

AI for Queensland SMEs

From LLMs to Agents and
Practical Business Workflows

Dr Taotao Cai

Senior Lecturer in Computing, UniSQ

Program Director, Master of Cyber Security



YJ LININGS
INTERIOR CONSTRUCTION SPECIALISTS

- WALL & CEILING SYSTEMS
- PARTITION & FRAMING SYSTEMS
- CLADDING SYSTEMS
- CARPENTRY & FIT-OUT WORKS



Today's Roadmap

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps

Now Moving To

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps

Dr Taotao Cai

Senior Lecturer in Computing | Program Director, Master of Cyber Security

University of Southern Queensland (UniSQ)



 **2017**
PhD commenced
The University of
Western Australia

1



Perth



2019–2021
Transferred to Deakin University (Mar 2019)
PhD completed (Nov 2020)
Associate Research Fellow (Jul 2020 – Feb 2021)

2

Melbourne



2021
Postdoctoral Research Fellow
Macquarie University

3

Sydney



4

Brisbane



 **2023–Present**
Continuing Academic Position
University of Southern
Queensland



Homepage

<https://taotao1992.github.io/>



WeChat

Research, Publications, and Patents



Research Interests

- **Graph learning** and dynamic network reasoning
- **Trustworthy AI** for social and information systems
- **Sustainable AI** for remote sensing, infrastructure, and smart built environments
- **Intelligent transportation** and operational decision support



Publications

- **30+** peer-reviewed outputs in **CORE A*** venues and high-impact journals
- Information Systems selected "*Community-Diversified Influence Maximization in Social Networks*" as one of its most influential papers of the past decade
- Recent work spans **TGRS, JSTARS, TKDE, ICDE**, and related AI venues



Patents

- **5** international patents in graph sampling, community search, and influence estimation

For the current publication and patent list, see my website, Google Scholar, and CV.



