

DESIGNING A HUMAN-AI INNOVATION EDGE FOR ENTERPRISES

**Strategy Use Cases
on Building "Trilingual Skills"
(Design x AI x Domain)
for the Future of Innovation**

Developed by



SINGAPORE UNIVERSITY OF
TECHNOLOGY AND DESIGN

Supported by



ACKNOWLEDGEMENTS

Crafting this set of Strategy Use Cases was a team effort. The publication was developed by the Singapore University of Technology and Design (SUTD), and supported by SGTech and SkillsFuture Singapore (SSG). As we write in the Executive Summary, it is the power of the community that made this possible.

SSG x SGTech x SUTD

SSG drives the national SkillsFuture movement to nurture a future-ready and resilient workforce. Their policies and programmes enable many of the upskilling and upskilling opportunities described in the Strategy Use Cases. SSG also provided guidance on the thrust and tone of this publication.

SGTech is the leading trade association for Singapore's tech industry. Representing over 1,400 member companies of all sizes, they brought keen insights into enterprise needs. SGTech also conceptualised the AI Impact Series mentioned in several Strategy Use Cases, that is now delivered in collaboration with SUTD. They too provided guidance on the thrust and tone of this publication.

SUTD developed the publication. SUTD is the world's first Design·AI university, building on its 17-year track record of nurturing design innovators, and over twelve years of research into the future of work and innovation. SUTD aims to nurture “trilingual talent”, who can combine Design x AI x Domain Expertise to innovate solutions that put humans back at the centre. The highly effective workshops run by the SUTD Academy are at the heart of many of the case studies.

Enterprises in Strategy Use Cases

We extend our deepest gratitude to all enterprises highlighted in the Strategy Use Cases, who have graciously and generously agreed to let us share their stories. Their stories are an inspiration to all of us.

AUTHORS AND CONTRIBUTORS

The publication was written by an internal team at SUTD together with a small group of external contributors for three of the use cases.

The internal SUTD team was drawn from its Future of Innovation and DesignZ endeavours, Design·AI Fab Lab/Building o, Lee Kuan Yew Centre for Innovative Cities, Business Excellence Office, and Office of Strategic Planning:

POON King Wang

Dinithi JAYASEKARA

HUANG Qian (Cathy)

LEE Ji-Eun

Candy MAK Yu Shan

Jose Paolo Tambagan BOQUIREN

Anieyrudh R

Titus Joseph Inquimboy RODRIGUEZ

Thijs WILLEMS

External contributors include:

Maria AIOLOVA (and her colleagues at SJ Group)

Kevin KHOO (Oral History Centre)

Azwar GHAZALI and Yean CHEONG (SGTech)

Contents

EXECUTIVE SUMMARY	1
<i>FUTURE OF ORGANISATIONS</i>	
NATIONAL: WORKING TOGETHER – HOW IT STARTED. HOW IT'S GOING.	8
SME: KHL PRINTING – UPSKILLING X INNOVATING FOR DAY-TO-DAY NEEDS	13
SME: VECTOR GREEN – DESIGNING AI THAT BUILDS ON DOMAIN EXPERTISE	17
LOCAL COMPANY: SIMPLYGO – EMBEDDING AI-DRIVEN INNOVATION	21
MNC: SJ GROUP – WHEN ARCHITECTS BUILT THEIR OWN AI DESIGN TEAM	24
MICRO-SME: ART LOFT – THE ART OF AI	29
GARDENS BY THE BAY – SEEDING SKILLS AND INNOVATIONS FOR A THOUSAND FLOWERS TO BLOOM	34
PUBLIC SECTOR: NLB ORAL HISTORY CENTRE – FROM STUDENT PROTOTYPE TO NATIONAL DEPLOYMENT	40
BUILDING AN ORGANISATION-WIDE AI AGENT AND CHATBOT ECOSYSTEM	45
NEW ROLE: THE FORWARD DEPLOYMENT DESIGNER (FDD)	50
LEADERSHIP: CAN AI REPLACE LEADERS?	56
STRATEGY: FORECASTING FAST VS SLOW	60

FUTURE OF INNOVATION

THE FUTURE IS BUILT, NOT TAUGHT	66
EJC: WHEN TEACHERS AND STUDENTS LEARNT AND INNOVATED TOGETHER	70
WSG X SUTD: INNOVATING WHAT'S NEXT FOR MATURE WORKFORCES	75
CITIZENS: ALL CAN INNOVATE	81
FUTURE OF CITIES: FROM CITIZEN PARTICIPATION TO CITIZEN CO-CREATION	87

Executive Summary

Introduction

Discussions about AI often focus on what AI can do. And what AI can do to people.

The Strategy Use Cases in this briefing focus instead on what **AI can do for people**. And what **people can do with AI**.

They show leaders ten decisions they can make to put humans at the heart of our age of AI.

1. Ensure no one is left behind, especially because all can innovate now.

From SMEs (e.g. KHL Printing; Vector Green) to large firms (e.g. SJ Group; Gardens by the Bay). From private sector to public sector (e.g. Oral History Centre). From the young to the not-so-young.

The 1,500 participants — across ages and career stages — in our workshops were all empowered to upskill and innovate with AI. Their AI adoption, proficiency, and confidence all rose by at least 30%. Creative self-efficacy rose in tandem.

How? By going beyond text prompting and chatbots. They experienced the full power of AI — from multi-modal AI to multi-agent solutions.

2. First let them learn to build. Then build to learn. *(And discover people's potential.)*

Using this full power, people who built AI solutions learned more, retained more, and changed their behaviour more than those taught about AI in a classroom. And once they could build, the building itself became the fastest way to keep learning.

Take Vector Green, an SME, for example. With a 41% gain in AI Confidence, the conversation about AI changed completely once employees saw they could build solutions to everyday problems. The motivation and plans to keep building followed naturally, and are already underway.



By building, we find better ways to learn about AI. By building, we discover what AI is fully capable of. By building, we discover how much more we are capable of.

3. Know when to work with AI, and when not to. (AI as Tool, Teammate, and/or Neither.)

Once we build with AI, we discover and develop a valuable new skill: judgement.

Judgement about when to use AI as a Tool (e.g. to execute a series of tasks). When to work with it as a Teammate (e.g. in a design process from conceptualisation to co-creation). And when to choose Neither — when the work demands we choose what humans value.

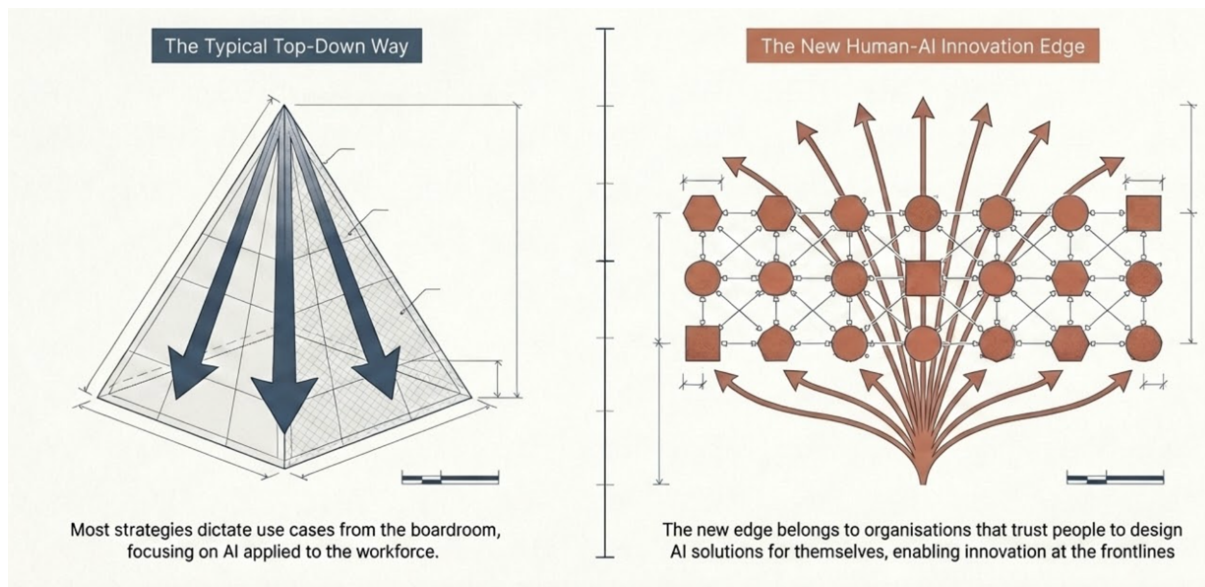
At KHL Printing, 67% of employees came to see AI as both Tool and Teammate, up from 56%. 22% saw it as Teammate/Collaborator. Only 11% still saw it as just a tool. That perceptual shift signals readiness for deeper AI integration.

And at Gardens by the Bay, employees developed the critical judgement to assess whether a use case is worth pursuing — and to challenge AI ideas that don't fit reality.

That judgement is what prevents wasted investment and raises ROI.

4. Trust the people closest to the work to innovate.

The employees who do the work every day know where the real problems are. Now that they are armed with innovation and judgement skills, they are better able than most to spot where AI helps, and to build solutions that work well.



At SimplyGo, employee-built prototypes led to automated compliance monitoring, resume screening, predictive IT operations, and expense tracking. They now have a formal AI Working Group governing the pipeline from prototype to production.

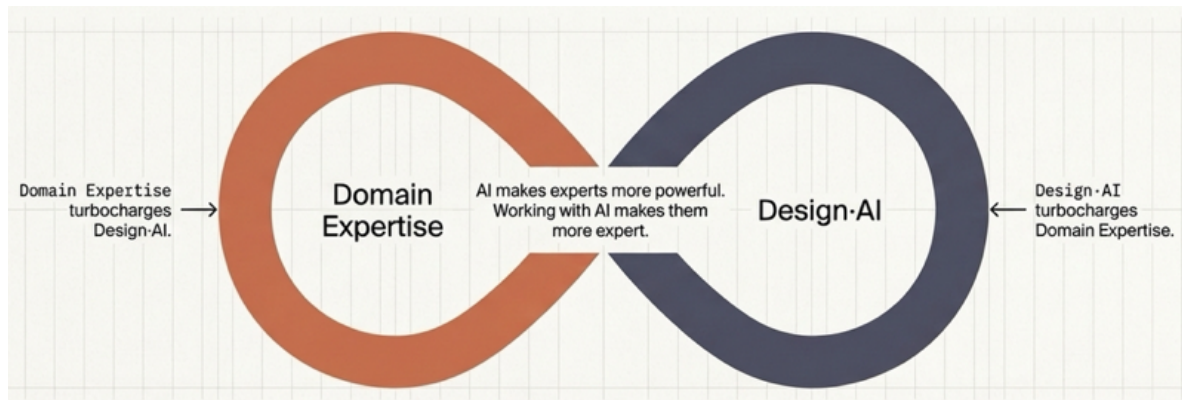
At SJ Group, the AI pipeline they built works because its builders understand architecture. From regulatory setbacks, block configurations, and site constraints, they know what is needed and what needs to be built.

Once we trust our people, innovation can come from anywhere and anyone.

5. Deepen domain expertise — because it multiplies Design-AI's power (and vice versa).

Trusting those who know the work best also means we must deepen their domain expertise.

Domain expertise turbocharges what people can do with Design and AI. Building on the SJ Group example earlier, it was the domain expertise that ensured each of the AI agents in their AI pipeline was accurate enough to trust.



What is true of MNCs is also true of micro-SMEs. Such as Art Loft. It was the founder-educator's decade of deep pedagogical practice that gave life and meaning to what AI can do for them. She has since designed six new AI-enhanced activities to improve the experience for students, teachers, and parents.

The reverse is also true. Design-AI turbocharges domain expertise just as much.

Designed well, AI makes experts more powerful. And working with AI makes them more expert. The two reinforce each other.

6. Scale confidence to act — because scaling skills without confidence stalls.

To scale such skills across an organisation and across society, we must also scale confidence.

Eunoia Junior College (EJC) first equipped a small core group of staff and students with AI skills. They then led the Learning Day sessions for the rest of the school. Throughout, care for confidence building was deliberate: a 1:5 instructor-to-learner ratio, platform-agnostic curriculum, and teachers and students learning from each other as peers.

As a result, staff rated their confidence to act at 5 out of 6. The high rating suggests the confidence-building worked; it also suggests the effort must continue.

The Lee Kuan Yew Centre for Innovative Cities at SUTD's Open Loop Career tools are another example. These tools were designed to empower Singaporeans for Open Loop Careers (i.e. continuous career transitions) via an SSG/Institute for Adult Learning research grant. They were tested and trialled with over 200 workers to ensure they built confidence.

In subsequent workshops for mature workers 50 years and older — What's Next: Reimagine Your Career With Design-AI — piloted in collaboration with WSG, 98% of the participants said AI was important for career exploration and transition. It signals confidence to act on AI, beyond just awareness of it.

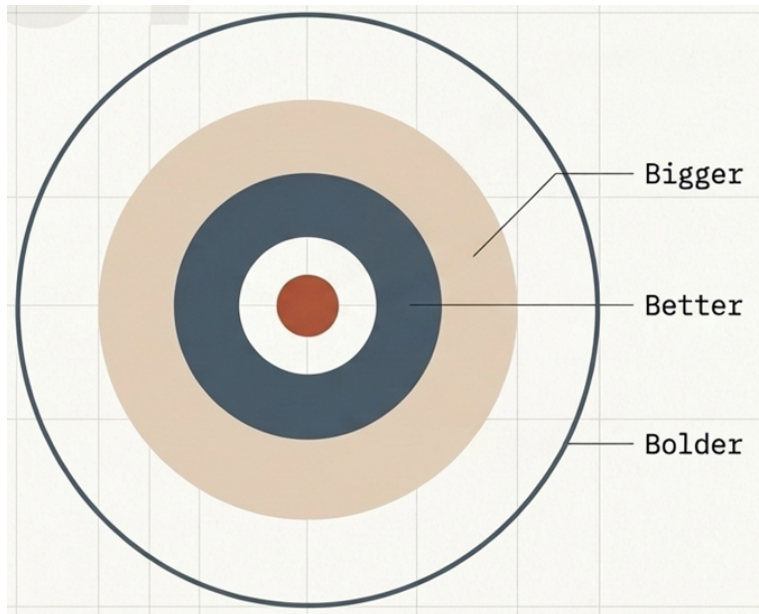
Care in scaling confidence is how we release the full power of skills.

