MBOT
LEMBAGA TEKNOLOGIS MALAYSIA
MALAYSIA BOARD OF TECHNOLOGISTS

TECHIES

Official Bulletin 35th Edition - KDN: PQ1780/J/187
www.mbot.org.my

THE 2 A.M. TUTOR: How a Generative AI Chatbot Walks a Student Through Additional Mathematics Problems

A close look at generative AI tutors in Malaysian classrooms. What they actually do, where they trip up, and why technologists must steer the wheel.

CONTINUED ON PAGE 02 >>

By Ts. Muhammad Suzarilshah, SWIFT

/what's inside

/editorial committee

ISSN 0128-1313
9 770128 131009

/chief editor's note

Artificial Intelligence (AI) has the potential to transform teaching and learning, making education more personalised, accessible, and effective. However, this transformation requires more than simply introducing new technologies into classrooms. We must remember that the goal is not to replace human intelligence with artificial intelligence, but to enhance human capabilities through responsible use of technology.

The integration of AI must remain aligned with curriculum objectives and learning outcomes. Students, educators, and technical experts should develop a deeper understanding of how AI systems work, including their capabilities, limitations, and potential risks, rather than using them blindly.

One of the most pressing challenges is assessment. As AI tools become increasingly capable, assessment methods must be revamped to ensure that students are genuinely evaluated on their knowledge, skills, creativity, and critical thinking. Assessments are designed to measure student learning, not the capabilities of a virtual AI brain.

I am confident that navigating this educational revolution requires expertise, commitment, and collaboration from all of us. Together, we can harness AI to strengthen education while preserving the values and purpose of learning.

Assoc. Prof. Dr. Mohamad Asmidzam Ahamat

CONTINUED FROM PAGE 01 >>

It is 2 a.m. Aisyah, a Form Five student in Ipoh, is two days away from her SPM Additional Mathematics trial examination. She is stuck on a single line of work in a differentiation problem. Her teacher is asleep. Her tuition group chat is silent. So she opens an app on her phone, types in the question, snaps a photo of her work, and asks, "Where did I go wrong?"

A generative AI tutor reads her photo, recognises the steps she wrote, and replies in conversational Bahasa Melayu. It does not give her the answer. Instead, it points to the second line and asks: "You applied the chain rule here. What is the derivative of the inside function?" Forty minutes later, Aisyah has worked through three similar problems on her own, and the chatbot has quietly logged where she still hesitates.

This is what a generative AI tutor looks like. Strip away the marketing, and what you have is a chatbot built on a large language model, or LLM. It is not a futuristic robot teacher. It is a piece of software, accessible using a smartphone, that can hold a back-and-forth conversation, explain ideas in plain language, and adapt to the student in front of it. For technologists working in Information and Communication Technology, this is one of the most consequential applications of AI we will deploy this decade.

Pulling back the curtain: what is a generative AI tutor?

A large language model is a program trained on enormous amounts of text. Textbooks, websites, forums, code, and even past examination papers. It studies all of that until it learns the statistical patterns of how humans write and explain things. When you type a question, the model does not "look up" an answer. It predicts, one word at a time, what a knowledgeable response would sound like. ChatGPT, Google's Gemini, and Anthropic's Claude are the best-known examples.

A "tutor" built on top of such a model is simply a chatbot wrapper with three things added on. The first is a system prompt that tells the model how to behave. Something like, "You are a patience Additional Mathematics tutor. Never give the final answer. Ask guiding questions." The second is a memory of the current conversation, so the model can refer back to what the student already tried. The third, sometimes, is a connection to a curriculum database, so the tutor stays anchored to the syllabus the student is actually sitting for.

The result is a system that does something a textbook cannot. It talks back. It does not just present the chain rule. It watches a student fumble through it and steps in at the exact line where understanding breaks down.

A real, specific application: Khanmigo and the Pandai generation

In May 2023, Khan Academy released Khanmigo, an AI tutor built on OpenAI's GPT-4 and offered to schools in the United States and a handful of international partners. Khanmigo refuses to give direct answers to homework. Instead, it asks the student what they have tried, suggests where they might check their work, and gives hints scaled to their level of stuckness. Early studies of Khanmigo, presented at ASU+GSV in 2024, suggested that students using the tool spent measurably more time on practice problems and were more willing to attempt harder questions, because the cost of being wrong had dropped.

Closer to home, Malaysian start-up Pandai has launched an AI-powered study companion targeted at SPM and UPSR students, with Bahasa Melayu and Mandarin interfaces. It uses an LLM under the hood, plus a layer of locally curated content tied to the Kurikulum Standard Sekolah Menengah (KSSM). According to Pandai's own reporting in 2024, the platform was used by more than one million Malaysian students, with the majority of usage occurring between 9 p.m. and 1 a.m. Those are precisely the hours when teachers are unavailable.

