PROJECT IDEAS · M.TECH · CSE SEMINAR TOPICS · 2026-27

M.Tech seminar topics for CSE, *with survey papers*

These 40 M.Tech seminar topics for CSE cover eight areas, from LLMs and computer vision to security, cloud, networks and data systems. Each topic comes with **one checked survey paper from 2023 to 2026** and an eight-slide seminar plan with a demo or table to show. Pick one close to your dissertation area.

Last updated 2026-11-17

A seminar is not a project. You read one recent survey paper closely, explain its area to your guide and panel, and answer their questions. We matched each survey paper below to its Crossref DOI record on 17 November 2026; titles, journals and years appear exactly as registered, and each card says only what the paper's abstract says. For topics you build rather than present, see our M.Tech project ideas and M.Tech machine learning projects with base papers.

KEY POINTS

- **40 M.Tech CSE seminar topics in eight groups,** each with one survey or review paper published from 2023 to 2026. Thirty-one appeared in IEEE Access and one in the IEEE Open Journal of the Computer Society; the rest are in ACM, Springer and MDPI journals.
- **Every card has a seminar plan:** eight slides in order, and one demo, figure or comparison table you can show.
- **Pick by three tests:** the topic is recent, a survey paper exists for it, and it leads into your dissertation area.
- **Present the survey's own structure.** Your contribution is a clear explanation, a fair comparison and one small demo, not a new result.
- **We quote no accuracy figures.** Take numbers from the papers themselves and cite them on your slides.

How do you pick an M.Tech seminar topic your panel will accept?

Many seminar topic lists online are copied from decks that are years out of date, and panels notice. A topic is easier to defend when it passes three tests.

- **It is recent.** The area should have moved in the last two or three years, so that you have something new to explain. Every topic on this page has a survey published from 2023 to 2026.
- **A survey paper exists.** One good survey gives you the structure of the talk, the main methods compared side by side and a reference list to follow. A topic with only blog posts and news articles behind it is hard to defend when the panel asks for sources.
- **It leads into your dissertation.** In many universities the seminar comes before or alongside Stage-I. If the topic sits in your dissertation area, the papers you read become your literature survey, and your guide sees that you already know the field.

Check two more things before you register the title. Ask your guide whether the department wants the seminar in the same area as the dissertation or prefers a different one. And check what your university asks for: the length of the talk, whether a written seminar report is required and its format all differ between universities, so follow your own department's current rules.

Where each seminar group leads for the dissertation

SEMINAR GROUP	DISSERTATION AREA IT FEEDS	BASE PAPERS TO BUILD ON
LLMs and NLP	NLP, LLM applications, Indian-language AI	M.Tech machine learning projects
Vision and multimodal	Image processing, detection, generative models	M.Tech image processing projects
Trustworthy AI	Explainable, fair or private machine learning	M.Tech machine learning projects
Security and privacy	Intrusion detection, cryptography, deepfakes	M.Tech cyber security projects
Cloud, edge and serverless	Scheduling, offloading, edge AI	IEEE base papers by domain
Networks, 5G/6G and IoT	SDN, network security, IoT	M.Tech IoT and embedded projects
Data systems	Retrieval, databases, data engineering	IEEE base papers by domain
SE and emerging areas	Software engineering, MLOps, quantum, blockchain	M.Tech project ideas

The middle column is our suggestion, not a rule. Your guide decides whether a seminar and a dissertation fit together.

Which seminar topics suit M.Tech CSE in 2026-27?

40 topics in eight groups. Each card says what the topic is and what to present. Open "The survey paper and a seminar plan" to see the paper, linked by its DOI, then an eight-slide outline and one thing to show the panel.

Topic ideas, not delivered projects

Each card names one survey or review paper, checked against its DOI record on Crossref, and describes only what the paper's abstract states. The seminar plan and the demo are our suggestions 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.

Large language models and NLP

The busiest area for 2026–27. Panels will ask what changed in the last two years, so build the talk on the survey's own taxonomy rather than on the history of the transformer.

TOPIC IDEA · M.TECH CSE · LLMS AND NLP

Retrieval-augmented generation (RAG) for large language models

Survey paper: ACM Computing Surveys, 2026

What it is: An LLM answers from what it learnt in training, so its answer can be out of date or simply made up. RAG searches an outside collection of documents at question time and gives the best passages to the model along with the question.

What to present: The RAG pipeline in the survey's four stages (pre-retrieval, retrieval, post-retrieval and generation), the main methods at each stage, how RAG systems are evaluated, and the open problems the survey lists.

− The survey paper and a seminar plan

Survey paper: A Survey on Retrieval-Augmented Text Generation for Large Language Models, Y. Huang and J. X. Huang, ACM Computing Surveys, 2026
Seminar plan: (1) Why LLMs give stale or made-up answers. (2) RAG in one diagram: question, retriever, document index, generator. (3–6) One slide for each of the survey's four stages, with two or three methods from the survey on each. (7) How RAG is evaluated, and the open problems the survey names. (8) Summary and one dissertation problem you could take up. Show: a live demo that answers questions on your syllabus PDFs with and without retrieval, using a small open model and a FAISS index.

TOPIC IDEA · M.TECH CSE · LLMS AND NLP

Hallucination in large language models: detection and mitigation

Survey paper: IEEE Access, 2025

What it is: LLMs sometimes write fluent text that is false or not supported by any source. In medicine, law or code, one confident wrong answer can do real harm, so finding and reducing hallucination is an active research area.

What to present: The review's four architecture types, from unaugmented generation to fully integrated detection and mitigation, how its 125 papers spread across nine application domains, and the three research priorities it names.

− The survey paper and a seminar plan

Survey paper: A Systematic Literature Review of Hallucinations in Large Language Models, C. Woesle et al., IEEE Access, 2025
Seminar plan: (1) What a hallucination is, with two examples you test yourself. (2) How the review was done: PRISMA, three databases, 125 papers. (3) The four architectures: unaugmented, post-hoc reactive validation, proactive detection-and-mitigation, fully integrated. (4) Which domains are well studied and which are not. (5) Detection methods. (6) Mitigation methods. (7) The review's three research priorities. (8) Summary. Show: a table of 10 factual questions put to one open model, each answer marked correct.

partly correct or hallucinated, before and after you add a simple verification step.

TOPIC IDEA · M.TECH CSE · LLMS AND NLP

LLM-based autonomous AI agents and agent protocols (MCP, A2A)

Survey paper: IEEE Access, 2026

What it is: An AI agent is an LLM that plans, calls tools and takes several steps to finish a task, instead of answering one prompt. Agents now connect to tools and to each other through shared protocols.

What to present: The review's taxonomy of about 60 benchmarks, the agent frameworks released from 2023 to 2025, applications in fields such as healthcare and software engineering, and the collaboration protocols ACP, MCP and A2A.