Grants, Teaching, and a Living AI Atlas



1 Grants



AEA Ignite highlight 1

IG250200014

SafeBuild AI: Intelligent Safety and Efficiency for Construction SMEs



2026–2027 | Lead Chief Investigator



Funding: approx. AUD 549,886



AEA Ignite highlight 2

IG240100414

Intelligent FCAS Agent for Renewable Energy Power Systems



2025–2026 | Chief Investigator



Lead: Dr Yanming Zhu, Griffith University



Funding: approx. AUD 508,442

2 Teaching



Teaching focus

CSC2220 Artificial Intelligence

Course focus: AI concepts, methods, tools, and applications

3 Living AI Atlas



AI Evolution Atlas

Interactive map of AI history, key papers, agents, videos, and future directions

Scan to visit



<https://taotao1992.github.io/ai-evolution-atlas/>

Now Moving To

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

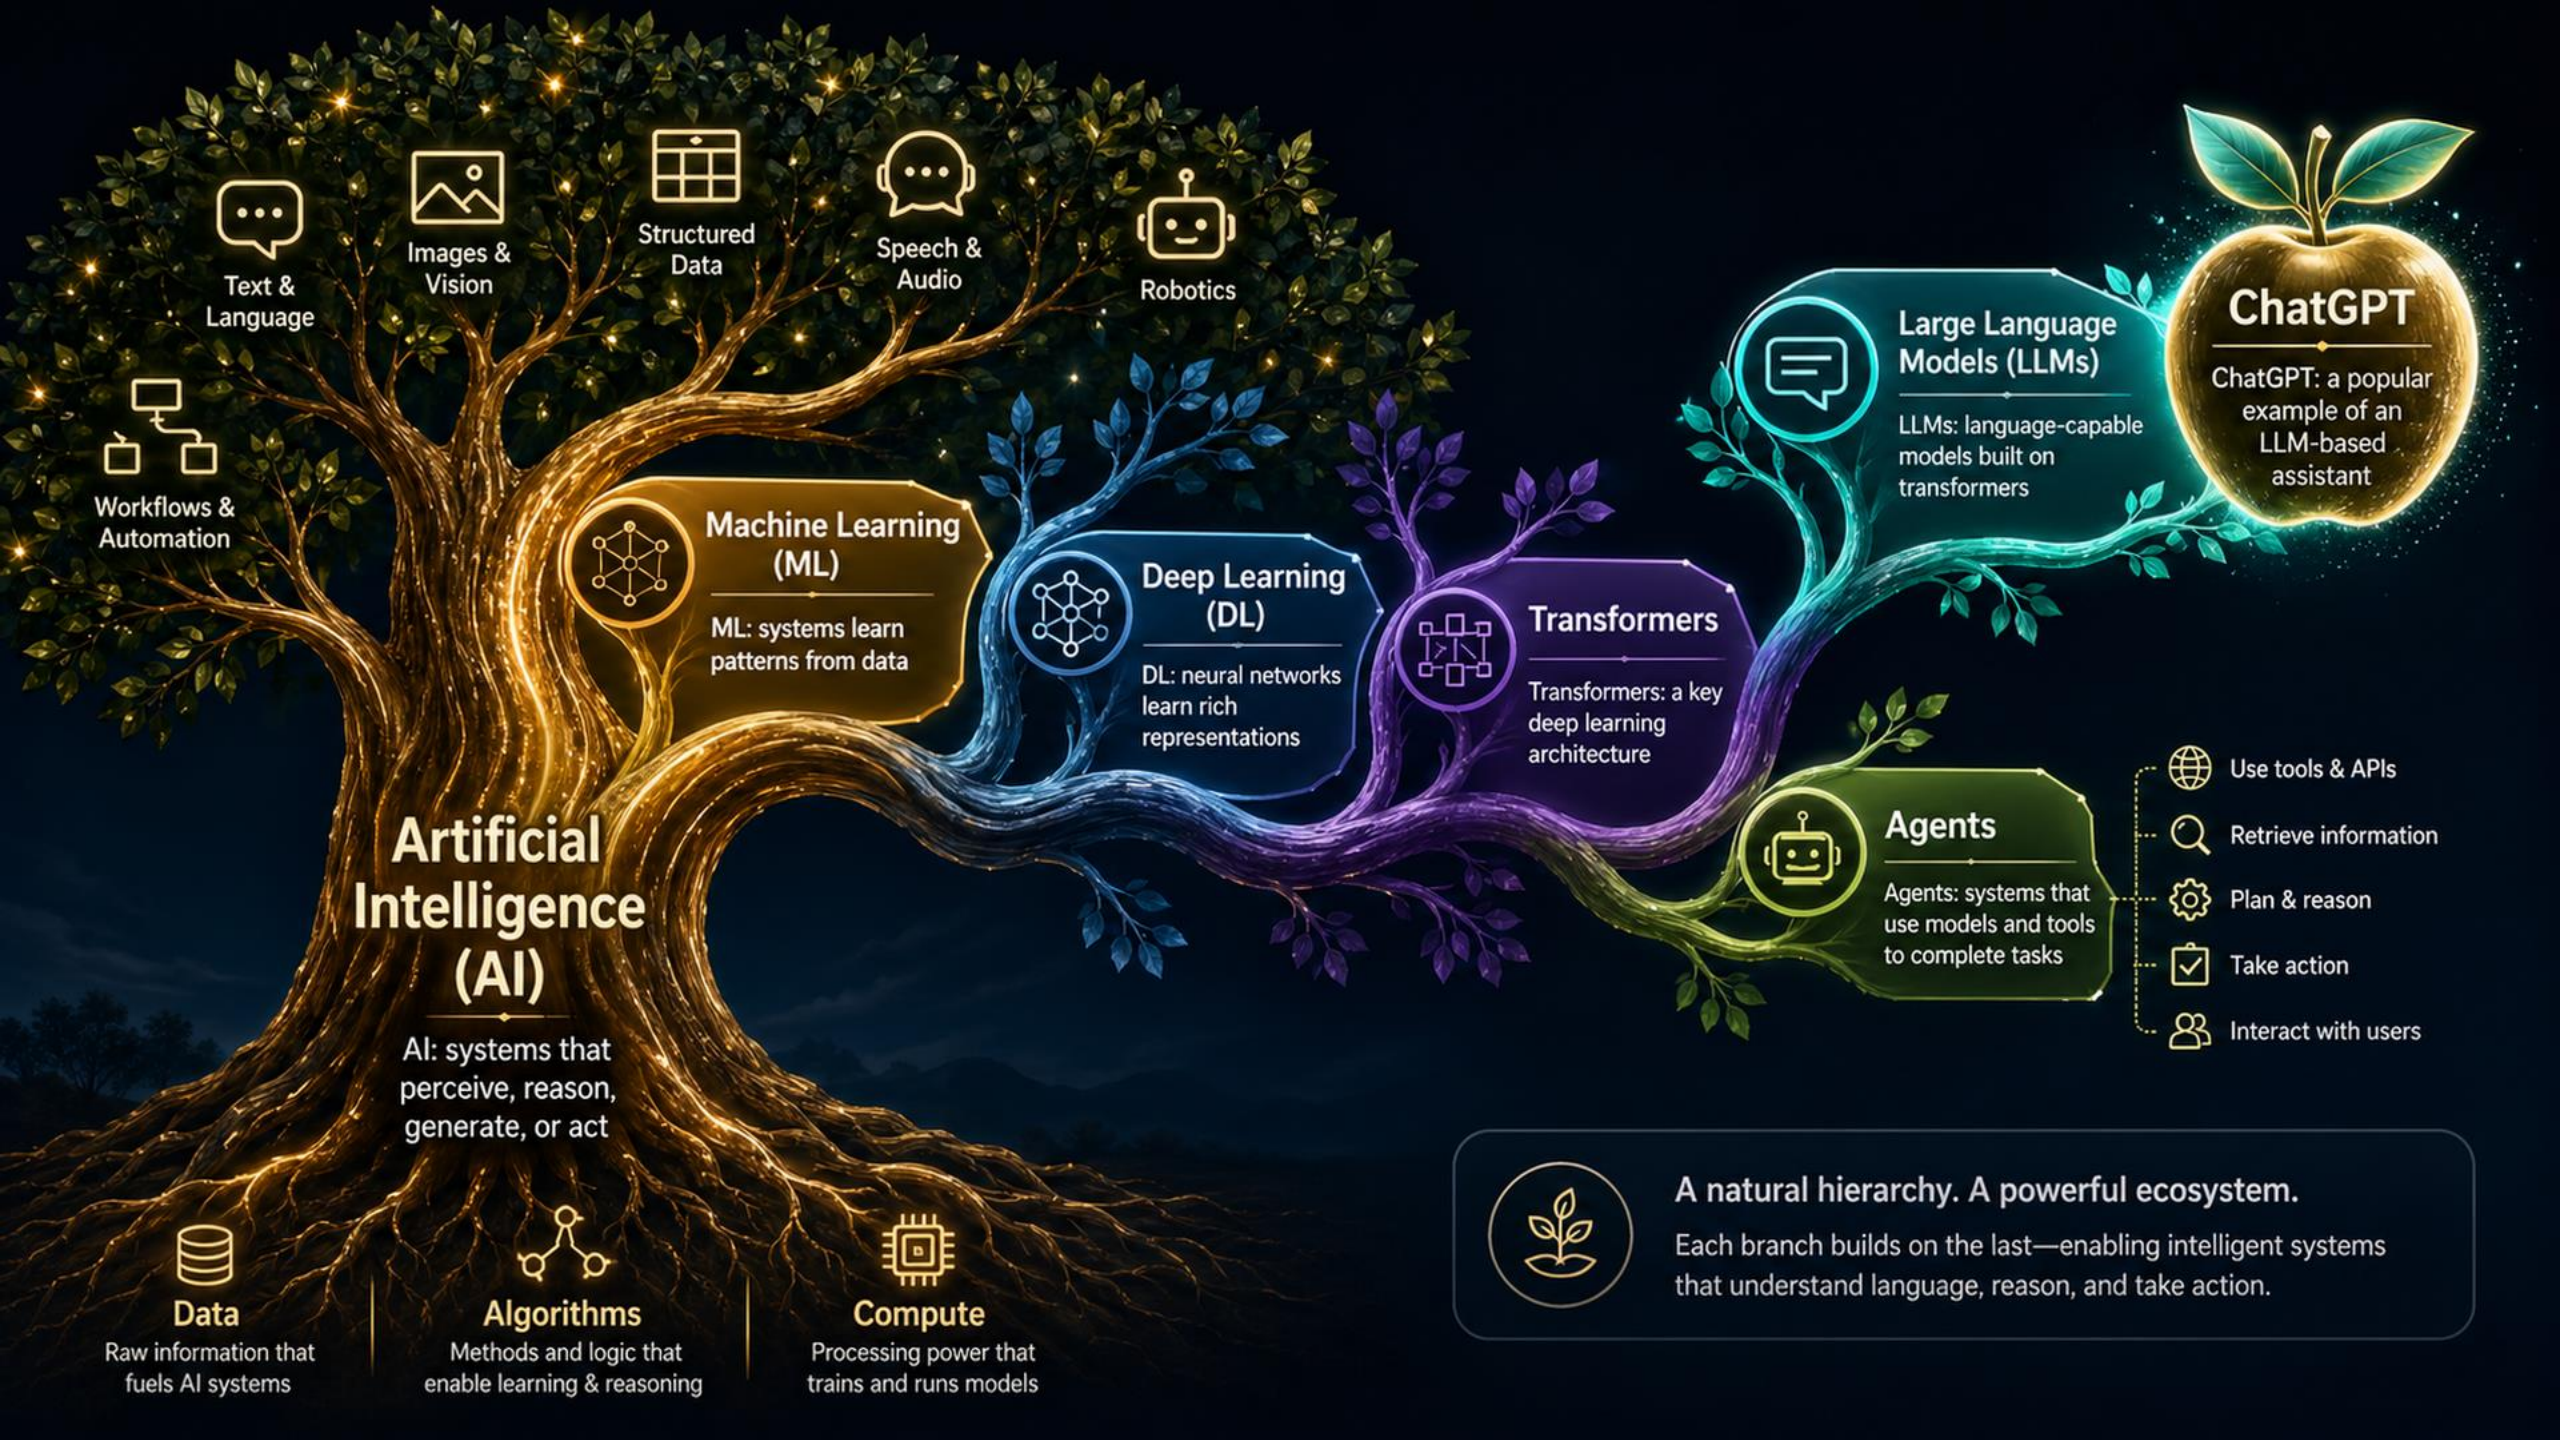
How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps



Artificial Intelligence (AI)

AI: systems that perceive, reason, generate, or act

Data
Raw information that fuels AI systems

Algorithms
Methods and logic that enable learning & reasoning

Compute
Processing power that trains and runs models

Workflows & Automation

Text & Language

Images & Vision

Structured Data

Speech & Audio

Robotics

Machine Learning (ML)
ML: systems learn patterns from data

Deep Learning (DL)
DL: neural networks learn rich representations

Transformers
Transformers: a key deep learning architecture

Large Language Models (LLMs)
LLMs: language-capable models built on transformers

Agents
Agents: systems that use models and tools to complete tasks

- Use tools & APIs
- Retrieve information
- Plan & reason
- Take action
- Interact with users

ChatGPT
ChatGPT: a popular example of an LLM-based assistant

A natural hierarchy. A powerful ecosystem.
Each branch builds on the last—enabling intelligent systems that understand language, reason, and take action.

From Text Completion to Deep Understanding: The Evolution and Early Applications of Large Language Models (LLMs)



Early Completion Models

Emergent Capabilities

Business Applications & Governance

1 Early Completion Models (Completion Models)

Limitations and breakthroughs: more like advanced text prediction. Good at completion, but not reliable for complex business logic.

2 Emergent Capabilities (Emergent Capabilities)

Large-scale training enabled zero-shot problem solving, and models began to show reasoning and transfer abilities.

3 Early Business Applications

Automated draft writing, large-scale document summarization, and basic customer service responses.

4 Hallucinations & Governance

Models better understand context and user intent, but accuracy, controllability, and governance still require balance.



Key takeaway: LLMs have evolved from statistical completion tools into knowledge work assistants with emerging reasoning ability.

The 2025/2026 LLM Landscape

Performance, Cost, and Best Fit for SMEs



CLOSED FRONTIER

 **Claude Opus 4.7**

-  **Strength:** Advanced coding and reasoning
-  **Metric:** Frontier-grade code and logic performance
-  **SME fit:** Best for complex business logic and mission-critical code

CLOSED FRONTIER

 **GPT-5.5**

-  **Strength:** Fast, versatile general intelligence
-  **Metric:** Strong low-latency and multimodal performance
-  **SME fit:** Best for real-time customer chat, drafting, and voice workflows

CLOSED FRONTIER

 **Gemini 3.1 Pro**

-  **Strength:** Long context and multimodal analysis
-  **Metric:** Strong large-context document and media handling
-  **SME fit:** Best for long meetings, reports, and mixed text-image analysis

OPEN-SOURCE DISRUPTOR

 **DeepSeek V4**

-  **Strength:** High capability with strong cost efficiency
-  **Metric:** Well-suited for scaled inference and batch workloads
-  **SME fit:** Best for high-volume data processing and cost-sensitive automation