The specific application here is narrow. It is an after-hours, syllabus-aligned conversational tutor that helps a student through a specific problem. It is not replacing the teacher. It is filling the gap between the bell ringing and the next class.

Why this matters for equity

For decades, the most reliable predictor of an SPM Add Maths grade in Malaysia has been simple. It is whether the family could afford private tuition. A good tutor in Petaling Jaya can charge RM150 an hour. A generative AI tutor on a smartphone costs roughly the same as mobile data. If the underlying tool is well-designed, this is the single largest cost reduction in personalised teaching in a generation.

Where generative AI tutors trip up

Three problems that demand the attention of every technologist deploying these systems are as follows:

Hallucinations

LLMs sometimes produce confident, fluent, and completely wrong answers. A model may insist that the integral of one over x-squared is one over x, with a flourish of reasoning to back it up. For a struggling student who does not yet know better, this is not a quirky bug. It is misinformation delivered by a trusted-sounding voice.

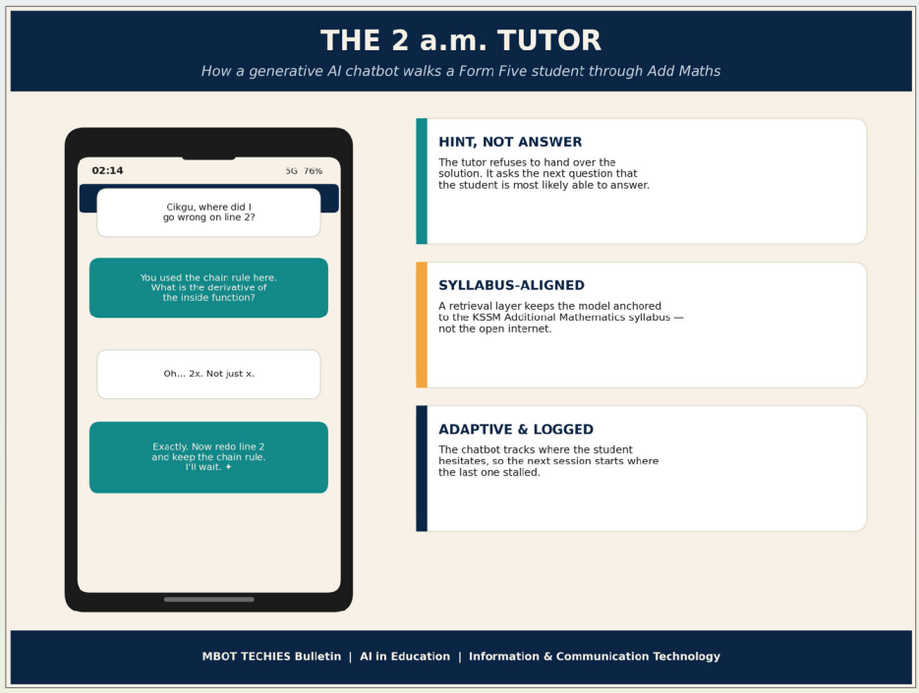

A generative AI tutor that coaches instead of answering. That single design choice decides whether students learn or lean.

Over-reliance

A tutor who answers too readily erodes the one skill that learning mathematics actually requires productive struggle. Designing a chatbot to be a coach rather than a calculator is harder than it sounds. Most consumer-oriented LLMs are designed to function as eager assistants by default.

Data and dignity

Students using these tools share photos of their homework, learning patterns, and location through their devices. Who stores that data? Who can buy it? Under the Personal Data Protection Act 2010, schools and vendors have clear duties to protect this information. However, those duties are seldom subject to judicial scrutiny and rarely audited.

The role of technologists

This is where MBOT-registered technologists, particularly those of us in ICT, must lead. Building a generative AI tutor responsibly is an engineering discipline, not a vibes exercise. It demands prompt design that enforces Socratic behaviour. It demands retrieval systems that ground answers in verified syllabus content. It demands evaluation harnesses that test the tutor against past SPM question papers. It demands consent flows that a fifteen-year-old can actually understand.

A teacher cannot inspect the weights of a language model. We can.

A parent cannot read a system prompt. We can.

The chatbot does not know it is talking to a child preparing for the most important exam of her life. We do.

Aisyah puts away her phone at 2:47 a.m. She has not finished her revision, but she understands the chain rule for the first time. Tomorrow, the AI tutor will not remember her. But she will remember the tutor.