The survey paper and a seminar plan

Survey paper: From LLM Reasoning to Autonomous AI Agents: A Comprehensive Review, M. A. Ferrag et al., IEEE Access, 2026
Seminar plan: (1) From a chatbot to an agent: plan, act, observe. (2) How agents are tested: the review's taxonomy of about 60 benchmarks. (3) Agent frameworks from 2023 to 2025 and what each adds. (4) Three applications from the review. (5) How agents talk to tools and to each other: ACP, MCP and A2A. (6) Failure modes and security risks in agent protocols. (7) Research directions from the review. (8) Summary. Show: a short screen recording of an open-source agent finishing a two-step task with one tool, with each step it took labelled.

TOPIC IDEA · M.TECH CSE · LLMS AND NLP

Efficient LLMs: compression, fine-tuning and inference on limited hardware

Survey paper: IEEE Access, 2025

What it is: Large models need memory and compute that most college labs and small companies do not have. Efficiency methods make it possible to fine-tune and run an LLM on one modest GPU.

What to present: The review's four categories of methods (parameter-centric, architecture-centric, training-centric and data-centric) and its five lifecycle tasks: compression for local deployment, faster pre-training, efficient fine-tuning, optimised inference and data preparation.

The survey paper and a seminar plan

Survey paper: Make Large Language Models Efficient: A Review, A. Mussa et al., IEEE Access, 2025
Seminar plan: (1) Why model size is a problem for small labs. (2) The four categories of efficiency methods. (3) Compressing a model for local deployment. (4) Faster pre-training. (5) Efficient fine-tuning on your own data. (6) Inference under tight memory. (7) Data preparation, and open problems. (8) Summary. Show: one small open model loaded at 16-bit and at 4-bit on a free Colab GPU, with memory used, tokens per second and its answers to 10 fixed questions in one table.

TOPIC IDEA · M.TECH CSE · LLMS AND NLP

Multilingual prompt engineering for Indian and other low-resource languages

Survey paper: IEEE Access, 2026

What it is: LLMs work best in English. Prompting methods can improve results in other languages without retraining the model, which matters for Hindi, Marathi, Tamil and the other Indian languages.

What to present: The survey's taxonomy of 39 prompting techniques across 30 multilingual NLP tasks, drawn from 36 papers with datasets spanning about 250 languages, and its findings by language family and by high- or low-resource language.

− The survey paper and a seminar plan

Survey paper: Multilingual Prompt Engineering in Large Language Models: A Survey Across NLP Tasks, S. Vatsal et al., IEEE Access, 2026
Seminar plan: (1) Why LLMs do worse outside English. (2) Prompting basics: zero-shot, few-shot and chain-of-thought. (3) The survey's taxonomy of multilingual prompting. (4) The NLP tasks and datasets it covers. (5) High-resource versus low-resource languages. (6) Findings across language families. (7) Gaps that suit an Indian-language dissertation. (8) Summary. Show: one task, such as sentiment or question answering, in English and in one Indian language, with three prompting methods side by side in

a table.

Computer vision and multimodal AI

These topics present well because you can show images and short videos. Keep one slide for how each method is evaluated, since panels ask about metrics. For full projects in this area, see our M.Tech image processing ideas.

TOPIC IDEA · M.TECH CSE · VISION AND MULTIMODAL

Vision transformers: design choices and trade-offs

Survey paper: IEEE Access, 2026

What it is: Vision transformers split an image into patches and use attention in place of, or alongside, convolutions. In a few years they have grown into a large family of models for classification, detection and multimodal systems.

What to present: The survey's six design pressures (representation, interaction, structure, inductive bias, efficiency and optimisation or generalisation), the six taxonomy axes they map to, and its point that lower theoretical cost does not always mean lower real latency or memory.

− The survey paper and a seminar plan

Survey paper: A Survey on Vision Transformers: Design Pressures, Architectural Taxonomy, and Future Outlook, C. P. Lee et al., IEEE Access, 2026
Seminar plan: (1) From CNNs to the original ViT: patches, tokens and attention. (2) The six design pressures. (3) Tokenisation and new attention operators. (4) Hierarchical and CNN-transformer hybrid designs. (5) Efficiency, compression and token reduction. (6) Nominal versus realised efficiency. (7) Where ViTs sit in multimodal and diffusion models, and open questions. (8) Summary. Show: attention maps from a pretrained ViT on five of

TOPIC IDEA · M.TECH CSE · VISION AND MULTIMODAL

Conditional diffusion models for controllable generation

Survey paper: IEEE Access, 2025

What it is: Diffusion models create images and other data by removing noise step by step. Conditioning on labels, signals or physical constraints lets the user control what gets generated.

What to present: Recent conditional diffusion architectures, their applications in computer vision, NLP, mechanics and healthcare, the main challenges the survey identifies and the open research problems it lists.

− The survey paper and a seminar plan

Survey paper: Conditional Diffusion Models: A Survey of Techniques, Applications, and Challenges, T. Panagiotakopoulos et al., IEEE Access, 2025
Seminar plan: (1) Generative models in one slide: GANs, VAEs and diffusion. (2) Forward and reverse diffusion in plain terms. (3) What conditioning means, and the kinds of side information. (4) Architectures from the survey. (5) Applications in vision and healthcare. (6) Applications in NLP and mechanics. (7) Challenges and open problems. (8) Summary. Show: a grid of images from one open text-to-image model for the same prompt at four guidance scales, with your note on what changes.

your own photos, beside Grad-CAM maps from a ResNet on the same photos.

TOPIC IDEA · M.TECH CSE · VISION AND MULTIMODAL

Ten years of YOLO: real-time object detection from v1 to v13

Survey paper: IEEE Access, 2025

What it is: YOLO detects objects in one pass over the image, which makes it fast enough for live video. Over ten years it has grown into one of the most influential families of real-time detectors.

What to present: The main versions from YOLOv1 to YOLOv13, the architectural trends across them, the main application areas, evaluation practice, ethical issues and the future directions the review discusses.

− The survey paper and a seminar plan

Survey paper: A Decade of You Only Look Once (YOLO) for Object Detection: A Review, L. T. Ramos and A. D. Sappa, IEEE Access, 2025
Seminar plan: (1) Detection basics: boxes, classes, IoU and mAP. (2) Two-stage versus one-stage detectors. (3) The early versions. (4) The middle versions. (5) The recent versions up to v13. (6) Application areas and evaluation practice. (7) Ethics and future directions. (8) Summary. Show: one short traffic video run through two YOLO versions, with frames per second and detections side by side.

TOPIC IDEA · M.TECH CSE · VISION AND MULTIMODAL

Gaussian splatting for 3D reconstruction and novel view synthesis

Survey paper: IEEE Access, 2024

What it is: Gaussian splatting builds a 3D scene from ordinary photos and renders it from viewpoints the camera never took. It is a fast-growing method for 3D capture.