OPEN / SELF-HOSTED

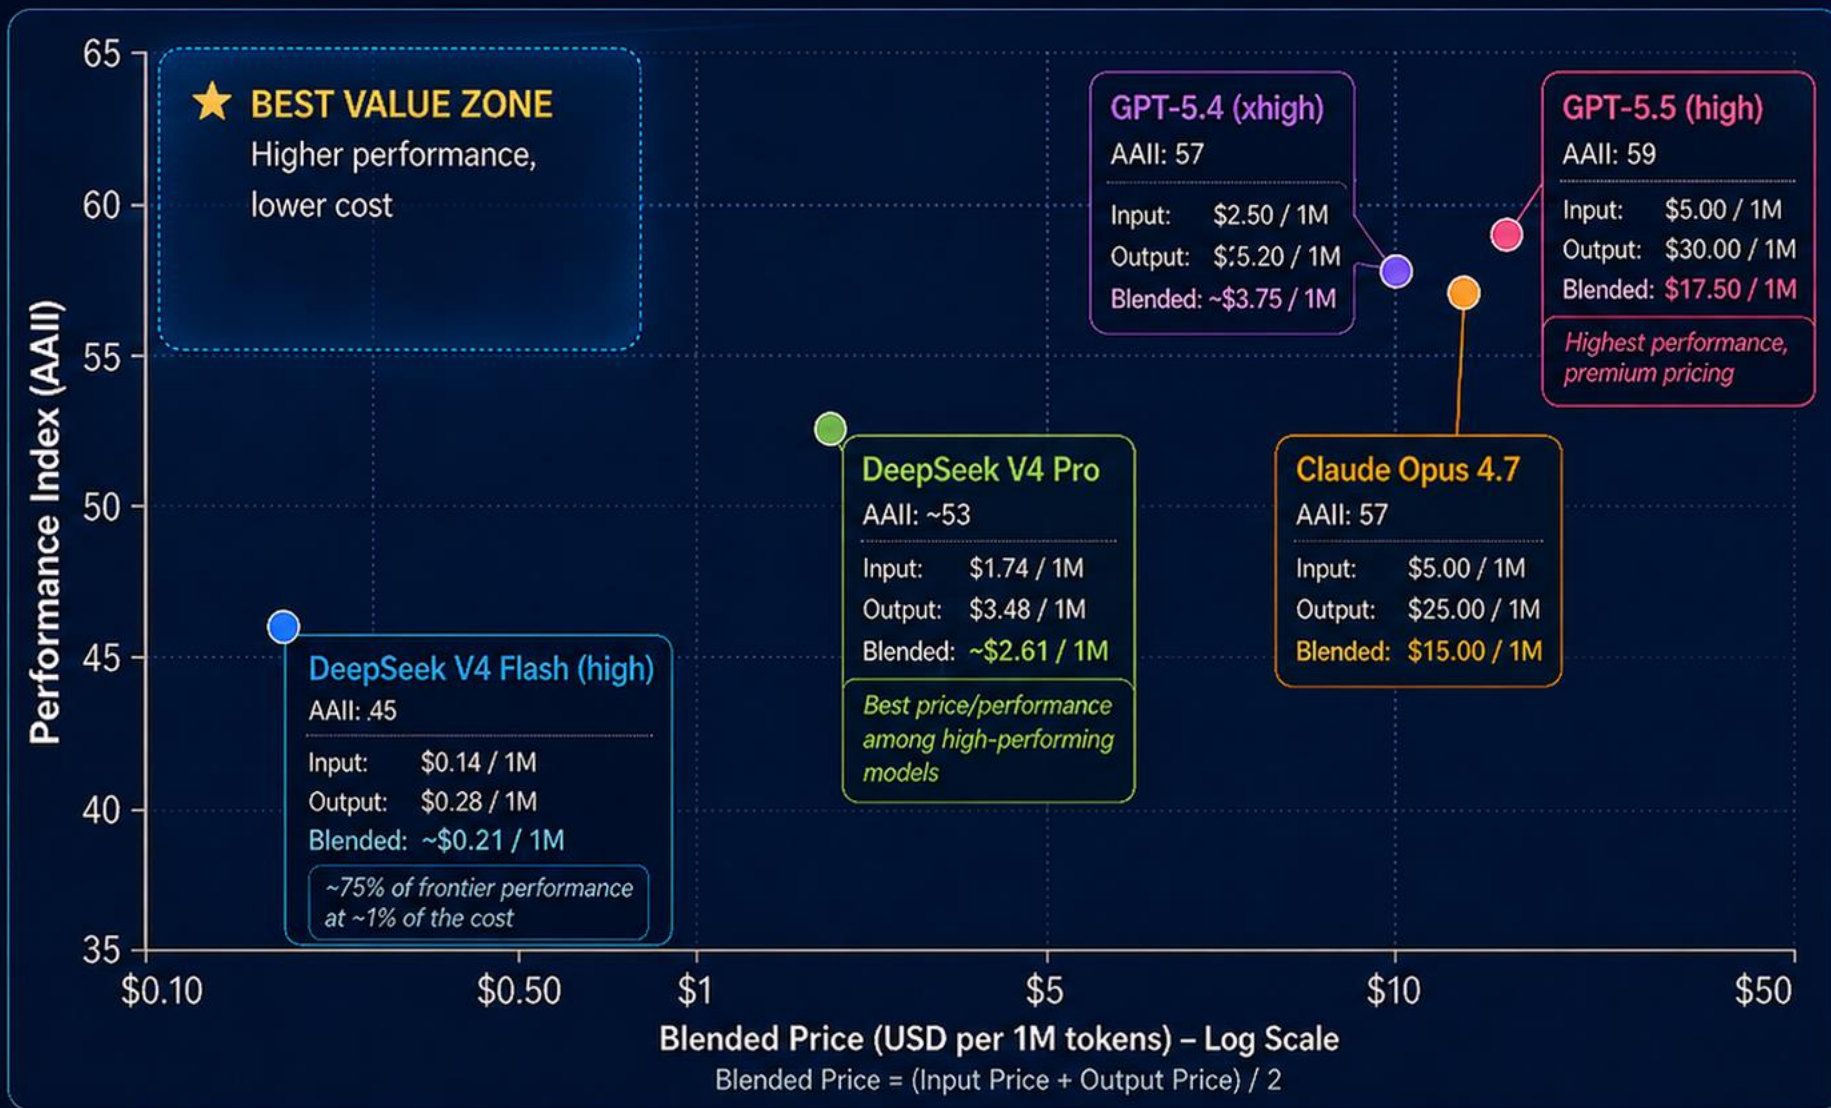
 **Llama 4 / Qwen 3**

-  **Strength:** Private deployment and no API dependence
-  **Metric:** Local deployment option
-  **SME fit:** Best for privacy-sensitive and compliance-heavy workflows

 **Decision rule for SMEs:** choose frontier closed models for the hardest reasoning and real-time experiences; choose open models when cost, scale, or privacy matter most. 