ABOUT THE AUTHOR

Ts. Muhammad Suzarilshah is Microsoft AI MVP, Docker Captain, a Microsoft Certified Trainer, and a technophile with a deep passion for technology, particularly in the realms of Docker, Cloud, IoT, and Networking. With a strong foundation in IT infrastructure, he also enjoys coding and developing software solutions.

By Ts. Dr. Raja Mohd Tariqi bin Raja Lope Ahmad, Fiscal Digest Sdn. Bhd. and
Raja Qatrun Nada binti Raja Mohd Tariqi, Fakulti Pendidikan, Universiti Kebangsaan Malaysia (UKM)

Bridging the Language Gap: Pre-trained NLP Models for Bahasa Melayu in University Learning

Introduction

A Computer Science lecturer teaching an Artificial Intelligence course ends class with a familiar question: *"Boleh faham tak topik tadi?"* Within minutes, forty written replies are collected.

Some are written in standard Bahasa Melayu, while others include Manglish, regional dialects, and code-switched expressions commonly used by Malaysian students.

The comments are uploaded into an AI-powered educational dashboard designed to analyse student sentiment and emotional reactions. Within seconds, the system reports that more than 95 per cent of the class felt positive. The lecturer hesitates. Something does not feel right.

Behind the dashboard are pre-trained transformer models originally trained on English-dominant corpora. Although these models perform well on global benchmarks, they often struggle to understand Malaysian educational language patterns, particularly dialects, Manglish, and culturally specific expressions.

The language gap in Malaysian educational AI

Large language models learn meaning from the data used during training. When training datasets are dominated by English-language content, the resulting models often fail to fully capture the linguistic diversity of Bahasa Melayu used in real Malaysian classrooms.

Expressions such as *"molek"*, *"berabe"*, *"siok bah"*, and *"kemain"* carry emotional meaning that is immediately understood by Malaysians but may be misinterpreted by generic multilingual AI systems. In many cases, the prediction is not completely wrong. It is simply too generic, reducing rich regional expressions into oversimplified positive or neutral signals. This creates a major challenge for universities attempting to use AI-based dashboards to evaluate student feedback, learning experiences, and classroom engagement.

Building a Malaysian educational NLP benchmark

To study this issue systematically, we developed MYUniDialectSentiment840, a Malaysian university feedback benchmark containing 840 student comments collected across fourteen Malaysian dialect groups, including Kelantan, Terengganu, Sabah, Sarawak, Kedah, Johor, Negeri Sembilan, Manglish, and KL slang.

Each feedback sample was manually labelled across four dimensions:

- Sentiment (positive, neutral, negative)
- Emotion (happy, love, surprise, sadness, anger, fear, or no clear emotion)
- Academic topic
- Learning context

This benchmark allowed the research team to identify where generic AI models failed and where locally adapted models could improve performance.

Technology stack

The prototype system was developed using lightweight and deployable open-source technologies suitable for Malaysian educational environments:

- Python
- Hugging Face Transformers
- ft-Malay-bert
- multilingual-emotion-classifier
- Streamlit Dashboard
- Apple Silicon MPS Acceleration
- Apache-2.0 Open-Source Deployment

Expressions such as "molek", "berabe", "siok bah", and "kemain" carry emotional meaning that is immediately understood by Malaysians but may be misinterpreted by generic multilingual AI systems.

What do the results show?

Initial testing using generic multilingual baselines revealed significant weaknesses in understanding Malaysian student feedback. Neutral comments were frequently classified as positive, while culturally specific emotional expressions were often missed entirely.

After fine-tuning locally adapted models using MYUniDialectSentiment840, the evaluation showed substantial improvements in sentiment and emotion classification performance.

Within the controlled benchmark environment:

- Sentiment classification achieved 99.32% accuracy
- Emotion classification achieved 97.96% accuracy
- Classification consistency improved across all fourteen dialect groups
- The “no_clear_emotion” category significantly reduced false positive emotional predictions

These results demonstrate that local adaptation is often more important than simply using larger models.

Educational AI governance and ethical considerations

As educational institutions increasingly adopt AI-powered analytics systems, governance and privacy considerations become critically important. The proposed approach focuses on anonymised educational feedback and lightweight deployable models that can operate within institutional governance environments. The system is designed to support educators rather than replace human judgement.

Responsible deployment also requires transparency, explainability, and sensitivity to cultural and linguistic diversity within Malaysian classrooms.