7. Start small, then build better, bigger, and bolder.

Every single Strategy Use Case here started small.

Some are small by default because of size — such as Art Loft, KHL Printing, and Vector Green. They all achieved big jumps in AI confidence, assessed outcomes, and actual impact.

Some are as small as the individual citizen — a parent, a patient, and a Pilates instructor — using AI on ordinary day-to-day tasks to do the extraordinary.



Others chose to start small.

The Oral History Centre built on a small student project to design Interview Pro, now a national institution’s research tool.

Similarly, SJ Group’s AI capability started as a team of four.

And SUTD’s pivot to be the world’s first Design-AI university started with small teams too. Now different teams are building faster, higher, and stronger.

What matters less is whether you are small, and more whether you start.

8. Redesign the work so that humans do what humans do best — and let AI do the rest.

“My domain expertise is what makes AI useful. I know what matters on the ground — the institutional context, the constraints, what our visitors actually need. AI doesn’t know any of that without me.” ~ Gardens by the Bay participant

Now that employees are equipped with the full power of AI, creativity, and confidence, and empowered with trust, expertise, and judgement — what is next?

Organisations can now reimagine and redesign work with humans at the centre.

This can be in how we redesign human-in-the-loop AI-native workflows, as seen in several cases.

It can also be for leadership decision-making, such as in the C-suite Digital Twin, where AI plays different roles depending on the type of decision. For simple Yes/No decisions, we can leave it to the AI. For complex ones, AI is more a devil’s advocate and even muse.

Or in co-forecasting with AI, where we combine AI’s speed with human insights to compress forecasting cycles from months into weeks to keep pace with accelerating AI advances.

It works too for individual citizens. Each of the ordinary citizens profiled has found new ways to incorporate AI into how they educate children, train adults, and manage medical conditions.

It is all about knowing what we can do that AI cannot. And then working together, with each other and with AI (as Tool, Teammate, and/or Neither), to achieve the outcomes we care about.

9. Create new tasks, new roles, and new value.

One outcome we care most about is creating new tasks, new roles, and new value.

Why? Because leading economics research shows that AI is only unambiguously pro-worker when it creates new tasks.

Leaders must thus strive to ensure that they, their employees, and their organisations do so. It is possible, based on the evidence across all our cases and in all our workshops, that all can innovate.

	Labor-Augmenting	Capital-Augmenting	Automation	Expertise-Leveling	New Task-Creating
Productivity	✓	✓	✓	✓	✓
Value of Human Expertise	~	~	✗	~	✓
Pro-Worker Status	Ambiguous	Ambiguous	No	Ambiguous	Unambiguously Yes

To further raise the odds and opportunities for creation, leaders can also consider working with Forward Deployment Designers (FDD). The FDD is a new role born from an enterprise need identified by SGTech and SUTD. Many projects stall at the gap between what AI can do and what people need. FDDs bridge that gap. They uncover what people need, prototype possible answers, and get to a working solution sooner than later. In one case, a stalled project went from stuck to beta testing in three months.

Piloted and tested by SGTech and SUTD, the FDD can be of any age and at any career stage. All they need is deep expertise in at least one domain, willingness to roll up their sleeves to design and prototype with AI, and the heart to always put humans at the heart of it all.

New tasks. New roles. New value. Creating them is hardly theory — it is happening in Singapore now.

10. Invest in community — because it's not just about "me"; it's also about "we."

Many of the breakthroughs in this briefing happened because of community. In EJC, teachers and students worked together. For both small and large enterprises, they worked together across their internal teams.

The Future of Cities case study shows how we can take working as a community further. In these two years, we showed how AI can turn citizen participation into citizen co-creation. Together, human facilitators and AI can:

- level up and create a common baseline of understanding across citizens,
- give depth to their discussions through medium to high fidelity visualizations and prototypes, and
- elevate the quality of outcomes and even create working solutions

At the national level, these breakthroughs were only possible because of the individual investments made by different stakeholders. Whether it is in policy by SSG; enterprise mobilisation by Trade Associations and Chambers such as SGTech; or research, training, education, and innovation by institutes of higher learning such as SUTD.

No single organisation could have produced these results alone. The ecosystem they built together is what makes it work. Especially when we collectively put humans at the heart of it.

Conclusion

With these ten decisions leaders can make, we strive closer to a human-centred future.

The age of AI is here. How human it becomes is up to us.

THE FUTURE OF ORGANIZATIONS

CASE STUDIES IN THIS SECTION

1. **NATIONAL: WORKING TOGETHER – HOW IT STARTED. HOW IT'S GOING.**
2. **SME: KHL PRINTING – UPSKILLING X INNOVATING FOR DAY-TO-DAY NEEDS**
3. **SME: VECTOR GREEN – DESIGNING AI THAT BUILDS ON DOMAIN EXPERTISE**
4. **LOCAL COMPANY: SIMPLYGO – EMBEDDING AI-DRIVEN INNOVATION**
5. **MNC: SJ GROUP – WHEN ARCHITECTS BUILT THEIR OWN AI DESIGN TEAM**
6. **MICRO-SME: ART LOFT – THE ART OF AI**
7. **GARDENS BY THE BAY – SEEDING SKILLS AND INNOVATIONS FOR A THOUSAND FLOWERS TO BLOOM**
8. **PUBLIC SECTOR: NLB ORAL HISTORY CENTRE – FROM STUDENT PROTOTYPE TO NATIONAL DEPLOYMENT**
9. **BUILDING AN ORGANISATION-WIDE AI AGENT AND CHATBOT ECOSYSTEM**
10. **NEW ROLE: THE FORWARD DEPLOYMENT DESIGNER (FDD)**
11. **LEADERSHIP: CAN AI REPLACE LEADERS?**
12. **STRATEGY: FORECASTING FAST VS SLOW**

Transforming Lives:
From Upskilling to Innovation

DESIGNING A HUMAN-AI INNOVATION EDGE FOR ENTERPRISES

STRATEGY USE CASE

National: Working Together – How It Started. How It's Going.

SSG x SGTECH x SUTD: Integrating Skills, Tasks, and Design-AI Innovation to Build a National Upskilling and Innovation Ecosystem

CONTEXT

The pathway to resilience in the future of work is via the Future of Innovation. The Future of Innovation is the widespread ability of ordinary workers to conceive, create, and deliver new value with AI. Singapore is already building this. Through a decade of skills ecosystem development by **SkillsFuture Singapore (SSG)**, enterprise mobilisation by **SGTech**, and pioneering research and training at the **Singapore University of Technology and Design (SUTD)**, the country has assembled a national approach to human-AI innovation that is ahead of the global research and practice curve. Their individual efforts have now converged — and the combined evidence shows this approach works.

THE IMPACT AT A GLANCE

	Reach	Depth	Signal
SSG	105,000+ individuals trained	1,600 AI-related courses supported	72% workers and jobseekers want more training and guidance on how AI applies to their field
SGTech	111 SMEs transformed through AI Impact Series	439 professionals trained across industries	92% would recommend programme to peers
SUTD Design-AI	1,158 C-suite/senior leaders trained	+30% confidence, +11% creative self-efficacy	92% would recommend programme to peers
Combined Workshop Impact	AI adoption 30% more effective than research literature baseline	AI proficiency 88% more effective than research literature baseline	92–98% satisfaction across all programmes and audiences
Singapore Edge	+11% creative self-efficacy	+30% confidence with AI	78% expected productivity (+13% above baseline)

* Data as of end-February 2026. SGTech figures cover the AI Impact Series (July 2025–February 2026). SUTD figures include the Executive Programme for Leaders and Design-AI professional workshops. The Combined and Edge rows are based on measurements from the SUTD Design-AI workshop conducted in January 2026.

HOW IT STARTED

Landmark economics research has shown that for most of the 20th century, technological displacement of workers was offset by the creation of new tasks that favoured human labour. But since the late 1980s, that balance has broken down: automation has accelerated while new human task creation has weakened, leading to slower employment growth and a falling labour share of income [1].

A major new study published in February 2026 warns that AI is continuing this pattern. Of five types of technology it identifies, only **new-task-creating technologies are unambiguously pro-worker**, because they generate demand for novel human expertise and expand both the variety and value of work [2]. Yet current AI investment overwhelmingly targets automation.

	Labor-Augmenting	Capital-Augmenting	Automation	Expertise-Leveling	New Task-Creating
Productivity	✓	✓	✓	✓	✓
Value of Human Expertise	~	~	✗	~	✓
Pro-Worker Status	Ambiguous	Ambiguous	No	Ambiguous	Unambiguously Yes

Leading research have long argued that mass flourishing depends on many people empowered to develop and market new ideas, instead of remaining the province of a few isolated geniuses [3]. Leading research on working with AI reinforces this: centralised, top-down approaches — through IT departments, consultants, or innovation councils — are too slow and too removed from actual work. Individual workers who understand their own problems and can experiment freely are far more likely to discover powerful uses for AI [4]. The implication is clear: **the way forward goes beyond training workers to use AI — it means empowering them to innovate and create with it.**

A DECADE OF SKILLS INFRASTRUCTURE

Singapore is uniquely positioned because of the skills infrastructure it has built over ten years. Since the SkillsFuture movement launched in 2015, SSG has developed **Skills Frameworks across 38 sectors** as part of Industry Transformation Maps — co-created with employers, unions, education institutions, and government. These frameworks provide a granular national vocabulary of job roles, career pathways, and emerging skills that few countries possess.

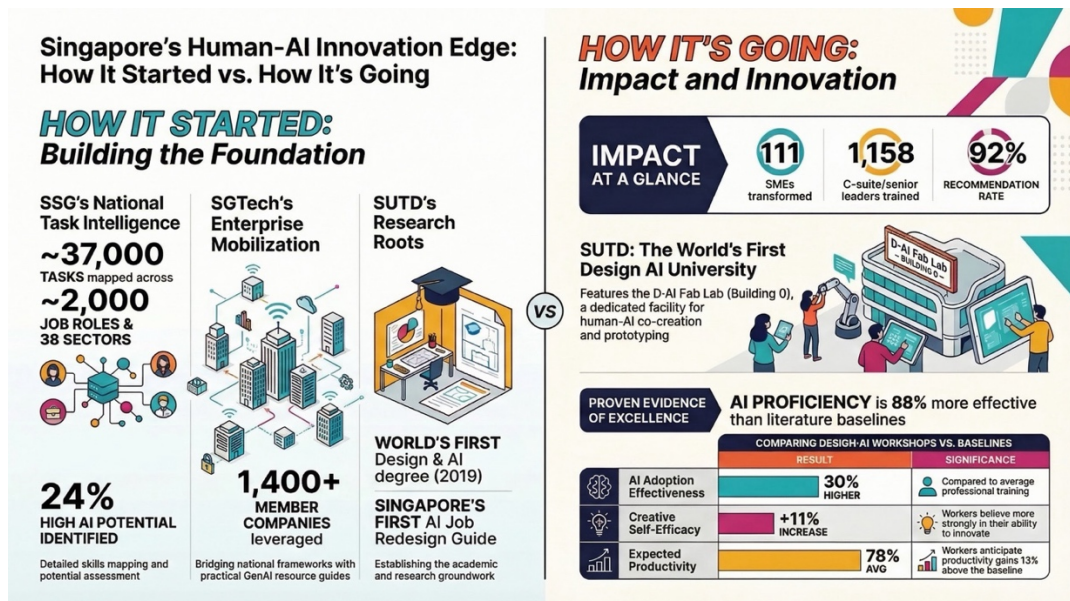
SSG's latest evolution takes this infrastructure from *skills* to *tasks*. A 2026 analysis of **~2,000 job roles** encompassing **~37,000 individual tasks** across all 38 frameworks found that roughly **24% of work tasks** have high potential for AI application. The new *AI Potential on Tasks* dashboard on the Jobs-Skills Portal provides task-level visibility into where AI can assist, and just as crucially, where human-centric tasks endure: decision-making, self-management, and learning agility. This shift from skills to tasks is exactly the analytical lens that the economics research calls for: understand displacement at the task level, then design reinstatement deliberately.

To further support Small and Medium-Sized Enterprises (SMEs) in navigating this transition, SSG has developed the SME AI Skills Launchpad in support of the National AI Impact Programme. The SME AI Skills Launchpad resource page, hosted on the SkillsFuture for Business Portal, provides a comprehensive collection of tools designed for every phase of AI adoption, including curated complimentary AI masterclasses developed by SSG's SkillsFuture QueenBee partners. These resources empower organisations to develop internal capabilities and accelerate their next stage of business transformation.

BUILDING THE ENTERPRISE BRIDGE AND LEADING BY EXAMPLE

SGTech, Singapore’s leading trade association for the tech industry with over **1,400 member companies**, was appointed by SSG as a **SkillsFuture Skills Development Partner (SDP)**. This role gave SGTech the mandate to champion skills development and drive skills adoption across the tech workforce. To accelerate AI readiness specifically, SGTech established the **AI Skills and Training Committee (AI-STC)**, a 26-member committee comprising government agencies, industry leaders, and Institutes of Higher Learning, working to drive and accelerate AI skills adoption across sectors.

In January 2024, responding to Singapore’s National AI Strategy 2.0, SGTech released the *Generative AI for the Tech Workforce* resource guide, developed in partnership with SSG and AI Singapore (AISG), with inputs from **30 companies** who pledged support for GenAI skills development. The guide was Singapore’s first publication articulating jobs and skills needs in GenAI, covering three business use cases, a tech talent profile with three tiers of GenAI expertise, and resources for business owners adopting GenAI. Targeted specifically at business owners planning to integrate GenAI into their transformation journey, it was a practical bridge between SSG’s national frameworks and what enterprises needed on the ground.



FROM RESEARCH AND EDUCATION TO PRACTICE AND INNOVATION

SUTD has been building the foundations for future of work and innovation well before the current AI wave. Since its establishment in 2009, SUTD has been nurturing talent in design innovation. And starting in the mid 2010s, it has been investing in future of work and innovation research. SUTD also started the world’s first Design and AI degree in 2019.