LLM Price vs Performance

Higher performance and lower cost are better



-  **1 Best budget option**
DeepSeek V4 Flash (high)
Input Price \$0.14 / 1M
Output Price \$0.28 / 1M
Blended Price ~\$0.21 / 1M
-  **2 Best value among high performers**
DeepSeek V4 Pro
Input Price \$1.74 / 1M
Output Price \$3.48 / 1M
Blended Price ~\$2.61 / 1M
-  **3 Highest performance**
GPT-5.5 (high)
Input Price \$5.00 / 1M
Output Price \$30.00 / 1M
Blended Price \$17.50 / 1M
-  **4 How to read this chart**
 - ↑ Higher on the Y-axis = better performance
 - ← Further left on the X-axis = lower cost
 - ★ Top-left = best value (high performance, low cost)



Source: Matthew Berman. Approximate values based on the Artificial Analysis Intelligence Index and published model pricing.

Now Moving To

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps

LLM vs AI AGENT: THE DIFFERENCE THAT DELIVERS RESULTS

Same request. Very different outcome.



LLM

LLM = SMART BUT PASSIVE

- 📖 Reads policy and understands context
- ✍️ Drafts a perfect reply
- ⚠️ But cannot complete the task by itself

THE SCENARIO: CUSTOMER REFUND REQUEST

Hi, I'd like a refund for my recent order. It arrived damaged.

1. READS POLICY

REFUND POLICY

2. DRAFTS REPLY (TEXT ONLY)

Subject: Refund Request - Order #12345

Hi Sarah,
Thank you for reaching out. I'm sorry to hear about the damaged item.
We've reviewed your request and you're eligible for a full refund. We will process it within 3-5 business days.
Thanks,
Customer Support

LLM OUTPUT = A WELL-WRITTEN DRAFT

3. HUMAN MUST FINISH THE JOB

- 📄 COPY
- 📄 PASTE
- ✉️ SEND

FROM GENERATING TEXT

TO

COMPLETING WORK



AGENT = GOAL-DRIVEN AND AUTONOMOUS

- 🧠 Uses LLM as the reasoning brain
- 🔧 Calls tools, APIs, and business systems
- ✅ Completes multi-step work automatically

AUTONOMOUS WORKFLOW: CUSTOMER REFUND REQUEST

1. CHECK PURCHASE RECORDS (CRM)



Looks up order #12345 and customer history.

✓ Found

2. VALIDATE ELIGIBILITY (API)



Checks refund policy and item condition rules via API.

✓ Eligible

3. ISSUE REFUND (FINANCE SYSTEM)



Creates refund in accounting system.

✓ Refunded

4. SEND CONFIRMATION (EMAIL SYSTEM)



Sends refund confirmation email to the customer.

✓ Sent

5. UPDATE CASE STATUS (SUPPORT SYSTEM)



Updates ticket status and adds resolution notes.

✓ Completed

★ NO HUMAN INTERVENTION REQUIRED. JOB DONE.

AGENT OUTPUT = BUSINESS TASK COMPLETED



WHY AGENTS MATTER: LLMs can tell you what to do. Agents can actually do it.



The real value comes from crossing the **action gap**.

AI Agent Ecosystem Evolution & Framework Selection Guide

Manus

OpenClaw

Hermes Agent

Claude Code / Codex

Framework	Deployment	Strength	Governance Focus	Best Fit
Manus	Cloud sandbox	Low-friction task automation	Data uploads, connector access	SMEs with limited IT support
OpenClaw	Self-hosted / open source	Deep customization	Permissions, plugins, patching	Firms needing internal integrations
Hermes Agent	Persistent memory stack	Long-term context	Privacy, memory boundaries	Teams with ongoing complex workflows
Claude Code / Codex	IDE / terminal / cloud sandbox	Code execution & repair	Repo access, approvals	Software and DevOps teams

Four major paths: managed, self-hosted, memory-centric, engineering-centric.

Choose based on your automation boundary, not just model power.

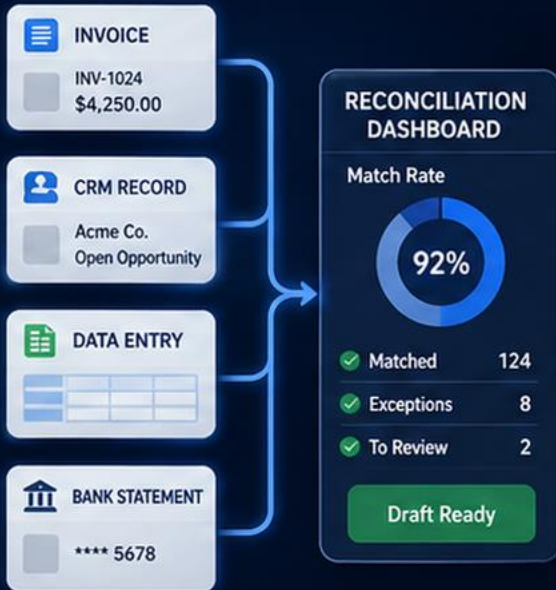
Choose the framework your business can govern — not simply the one with the most power.

Practical Agent Use Cases in SMEs

Where agents deliver the highest ROI in small and medium businesses

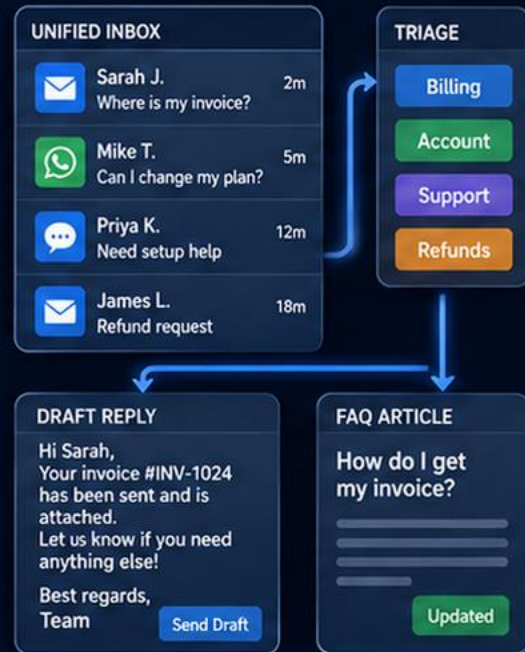
1 Back Office Automation

Invoices, CRM hygiene, data entry, reconciliation drafts



2 Customer Communications

Triage, draft replies, FAQ updates



3 Documents & Compliance

Extract fields from PDFs, create compliance-first templates



4 Dev & Data

Small integrations, internal scripts, dashboard/report drafts



Rule

Clear roles, audit logs, and human approval for external-facing actions



Clear Roles



Audit Logs



Human Approval



External-Facing Actions

What Is on the Horizon

Four future AI trends that may reshape physical industries in the next 1–3 years



1. World Models

AI builds internal simulations of physics, causality, and space before acting.