What to present: Learning-based 3D reconstruction and the recent work on Gaussian splatting: input types, model structures, output representations and training strategies, with the unresolved challenges and future directions the review gives.

− The survey paper and a seminar plan

Survey paper: Gaussian Splatting: 3D Reconstruction and Novel View Synthesis: A Review, A. Dalal et al., IEEE Access, 2024
Seminar plan: (1) The 3D reconstruction problem. (2) Earlier learning-based approaches, in one slide. (3) What a 3D Gaussian is and how splatting renders it. (4) Input types and model structures. (5) Output representations and training strategies. (6) Applications. (7) Challenges and future directions. (8) Summary. Show: an object you film with a phone, reconstructed with an open-source Gaussian splatting implementation and played as a short fly-around video.

TOPIC IDEA · M.TECH CSE · VISION AND MULTIMODAL

Vision-language-action models for robots

Survey paper: IEEE Access, 2025

What it is: VLA models join vision, language and robot action data in one model, so that a robot can follow an instruction such as "pick up the red cup", with the aim of handling new tasks with little or no extra training data.

What to present: The survey's full-stack view: how VLA architectures developed, their building blocks, processing for each modality and learning paradigms, plus robot platforms, data collection, public datasets, data augmentation and evaluation benchmarks.

− The survey paper and a seminar plan

Survey paper: Vision-Language-Action Models for Robotics: A Review Towards Real-World Applications, K. Kawaharazuka et al., IEEE Access, 2025
Seminar plan: (1) From LLMs and VLMs to VLAs. (2) The shift in strategy and architecture the survey traces. (3) Building blocks and processing for each modality. (4) Learning paradigms. (5) Robot platforms and data collection. (6) Public datasets and benchmarks. (7) What still blocks real-world use. (8) Summary. Show: a comparison table of four public VLA datasets or benchmarks taken from the survey's project website, with the robot, the kind of tasks and the licence of each.

Trustworthy AI: explainability, fairness, privacy and robustness

These topics connect to almost any applied dissertation, because any model you build can be questioned on bias, privacy or robustness. For base papers to build on, see our M.Tech machine learning projects.

TOPIC IDEA · M.TECH CSE · TRUSTWORTHY AI

Trustworthy and explainable AI (XAI)

Survey paper: IEEE Access, 2023

What it is: Many AI models are black boxes. Trustworthy AI asks whether a system is reliable, fair and secure, and explainable AI gives reasons for its decisions that people can check.

What to present: The components of trustworthy AI and the biases that make systems unreliable, why banking, healthcare, autonomous systems and IoT need it, transparent and post-hoc explanation models with their pitfalls, and policies for safe autonomous vehicles.

The survey paper and a seminar plan

Survey paper: A Review of Trustworthy and Explainable Artificial Intelligence (XAI), V. Chamola et al., IEEE Access, 2023
Seminar plan: (1) Why black-box models are a problem. (2) The components of trustworthy AI. (3) Sources of bias and unreliability. (4) Transparent models versus post-hoc explanations. (5) Two explanation methods in detail, such as LIME and SHAP. (6) Use in banking, healthcare and autonomous vehicles. (7) Pitfalls of XAI and open problems. (8) Summary. Show: SHAP and LIME explanations for the same three predictions of one tabular model, such as a loan-approval or diabetes model, side by side.

TOPIC IDEA · M.TECH CSE · TRUSTWORTHY AI

Sources of bias in machine learning pipelines

Survey paper: IEEE Access, 2026

What it is: Bias is a systematic error that can enter at any stage, from collecting data to reporting results. Adjusting a fairness metric at the end does not help if the cause sits earlier in the pipeline.

What to present: The survey's taxonomy of more than forty potential sources of bias in data and models along the machine learning pipeline, with an example for each, and how knowing the source helps detection and mitigation.

The survey paper and a seminar plan

Survey paper: A Survey on Bias in Machine Learning Research, A. Mikołajczyk-Bareła and M. Grochowski, IEEE Access, 2026
Seminar plan: (1) Bias as systematic error, not only unfairness. (2) The machine learning pipeline as a map. (3) Bias in data collection and labelling. (4) Bias in splitting and preprocessing. (5) Bias in models and training. (6) Bias in evaluation and reporting. (7) Detection and mitigation. (8) Summary. Show: one public dataset where you report accuracy separately for two groups, such as two age bands, before and after one mitigation step.

TOPIC IDEA · M.TECH CSE · TRUSTWORTHY AI

Differential privacy in federated learning

Survey paper: IEEE Access, 2025

What it is: Federated learning trains one model across hospitals, banks or phones without pooling their data. Differential privacy adds noise so that the shared updates reveal even less about any single record.

What to present: The two families the survey separates, central and local differential privacy federated learning, the trade-off between privacy and accuracy, the challenges of combining the two methods, and the three research ideas the survey proposes.

− The survey paper and a seminar plan

Survey paper: A Survey of Differential Privacy Techniques for Federated Learning, X. Wang et al., IEEE Access, 2025
Seminar plan: (1) Data silos and why data cannot simply be pooled. (2) Federated learning in one diagram. (3) Differential privacy and the privacy budget. (4) Central differential privacy in federated learning. (5) Local differential privacy, including the shuffle model. (6) Privacy versus accuracy and communication cost. (7) The survey's three future research ideas. (8) Summary. Show: a plot of test accuracy against the privacy budget for a small federated model on MNIST with 10 simulated clients, built with Flower or a similar open-

TOPIC IDEA · M.TECH CSE · TRUSTWORTHY AI

Adversarial attacks and defences in deep learning

Survey paper: IEEE Access, 2024

What it is: A small, carefully chosen change to an input can make a deep network give a wrong answer with high confidence. That is a risk for self-driving cars, malware detection and healthcare.

What to present: Recent attacks on object recognition grouped by the attacker's knowledge and capacity, defences grouped by strategy, how vision transformers compare with CNNs, and the common datasets and metrics. The survey also lists attacks and defences that have code available.

− The survey paper and a seminar plan

Survey paper: How Deep Learning Sees the World: A Survey on Adversarial Attacks & Defenses, J. C. Costa et al., IEEE Access, 2024
Seminar plan: (1) What an adversarial example is, with one picture. (2) Threat models: white-box and black-box. (3) The main attacks. (4) The main defences. (5) Vision transformers versus CNNs. (6) Datasets, metrics and available code. (7) Open issues. (8) Summary. Show: clean and attacked accuracy of one CIFAR-10 classifier under FGSM and PGD at three perturbation sizes, in a table with one example image for each.

source framework.

TOPIC IDEA · M.TECH CSE · TRUSTWORTHY AI

Machine unlearning: removing data from a trained model

Survey paper: IEEE Open Journal of the Computer Society, 2025

What it is: Machine unlearning removes the influence of chosen training data from a model, for privacy, copyright or data-quality reasons, without training again from scratch. Unlearning systems also open new ways to attack a model.