In 2020, SUTD’s Lee Kuan Yew Centre for Innovative Cities (LKYCIC) collaborated with IMDA and PDPC to publish Singapore’s first *Guide to Job Redesign in the Age of AI* – an industry-agnostic framework helping organisations redesign job roles to harness AI while increasing the value of human work. The Centre’s broader Future of Work research was recognised in Singapore’s National AI Strategy, and cited in SSG’s *Skills Demand for the Future Economy Report 2025* (p. 62) [5].

LKYCIC has also pioneered innovative career development tools, including an AI-powered tasks-skills stack that analyses tasks across sectors and generates job transition pathways. Their work in the last 5 years has extended into the future of innovation and creativity itself, focusing on strategies on how humans can flourish. For example, research has investigated how people and AI can co-create through “inside”,

“on”, and “outside the box” thinking — a practical framework for the kind of mass flourishing that the economics research envisions, now powered by AI.

In January 2025, SUTD announced a \$50 million pivot to become the **world’s first Design AI university**, specialising in design and AI across education and research for both undergraduates and postgraduates. Central to this pivot is the view that AI should no longer be treated as a tool but as a teammate – a concept SUTD calls “innovation beyond human imagination” – grounded in LKYCIC research showing that mastery in the AI age comes from knowing how to work with many people and machines. Amongst multiple initiatives, SUTD is weaving AI into all three terms of its first-year curriculum through a new course, *Innovating with Design and AI (iDeA)*, launching Singapore’s first **MSc in Design AI for mid-career professionals (DAI-E)**, and has established the **D·AI Fab Lab** (also affectionately called ‘Building o’) – a newly established facility dedicated exclusively to Design·AI objectives, where students and faculty explore, experiment, and shape the future of innovation. Distinct from SUTD’s existing fabrication labs, the D·AI Fab Lab is purpose-built for human-AI co-creation and rapid prototyping of AI-augmented solutions. Through an MOU with SJ Group, SUTD has also launched the **D·AI Fab Lab @ SJ**, Singapore’s first living lab in Design and AI bridging academia and industry, bringing together “trilingual” talent skilled in combining design, AI, and domain expertise to tackle complex urban, infrastructure, and environmental challenges. A second MOU was also signed to establish a **D·AI Fab Lab with RSAF**.

HOW IT’S GOING

MOBILISING ENTERPRISE

With the GenAI resource guide laying the groundwork, SGTech moved from advocacy to action. In July 2025, it partnered with SUTD to launch the **AI Impact Series** — a demand-led programme helping enterprises define practical AI business use cases and identify the skills their workers need. Where the guide had given business owners a map, the AI Impact Series gave them a vehicle: structured workshops that translate SSG’s national task-level data into concrete enterprise AI strategies.

Critically, as a Trade Association and Chamber, SGTech led by example by putting its own teams through the AI Impact Series workshops and developing organisation use cases that it is now working on with SUTD. To date, the AI Impact Series has made the following progress:

- **111 SMEs transformed**
- **439 professionals trained**
- **92% of the above would recommend the workshops to their peers**

THE CONVERGENCE: SUTD’S DESIGN·AI STRATEGY

These individual efforts — SSG’s national infrastructure, SGTech’s enterprise mobilisation, SUTD’s strategies — have now converged in SUTD’s Design·AI strategy. Design·AI goes further than treating AI as a tool for efficiency; it trains people to work with AI across a deliberate spectrum: **as a tool, as a teammate, and/or choosing not to use it**. This design-led approach, integrating proprietary research-grounded methods that blend Design × AI × Domain expertise, focused on what the research says matters most: new task creation — empowering anyone to innovate 10× to 100× (or more) through democratised creation with AI.

The results speak across for themselves (see AI Impact Series progress above). In addition, the SUTD Academy has also made strong progress in executive education:

- **1,158 C-suite/senior leaders and their teams trained**
- **92% of the above would recommend the Executive Programme to their peers**

THE EVIDENCE: DESIGN-AI PROFESSIONAL WORKSHOP OUTCOMES (JAN 2026)

METRIC	RESULT	SIGNIFICANCE
AI Adoption Effectiveness	30% more effective than literature baseline	Vs. average professional training effectiveness from overseas research
AI Proficiency Effectiveness	88% more effective than literature baseline	Effect sizes (Cohen's $d \geq 0.8$ = large) compared across training interventions
Creative Self-Efficacy	+11% increase	Workers believe more strongly in their own ability to create and innovate
Confidence with AI	+30% increase	Reduces the adoption friction that research warns slows progress
Expected Productivity	78% avg (+13% above baseline)	Workers anticipate even more meaningful productivity gains with AI

WHAT THIS MEANS

Workers who are 30–88% more effective at adopting AI, 30% more confident, and 11% more creative are precisely the workers equipped to create the new tasks that keep AI pro-worker. It is how the collective efforts of everyone puts humans at the heart of the ecosystem.

REFERENCES

- [1] Acemoglu, D. & Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- [2] Acemoglu, D., Autor, D. & Johnson, S. (2026). Building Pro-Worker Artificial Intelligence. The Hamilton Project, Brookings Institution. NBER Working Paper No. 34854.
- [3] Phelps, E. (2013). *Mass Flourishing: How Grassroots Innovation Created Jobs, Challenge, and Change*. Princeton University Press.
- [4] Mollick, E. (2024). *Co-Intelligence: Living and Working with AI*. Portfolio/Penguin.
- [5] SkillsFuture Singapore (2025). *Skills Demand for the Future Economy Report 2025*, p. 62. See also: SUTD LKYCIC & IMDA/PDPC (2020). *A Guide to Job Redesign in the Age of AI*.

SME: KHL Printing – Upskilling x Innovating for Day-to-Day Needs

KHL Printing x AI Impact Series: How 22 Employees Built the Case for Organisational AI Adoption

EXECUTIVE SNAPSHOT

61%

Average Improvement Assessed
Across productivity, creativity,
quality, and time savings

22

Employees Trained
Across two corporate runs
of the AI Impact Series

63%

Peak Time Savings Assessed
Reduction in repetitive task time

CONTEXT

KHL Printing Co Pte Ltd, a printing and packaging solutions provider, partnered SGTech and SUTD under the AI Impact Series to strengthen workforce AI capabilities and identify high-impact operational use cases. Twenty-two employees across two corporate runs designed and prototyped AI solutions aligned to day-to-day operational challenges using AI tools such as ChatGPT, Gemini, and Copilot.

KEY FINDINGS

- 61% average improvement assessed across productivity, creativity, work quality, and time savings
- 63% reduction in repetitive task time
- Post-workshop, 67% view AI as both tool and teammate; 22% see it as a true collaborator
- Multiple AI use cases prototyped spanning front, middle, and back offices; employees demonstrated increased confidence in applying AI both as a productivity tool and as a collaborative work partner
- Post-workshop deployment now underway with SUTD

The KHL Printing × SGTech × SUTD AI Impact Series represents a shift in how SMEs can approach AI adoption. Beyond treating AI as a centralised technology rollout or upskilling initiative alone, the programme empowered employees to design and prototype AI solutions directly applicable to their daily workflows. Twenty-two participants worked with the SUTD team across two corporate workshop runs to strengthen workforce AI capabilities and build functional prototypes targeting high-impact operational use cases.

Across the assessed use cases, participants reported an average improvement of approximately 61% across productivity, creativity, work quality, and time savings. These gains were assessed based on the

employees’ practical domain expertise of the existing workflows. Most solutions were built using AI platforms such as ChatGPT, Gemini, and Copilot, demonstrating that transformation does not require advanced AI capability – it just requires design experimentation and workflow insight.



Perceived AI Impact at Work

Survey scores on a 0–100 scale

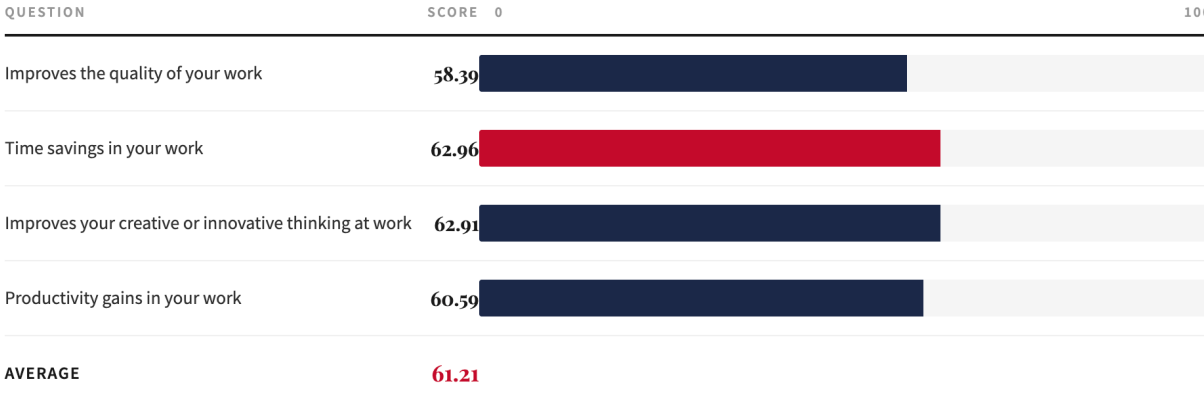


Figure 1: Perceived Impact of AI Prototypes on Work Quality, Productivity, and Creativity

TRANSFORMATION IN ACTION

Use Cases and Workflow Transformations

The programme significantly strengthened AI literacy and practical innovation capability among employees, enabling them to translate learning directly into workplace applications. The enterprise gains spanned multiple different functions.

Table 1: AI Prototypes Built by KHL Printing Employees — Workflow Shifts and Human Roles Retained

Function	Assessed/Potential Impact
Data Processing	91% quality, 60% time savings
Inbound Call Handling	Up to 88% time savings, 86% productivity
Sales Pipeline	75–90% quality, 85–92% creativity
Collections & Enquiries	86% quality, 70% time savings
Procurement	70% quality, 77% time savings
Payment Processing	70% quality, 71% time savings
Product Listings	50% quality, 49% time savings
Knowledge Sharing	80% time savings, 80% creativity

Table 1 presents the AI prototypes built by KHL Printing employees. The details have been summarised and/or redacted to protect confidential information. The diversity of use cases — spanning the front, middle, and back offices — reflects the breadth that AI can address when employees themselves identify the innovation possibilities. All prototypes retained a human-in-the-loop for decision-making, quality assurance, and judgement calls.

SHIFTING MINDSETS

Changes in Perception: From Tool to Teammate

The most significant finding was the measurable shift in how employees perceived AI's role in their work. Using the framework of viewing AI as a Tool, Teammate, and/or Neither, the programme tracked attitudinal changes across both workshop runs.

Prior to the second workshop, 56% of participants viewed AI as both a tool and a teammate, 33% perceived AI only as a tool, and 11% were unsure of its role. After Phase 2, perceptions shifted markedly towards more collaborative uses of AI: two-thirds (67%) now saw AI as both a tool and a teammate, and 22% viewed it as a true collaborator. Only 11% continued to see AI strictly as a tool.

How Do You See AI?

Perception shift from pre-workshop to post-workshop



Figure 2: Changes in Perceptions of Generative AI Before (left) and After (right) the Workshop

This perceptual shift matters because it signals readiness for deeper AI integration. When employees see AI not merely as a productivity shortcut but as a collaborative work partner, the organisation is better positioned for sustained adoption and cultural embedding of AI-enabled workflows.

THE TRILINGUAL FRAMEWORK

Building 'Trilingual' Skills and Tasks (Design x AI x Domain)

The gains observed were both the result upskilling in AI tools, as well as innovating with them, by combining AI with domain expertise and design. This "two-for-one" model – where training in skills and innovating via operational value creation occur simultaneously – is particularly powerful for SMEs with tight resources. The AI training programmes thus becomes embedded redesign laboratories where learning directly produces operational value through the innovations built.

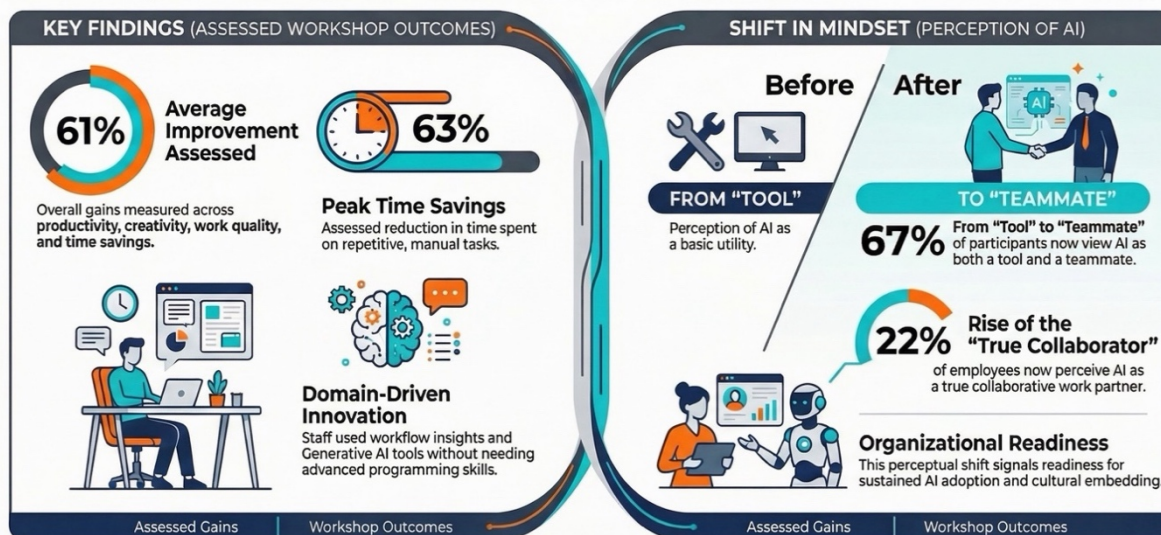
Post-programme, KHL and SUTD are working together to explore AI solutions to redesign multiple operational workflows, including automation of manual data processing and reporting tasks, AI-enabled call handling and customer enquiry routing, sales follow-up automation and lead-capture workflows, as well as using AI to improve administration work efficiency.