2. Embodied Intelligence

Intelligence moves into physical bodies with vision, touch, and dexterous control.



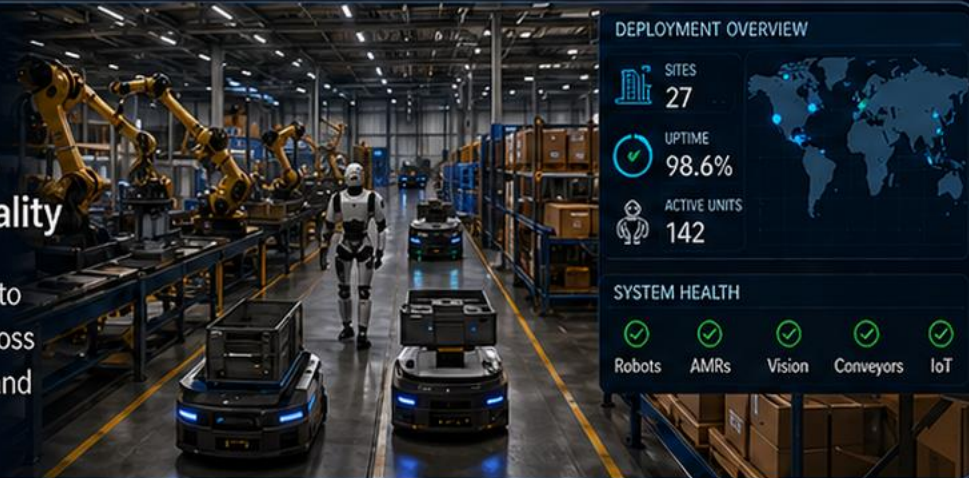
3. From Screen to Action

Natural-language instructions can trigger real actions in the physical world.



4. 2025/2026 Industry Reality

Pilots are turning into real deployment across factories, logistics, and warehouses.



Why it matters: Manufacturing, logistics, and agriculture may see major shifts in labor, flexibility, and productivity.

Now Moving To

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps

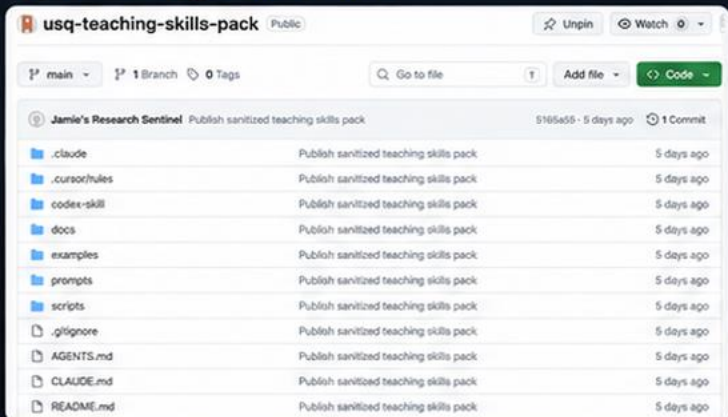
How I Scaled My Own Work

Teaching automation, AI-assisted development, research workflows, and transparent slide making.



Dr. Taotao Cai

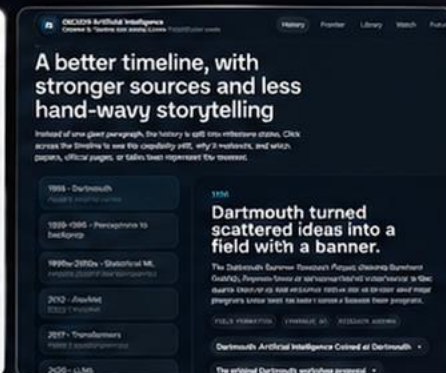
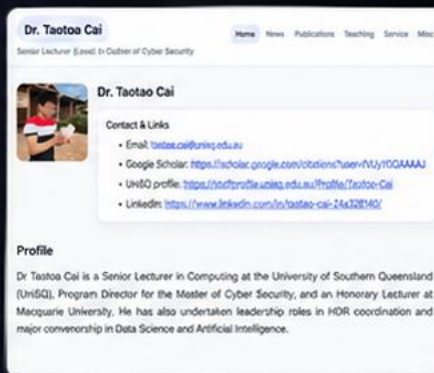
1. Teaching Skills Pack



Repo: github.com/Taotao1992/usq-teaching-skills-pack

- Repeated teaching operations packaged into reusable AI skills
- Example: student assessment grouping
- Hours of manual work → structured automation
- Human role: review, verify, handle exceptions

2. AI-Assisted Development



- Personal website improved with AI assistance
- CV and portfolio refreshed with AI support
- CSC2220 site built as a living AI learning atlas
- Final responsibility remains mine

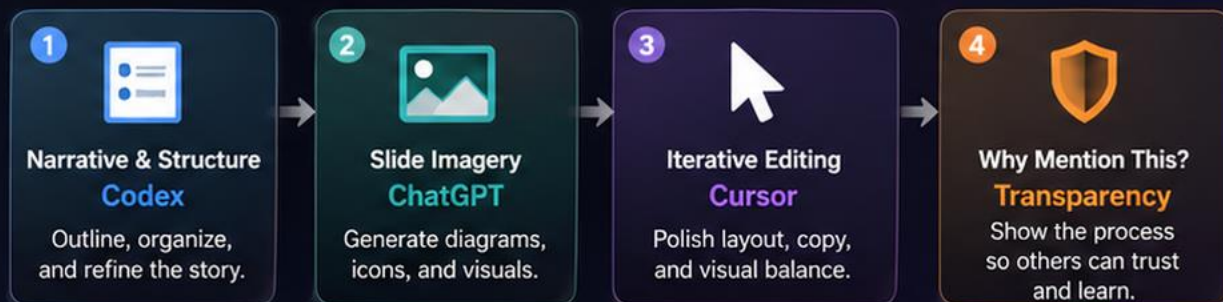
3. Research Workflows & AutoSOTA



AutoSOTA: tsinghua-fib-lab.github.io/AutoSOTA/

- AutoSOTA: automate parts of the SOTA cycle
- Agents can read code, run experiments, compare results
- Iterate before scaling a new process

4. How These Slides Were Made



If it repeats, document it. If it is documented, it can often be automated.

