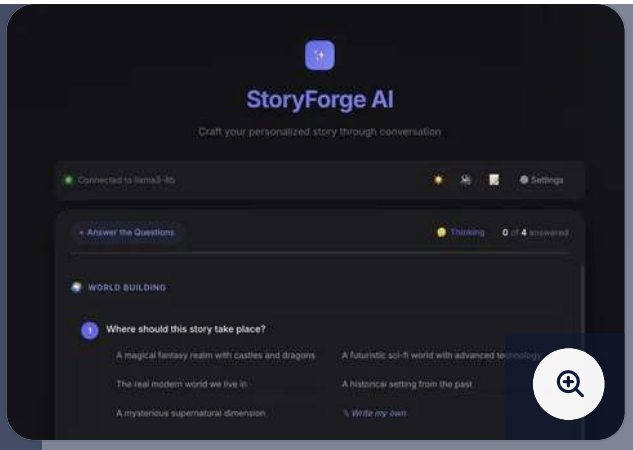
What does a delivered M.Tech AI project look like?

Here are four of the M.Tech AI projects we delivered, all linked from the tables above. Each came with working software, a thesis or report and an IEEE-format paper.

- Parkinson's disease detection using deep learning: two EfficientNetB0 models with Grad-CAM; ROC-AUC 0.930 on a 300–image test set.
- Diabetic retinopathy detection using deep learning: two public ViT-B/16 models ensembled without retraining, with attention rollout.
- Interactive storytelling AI using LLMs: guided questions, illustrations and narration, on local or hosted models.
- Text to video generator project (AI video storyteller): LLM script, diffusion scenes and text-to-speech combined into a narrated 1080p video.



Web app Video storyteller: topic, genre, tone and scene count in, narrated video out.



Web app Interactive storytelling: the reader shapes the story world through guided questions.

Shown with student, guide and institute details removed.

Looking beyond AI? See 53 final year project ideas for CSE across eight domains, IoT project ideas with the hardware named, or all our final year project ideas, sorted by branch.

Frequently asked questions

Which AI project is best for a final-year B.Tech student?

One that uses a public dataset you can download today, trains on a laptop or a free cloud GPU, and ends in a working demo. Image and text classifiers with an explanation screen are easy to demonstrate. RAG chatbots are popular in 2026, but they need a measured evaluation, not just a chat window.

Should I choose machine learning or deep learning for my project?

Let the data decide. On tabular data such as customer or sensor records, gradient boosting is usually strong and easier to explain. For images, audio and text, fine-tune a pretrained deep learning model. Many good projects compare one of each on the same test set.

Can I build an LLM project without a paid API?

Yes. Small open-weight models run locally through tools such as Ollama or on a free cloud GPU, and QLoRA fine-tunes a small model in 4-bit precision. Our delivered interactive storytelling project can use a local Ollama model as well as hosted APIs.

What makes an AI project good enough for M.Tech?

A contribution you can measure. Reproduce a recent base paper on the same data, add one change, and prove it with an ablation study and, ideally, a test on a second dataset. The M.Tech ideas above are scoped that way.

Where do I find datasets for machine learning projects?

Start with the UCI Machine Learning Repository for tabular and sensor data, Kaggle for datasets that come with a fixed metric, Hugging Face Datasets for text and speech, PhysioNet for medical signals, AI4Bharat for Indian languages and data.gov.in for local problems. Check each dataset's licence and download it before you commit to the idea.

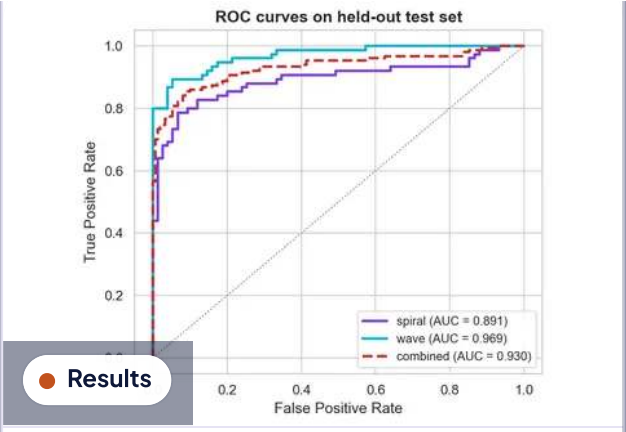
Can you build one of these ideas with me?

Yes. We build B.Tech and M.Tech AI projects with code, results, report and PPT, and walk you through the method for your viva. Share the idea, your level and your deadline for a free consultation, and follow your institution's rules on outside help.

DELIVERED WORK

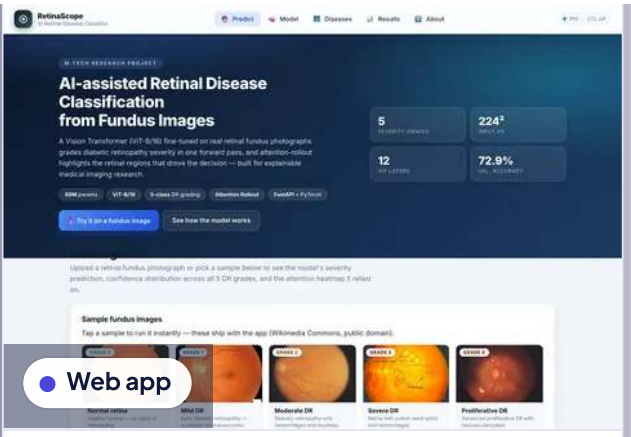
AI projects we delivered

M.Tech projects that started as ideas like the ones above. Shown with names and institute details removed.



M.Tech Medical AI · Explainable AI

Parkinson's disease detection from spiral drawings



M.Tech Medical imaging · Deep learning

Diabetic retinopathy detection with a ViT ensemble



M.Tech Generative AI · NLP

Interactive storytelling AI using LLMs

[See all 8 case studies →]

RELATED SERVICES

Help with your AI project

Final year projects (CSE and IT)

B.Tech builds with a demo, report and PPT.

M.Tech projects

Base paper, a measured extension, thesis and paper.

Dissertation, report & PPT

Chapters, black books and review slides.