Table 2: Building 'Trilingual' Skills and Tasks

Skills and Tasks	Before (Pre-Redesign)	Now (AI-Embedded)
Design	Only a small number innovate.	All can innovate. Workflow-level redesign.
AI	Primarily manual execution e.g. human-driven data extraction, for manual, step-heavy processes in operations. Repetitive tasks embedded in front, middle, and back-office workflows.	Shift from manual completion to AI-assisted orchestration. Tool Mode: automation of processes; benefits typically more quantitative. Teammate Mode: strategy, analysis, and creative assistance. AI accelerates; benefits are more qualitative. Humans hold authority.
Domain Knowledge	Employees manually manage both execution and judgement. Knowledge often tacit and individual.	AI handles structured execution. Humans retain interpretive authority over critical client, relationship, and negotiation processes and decisions.

SME AI Innovation: Assessing the Impact of Employee-Led Redesign

KHL Printing partnered with SGTech and SUTD to train 22 employees to design AI prototypes for daily workflows, demonstrating how domain expertise, rather than coding skills, drives innovation.



SME: Vector Green – Designing AI That Builds on Domain Expertise

Vector Green x AI Impact Series: From Consumers of AI to Builders of AI as Part of Organisational Transformation

EXECUTIVE SNAPSHOT

41%

AI Confidence Improvement
After Phase 1

82%

Improved in AI
After Phase 2

100%

Transfer Intent
For use cases; and all reported
concrete plans to use AI at work

CONTEXT

Something Shifted in How They Talked About Work

Vector Green is a Singapore-based green energy company. Founded in 2020, they design and deploy solar photovoltaic systems, build digital energy monitoring platforms, and manage utility infrastructure for large-scale facilities. Their work spans engineering, project management, commercial operations, and compliance. It is technically complex, operationally demanding, and spread across a lean team of professionals who each carry deep but specialised knowledge.

In early 2026, as part of their organisational transformation, a group of those professionals went through the SGTech x SUTD AI Impact Series workshops at SUTD. The first built their comfort with designing with AI. The second challenged them to build working AI use cases for real business problems. By the end of the programme, something fundamental had changed:

“Once you see your own people build something that solves a problem they face every day, the conversation about AI changes completely.”

PHASE 1:

Building the Baseline

The first workshop raised the floor. Many participants had used AI casually, asking ChatGPT to draft an email, looking up quick answers. Few had a deeper understanding of what AI could do for their specific roles.

SUTD’s approach was hands-on from the start. Combining research- and evidence-based insights on how design and AI can amplify domain expertise (and vice-versa), the workshop guided participants through designing and building actual applications.

The results told a clear story:

- Confidence in using generative AI rose by **41%**
- Creative self-efficacy rose by **15%**
- Expected productivity with AI rose from **65% to 73%**.

Phase 1 gave a diverse group of professionals — engineers, project managers, commercial leads, HR, safety officers — a common language for thinking about AI and enough hands-on experience to know it was worth pursuing further.

PHASE 2:
From Baseline to Business Use Cases

The company returned for a second workshop. This session had a tougher mandate: take the baseline built in Phase 1 and turn it into business use cases. Participants identified business problems from their experience, and designed working AI prototypes by the end of the day. The improvements across four major skill domains (measured with matched pre- and post-survey data):

Table 1: Phase 2 Competency Gains Across Four Domains

Skills and Tasks	% Improved
Design	55%
AI	82%
Domain Knowledge	45%
Integration/Critical Thinking (i.e. combining Design x AI x Domain)	73%

AI competency continued to show strong gains. Specifically, the biggest single gain across all survey items within each of the four major domains was on “*understanding which business problems AI can solve*”, the most critical skill for anyone trying to apply AI at work.

For Design Thinking and Domain competencies, uplifts were also seen because the workshop activated existing knowledge, building on them by connecting them to the application of AI.

Integration/Critical Thinking skills — the ability to combine design thinking, AI knowledge, and domain expertise into workable solutions — saw the second highest improvement. The most notable gain was in the survey item “*spotting AI ideas that are technically interesting but not practically useful*”. Such critical judgement helps organisations assess AI utility versus AI hype.

It is worthwhile to note that **Integration/Critical Thinking are also Design skills**. Design researchers have a name for this. Herbert Simon’s foundational definition of design centres on exactly this kind of synthesis — combining diverse inputs into coherent, practical solutions. Roger Martin’s work on “integrative thinking” frames it as a core design competency. What the data shows is that participants improved in both explicit problem framing (the Design Thinking category) and this higher-order integration skill — the ability to bring disparate considerations together into viable solutions.

This matters for organisations because integration is the skill that turns AI knowledge into business value. An employee who understands AI but cannot connect it to real user needs and business constraints will

generate ideas that sound good but go nowhere. The workshops' design-led approach builds both capabilities simultaneously — and the data suggests the integration skill is where the most growth happens.

TRANSFORMATION IN ACTION

What They Built

During Phase 2, participants built working AI prototypes that addressed problems they encounter in their day-to-day roles:

- **A lead generation agent**
- **A contract review tool**
- **A pricing manager**
- **An operational dashboard**
- **A form automation system**

100% of the builds were solutions the participants themselves believed could be deployed.

THE REAL SHIFT:

From Consumer to Builder

Pre-workshop, AI use was passive and generic. Participants described using AI to draft emails, look up quick answers, and check the language in proposals. The prevailing mode was consumption: AI as a slightly smarter search engine.

Post-workshop, the language changed fundamentally. Participants described building systems, beyond using tools. They talked about training custom AI with their domain knowledge, automating workflows, and creating agentic applications. One participant noted that his inputs into could “help build the thinking.” Another recognised that domain knowledge “provides the guidelines and constraints that the tools may not know.”

This framing of domain expertise as an asset that makes AI more useful, not a skill made redundant by it, is the difference between asking AI a question and designing an AI-powered solution.

“Probably through the use of AI Studio, or agentic AI to automate some processes that are non-confidential and just time consuming.”

— Participant (on post-workshop AI integration plans)

What Comes Next

Post-workshop engagement indicated that the prototypes developed during the session will likely be used to support the redesign of multiple operational workflows at Vector Green. The programme has raised their AI baseline. The barrier to building with AI has dropped. The company now has a team of people who have upskilled in and innovated with AI. They understand that their own domain expertise is what makes AI solutions useful. The next step is to ensure organisation is set up to let them keep building.

From Consumers to Builders: Vector Green's AI Transformation



THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

The table below elaborates on this gain in ‘trilingual’ skills and tasks.

Table 2: Building ‘Trilingual’ Skills and Tasks

Skills and Tasks	Before (Pre-workshop)	Now (AI-embedded)
Design	Problem framing relied on intuition. Challenges described broadly without structured decomposition into testable use cases.	All can innovate. Participants translate challenges into testable AI use cases. Design Thinking reframed as iterative, flexible mindset. Prototyping used to explore ideas early.
AI	AI use was passive: drafting emails, asking questions, ad hoc research.	Participants designed and built working AI prototypes for their own workflows, moving from basic AI consumption to hands-on co-creation with no/low-code AI.
Domain Knowledge	Employees possessed tacit domain expertise but had not applied it in an AI context. Knowledge was individual and unstructured.	AI handles structured execution. Humans retain interpretive authority over ambiguity, exceptions, and stakeholder implications.

Local Company: SimplyGo – Embedding AI-Driven Innovation

SimplyGo x AI Impact Series: What Employee-Built Prototypes Reveal About Strengthening Future-Ready Workforce Capabilities

EXECUTIVE SNAPSHOT

**Productivity
Gain**

Across assessed use cases

**28
Prototypes**

By non-technical employees, low-code tools

**Peak Time
Savings**

Across assessed tasks

CONTEXT

As part of the SimplyGo x SGTech x SUTD AI Impact Series, Phase 2 placed AI capability directly in the hands of 28 employees. Participants were given time, facilitation, and low-code tools to prototype AI solutions for their own workflows, then present working solutions within the programme.

KEY FINDINGS

- On average, participants generally perceived that AI tools improved productivity and reduced the time needed to complete assessed tasks.
- Most prototypes built using low-code conversational AI — no programming required
- Fostered a strong culture of participant-led experimentation, while encouraging broader adoption of AI tools across departments—reflected in over 85% of participants indicating regular use of AI platforms.
- By combining workforce upskilling with rapid prototyping, AI-driven innovation is embedded as a core capability together with strengthening future-ready workforce capabilities.

SimplyGo provides integrated payment and ticketing solutions supporting large-scale transport and digital payment ecosystems. To accelerate workforce AI readiness and operational efficiency, SimplyGo participated in two corporate runs of the SGTech x SUTD AI Impact Series programme, where 28 employees designed AI prototypes addressing operational challenges such as report automation, customer query handling, compliance monitoring, and workflow automation using conversational AI and low-code tools.

The programme enabled participants to explore how AI could be applied to real workplace scenarios. Participants reported that AI tools helped them work more productively and complete assessed tasks more efficiently during the workshop. Beyond these perceived gains, the initiative fostered a strong culture of

participant-led experimentation, while encouraging broader adoption of AI tools across departments—reflected in over 85% of participants indicating regular use of AI platforms.

Most solutions were built using low-code conversational AI platforms, demonstrating that transformation can be achieved with them when paired with structured experimentation and workflow insight.

TRANSFORMATION IN ACTION

Use Cases and Workflow Transformations

Table 1: AI prototypes built by SimplyGo employees in Phase 2, showing workflow shifts post-redesign. The details have been large redacted to protect confidential information.

Business Function	Assessed/Potential Impact
Reporting & Analytics	50–80% perceived time savings (in structured cases)
Contract Review	Significant reduction in review time
Recruitment Screening	Faster candidate shortlisting
Customer Enquiry Routing	Faster routing, reduced response time
Developer Progress Tracking	Reduced monitoring overhead
Incident Reporting	Improved visibility, reduced reporting time
Compliance Monitoring	Improved consistency, qualitative enhancement
Vulnerability Management	Faster issue identification
Project Forecasting	Improved foresight (longer validation needed)
Data Cleaning	Approximately, 30–50% perceived time savings
Transaction Monitoring	Faster anomaly detection
Knowledge Management	Reduced search time

THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

The gains observed were the result of both upskilling in AI tools, and aligning AI with domain expertise and design to innovate simultaneously.

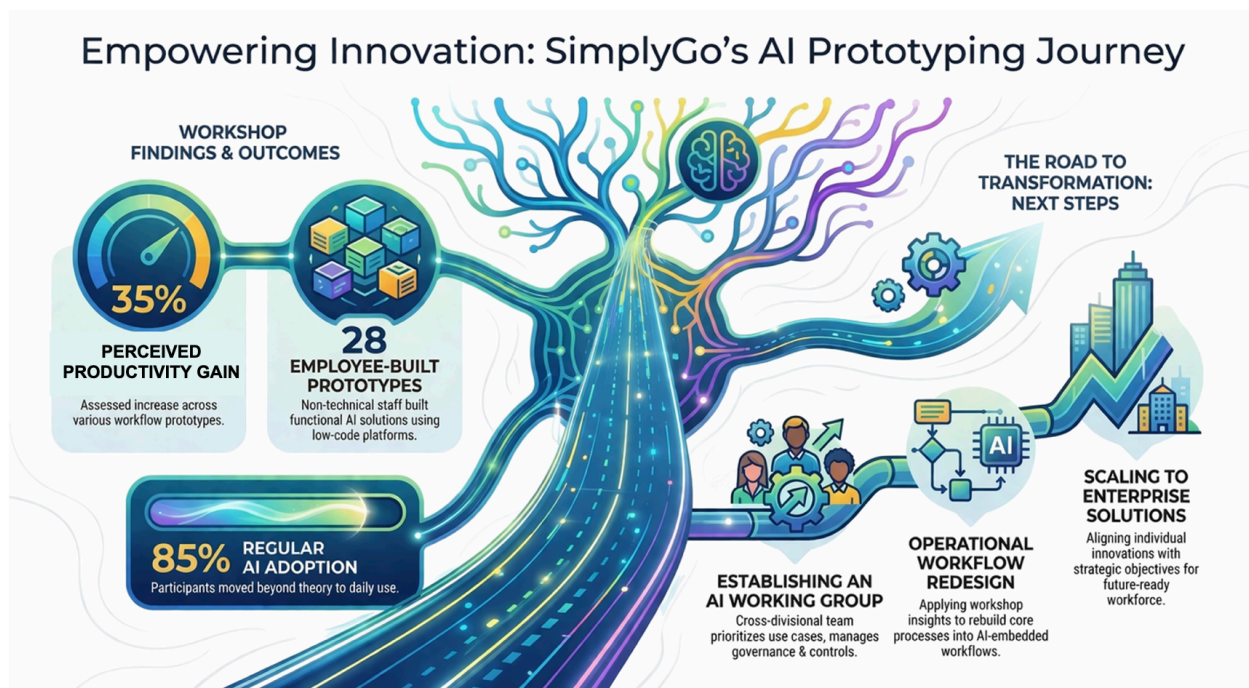
Post-programme engagement indicated that employee-developed AI use cases will likely be used to support the redesign of multiple operational workflows. These include automated compliance monitoring and transaction review processes, resume screening and HR workflow automation, predictive/agentive IT system operations as well as automated expense tracking, reporting, and analytics workflows.

To further strengthen these efforts, SimplyGo has established an AI Working Group. This AI Working Group is made up of representatives from all divisions, to capture and prioritise AI use cases with the appropriate governance and controls. This group serves as a focal point for aligning workforce-driven innovation with strategic business objectives, ensuring that promising prototypes can be scaled into enterprise-wide solutions.