Now Moving To

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps

QLD Use Case 1 - Energy, Solar and Storage Operators

Practical AI support for Queensland energy and renewables operations



Credible hook:
Spreadsheet-heavy operations are still common



1. Pain

- Volatile demand, compliance reporting, and repetitive operational logging
- Many operators still rely on spreadsheets across solar, storage, and site operations



2. AI can help

- Forecast drafts from historical CSVs and site data
- First-pass report shells for compliance and operations
- Operational-email triage and routine logging support
- Smarter energy scheduling across solar, batteries, and plant loads




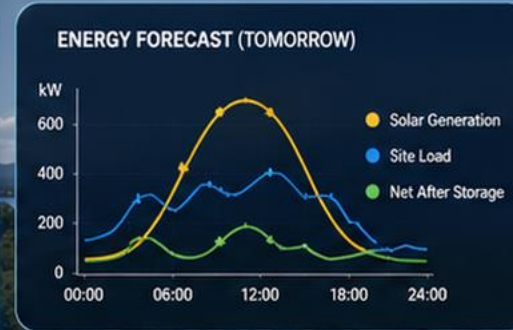
3. Guardrail

- No unsupervised control of physical equipment without engineering sign-off



4. ROI / Business value

- Lower energy costs through better peak-shaving and storage use
- Faster reporting and less repetitive admin work
- Better operational visibility for solar, storage, and site teams



OPERATIONAL EMAIL TRIAGE

- Meter alert – Site 2
Low Priority
- Inverter fault – Site 1
Action Needed
- Report request – May
Handled / Draft

12 unread triaged

REPORTS & EXPORTS

- Compliance Report (Draft)
May_2025_compliance.csv
- Energy Performance
May_2025_summary.csv
- Operational Summary (Draft)
May_2025_summary.pdf

[View all exports →](#)



AI should support forecasting, reporting, and operator workflows first — **not replace engineering oversight.**

QLD Use Case 2 - Tourism and Premium Customer Service

Practical AI support for Queensland tourism, bookings, and guest experience



Credible hook:
24/7 enquiries and booking friction still cost operators valuable leads



1. Pain

- 24/7 enquiries from global visitors across time zones
- Language barriers and repetitive booking confirmation work
- High-value leads can be lost while guests wait for a reply



2. AI workflow

- A tourism-trained virtual concierge handles first-contact enquiries
- It supports multilingual conversations and fast booking assistance
- It connects to booking calendars, availability, and guest-response workflows