What to present: Why unlearning matters for AI governance and compliance with privacy regulations, unlearning offered as a service alongside machine learning as a service, and the survey's taxonomy of threats, attacks and defences in unlearning systems.

− The survey paper and a seminar plan

Survey paper: Threats, Attacks, and Defenses in Machine Unlearning: A Survey, Z. Liu et al., IEEE Open Journal of the Computer Society, 2025
Seminar plan: (1) Why a model may need to forget. (2) What unlearning has to achieve. (3) Exact versus approximate unlearning. (4) Unlearning as a service. (5) Attacks on unlearning systems. (6) Defences. (7) Open problems from the survey. (8) Summary. Show: a small experiment that makes a trained MNIST model forget one digit by fine-tuning, with accuracy on that digit and on the rest, compared with full retraining.

Security and privacy

Each of these has a clear threat, a defence and a way to measure it, which makes the seminar easy to structure. For full projects in this area, see our M.Tech cyber security ideas.

TOPIC IDEA · M.TECH CSE · SECURITY

Zero trust architecture and trust scoring

Survey paper: IEEE Access, 2026

What it is: Zero trust drops the idea of a trusted internal network and checks every access request. Trust quantification gives users, devices, sessions and behaviour a score that drives each access decision.

What to present: The survey's three-stage framework for trust quantification (data collection, data pre-processing and the trust model), its evaluation criteria taken from the core tenets of zero trust, and the limits of current trust-scoring research.

− The survey paper and a seminar plan

Survey paper: Trust Quantification on Zero Trust Architecture: A Survey, S.-W. Nam et al., IEEE Access, 2026
Seminar plan: (1) Why perimeter security fails. (2) Zero trust principles: never trust, always verify. (3) The parts of a zero trust architecture. (4) Trust quantification and where it fits. (5) The three stages of the survey's framework. (6) How existing methods compare on the zero trust criteria. (7) Limits and research directions. (8) Summary. Show: a flow diagram of one login request passing through a zero trust system, with the signals that raise or lower its trust score.

TOPIC IDEA · M.TECH CSE · SECURITY

NIST post-quantum cryptography standards

Survey paper: IEEE Access, 2026

What it is: A large quantum computer could break RSA and elliptic-curve cryptography. NIST has standardised new algorithms meant to resist quantum attacks, and systems will have to migrate to them.

What to present: The three standardised algorithms (module-lattice key encapsulation, module-lattice signatures and stateless hash-based signatures), the FALCON and HQC schemes, and the survey's comparison of software and FPGA or ASIC implementations, from IoT devices to servers.

− The survey paper and a seminar plan

Survey paper: NIST Post-Quantum Cryptography Standards: A Comprehensive Review of Theoretical Foundations and Implementations, Q. D. Truong et al., IEEE Access, 2026
Seminar plan: (1) Why quantum computers threaten RSA and ECC. (2) Lattice problems in plain terms. (3) The module-lattice key encapsulation standard. (4) The two signature standards. (5) FALCON and HQC. (6) Software and hardware implementations compared. (7) Migration challenges. (8) Summary. Show: key sizes and timings for RSA-2048, ECDH and ML-KEM measured on your laptop with OpenSSL and the open-source liboqs library, in

one table.

TOPIC IDEA · M.TECH CSE · SECURITY

Deepfake detection: architectures, methods and datasets

Survey paper: IEEE Access, 2024

What it is: New generators such as DALL-E 3 and Stable Diffusion make fake images and video harder to spot. Detectors that score well in the lab often fail on new data from the wild.

What to present: The three parts of a deepfake detector the review studies (architecture, detection method and dataset), and its main finding that a lack of diversity in datasets limits how well detectors generalise to unseen data.

− The survey paper and a seminar plan

Survey paper: A Review of Deepfake Techniques: Architecture, Detection, and Datasets, P. Edwards et al., IEEE Access, 2024
Seminar plan: (1) What deepfakes are and how they are made. (2) Why detection gets harder every year. (3) Detector architectures. (4) Detection methods. (5) Datasets and their gaps. (6) Generalising to unseen data. (7) Ethics, privacy and future directions. (8) Summary. Show: one open-source pretrained detector tested on images from its own benchmark and on images you generate yourself, with the accuracy on each set in a table.

TOPIC IDEA · M.TECH CSE · SECURITY

Deep learning for network intrusion detection

Survey paper: IEEE Access, 2025

What it is: A network intrusion detection system watches traffic for attacks. Deep learning can learn attack patterns from packets or flows without hand-built features.

What to present: The survey's four-way taxonomy (reconstruction and generative models, transformer sequence models, CNN and DNN classifiers, and hybrid deployment-oriented systems), the datasets and features of 31 studies from 2020 to 2025, and the open challenges.

− The survey paper and a seminar plan

Survey paper: A Survey on the Applications of Deep Learning in Network Intrusion Detection Systems to Enhance Network Security, F. M. Anis et al., IEEE Access, 2025
Seminar plan: (1) Attacks and the types of IDS. (2) Network traffic as data: packets and flows. (3) Generative and reconstruction models for anomaly detection. (4) Transformer sequence models. (5) CNN and DNN classifiers. (6) Hybrid and deployable systems. (7) Datasets and open challenges. (8) Summary. Show: an autoencoder and a random forest trained on one public flow dataset such as CIC-IDS2017, with precision and recall for each attack class in a table.

TOPIC IDEA · M.TECH CSE · SECURITY

Fully homomorphic encryption for privacy-preserving machine learning

Survey paper: IEEE Access, 2026

What it is: Fully homomorphic encryption lets a server compute on encrypted data without decrypting it. It could let hospitals or banks use cloud AI without exposing their records, but it is still slow.

What to present: The review's four areas: hardware acceleration on GPUs, FPGAs and ASICs, algorithmic speed-ups such as SIMD batching, privacy-preserving machine learning with non-linear functions, and federated learning with secure aggregation, plus why encrypted inference is easier than encrypted training.

− The survey paper and a seminar plan

Survey paper: Fully Homomorphic Encryption and Artificial Intelligence: A Review of Architectures, Optimizations, and Distributed Privacy, N. Shukla et al., IEEE Access, 2026
Seminar plan: (1) Computing on encrypted data: the idea. (2) FHE basics: ciphertexts, noise and bootstrapping. (3) Why FHE is slow. (4) Hardware acceleration. (5) SIMD batching and other algorithmic speed-ups. (6) Encrypted inference versus encrypted training. (7) FHE with federated learning, and open problems. (8) Summary. Show: the time for one logistic-regression prediction in plain form and under

CKKS encryption, measured with an open-source library such as TenSEAL.