Table 2: Building 'Trilingual' Skills and Tasks

Skills and Tasks	Before (Pre-Redesign)	Now (AI-Embedded)
Design Thinking	Only a small number innovate.	All can innovate. Workflow-level redesign.
AI	Primarily manual execution e.g. human-driven data extraction, review, comparison, and synthesis. Repetitive tasks embedded inside broader workflows.	Shift from manual completion to AI-assisted orchestration. Tool Mode: automation of structured, rule-based tasks; benefits typically more quantitative. Teammate Mode: compliance, advisory, predictive monitoring, forecasting etc; benefits are more qualitative. AI accelerates structured cognition but does not hold final authority.
Domain Knowledge	Employees manually manage both execution and judgement. Knowledge often tacit and individual. Cognitive load concentrated on staff.	AI handles structured execution. Humans retain interpretive authority over ambiguity, exceptions, ethical trade-offs, and stakeholder implications.

The above efforts have already demonstrated strong organisational value creation by supporting scalable AI-enabled workflows and operational efficiency improvements across functions. Together with other initiatives, they are expected to further accelerate SimplyGo's internal AI adoption journey. By combining workforce upskilling with rapid prototyping, AI-driven innovation is embedded as a core capability together with strengthening future-ready workforce capabilities.



MNC: SJ Group – When Architects Built Their Own AI Design Team

How Three SUTD Graduates Inside SJ Group Created AI Agents That Compress Days of Design Work into Seconds

EXECUTIVE SNAPSHOT

30 seconds

Per Design Iteration
Half-day tasks reduced to seconds in internal testing

5 AI Agents

Full Pipeline
Setbacks, blocks, site layout, narratives

< 1 Year

Student Project to Plugin
From SUTD academic project to SJ Group's AI Studio Plugin v1.3

CONTEXT

In Singapore's residential development sector, early-stage design is where the most consequential decisions are made — and where the most time is lost. When a housing developer issues a design brief, the architect typically has two weeks to analyse the site, calculate setbacks, generate unit configurations, produce site layouts, assess climatic performance, and prepare visualisations. Each stage is labour-intensive, iterative, and tightly interdependent. If the developer requests changes to the unit mix or orientation, much of the work starts again.

The tools architects rely on — Revit, AutoCAD, manual spreadsheets — were not designed for rapid iteration at the conceptual stage. Architects spend the majority of their early-stage time on tasks that are calculable, repetitive, and rule-governed, leaving too little time for the creative judgment and design thinking that differentiate one proposal from another.

THE ORIGIN STORY

From SUTD Classroom to SJ Group's AI Design Team

The work began under the SUTD–SAA DAI Programme, where an initial group of four DAI team members, along with SAA's mentorship, investigated how AI might automate early-stage architectural design for Singapore's public housing sector. Within four months, the team delivered working proof-of-concepts for the first stages of the design pipeline. Although the complete workflow could not be realised within the programme's timeframe, the results were compelling.

After the proof of concept, the effort was consolidated into a dedicated AI team of three — a software engineer, a machine learning engineer, and a computational designer — who joined SAA Architects officially to continue development beyond the academic setting. The Senior Executive Director of SAA Architects played a key role in defining the scope: SAA already had five established design teams. This would be the

sixth — but unlike the others, Team 6 would not compete for projects using human designers alone. It would develop AI-powered tools to augment every team’s capability, starting with early-stage design and feasibility studies.

A fourth team member — a backend engineer — was added to build the infrastructure and APIs. Today the team of four blends machine learning, backend engineering, computational design, and product management. What makes the team distinctive is that its founding members studied the problem at university and now sit inside the firm whose workflows they are trying to improve.

“The goal is to automate the mundane aspects of architectural work, freeing up time for architects to focus on creative and high-value tasks.”

WHAT THEY BUILT

The Team 6 Plugin

The AI Team developed a suite of modular AI agents that automate sequential stages of early-stage residential design. Each agent is named after a renowned architect — a concept introduced by the Senior Executive Director of SAA Architects to personify the tools and make their functions intuitive to design teams. The agents, collectively known as the Team 6 Plugin, cover four core stages of the workflow, plus two cross-cutting capabilities:

Alvar (Development Controls) fetches site data automatically and calculates setback envelopes from regulatory parameters. In internal testing, it replaces a manual process that typically consumed half a day with a 30-second automated computation.

Louis (Unit to Block Configuration) takes user-defined unit types and parameters — apartment sizes, mix ratios, corridor widths — and generates compliant block configurations. Produces in three minutes what previously required half a day of manual layout work.

Kenzo (Site Configuration) places configured blocks onto the site, respecting setbacks, orientation, and spacing requirements. Generates viable site arrangements in three to five minutes per iteration.

Antoni (Concept Narratives) generates written design narratives and concept descriptions to accompany proposals. Reduces narrative preparation from half a day to approximately 30 seconds. Uses SAA's Design Playbook to better structure its responses. [Design Playbook | SAA Architects](#)

Mies (Concept Visualisation) produces AI-powered architectural renders, a capability that did not previously exist at the early feasibility stage.

The team also built Rotom, a Photoshop plugin leveraging AI for rapid image editing and facade generation, which became one of the most widely adopted tools due to its immediate, transparent utility in fast-paced design tasks.

IMPACT

Early Results from Internal Testing

The tools are still in the experimentation stage and have not yet been deployed in live competition bids or client-facing projects. However, internal benchmarking shows significant potential. Each pipeline stage that previously consumed half a day now completes in seconds to minutes per iteration. The significance

lies not just in the speed of any single run, but in the ability to generate multiple design options in the time it previously took to produce one.

Table 1: Estimated Impact

Design Stage	Manual	AI-Assisted	Reduction	Agent
Development Controls	0.5 day	30 sec*	>99%	Alvar
Unit to Block Config	0.5 day	3 min*	>99%	Louis
Site Layout Config	0.5 day	3–5 min*	>99%	Kenzo
Concept Narratives	0.5 day	30 sec*	>99%	Antoni
Concept Visualisation	N/A	30 sec*	New	Mies

* Per iteration. Estimated based on internal benchmarking by Team 6. Not yet validated in live project conditions.

An architect using the full pipeline could explore ten or fifteen site configurations in a single morning, comparing climatic performance and unit mix across each, rather than committing to a single option early and refining it linearly. The next phase of validation will involve applying the tools to real project timelines and measuring impact under production conditions.

HOW THE PIPELINE WORKS

The system reflects a deliberate design principle: modularity with human-in-the-loop control. Each AI agent operates independently, meaning architects can engage at any stage without being forced through the full sequence. At each handoff point, the architect reviews and adjusts the output before passing it to the next agent. This is not a concession to caution — it is a design requirement. Architectural judgment is needed precisely at the boundaries: Is this setback interpretation correct for the site? Does this block configuration reflect the developer’s priorities? Is this site arrangement viable given features the algorithm cannot see?

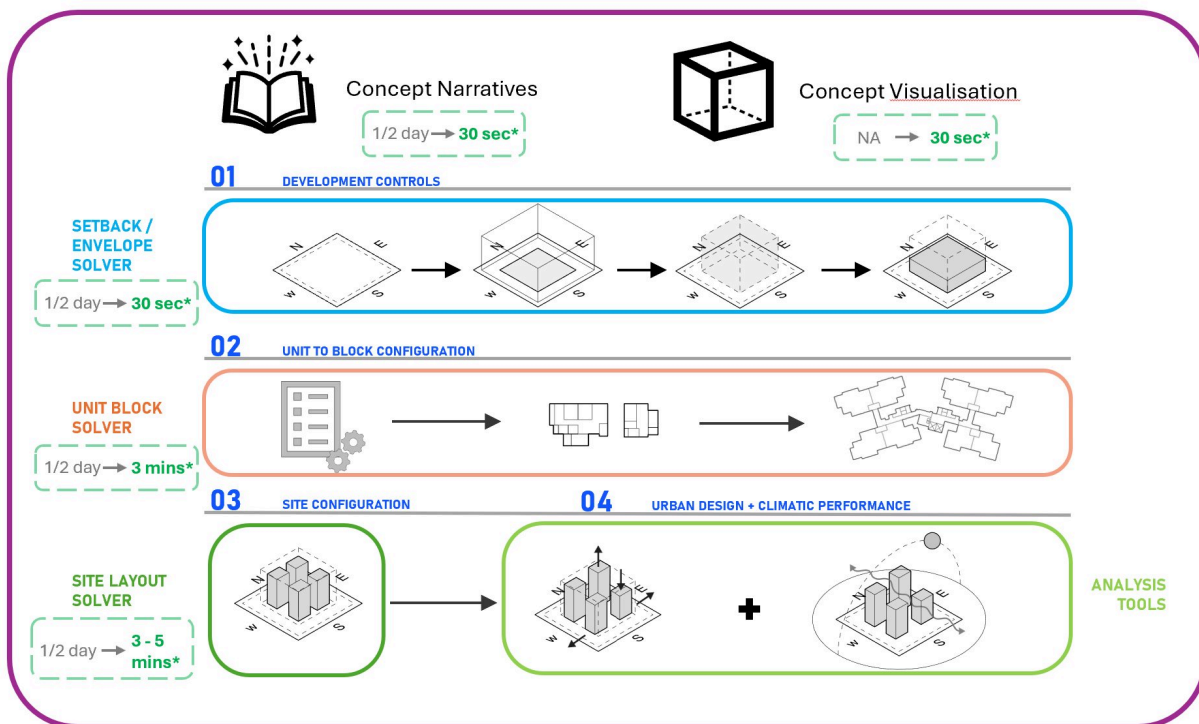


Figure 1: The SAA AI Studio Plugin pipeline — four sequential modules with estimated time savings per iteration. Each module allows human review before progressing to the next stage. All figures based on internal testing.

ADOPTION AND TRUST

Building the tools was one challenge. Getting architects to use them is another and work on this front is ongoing. Initial reactions from SAA's design teams were positive: the speed of iteration and quality of generated options were immediately compelling. But adoption has been uneven, and the reasons are instructive.

The strongest uptake came for tools with transparent, task-based outputs. Rotom, the Photoshop plugin, saw the highest adoption because its input-output relationship was immediately visible. The more complex agents — particularly those involving algorithmic site layout generation — faced more resistance. The concern was not about capability but about transparency: when an algorithm generates a site configuration, the architect needs to understand why that arrangement was proposed in order to defend it to a client or regulatory authority. This finding is shaping Team 6's roadmap: future iterations will prioritise explainability — surfacing which parameters drove each output.

“Adoption was highest for tools with transparent, task-based outputs. When architects could see the reasoning, they trusted the results.”

WHAT THIS MEANS

Team 6's origin carries a message for every architecture school and AEC firm: the barrier to building domain-specific AI tools has dropped. The team's founding members were DAI students when they built the first proof-of-concepts. The strategic question is not whether AI can handle early-stage design tasks — Team 6's internal experiments suggest that it can — but whether firms will build this capability internally, where it can be shaped by practitioners who understand the work, or outsource it to external vendors who do not sit inside the workflow.

For SJ Group, the next step is moving from experimentation to live project deployment — applying these tools under real competition timelines and measuring impact on win rates, proposal quality, and design team satisfaction. The research collaboration with SUTD, through the SJ–SUTD partnership on mapping AI-ready workflows in AEC, will provide the structured methodology needed to track these outcomes rigorously.

What started as a four-month university project under the SUTD–SAA DAI Programme became SJ Group's dedicated AI design capability. The tools are still experimental — but the team, the pipeline, and the proof of concept are real. As SJ Group moves from experimentation to deployment, the knowledge lives inside the organisation, not inside a vendor contract.

THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

Team 6's distinctive strength lies beyond any single skillset, and in the combination: people who understand both the architecture and the algorithms, who studied the problem at university and now sit inside the firm whose workflows they are trying to improve. The three dimensions that made this work:

Table 2: Building 'Trilingual' Skills and Tasks

Skills and Tasks	Without AI Integration	With AI Integration
Design	Sequential, linear refinement. Single option pursued. Rapid iteration limited by tool constraints (Revit, AutoCAD).	Multiple configurations explored in parallel. Human architects focus on creative and high-value judgment, not repetitive tasks.
AI	Manual execution of calculable, rule-governed tasks. Half-day work embedded in each design iteration.	Modularity with human-in-the-loop: AI agents handle routine computation; architects retain interpretive authority and control.
Domain Knowledge	Deep domain expertise required upfront to even propose a design. Expertise gatekeeps participation in early-stage ideation.	Domain expertise shifts from prerequisite to evaluative lens. Architects focus on validation and judgment: Is the setback interpretation correct? Does the config reflect priorities?

DESIGNING A HUMAN-AI INNOVATION EDGE FOR ENTERPRISES

STRATEGY USE CASE

Micro-SME: Art Loft – The Art of AI

Art Loft x AI Impact Series: Drawing on a Decade of Domain Expertise to Design AI That Is Meaningful for Children, Parents, Teachers

EXECUTIVE SNAPSHOT

< 10

People, One Studio

A micro-enterprise with limited resources

10+ Years

Domain Expertise

That the foundation AI was built on

See-Think-Wonder

A Framework That Works from Experience

Used to co-design with AI

HOW TO READ THIS BRIEFING

This briefing takes a different format from the others: it dives deep into the human experience. The left side from the second to fifth pages carries the actual conversation between a researcher and founder of a micro-enterprise (and experienced art educator) as they worked together to explore what AI could and could not do for a micro-SME. The *right margin annotations in red* draw out the larger themes — about foundations, and what makes AI meaningful for professionals and small enterprises.

CONTEXT

Who Is in the Room

Art Loft Singapore is a micro-enterprise. Its founder, Wong Seet Fun, is an artist and early childhood educator. For over a decade, she has built an art education practice grounded in inquiry: asking children open-ended questions, guiding them to observe closely, and treating the creative process as more important than the finished product.

Over several months, a researcher from SUTD's Future of Innovation Lab and the Lee Kuan Yew Centre for Innovative Cities worked with her through museum visits, observations, volunteer guiding sessions, and iterative AI design. Their goal was to explore whether AI could help micro-SMEs such as training museum educators to ask better questions during guided encounters with artworks.

What follows is drawn from that collaboration. It is a deep dive into the human experience of what two people discovered when they put deep domain expertise and AI in the same room.

01 — What AI Seems to Do Better

A quiet table in a museum café after a guided tour. The researcher opens a laptop and shows the educator a prototype of the custom GPT.

RESEARCHER

We've been experimenting with this tool for a while now. What surprised you the most?