Key Lessons for Malaysian Higher Education

Three important lessons emerged from this work:

1. Local benchmarks matter: AI systems cannot improve without datasets that accurately reflect Malaysian linguistic diversity.
2. Neutral student feedback is meaningful: Calm and descriptive responses should not automatically be interpreted as positive emotional reactions.
3. Local adaptation is practical and affordable. Lightweight transformer models can deliver strong results without requiring expensive frontier-scale infrastructure.

A Path Forward

The next challenge for Malaysian educational AI is not simply larger models, but broader and more representative Malay-language datasets. A collaborative national academic-feedback benchmark involving Malaysian universities could help create more inclusive and culturally aware educational AI systems. Such efforts would support better learning analytics, more accurate institutional insights, and fairer educational technologies for diverse student populations.

For MBOT-aligned technologists, this work demonstrates that meaningful AI innovation does not always require massive infrastructure investments. Sometimes, the most impactful innovation begins with understanding how people actually speak.

OUR IT STAFF ARE NOT READY FOR AI, AND THAT'S OKAY. FOR NOW.

By Ts. Elizabeth Matthew, Mindvalley Labs Sdn Bhd

How Malaysian universities can help their campus IT teams grow alongside the AI tools they are being asked to support

Let's start with an honest admission: when most Malaysian universities decided to adopt AI tools, the conversation was almost entirely about the technology. Which platform to buy? How much does it cost? What can it do for students and lecturers?

Almost nobody asked: "Is our IT team ready for this?"

That question matters more than most people realise. Because behind every AI-powered student portal, every smart learning management system, and every automated academic chatbot, there is a group of IT staff who are expected to install it, run it, fix it when it breaks, and make sure it stays secure. And right now, many of those teams are being asked to support technology they were never trained for.

This is not a criticism; it is simply where we are. The good news is that it is entirely fixable if universities are willing to invest in their people as thoughtfully as they invest in their systems.

So, what exactly has changed?

Traditional IT work in a university is quite predictable. They manage the network, support the student information system, make sure computers in labs are working, and ensure that the Wi-Fi reaches the lecture hall. These are well-understood problems with well-understood solutions.

AI systems are different. They do not just sit still and wait for someone to use them. They process data constantly, learn from patterns, and can behave in unexpected ways. An AI chatbot might give a student the wrong answer. A predictive analytics tool might flag a student as "at risk" based on flawed data. A facial recognition attendance system might fail because the lighting in a new building is slightly different from the environment in which it was trained.

These are not the kind of problems you solve by restarting the server. They require a new kind of thinking and a new set of skills.

"The IT staff who keep our university systems running are some of the most important people in any AI adoption story. But they are also often the most overlooked. That needs to change."

The gap in practice

The skill gap in campus IT is not one big problem; it shows up differently depending on who you are talking to:

- **For system administrators and network engineers,** the challenge is that most AI platforms run on cloud infrastructure, use container technologies, and communicate through APIs. These are relatively new concepts for teams that were trained primarily on on-premise systems and traditional networking.
- **For helpdesk staff,** the difficulty is that AI problems do not look like normal IT problems. When an AI tool gives a wrong answer, produces a strange output, or behaves inconsistently, there is no standard troubleshooting guide to follow. You need to understand a little about how the AI works before diagnosing the issue.

- **For IT managers,** the gap is often at the decision-making level. When a vendor pitches a new AI tool, do you know the right questions to ask about data privacy, model reliability, or long-term maintenance costs? Many IT leaders find themselves in unfamiliar territory during these conversations, and that is completely understandable.
- **For cybersecurity teams,** AI introduces entirely new types of threats that traditional security training does not cover, such as attackers trying to manipulate an AI model's outputs, or sensitive training data being exposed through clever queries.

None of these gaps is the fault of the individuals involved. They are the natural result of technology outpacing training programmes.

No need to turn everyone into a data scientist

Here is the part that often gets lost in these conversations: closing the AI skill gap does not mean sending your entire IT department back to university. It does not mean everyone needs a machine learning certification or a degree in data science.

What it means is giving each person the level of AI knowledge that they actually need to do their job well. Think of it in three simple layers:

Layer 1: The basics for everyone

Every single person in your IT department should understand, at a basic level, what AI is, what it can and cannot do, and what kinds of problems it tends to create. There are short, free online courses from Google, Microsoft, and local providers such as MDEC that cover AI literacy in under 20 hours. The goal is not technical expertise; it is informed common sense. An IT support officer who understands that AI systems can produce confident but incorrect outputs will handle user complaints far more effectively than one who treats an AI glitch like a corrupted file.

Layer 2: Hands-on skills for technical staff

Staff who actually deploy, configure, and maintain AI systems need more specific training. The most useful areas to focus on are cloud platforms (AWS, Azure, and Google Cloud all offer AI-focused certifications), working with APIs, and basic data management. The great news is that most of the major AI platform vendors offer their own training and certification programmes, so if your university has already chosen a platform, there is likely a structured learning path available.