Cloud, edge and serverless computing

Most of these can be demonstrated on a laptop with Docker, a free cloud tier or a Raspberry Pi. A measured latency or cost makes a stronger slide than an architecture diagram alone.

TOPIC IDEA · M.TECH CSE · CLOUD AND EDGE

Serverless computing for HPC, AI and big data

Survey paper: IEEE Access, 2025

What it is: Serverless platforms run code as short functions and scale them automatically, so the user does not manage servers. Researchers are now testing whether this model suits heavy parallel jobs, not only web back-ends.

What to present: The review of 122 articles from 2018 to early 2025, its taxonomy of eight research directions and nine use-case domains, and publication trends in serverless computing for compute-intensive work across cloud, HPC and hybrid environments.

− The survey paper and a seminar plan

Survey paper: High-Performance Serverless Computing: A Systematic Literature Review on Serverless for HPC, AI, and Big Data, V. Besozzi et al., IEEE Access, 2025
Seminar plan: (1) From servers to virtual machines, containers and functions. (2) How a serverless platform works: triggers, cold starts and scaling. (3) Why HPC and AI researchers are interested. (4) How the review was done. (5) The eight research directions. (6) The nine use-case domains. (7) Open challenges. (8) Summary. Show: the same image-resizing job run as a serverless function on a free tier and on one virtual machine, with cold-start time, run time and cost per 1,000 calls in

TOPIC IDEA · M.TECH CSE · CLOUD AND EDGE

The edge-cloud continuum: architectures and platforms

Survey paper: IEEE Access, 2026

What it is: Some IoT and real-time applications cannot wait for a distant data centre. The edge-cloud continuum spreads work across devices, edge nodes and the cloud and manages them as one system.

What to present: The survey's developer-oriented framework: architectural models, performance metrics, paradigms for computation, communication and deployment, enabling technologies, edge-to-cloud platforms, testing tools, and applications in smart cities, healthcare and Industry 4.0.

− The survey paper and a seminar plan

Survey paper: Navigating the Edge-Cloud Continuum: A State-of-Practice Survey, L. Belcastro et al., IEEE Access, 2026
Seminar plan: (1) Why the cloud alone is not enough. (2) Device, edge and cloud layers. (3) Architectural models. (4) Performance metrics. (5) Deployment paradigms and platforms. (6) Tools for testing and experiments. (7) Applications and open challenges. (8) Summary. Show: round-trip latency for one small inference request served from a Raspberry Pi on your LAN and from a cloud region, measured over 100 requests and shown as a box plot.

a table.

TOPIC IDEA · M.TECH CSE · CLOUD AND EDGE

Kubernetes scheduling algorithms

Survey paper: Journal of Cloud Computing, 2023

What it is: Kubernetes decides on which machine each container runs. Better scheduling can cut cost and delay in a data centre, and it is an active research area.

What to present: The survey's four groups of Kubernetes schedulers (generic, multi-objective optimisation, AI-focused and autoscaling-enabled) and the gaps it identifies for further research.

− The survey paper and a seminar plan

Survey paper: A survey of Kubernetes scheduling algorithms, K. Senjab et al., Journal of Cloud Computing, 2023
Seminar plan: (1) Containers and why they need orchestration. (2) Kubernetes architecture: control plane, nodes and pods. (3) How the default scheduler picks a node. (4) Generic scheduling. (5) Multi-objective optimisation. (6) AI-focused and autoscaling-enabled scheduling. (7) Gaps from the survey. (8) Summary. Show: a three-node local cluster in kind or minikube where you deploy the same workload with default settings and with node affinity rules, showing where the pods land.

TOPIC IDEA · M.TECH CSE · CLOUD AND EDGE

Computation offloading in mobile edge computing

Survey paper: IEEE Access, 2025

What it is: Phones, IoT nodes and drones have little battery and processing power. Offloading sends heavy tasks to edge servers or the cloud, trading energy and time against network delay.

What to present: Architectural models, full versus partial offloading, decision strategies using supervised, unsupervised, deep and reinforcement learning, the role of 5G and edge intelligence, and the simulation models used to evaluate offloading schemes.

− The survey paper and a seminar plan

Survey paper: A Survey on Computation Offloading and Current Trends, A. Kumar and A. Pal, IEEE Access, 2025
Seminar plan: (1) Why offload at all. (2) Where to: edge or cloud. (3) Full versus partial offloading. (4) The decision problem: delay, energy and resources. (5) Machine learning and reinforcement learning for offloading decisions. (6) 5G, simulation models and applications. (7) Open challenges. (8) Summary. Show: a small Python or iFogSim simulation comparing local execution, full offloading and a simple threshold policy on delay and energy.

TOPIC IDEA · M.TECH CSE · CLOUD AND EDGE

TinyML: machine learning on microcontrollers

Survey paper: IEEE Access, 2024

What it is: TinyML runs trained models on small, battery-powered microcontrollers with very little memory, so a sensor can make decisions without sending its data to the cloud.

What to present: The three workflows for building a TinyML system (ML-oriented, hardware-oriented and co-design), the survey's taxonomy of model optimisation and learning techniques, and current hardware devices and software tools.

− The survey paper and a seminar plan

Survey paper: A Machine Learning-Oriented Survey on Tiny Machine Learning, L. Capogrosso et al., IEEE Access, 2024
Seminar plan: (1) Why run ML on a microcontroller. (2) Hardware limits: memory, power and compute. (3) The three workflows. (4) Model optimisation families from the survey. (5) Learning techniques under the TinyML lens. (6) Boards and software tools. (7) Challenges and future directions. (8) Summary. Show: a keyword-spotting or gesture model on an ESP32 or Arduino Nano 33 BLE Sense, with model size, RAM use and inference time before and after 8-bit quantisation.

Networks, 5G/6G and IoT

A networking seminar stands out when it has one simulation or measurement. Mininet, ns-3 and public traffic traces are free. For hardware projects, see our M.Tech IoT and embedded ideas.

TOPIC IDEA · M.TECH CSE · NETWORKS

AI, machine learning and LLMs in 5G and 6G networks

Survey paper: IEEE Access, 2025

What it is: 6G is expected to carry real-time telemedicine, extended reality and autonomous systems. The survey looks at how AI, ML and LLMs can manage, secure and repair these networks automatically.

What to present: The survey's roadmap for bringing AI, ML and LLMs into 6G for network management, resource allocation and threat mitigation, its frameworks for AI-enhanced security and self-healing networks, federated learning for privacy, and the ethics issues it raises.