FOUNDER AND MUSEUM EDUCATOR

The speed. When I upload an artwork image and ask the AI to generate inquiry questions, it produces ideas instantly. Sometimes ten different prompts in seconds.

RESEARCHER

So, it's useful for brainstorming.

FOUNDER AND MUSEUM EDUCATOR

Exactly. For example, when I asked it to generate questions using the See–Think–Wonder framework, it quickly produced prompts like: What do you notice in this painting? What might be happening in this scene? What do you wonder about the artist's intention?

RESEARCHER

That must make preparation easier.

FOUNDER AND MUSEUM EDUCATOR

Yes. It helps me start thinking. It's like having a research assistant that never gets tired. But it's only the starting point. Good guiding still requires human judgement.

It's like getting a rough sketch. I still need to refine it before using it with visitors.

02 — What Humans Do Better

Later, the researcher joins the educator in a gallery during a guided tour. Visitors gather in front of a painting.

RESEARCHER

What about the things AI cannot do?

FOUNDER AND MUSEUM EDUCATOR

Teaching is relational. When I guide visitors, I watch their expressions. I listen to their responses. Sometimes people hesitate before speaking. Sometimes someone notices something surprising. A guide has to respond to that moment.

RESEARCHER

So guiding is not just asking questions.

AI's advantage is volume and pace — but notice she says "ideas," not "answers."

See–Think–Wonder is widely used in art education to encourage careful observation. Her knowledge that this framework work came from her deep practice, not from the AI.

This is the prerequisite argument in miniature: AI can generate, but only a practitioner with years of experience can judge.

An art metaphor from an art educator. The sketch needs the artist.

This is the domain expertise that ten years of practice built — and that no prompt can replicate.

FOUNDER AND MUSEUM EDUCATOR

Exactly. It's about sensing the room. A question that works for one group may not work for another. Sometimes I change my approach completely depending on the audience. Children might focus on colours and stories. Adults might ask about symbolism or history.

RESEARCHER

That kind of flexibility must be difficult to encode into AI prompts.

MUSEUM EDUCATOR

Very difficult. Many guiding decisions rely on tacit knowledge built through experience. When to allow silence. When to encourage a quieter visitor. When to ask a follow-up question. AI can suggest questions. But it cannot fully capture the human dynamics of a gallery conversation.

AI can help with the what. But humans understand how it feels in the moment.

This is tacit knowledge — the kind built through experience, difficult to formalise, and impossible to shortcut. It is why the foundation matters.

This distinction — the what versus the how-it-feels — is the clearest statement of what AI cannot replace.

03 — What Changed When They Designed Together

After several rounds of testing, the researcher and educator review the evolving AI system.

RESEARCHER

Let's talk about the design process. How did we improve the system?

FOUNDER AND MUSEUM EDUCATOR

At first, the AI produced very generic questions. Things like: "What colours do you see?" "What is happening in this painting?" Correct questions — but not very engaging.

RESEARCHER

So we refined the prompts.

FOUNDER AND MUSEUM EDUCATOR

Yes. We added examples from real museum tours and structured the prompts using the See–Think–Wonder framework. Gradually the questions became better.

RESEARCHER

What did you learn from that process?

FOUNDER AND MUSEUM EDUCATOR

Designing the AI actually helped me reflect on my own teaching. I had to explain things I normally do automatically. Why some questions invite deeper discussion. How wording influences visitor responses. How curiosity develops through layered questioning.

Generic output is what you get when AI has no foundation to draw on. The improvement came from feeding it expert knowledge.

The AI improved because the educator's ten years of practice gave it something meaningful to learn from. More data would not have done this. Domain expertise did.

The virtuous cycle: the foundation made the AI useful, and the AI made the foundation visible. Neither works without the other.

RESEARCHER

So the AI became a tool for reflection.

FOUNDER AND MUSEUM EDUCATOR

Exactly. It helped externalise aspects of professional expertise that usually remain invisible. But it also showed us something important. Some parts of teaching cannot be fully formalised. They remain human.

This is not a failure of AI. It is a clear-eyed assessment of where human expertise begins and where AI support ends.

04 — The Questions We Keep Asking

RESEARCHER

Let's return to our three big questions. First: what can AI do better than humans?

FOUNDER AND MUSEUM EDUCATOR

Speed, variety, and information synthesis. It can generate many question ideas quickly and structure teaching frameworks efficiently.

AI as a productivity tool — useful, but not transformative on its own.

RESEARCHER

What can humans do better than AI?

FOUNDER AND MUSEUM EDUCATOR

Human presence. Understanding the audience. Responding to emotions, curiosity, and unexpected interpretations. Teaching is relational.

She keeps returning to this word: relational. It is the core of what ten years of practice taught her.

RESEARCHER

And what can humans and AI do better together?

FOUNDER AND MUSEUM EDUCATOR

When AI suggests ideas, educators can refine them using experience and judgement. AI becomes a learning partner, helping educators practise questioning and reflect on their teaching.

A learning partner — not a replacement, not even a tool. A partner that is only as good as the expertise it is paired with.

RESEARCHER

So the goal is not automation.

FOUNDER AND MUSEUM EDUCATOR

Exactly. It's augmentation. AI can support educators, but the spark of curiosity between people remains human.

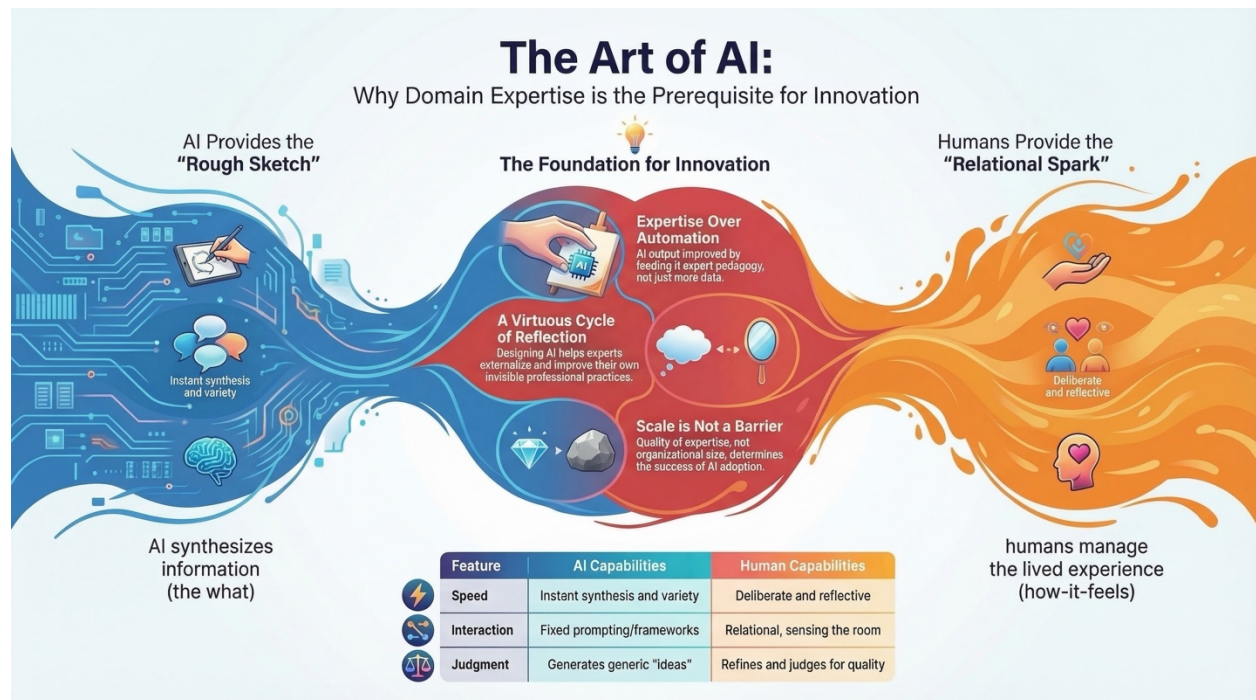
The spark of curiosity between people. That is what a decade of inquiry-based teaching was built to protect.

Art Loft is now using AI beyond museum education. Using her deep experience, she has designed 6 new activities and techniques that she is carefully and conscientiously iterating. These span possibilities in documenting students' learning journeys, transforming static records into personalised narratives that

capture creative processes over time, and tracking progress more precisely. Parents can also see more than just what their child made; they can also see how their child thought through it.

All these are only possible because the underlying pedagogy generates observations rich enough for AI to meaningfully process. A studio that taught children to copy paintings would have much less room for innovation.

The founder-educator herself is excited and energised as a leader. She knows she must take it upon herself to role model the possibilities bit by bit. It is a process that is both slow and exciting, and also both new yet familiar.



WHAT THIS GROUND CONVERSATION TELLS US

The foundation is the prerequisite. AI’s output improved not because of more data, but because one educator’s decade of practice gave it something meaningful to work with. Without the pedagogy, the See–Think–Wonder framework, and the years of reading visitors’ faces in galleries, the AI would have produced generic questions and no one would have known the difference. Organisations that skip the foundation and go straight to the tool may find that AI simply automates what was already shallow.

Small organisations are better positioned than they think. Art Loft has fewer than ten people, no tech department, and no innovation budget. Yet the depth of its practice gave AI much to work with. The barrier to meaningful AI adoption is less about scale, and more about the quality of the expertise you already have. And that possibility is open to everyone, not just large institutions, including a small studio. Where one educator spent a decade learning how to ask children the right questions. And then taught an AI to help.

REFERENCES

Huang, Q., & Wong, S.F. (2026 June). Co-designing Human-AI Collaborative Training for Museum Educators: An Autoethnographic Inquiry, Redesigning Pedagogy International Conference (RPIC) 2026, National Technological University

Gardens by the Bay – Seeding Skills and Innovations for a Thousand Flowers to Bloom

AI Capability Building: Evidence from Gardens by the Bay

EXECUTIVE SNAPSHOT

2×

Industry Average

Effect sizes double the published benchmark for workplace training

46%

AI Literacy Jump

From 2.6 to 3.8 on a 5-point scale in a single day

84%

Will Apply in 3 Months

Participants committed to using new AI skills at work

CONTEXT

Most AI training leave employees feeling inspired but rarely equipped. This study examines what happens when you flip that model: instead of teaching people about AI, you put them in a room with real business problems; integrate design, AI, and domain knowledge; and ask them to build AI solutions from scratch.

Across two workshop sessions at Gardens by the Bay, the employees did exactly that. The measurable impact was roughly double what overseas published research says typical workplace training achieves.^{1,2}

The question is no longer whether to train employees in AI, but whether the training formats being used are effective enough to justify the investment. This case study provides evidence that a specific format — hands-on, problem-based, integrating design with AI and domain knowledge — produces results well above the industry norm.

“Before the workshop, AI was basically ChatGPT to me. After, I could see there’s a whole toolkit — and the real question isn’t ‘how do I use AI?’ but ‘which AI tool fits this problem?’”

— Participant reflection, synthesised from March session open-ended responses

¹ Arthur et al. (2003) meta-analysis of 15,000+ participants found typical organisational training produces effect sizes of $d = 0.60$ – 0.63 . Cohen’s d is a standard measure of how large an effect is; values above 0.8 are considered large. Both workshop sessions exceeded this benchmark substantially.

² Hattie (2009) uses $d = 0.40$ as the threshold for a meaningful educational intervention. Lacerenza et al. (2017) found $d = 0.63$ – 0.82 for well-designed leadership training.

KEY FINDINGS

- **AI skills jumped significantly in both sessions.** AI Literacy rose 46%³ and AI Confidence rose 31%.⁴ Both had very large effects by research standards.
- **The impact was roughly double** what typical workplace training achieves, compared against the largest published meta-analysis of organisational training.
- **Integration/Critical thinking improved alongside enthusiasm.** Employees got better at telling the difference between an AI idea that sounds good and one that would work.⁵ These integration skills — combining user needs, AI, and business constraints — are a core design competency, and showed the second-largest gains.

'TRILINGUAL' SKILLS (DESIGN x AI x DOMAIN)

Design: Framing real business problems before reaching for AI. Translating challenges into testable use cases. Crucially, the integration/critical thinking skill is itself a design competency.

AI: Understanding what AI can and cannot do. Writing effective prompts, fact-checking outputs, knowing which tools fit which problems. Literacy scores rose.

Domain Knowledge: Applying professional expertise to judge whether an AI solution works in practice. Participants described their domain knowledge as providing fact-checking ability, practical judgment, and the institutional context AI needs to produce useful outputs.

Across two sessions in January and March 2026, Gardens by the Bay employees went through SUTD's Design·AI Workshop. Each session used a different survey instrument, giving us two independent lenses on the same question: does this workshop change how employees think about and use AI?

THE CORE FINDING

Two Surveys, One Consistent Story

The survey results showed that:

- **AI Confidence rose 31%** (January) — the practical ability to use generative AI tools, from writing effective prompts to fact-checking outputs⁴
- **AI Literacy rose 46%** (March) — understanding what AI can do, which problems it suits, and what data it needs³
- **Both groups started in the mid-2s** on a 5-point scale and **moved well past the midpoint** — from below “Neutral” to near “Agree”

How large are these gains? Researchers used a standardised measure called an effect size to compare training programmes across different studies. The chart below shows the GBTB Design·AI Workshop's results against overseas benchmarks:

	Overseas Benchmark	GBTB Design·AI Workshop
Effect size range	0.60–0.82	1.23–1.86
Relative performance	1×	~2×

³ AI Literacy: Cohen's d = 1.38–1.86 across three items, all p < .001.

⁴ AI Confidence: Cohen's d = 1.23, t(19) = 7.09, p < .001. Matched pre–post pairs, n = 20.

⁵ Integration/Critical Thinking: Cohen's d = 1.25–1.66 across six items, all p < .001.