Layer 3: Strategic understanding for IT leaders

Managers and IT directors do not need to know how to build an AI model. But they do need to be able to sit in a meeting with an AI vendor and ask smart questions. They need to know when a proposed AI system poses a data privacy risk. They need to understand what "model drift" means and why it matters for long-term system reliability. Programmes offered through MBOT, ISACA, and other professional bodies are increasingly including AI governance and leadership modules specifically designed for this audience.

Making learning part of the job

Training budgets are tight in most university IT departments. And even when the funding is there, finding time to learn while keeping systems running is a real challenge. Here are a few approaches that work even under these constraints:

- **Monthly knowledge-sharing sessions:** Set aside one hour a month for IT staff to share something new they have learned like a tool they tried, a problem they solved, or an article they found useful. It costs nothing and builds collective knowledge over time.
- **Certifications with recognition:** If a staff member earns an AI or cloud certification, acknowledge it in their performance review, in an internal announcement, or even just in a team meeting. Recognition sends a clear signal that learning is valued.
- **Involvement in AI projects:** Include IT staff in AI project planning from the very beginning, not just at the deployment stage. When technical staff understand why a system is being built and what it is supposed to achieve, they become far more effective at supporting it.
- **Short attachments or site visits:** If another university or a technology company nearby has already implemented AI tools successfully, arrange a visit. Seeing how a more experienced team operates is often worth more than any online course.

A message to MBOT-registered technologists

If you are a technologist registered with MBOT and you work in campus IT, you are in a genuinely important position right now. The 24 fields of technology recognised by MBOT touch almost every aspect of AI deployment in education from ICT infrastructure to systems engineering to cybersecurity.

Your CPD requirements are not just a box to tick. They are an invitation to document your growth during one of the most interesting periods in the history of technology. AI-related training, certifications, and project experience all count towards your professional development record.

More importantly, they make you better at what you do and better placed to guide your institutions through a transition that will only accelerate.

Final thought: Technology is only as good as the team behind it

We talk a lot about AI changing education. And it will. But the change will only be as good as the people making it happen including the IT staff working behind the scenes to keep everything running.

Investing in their skills is not a nice-to-have. It is the difference between an AI strategy that delivers on its promise and one that quietly falls apart six months after launch. Malaysian universities have a real opportunity here to build IT teams that are not just keeping up with AI, but genuinely helping to shape how it is used. That starts with taking their training as seriously as we take the technology itself.

/mbot news

SEMINAR ON DIGITAL EVOLUTION IN THE MANUFACTURING SECTOR

On 9 June 2026, the Digital Evolution in the Manufacturing Sector Seminar was successfully organised by the Malaysia Board of Technologists (MBOT) and the Japan-Malaysia Technical Institute (JMTI). The seminar took place at JMTI, Simpang Ampat, Penang. It was attended by 115 participants from industry, academia, and the technology community.

The event was graced by Datuk Rospiagos bin Taha, Director General of the Department of Human Resources (JTM). Three insightful presentations were delivered by:

- Mr. Zambri bin Abdul Hamid
  Vocational Training Officer (VCAT-II Vibration Analyst), JMTI
- Ts. Mohammad Syukur Md. Noh
  Director of Production and Engineering, B. Braun Medical Industries Sdn. Bhd.
- Ts. Ir. Dr. Shamsul Kamar Abu Samah
  Chief Executive Officer, National Aerospace Industry Corporation Malaysia (NAICO)

With the theme "Navigating Transformation Towards Sustainability and Social Values in the Industry 5.0 Era", the seminar aimed to explore the latest developments in digital transformation, workforce development, and sustainable industrial practices in the Industry 5.0 era.

The seminar reflects the commitment of MBOT and JMTI to strengthening the nation's technology ecosystem and nurturing a highly skilled workforce capable of succeeding in a human-centric and sustainable Industry 5.0 environment.

MBOT extends its sincere appreciation to all speakers, participants, and partners for their valuable contributions to the success of this seminar.

> *The seminar reflects the commitment of MBOT and JMTI to strengthening the nation's technology ecosystem and nurturing a highly skilled workforce capable of succeeding in a human-centric and sustainable Industry 5.0 environment.*

/mbot registration

68,922	15,798	28,679	4,209
Graduate Technologists	Qualified Technicians	Professional Technologists	Certified Technicians

117,608 Total MBOT Registrants (As of June 2026)

TECHIES is published by the Publication Committee of the Malaysia Board of Technologists.