− The survey paper and a seminar plan

Survey paper: AI, ML, and LLM Integration in 5G/6G Networks: A Comprehensive Survey of Architectures, Challenges, and Future Directions, Y. Usman et al., IEEE Access, 2025
Seminar plan: (1) From 5G to 6G: what changes. (2) The applications 6G must support. (3) AI for network management and resource allocation. (4) AI for security and threat mitigation. (5) Self-healing networks and LLMs in telecom. (6) Federated learning and ethics. (7) Research directions. (8) Summary. Show: a table of where AI acts in a 6G network (management, resource allocation, security, self-healing), with one method from the survey in

TOPIC IDEA · M.TECH CSE · NETWORKS

Machine learning for network slicing

Survey paper: IEEE Access, 2023

What it is: Network slicing splits one physical 5G network into virtual networks with different guarantees, for example one for video and one for factory robots. Running slices in practice is hard, and ML is being used to manage their resources automatically.

What to present: The survey's categories and sub-categories of ML contributions to network slicing, the lessons it draws and the open research challenges, with possible solutions.

− The survey paper and a seminar plan

Survey paper: Machine Learning in Network Slicing—A Survey, H. P. Phyu et al., IEEE Access, 2023
Seminar plan: (1) One network, many services. (2) What a slice is: eMBB, URLLC and mMTC. (3) The slice lifecycle. (4–5) The survey's main categories of ML use, with one example each. (6) Lessons learned. (7) Open challenges and possible solutions. (8) Summary. Show: a toy simulation in which a reinforcement learning agent shares bandwidth between two slices, plotted over time against a fixed split.

eacn row.

TOPIC IDEA · M.TECH CSE · NETWORKS

DDoS attack detection in software-defined networks

Survey paper: IEEE Access, 2024

What it is: SDN moves network control into a central controller. That controller becomes a single point of failure and the main target of DDoS attacks.

What to present: The SDN architecture and the DDoS attacks on it, common datasets and evaluation measures, data preprocessing and dimensionality reduction, and the strengths and weaknesses of statistical, machine learning and deep learning detectors.

− The survey paper and a seminar plan

Survey paper: Overview of DDoS Attack Detection in Software-Defined Networks, H. Wang and Y. Li, IEEE Access, 2024
Seminar plan: (1) SDN: control plane and data plane. (2) Why the controller is a target. (3) DDoS attacks on SDN. (4) Datasets and evaluation measures. (5) Preprocessing and dimensionality reduction. (6) Statistical, ML and deep learning detectors compared. (7) Challenges and future work. (8) Summary. Show: a Mininet topology with a Ryu controller in which one host launches a SYN flood and an entropy-based detector flags it, with the detection time on screen.

TOPIC IDEA · M.TECH CSE · NETWORKS

Security, privacy and trust in IoT systems

Survey paper: IEEE Access, 2024

What it is: IoT devices collect personal data and control physical things, often with little computing power to spare for security. The survey names security, privacy and trust as critical barriers to wider IoT adoption.

What to present: IoT security, privacy and trust organised by a three-layer IoT architecture, the security requirements and challenges in that architecture, recent research trends and methods, and the unresolved challenges with possible solutions.

− The survey paper and a seminar plan

Survey paper: A Survey on Security, Privacy, Trust, and Architectural Challenges in IoT Systems, M. Adam et al., IEEE Access, 2024
Seminar plan: (1) The IoT and why it is hard to secure. (2) The three-layer architecture. (3–5) One slide per layer: its threats and its defences. (6) Privacy and trust methods. (7) Open challenges. (8) Summary. Show: a Wireshark capture of an unencrypted MQTT message next to the same message sent over TLS.

TOPIC IDEA · M.TECH CSE · NETWORKS

Network traffic forecasting: from ARIMA to LSTM

Survey paper: IEEE Access, 2023

What it is: Operators forecast traffic to plan capacity and to spot anomalies. Both classical time-series models and neural networks are used, and the choice affects computing cost as well as accuracy.

What to present: The mathematics of ARMA, ARIMA and SARIMA and of RNN, LSTM, GRU and CNN forecasters, and the paper's own experiments on real data comparing fit and computing cost. The authors have released their code.

− The survey paper and a seminar plan

Survey paper: Forecasting Network Traffic: A Survey and Tutorial With Open-Source Comparative Evaluation, G. O. Ferreira et al., IEEE Access, 2023
Seminar plan: (1) Why forecast traffic. (2) Traffic as a time series: trend and seasonality. (3) ARMA, ARIMA and SARIMA. (4) RNN, LSTM and GRU. (5) CNN forecasters. (6) The paper's comparison of fit and computing cost. (7) Open problems. (8) Summary. Show: the authors' released code run on one traffic trace, with a plot of actual against forecast traffic for SARIMA and LSTM, and their error and training time in a table.

Data systems: vector databases, data lakes and streams

These suit students headed for data engineering, and each can be demonstrated with free, open-source software.

TOPIC IDEA · M.TECH CSE · DATA SYSTEMS

Vector database management systems

Survey paper: The VLDB Journal, 2024

What it is: LLM applications store text, images and code as embedding vectors and search them by similarity. The survey counts more than 20 commercial vector databases, all built in the five years before it was written.

What to present: The survey's five obstacles to vector data management, techniques for query processing, storage and indexing (quantisation and partitioning) and query optimisation for hybrid queries, native and extended systems, and benchmarks.

The survey paper and a seminar plan

Survey paper: Survey of vector database management systems, J. J. Pan et al., The VLDB Journal, 2024
Seminar plan: (1) Embeddings and similarity search. (2) The five obstacles from the survey. (3) Similarity scores and query types. (4) Storage and indexing: quantisation and partitioning. (5) Hybrid queries on vectors and attributes. (6) Native versus extended systems, and benchmarks. (7) Research challenges. (8) Summary. Show: the same 100,000 embeddings indexed with flat, IVF and HNSW indexes in FAISS, with recall and query time for each in a table.

TOPIC IDEA · M.TECH CSE · DATA SYSTEMS

Data lake architectures

Survey paper: Computers, 2024

What it is: A data lake keeps raw data of every type until it is needed, while a data warehouse holds cleaned, structured data for reporting. How a lake is organised decides whether anyone can find and trust its data later.

What to present: The major types of data-lake architecture the review identifies and how they evolved, its comparison diagrams, the differences between data warehouses and data lakes, the main implementation challenges and future directions.

The survey paper and a seminar plan

Survey paper: Data Lakes: A Survey of Concepts and Architectures, S. Azzabi et al., Computers, 2024
Seminar plan: (1) Data warehouse versus data lake. (2) Why organisations adopted lakes. (3–4) The architecture types from the review, redrawn as your own diagrams. (5) Metadata and governance. (6) Implementation challenges. (7) Future directions. (8) Summary. Show: a small lake on your laptop with raw, cleaned and curated folders of Parquet files, queried with DuckDB, beside a diagram of the folders.

TOPIC IDEA · M.TECH CSE · DATA SYSTEMS

LLMs for text-to-SQL

Survey paper: ACM Computing Surveys, 2025

What it is: Text-to-SQL turns a question in plain English into a database query, so that people who do not know SQL can still query data. LLMs have changed how well this works.