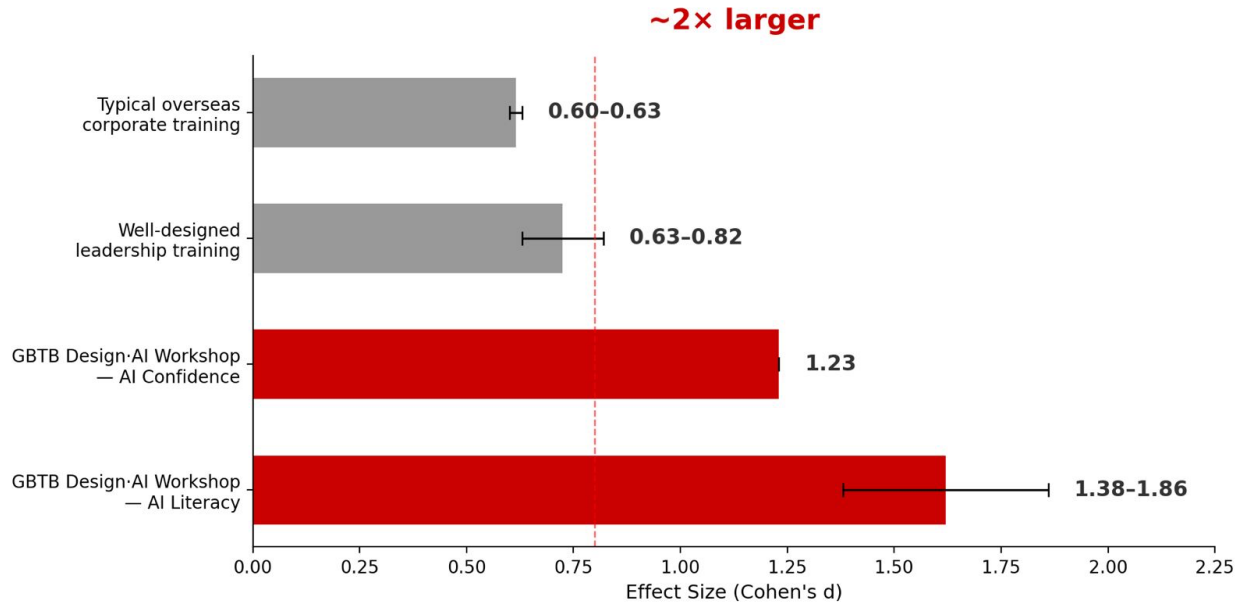


Figure 1: Effect sizes compared — GBTB Design-AI Workshop vs. overseas training benchmarks.

The GBTB Design-AI Workshop’s effects are roughly double the overseas industry benchmark. Both sessions cleared these benchmarks by a wide margin.

WHERE THE REAL GROWTH HAPPENED

Integration Is a Design Skill — and It Showed the Biggest Gains

The March survey measured 15 competency items across four categories: Design Thinking, AI Literacy, Domain Knowledge, and Integration/Critical Thinking. AI Literacy showed the largest absolute jump. But the second-largest gains — and arguably the most strategically significant — came from Integration/Critical Thinking.

Integration/Critical Thinking describe a specific set of abilities: connecting user needs with AI capabilities and business constraints. The March session showed large gains – the employees became more than AI enthusiasts, they developed judgement. Challenging an AI idea when it doesn’t fit reality. Assessing whether a use case is worth pursuing. Telling when something is technically interesting but not practically useful.⁵ Identifying implementation risks before committing resources. These are skills that prevent wasted investment.

Design researchers have a name for this. Herbert Simon’s foundational definition of design centres on exactly this kind of synthesis — combining diverse inputs into coherent, practical solutions. Roger Martin’s work on “integrative thinking” frames it as a core design competency. What the data shows is that participants improved in both explicit problem framing (the Design Thinking category) and this higher-order integration skill — the ability to bring disparate considerations together into viable solutions.

This matters for organisations because integration is the skill that turns AI knowledge into business value. An employee who understands AI but cannot connect it to real user needs and business constraints will generate ideas that sound good but go nowhere. The workshop’s design-led approach builds both capabilities simultaneously.

“I used to think of AI as something that does things for you. Now I see it as something you think with — a partner, not a replacement.”

— Participant reflection, synthesised from March session open-ended responses

“My domain expertise is what makes AI useful. I know what matters on the ground — the institutional context, the constraints, what our visitors actually need. AI doesn’t know any of that without me.”

— Participant reflection, synthesised from March session open-ended responses

WHAT THEY BUILT

The Proof Is in the Prototypes

This was beyond theory. The participants also created a diverse range of working AI applications.

“I didn’t know you could build a working app without being a programmer. We built one in a day — that changes what’s possible.”

— Participant reflection, synthesised from January session feedback

The numbers back up the energy in the room:

- **84%** of March participants said they would apply their learnings within three months
- **55%** estimated improvement from the AI solutions participants built (March average)
- **61% → 77%** expected AI productivity gain (January, pre to post workshop)

Creative self-efficacy improved significantly. Participants’ belief in their own creative problem-solving abilities rose across the board, with the largest gain in generating imaginative and original ideas.⁶ This is unusual for a skills training programme. It shows the workshop goes beyond teaching tools; it also unlocks creative confidence, the seed from which a thousand flowers grow.

Going forward, concrete plans for applying AI included presentation creation, document summarisation, data analysis, chatbot building, and fostering AI literacy within their teams.

THE TRILINGUAL FRAMEWORK

Why It Works: Design × AI × Domain, Not AI Alone

The workshop’s effectiveness is more than AI training alone. It integrates three skills that reinforce each other as shown above. The table below shows how each domain shifted from before to after the workshop.

⁶ Creative self-efficacy: +0.20 points, Wilcoxon $r = .47$ (medium-to-large effect), $p = .003$. The modest absolute gain likely reflects a ceiling effect — participants already rated their creativity above the midpoint before the workshop.

Table 1: Through the ‘Trilingual’ Lens — Design × AI × Domain Knowledge.

Skills and Tasks	Before Workshop	After Workshop
Design	Design described in traditional terms: aesthetics and user-friendliness. Problem framing seen as someone else’s job.	Design understood as a problem-solving framework. Participants translate challenges into testable AI use cases, with a large, statistically significant improvement. ⁷ Design extends beyond visuals to creating solutions, processes, and experiences.
AI	AI seen as a single tool (typically ChatGPT). Limited understanding of what AI suits, what data it needs, or how to verify outputs. Scores in the mid-2s on a 5-point scale.	Broader view: different AI tools serve different purposes, verification is essential, AI goes beyond text generation. Literacy scores near “Agree” (3.8+). ³ Participants build working applications — image recognition, chatbots, automated video tools.
Domain Knowledge	Professional expertise not connected to AI. Employees uncertain how their knowledge relates to AI adoption. Limited confidence in judging whether AI solutions would work in their context.	Domain knowledge recognised as what makes AI solutions practical and reliable. Employees describe it as providing fact-checking, better prompting, practical judgment, and institutional context. ⁸

“Design isn’t just about how something looks. It’s about framing the right problem first — and that changes everything about how you approach AI.”

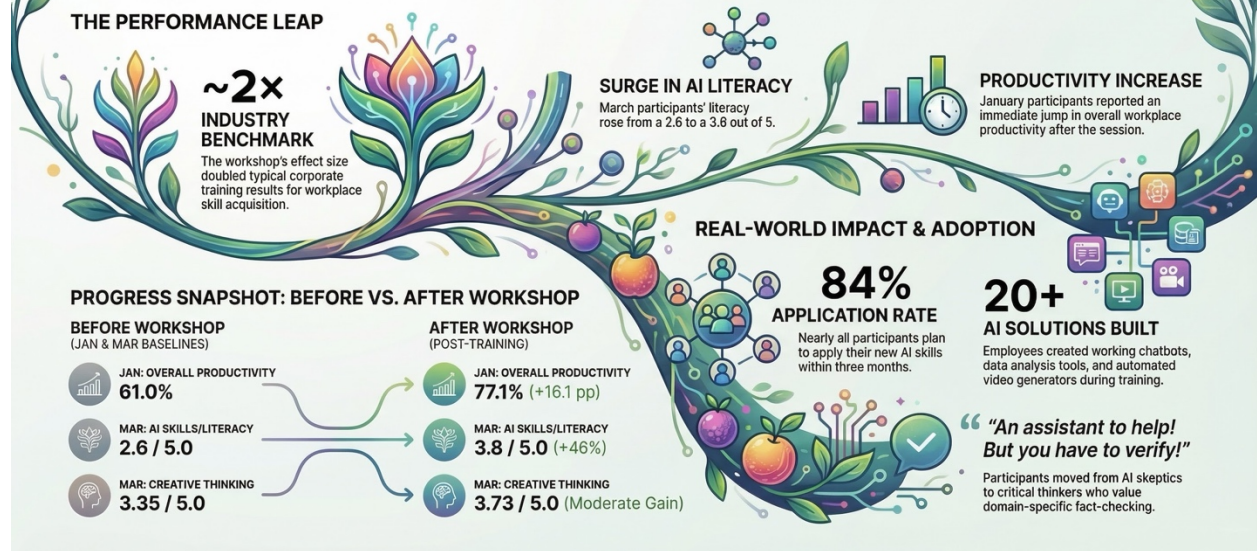
— Participant reflection, synthesised from March session open-ended responses

⁷ Translating challenges into testable use cases: +0.60 points, Cohen’s d = 1.08, p < .001

⁸ Judging if a solution works in practice: +0.61 points, Cohen’s d = 0.95, p = .001.

Gardens by the Bay: The Design-AI Transformation

Across two workshop sessions in early 2026, 49 Gardens by the Bay employees achieved skill gains doubling industry benchmarks. The training moved participants from a 'below-neutral' baseline to 'clearly confident,' resulting in dozens of working AI applications and significant projected productivity increases.



REFERENCES

- Arthur, W., Bennett, W., Edens, P. S., & Bell, S. T. (2003). Effectiveness of training in organizations: A meta-analysis of design and evaluation features. *Journal of Applied Psychology*, 88(2), 234–245.
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Lawrence Erlbaum Associates.
- Hattie, J. (2009). *Visible Learning: A Synthesis of Over 800 Meta-Analyses Relating to Achievement*. Routledge.
- Lacerenza, C. N., Reyes, D. L., Marlow, S. L., Joseph, D. L., & Salas, E. (2017). Leadership training design, delivery, and implementation: A meta-analysis. *Journal of Applied Psychology*, 102(12), 1686–1718.

Public Sector: NLB Oral History Centre – From Student Prototype to National Deployment

NLB Oral History Centre x SUTD: How a Classroom AI Tool Became the Oral History Centre's Interview Pro

EXECUTIVE SNAPSHOT

<h3>1 Month</h3> <p>Prototype to Demo Six students built Interview Companion from concept to showcase</p>	<h3>3 Months</h3> <p>Demo to Pro-version OHC extended prototype to a 4-function research tool</p>	<h3>Low Code, High Performance</h3> <p>Required to Build Both versions created as custom GPTs without programming</p>
--	--	--

CONTEXT

Interview skills are commonly taught through conceptual frameworks, written guides, and occasional role-plays. Students can often describe what good interviewing looks like before they can reliably perform it live. Authentic oral rehearsal is time-intensive, requires facilitation, and depends heavily on instructor availability. The result is a persistent Knowledge-Application Gap: students understand textbook principles yet feel underprepared when conducting real interviews. This gap is even more acute in specialised domains like oral history, where interviewers must navigate memory, emotion, silence, and cultural sensitivity – skills that cannot be developed from written guides alone.

KEY FINDINGS

- Students with no prior AI experience built a working voice-enabled practice tool in four weeks through learning-by-building
- The Oral History Centre (OHC) independently adopted and expanded the concept into OHC Interview Pro – validating the prototype's real-world transferability
- OHC's version restructured the instructional architecture: adding function-drift prevention, structured analytical frameworks, and domain-specific oral history

THE CORE INSIGHT

From Practice Tool to Research Infrastructure

The original system helped researchers conduct interviews better. The Pro version helps researchers design better research, conduct better interviews, and interpret the stories they collect. The upgrade was not about adding features – it was about restructuring the intelligence of the system

guidance

- The upgrade expanded from interview practice to the full oral history workflow: topic feasibility, outline review, mock interviewing, and transcript analysis
- The tool demonstrates a replicable adoption-adaptation pathway: classroom prototype → institutional demo → domain-specific deployment

KEY TAKEAWAY

“The most meaningful use of AI in skill development may not be content generation, but the creation of repeatable, human-interaction-focused practice systems that make high-quality rehearsal feasible at scale.”

THE PROTOTYPE

Building Interview Companion in the Classroom

Interview Companion was built over four weeks in March and April 2025 by six students enrolled in a graduate-level elective, Creating the Frontiers of a No Code Smart City, under the Master of Science in Urban Science, Policy and Planning programme at SUTD. Despite having limited experience with AI platforms and custom GPTs, the students’ skill development was accelerated through learning-by-building.

Two early design decisions guided the build. First, live mode voice role-play became the primary practice format, because voice rehearsal better approximates real interviews and reveals delivery issues that transcripts alone cannot. Second, qualitative coaching replaced scores and grades —the feedback focused on what to improve next rather than optimising a numeric rating.