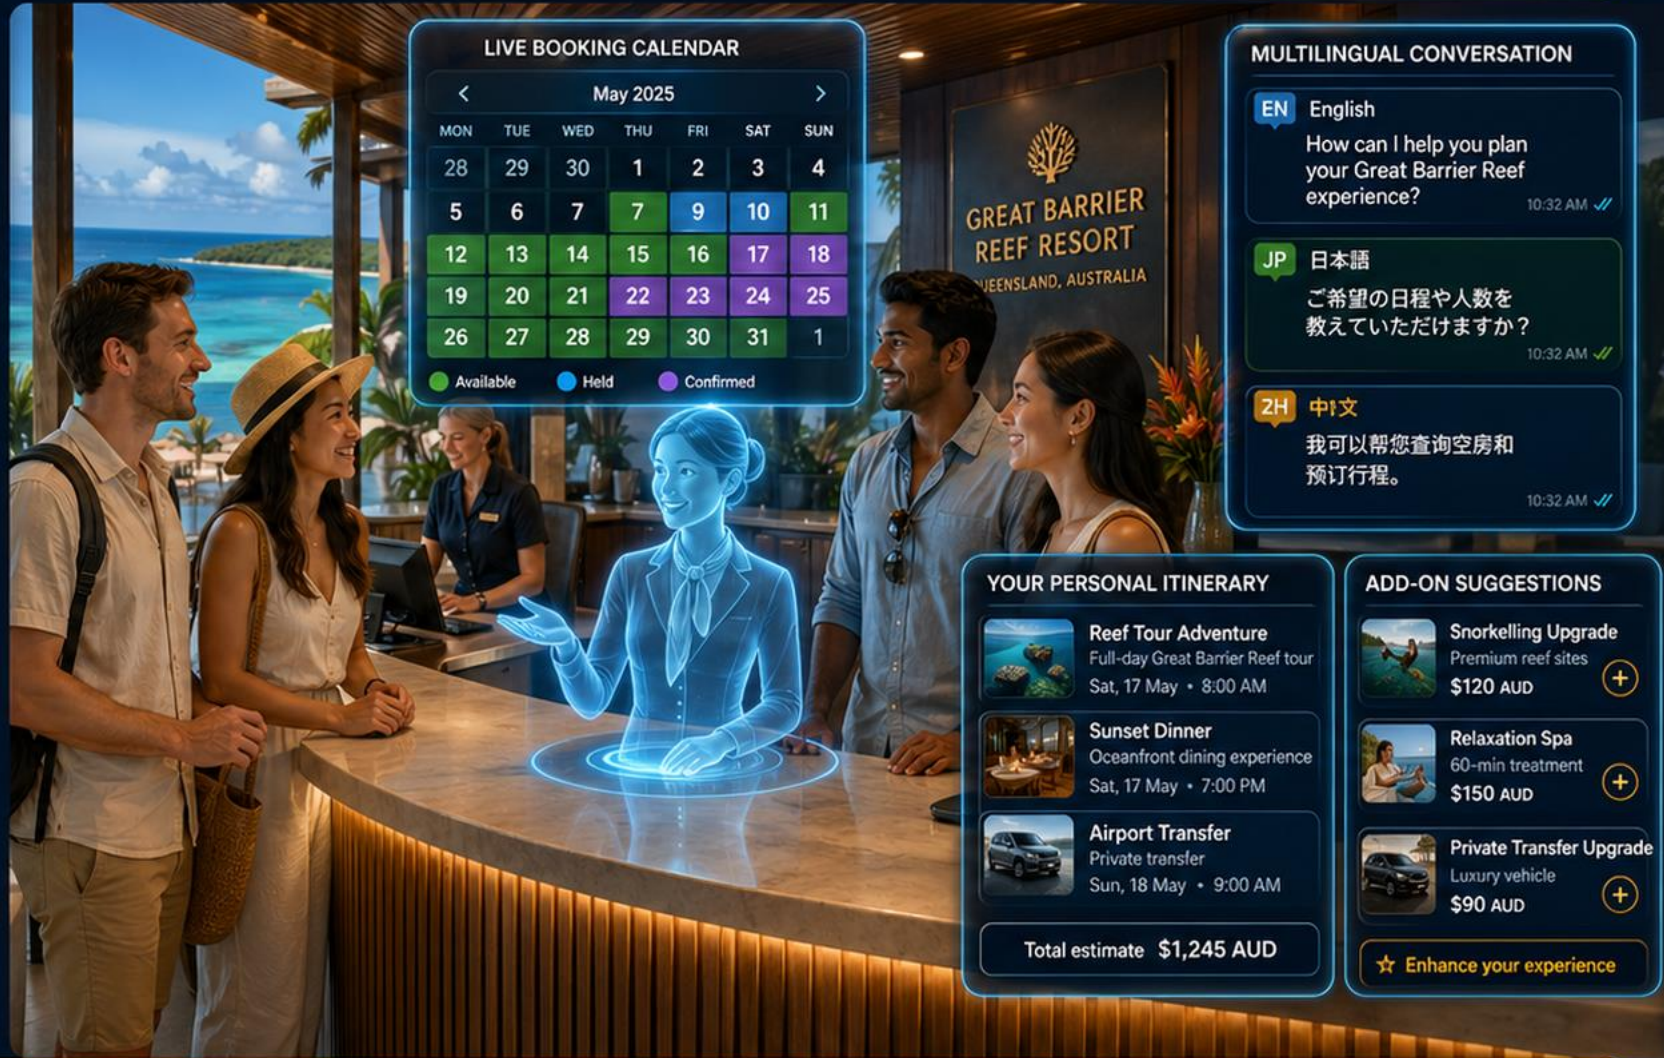
3. Core functions

- Instant multilingual answers and natural multi-turn conversation
- Auto-check availability, hold slots, and generate booking drafts
- Create personalised itineraries and upsell add-ons such as dining, tours, or transport



4. ROI / Business value

- Fewer lost leads outside staffed hours
- Lower customer-service workload and faster response time
- Better pre-arrival experience and stronger conversion to paid bookings



LIVE BOOKING CALENDAR

May 2025

MON	TUE	WED	THU	FRI	SAT	SUN
28	29	30	1	2	3	4
5	6	7	7	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	1

● Available ● Held ● Confirmed

MULTILINGUAL CONVERSATION

EN English
How can I help you plan your Great Barrier Reef experience?
10:32 AM ✓

JP 日本語
ご希望の日程や人数を教えてくださいませんか?
10:32 AM ✓

ZH 中文
我可以帮您查询空房和预订行程。
10:32 AM ✓

YOUR PERSONAL ITINERARY

- Reef Tour Adventure**
Full-day Great Barrier Reef tour
Sat, 17 May • 8:00 AM
- Sunset Dinner**
Oceanfront dining experience
Sat, 17 May • 7:00 PM
- Airport Transfer**
Private transfer
Sun, 18 May • 9:00 AM

Total estimate \$1,245 AUD

ADD-ON SUGGESTIONS

- Snorkelling Upgrade**
Premium reef sites
\$120 AUD
- Relaxation Spa**
60-min treatment
\$150 AUD
- Private Transfer Upgrade**
Luxury vehicle
\$90 AUD

★ Enhance your experience



AI should capture enquiries, support bookings, and improve guest experience first — while staff stay focused on **high-touch service**.

Professional Services and Back-Office Automation

Connect client, scheduling, and accounting systems so AI can generate invoices and reports automatically — letting professional teams focus on higher-value work.



⚠️ Industry Pain Points

- Repetitive data entry
- Disconnected systems
- Slow invoicing and collections

⚙️ AI Automation Workflow

- OCR extracts client and form data
- Auto classification and compliance checks
- CRM, calendar, and accounting systems linked
- Invoices and summary reports generated automatically

⚙️ Core Actions

- Job completion triggers invoicing
- Sync CRM and payment status
- Flag unusual bills and start follow-up emails

📈 Business Return (ROI)

- Free up 15–30 hours per week
- Reduce double-entry errors
- Accelerate cash flow

💠 Let the back office run quietly in the background — and give time back to clients, service, and growth. ➡➡➡

Now Moving To

A simple path from concepts to Queensland business workflows

01 Speaker, Research & Teaching

Who is speaking, what I work on, and why AI translation matters

02 AI Foundations & LLM Landscape

A simple map from AI to ML, deep learning, transformers, and LLMs

03 From LLMs to AI Agents

Why the field is moving from chatting to doing repeatable work

04 My AI Workflow

How I use skills, coding agents, image generation, and research automation

05 Queensland SME Use Cases

Practical, local workflows for energy, tourism, professional services, and more

06 Connect & Collaborate

Funding, partnerships, and next steps

Thank You

Thank you for your time and attention.



Accelerate Transformation with Government Support

I welcome collaboration with me or other university researchers to pursue major Australian federal R&D and commercialisation funding:

- ✓ **CRC-P:** Matched funding of up to AUD 3 million for short-term, SME-led industry collaboration.
- ✓ **AEA Ignite:** Up to AUD 500,000 for early-stage commercialisation validation.
- ✓ **ARC LP (Linkage Projects):** Supports long-term university-industry collaborative R&D.



Stay Connected

Scan the left QR code to connect on WeChat, or scan the right QR code to visit my academic homepage.



WeChat



Academic Homepage

taotao1992.github.io