What to present: Classic text-to-SQL benchmarks and metrics, the survey's taxonomy of the two main approaches (prompt engineering and fine-tuning), its analysis of models on well-known datasets, and the challenges and future directions.

− The survey paper and a seminar plan

Survey paper: A Survey on Employing Large Language Models for Text-to-SQL Tasks, L. Shi et al., ACM Computing Surveys, 2025
Seminar plan: (1) The text-to-SQL task, with one example. (2) Benchmarks and metrics. (3) Prompt engineering methods. (4) Fine-tuning methods. (5) How models compare on standard datasets. (6) Challenges. (7) Future directions. (8) Summary. Show: 15 questions on a sample college database (students, courses, marks) answered by one open LLM with a zero-shot and a few-shot prompt, scored by whether each query returns the right rows.

TOPIC IDEA · M.TECH CSE · DATA SYSTEMS

Graph retrieval-augmented generation (GraphRAG)

Survey paper: ACM Transactions on Information Systems, 2025

What it is: Plain RAG retrieves passages of text. GraphRAG retrieves from a graph of entities and their relations, which helps with questions that need several connected facts.

What to present: The GraphRAG workflow as the survey formalises it (graph-based indexing, graph-guided retrieval and graph-enhanced generation), the methods and training at each stage, tasks, application domains, evaluation and industrial use.

− The survey paper and a seminar plan

Survey paper: Graph Retrieval-Augmented Generation: A Survey, B. Peng et al., ACM Transactions on Information Systems, 2025
Seminar plan: (1) RAG recap and its limits. (2) Knowledge graphs: entities, relations and triples. (3) Graph-based indexing. (4) Graph-guided retrieval. (5) Graph-enhanced generation. (6) Tasks, domains and evaluation. (7) Future directions. (8) Summary. Show: the same five multi-hop questions answered by plain vector RAG and by retrieval over a small knowledge graph you build in Neo4j or NetworkX, marking which answers each got right.

TOPIC IDEA · M.TECH CSE · DATA SYSTEMS

Networked data streaming with Apache Kafka

Survey paper: IEEE Access, 2023

What it is: Apache Kafka carries streams of events, such as clicks, sensor readings and payments, between systems in real time. It has become a popular choice for streaming data, and researchers work on its algorithms, networking and security.

What to present: The survey's five macro areas of Kafka research (algorithms, networks, data, cyber-physical systems and security), the optimisation aspects in each, cross-domain applications and the open research challenges.

− The survey paper and a seminar plan

Survey paper: A Survey on Networked Data Streaming With Apache Kafka, T. P. Raptis and A. Passarella, IEEE Access, 2023
Seminar plan: (1) Batch versus stream processing. (2) Kafka basics: topics, partitions, brokers and consumer groups. (3) Research on algorithms. (4) Networks and data. (5) Cyber-physical systems. (6) Security. (7) Research challenges. (8) Summary. Show: a single-node Kafka on Docker with a producer sending simulated sensor readings, and a chart of throughput and end-to-end latency as you change the number of partitions.

Software engineering and emerging areas

These suit students whose dissertation is in software engineering, or who want a topic few classmates will pick.

TOPIC IDEA · M.TECH CSE · SE AND EMERGING

Large language models for software engineering

Survey paper: ACM Transactions on Software Engineering and Methodology, 2024

What it is: LLMs now write code, tests and documentation and help fix bugs. Researchers are measuring where they help in software engineering and where they fall short.

What to present: The review of 395 articles from January 2017 to January 2024, organised around four research questions: which LLMs are used, how software engineering data is collected and prepared, how performance is optimised and evaluated, and which tasks LLMs have helped with.

− The survey paper and a seminar plan

Survey paper: Large Language Models for Software Engineering: A Systematic Literature Review, X. Hou et al., ACM Transactions on Software Engineering and Methodology, 2024
Seminar plan: (1) LLMs meet software engineering. (2) How the review was done. (3) Which LLMs are used. (4) Datasets and preprocessing. (5) Optimisation and evaluation. (6) Software engineering tasks where LLMs help. (7) Gaps and future work. (8) Summary. Show: ten programming tasks from a public benchmark such as HumanEval given to one open code model, with how many pass their tests on the first try.

TOPIC IDEA · M.TECH CSE · SE AND EMERGING

MLOps: tools, processes and a reference architecture

Survey paper: IEEE Access, 2025

What it is: Training a model in a notebook is only the start: it then has to be deployed, monitored and retrained. MLOps brings DevOps practice to machine learning: versioning, pipelines, deployment and monitoring.

What to present: The classes of tools in MLOps environments and their most popular implementations, the components unique to MLOps (model orchestrators, feature stores and model repositories), the review's reference architecture and production stages, and its finding that no satisfactory metrics for MLOps effectiveness were found.

− The survey paper and a seminar plan

Survey paper: MLOps Components, Tools, Process, and Metrics: A Systematic Literature Review, A. P. Woźniak et al., IEEE Access, 2025
Seminar plan: (1) Why ML projects stall after the notebook. (2) DevOps versus MLOps. (3) Tool classes from the review. (4) Components unique to MLOps. (5) The reference architecture. (6) The stages of model production. (7) The missing metrics, and open problems. (8) Summary. Show: one model tracked with MLflow over three training runs, registered and served as a REST endpoint, with a screenshot of each step.

TOPIC IDEA · M.TECH CSE · SE AND EMERGING

Quantum machine learning for computer scientists

Survey paper: Entropy, 2023

What it is: Quantum machine learning asks whether quantum computers can speed up or improve learning. This review is written for readers without a physics background, which suits CSE students.

What to present: Quantum machine learning explained from the view of conventional techniques, its basic algorithms, and the paper's case studies: quanvolutional networks against CNNs on handwritten digits, QSVM against SVM on breast cancer data, and a variational quantum classifier on Iris.

− The survey paper and a seminar plan

Survey paper: Quantum Machine Learning: A Review and Case Studies, A. Zeguendry et al., Entropy, 2023
Seminar plan: (1) Why ML researchers look at quantum computing. (2) Qubits, superposition and gates without heavy physics. (3) Basic QML building blocks. (4) Quanvolutional neural networks. (5) Quantum SVM. (6) Variational quantum classifiers. (7) Limits of today's hardware. (8) Summary. Show: a variational quantum classifier on the Iris dataset in Qiskit or PennyLane on a simulator, beside a classical SVM, with accuracy and training time in a table.

TOPIC IDEA · M.TECH CSE · SE AND EMERGING

Blockchain for supply chain management

Survey paper: IEEE Access, 2024

What it is: A supply chain links many companies that do not fully trust each other, and a central system can be attacked or fail. A shared, tamper-evident ledger can track goods from source to customer.

What to present: Blockchain basics for supply chain management, industry applications, persistent challenges and research opportunities from the past six years, and the survey's list of available blockchain solutions for global supply chains.

− The survey paper and a seminar plan

Survey paper: Blockchain Technology for Global Supply Chain Management: A Survey of Applications, Challenges, Opportunities and Implications, P. Dudczyk et al., IEEE Access, 2024
Seminar plan: (1) How a supply chain works and where it breaks. (2) Blockchain basics: ledger, consensus and smart contracts. (3) Public versus permissioned chains. (4) Industry applications. (5) Available platforms from the survey's list. (6) Persistent challenges. (7) Research opportunities. (8) Summary. Show: a small Solidity smart contract on a local test chain such as Hardhat that records each hand-over of a batch of medicines, with the transaction history displayed.

TOPIC IDEA · M.TECH CSE · SE AND EMERGING

AI-driven migration from monoliths to microservices

Survey paper: IEEE Access, 2026

What it is: Many companies want to split a large monolithic application into microservices for the cloud. Deciding where to cut is hard, and AI methods are being used to suggest the split.

What to present: The review of 47 studies from 2020 to 2025, classified by analysis type, automation level and decomposition strategy; AI methods such as clustering, graph neural networks and reinforcement learning; reverse migration back to monoliths; and the proposed four-stage Adaptive Migration Framework.

The survey paper and a seminar plan

Survey paper: **A Systematic Literature Review on AI-Driven Migration From Monolithic Architectures to Microservices**, M. Labaalli et al., IEEE Access, 2026
Seminar plan: (1) Monolith versus microservices. (2) Why migration is hard. (3) The review's three classification dimensions. (4) AI methods: clustering, graph neural networks, reinforcement learning. (5) Reverse migration. (6) The Adaptive Migration Framework. (7) Open questions. (8) Summary. Show: the class dependency graph of an open-source monolith such as Spring PetClinic, clustered into candidate services and drawn as before-and-

after diagrams.

Choosing a seminar topic?

Send your specialisation, your dissertation area and your seminar date. The consultation is free.

Message before 7 pm for a reply today.

Free consultation →

WhatsApp

How do you prepare the seminar report and PPT?

Work from the survey outwards. The survey gives you the structure; your job is to explain it clearly, check it against a few of the papers it cites and add one thing of your own to show.

1

Read the survey twice.

The first time for the big picture, the second time with notes on its taxonomy,

its comparison tables and its list of open problems.

2 **Pick six to ten papers it cites** and read their abstracts and results. These become the rows of your comparison table and the reference list of

your report.

3 Build one comparison table:

method year, data, metric and the main limitation, with each row cited. Our [literature survey guide](#) shows the table format.

4 Write the report like a short review paper:

introduction, background, the methods [review paper format guide](#)

grouped by the survey's taxonomy, the comparison table, applications, open challenges and a conclusion. Our gives the section order and shows how to build the comparison table.

5

Make the slides from the report,

using the eight-slide plan on your card, then add title, outline and reference slides. Our [M.Tech dissertation PPT guide](#) has the slide design rules, and they apply to a seminar too.

6

Prepare the demo or table you will show,

and recor

d a scree n video of it in case the hall comp uter or the intern et fails.

7 **Rehearse aloud with a timer,** twice, and cut slides until the talk fits the slot.

A seminar report, chapter by chapter, and where each part comes from

CHAPTER	WHAT GOES IN IT	WHERE IT COMES FROM
Introduction	The problem, why it matters now, and the scope of the report	The survey's introduction, in your own words
Background	The basic ideas a reader needs first	Textbooks and the survey's background section
Methods	The main approaches, grouped by the survey's taxonomy	The survey, checked against the papers it cites
Comparison	One table of methods, data, metrics and limitations	The six to ten papers you read

CHAPTER	WHAT GOES IN IT	WHERE IT COMES FROM
Applications and demo	Where the methods are used, and what you ran yourself	The survey, and your own demo
Challenges and conclusion	Open problems and what you would study next	The survey's open problems, and your view
References	Every source cited in the text, usually in IEEE style	The survey and the papers you read

Page limits, front pages and the font and spacing rules differ between universities. Follow your department's template where it has one.

What questions do panels ask after an M.Tech seminar?

Expect questions that check whether you understood the survey or only copied its slides. Prepare a short answer to each of these.

- Why did you choose this topic, and how is it connected to your dissertation?
- Explain the main method on one slide again, without the slide.
- What is the difference between the two approaches in your comparison table?
- Which dataset and metric do papers in this area use, and why those?
- What is the biggest open problem the survey names, and how would you start on it?
- What did you run or measure yourself?

For the questions examiners ask at a project viva, see our viva questions guide.

How can we help with your M.Tech seminar?

We help M.Tech students in CSE and IT choose a seminar topic that leads into their dissertation, go through the survey paper with them, and plan the report, the slides

and the demo. When the seminar is done, the same reading becomes the literature survey for Stage-I, and we can carry on with the base paper and the dissertation.

Start with our dissertation, report and PPT page, or see M.Tech projects for CSE and AI for the full dissertation service. The consultation is free, and you get a written plan and a fixed quote.

Frequently asked questions

What is the difference between an M.Tech seminar and an M.Tech project?

A seminar is a talk and a short report on an area you have studied, usually from one or more survey papers. You explain the area, compare the main methods and answer questions; you are not expected to build a new system or claim new results. A project or dissertation starts from a base paper that you reproduce and then extend. For those, see our M.Tech project ideas.

How many slides should an M.Tech seminar PPT have?

Your department's time slot decides it. A common rule is about one content slide per minute, plus title, outline and reference slides. The eight-slide plans on this page cover the core content; add one slide for your demo or comparison table. Our M.Tech dissertation PPT guide covers slide design rules that apply to seminars too.

Can I take my seminar and dissertation in the same area?

Many students do, because the papers you read for the seminar then feed the literature survey of your Stage-I report. Whether the seminar and the dissertation may share a title is a department rule, so ask your guide. Our literature survey guide shows how to turn seminar reading into a survey table.

Where can I read these survey papers for free?

Open the DOI link on each card. IEEE Access, the IEEE Open Journal of the Computer Society, the Journal of Cloud Computing, Computers and Entropy are open-access journals. For the ACM and VLDB Journal papers, Semantic Scholar listed a free copy in November 2026, either on the publisher's site or as the authors' preprint on arXiv.

How recent should the survey paper for an M.Tech seminar be?

As recent as you can find, ideally from the last three years, so that the methods you present are still current. Every paper on this page is from 2023 to 2026. If your topic has no survey that recent, a newer research paper can bring the older survey up to date. Our review paper format guide explains how a survey is organised.

Can you help me prepare the seminar report and PPT?

Yes. Send the topic, your dissertation area and your seminar date. We help you plan the report and slides around the survey, set up the demo and prepare for questions, and the consultation is free. See our report and PPT help.

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. →

Report & PPT
The report, black book or dissertation, and review slides. →


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