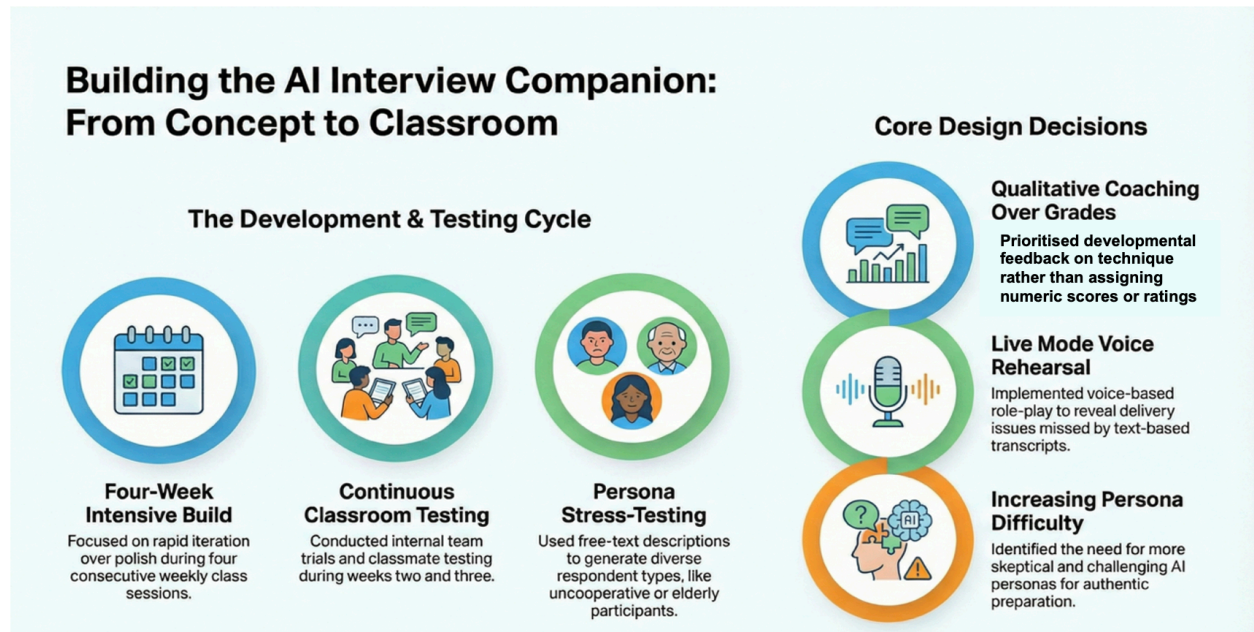


Figure 1. The design process: from concept to classroom deployment in four weeks.

The demo-ready version facilitated interview practice in three modes, each anchored to the user’s research question so feedback remains contextual rather than generic:

Interview Companion: 3 Ways to Master Your Skills

AI-powered tools for qualitative research excellence



Qualitative Transcript Feedback

Receive coaching on interview transcripts aligned with your specific research questions and guides.



Real-Time Live Role-Play Coaching

The GPT listens to live voice conversations between two humans and provides technique and delivery feedback.



Interactive Mock Interview Practice

Practice as the interviewer with the GPT as your interviewee to receive performance critiques and probe suggestions.

Figure 2. Interview Companion's three modes: transcript feedback, live coaching, and mock interview practice.

Iteration also revealed important limitations. Persona-driven practice increased learning value but exposed a ceiling: AI interviewees often remained polite and structured, while real interviews can be ambiguous, messy, or emotionally challenging. This became a key design insight for the next stage of development.

FROM PROTOTYPE TO PUBLIC SECTOR

The Oral History Centre Adopts and Transforms the Concept

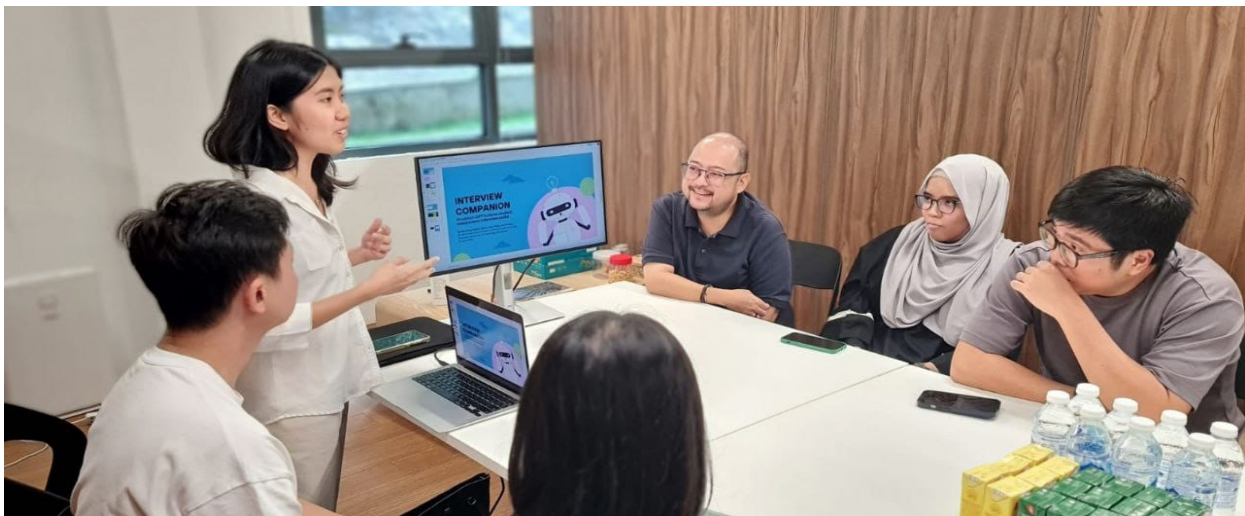


Figure 3. Students demonstrating Interview Companion to the National Library Board's Oral History Centre.

Interview Companion was subsequently made available under a Creative Commons Attribution licence and adopted by the National Archives of Singapore's Oral History Centre. What makes this case significant is what happened next: OHC adopted the tool, and went beyond, developing their own custom GPT, **OHC Interview Pro**, tailored to their institutional methodology and oral history practice. This next stage was made possible by the domain expertise of lead developer Grace Gui Ge, a postgraduate intern from the NUS

Department of History, together with Kevin Khoo, co-developer and project sponsor, whose experience as an Archivist at OHC helped bring the tool to a more advanced level.

As Kevin Khoo explained, the original system focused on improving interview techniques, helping researchers ask better questions, probe deeper, and maintain a professional conversational tone. But the OHC team recognised that interviewing is only one stage in a larger oral history process. Researchers often struggle earlier – when defining a research topic, structuring an interview outline, or interpreting the narrative that emerges from interviews.

THE UPGRADE

From Interview Practice to Full Research Workflow

FUNCTION	INTERVIEW COMPANION (STUDENT)	OHC INTERVIEW PRO (NLB)
Topic Feasibility	Not available	Evaluates oral history topics: academic framing, archive comparison, ethical challenges, scope refinement
Outline Review	Not available	Analyses narrative structure, flow, tone, cultural sensitivity, translation issues, and ethical red flags
Mock Interview	3 modes: transcript feedback, live coaching, mock interviewee with persona variation	Enhanced: memory-centric prompting, emotional context, strict character consistency, oral history interviewing guidance
Transcript Analysis	Basic transcript feedback aligned to research questions	Deep analysis: narrative structure, emotional turning points, positionality, historical themes, keyword extraction, bias detection
Architecture	Single-mode switching based on user requests	Structured instructional architecture with function-drift prevention, consistent analytical frameworks, reproducible analysis

WHY IT MATTERS

A Replicable Adoption-Adaptation Pathway

The journey from Interview Companion to OHC Interview Pro demonstrates a replicable pattern for how AI tools can move from education to industry deployment. The pathway has three stages: classroom prototype built through learning-by-building, institutional demonstration that sparks domain-specific interest, and independent adaptation that transforms the concept into a professional tool.

For educators, the instructor’s role shifts from being the primary source of feedback to being the architect of practice: defining rubrics, setting ethical boundaries, and calibrating what “good” looks like. AI makes repetition feasible. For institutions, custom GPTs and other AI artefacts become modular training assets that can be adapted to local contexts and professional needs.

The broader contribution goes beyond a single product. It is a transferable approach to designing AI-enabled rehearsal infrastructure. Similar applications may emerge for customer service training, HR interview preparation, and media rehearsal, and any domain where performance-based skills require repetitive, high-quality practice at scale.

A student prototype built in four weeks became a national institution’s AI research tool. The Oral History Centre did not just adopt Interview Companion — they transformed it. This is the adoption-adaptation pathway in action: what begins as a classroom exercise can become public sector infrastructure when the design is transferable and the domain expertise is encoded.

THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design × AI × Domain)

Skills and Tasks	Before AI	OHC Interview Pro (NLB)
Design	Interview skills taught through lectures, written guides, and occasional role-plays. Practice bottlenecked by instructor availability. Limited oral rehearsal opportunities.	Design extends beyond interview practice to the full oral history research workflow. Structured instructional architecture with function-drift prevention ensures each mode serves a distinct research stage. The system is designed as reusable institutional infrastructure, not a one-off tool.
AI	No AI support for interview training. Feedback depended on instructor availability. No way to rehearse live interview dynamics on demand.	Expanded to four functions covering the full research workflow: topic feasibility, outline review, enhanced mock interview (memory-centric prompting, emotional context, strict character consistency), and deep transcript analysis (narrative structure, bias detection, keyword extraction). Consistent analytical frameworks with function-drift prevention ensure reproducible results.
Domain Knowledge	Oral history expertise held by specialist archivists and senior researchers. Interview methodology learnt through apprenticeship, mentorship, and experience — not systematically transferable on demand.	Specialist oral history methodology encoded directly into the system’s instructional architecture. Domain experts (OHC archivists) shaped topic feasibility criteria, ethical red flags, cultural sensitivity protocols, narrative analysis frameworks, and memory-centric prompting. Expert authority preserved for all consequential decisions.

Building An Organisation-Wide AI Agent and Chatbot Ecosystem

Transformation: How SUTD Built a Portfolio of 15+ Purpose-scoped AI Spanning Every Major Function

EXECUTIVE SNAPSHOT

15+

AI Agents & Chatbots Deployed
Covering every major operational domain across SUTD — and growing

1

Unified Platform
Microsoft Copilot Studio + Teams, within SUTD's secure Microsoft tenant

∞

Portfolio by Design
Purpose-scoped, modular architecture — not one all-encompassing assistant

CONTEXT

SUTD has built one of the most comprehensive institutional AI agents and chatbot ecosystem in Singapore higher education. Rather than deploying a single all-encompassing assistant, the university has deliberately constructed a portfolio of 15+ purpose-scoped AI agents and chatbots — each aligned to a specific operational domain, each grounded in approved internal knowledge, and each governed by a named owner. This case documents that architectural strategy, spotlights SUTD Wiki as a representative example, and surfaces the design principles that have made the model replicable and sustainable.

KEY FINDINGS

- Shaping AI to fit institutions: rather than asking staff to adapt to AI, SUTD shaped AI tools to fit institutional work — reducing friction, improving consistency, and preserving human judgement at scale.
- Innovating new development models: The SUTD Wiki demonstrates a reusable student–staff co-creation model: students prototype, Business Excellence governs and deploys.
- Expanding with needs: The portfolio continues to expand — new domains are added as staff and departments identify use cases.

'TRILINGUAL' SKILLS (DESIGN x AI x DOMAIN)

Design: Innovating use cases with new AI skills.

AI: Tool mode throughout the portfolio. AI agents retrieve, surface, and structure institutional knowledge. No AI agent approves, decides, or acts autonomously.

Domain Knowledge: Institutional knowledge externalised into bot training by named domain owners. Expert authority preserved for all consequential decisions and escalations.

By 2024, SUTD's Microsoft Teams workspace had become a live directory of more than fifteen AI agents and chatbots, each purpose-built for a specific institutional function, each accessible to staff and students in the same platform they already used every day. There were bots for procurement questions, for campus wayfinding, for chemical compliance, for research ethics applications, for finance claims, for initiative reviews, for marketing support, for IT policies. And the list continued to grow.

This was not the result of a single top-down AI transformation initiative. It was the deliberate outcome of an architectural philosophy: build a portfolio of narrow, purpose-scoped agents and chatbots rather than a single all-encompassing assistant. SUTD's Business Excellence Office, which led the programme, recognised early that the challenge was not AI capability; it was institutional fit. What staff needed were tools that could respond to common queries using approved internal knowledge, within existing governance structures, without requiring new systems or significant change management.

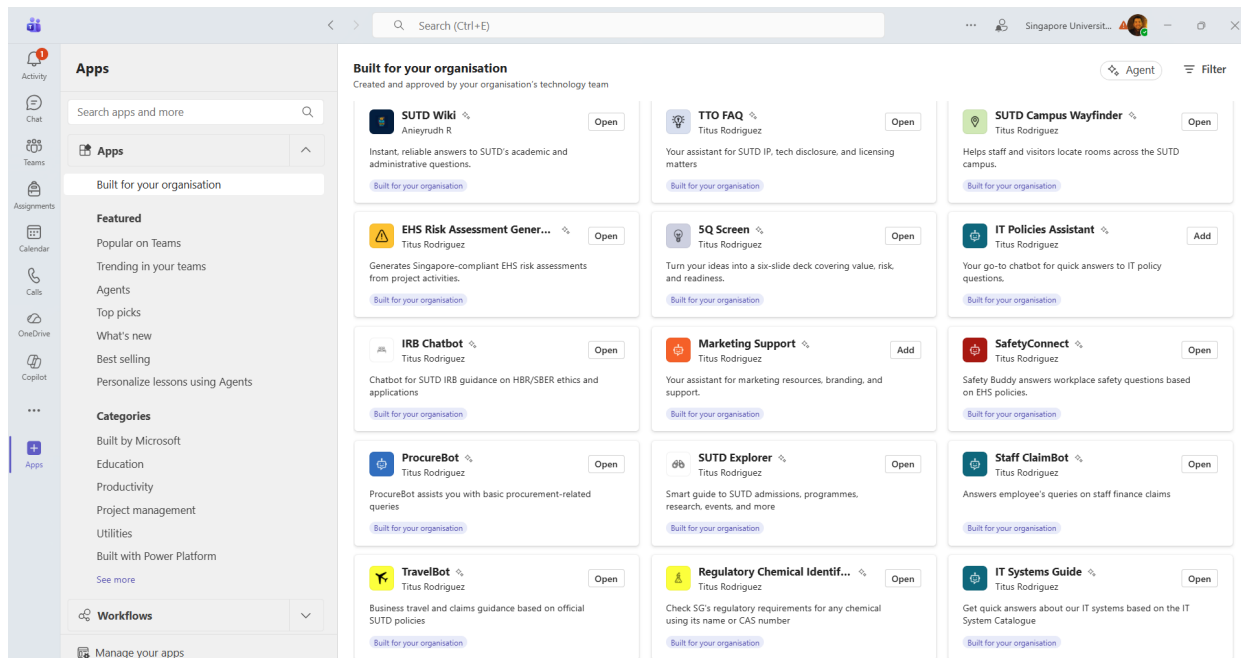


Figure 1: SUTD's AI agent and chatbot ecosystem in Microsoft Teams — the unified interface for 15+ AI agents spanning every major university function.

All agents and chatbots are built on Microsoft Copilot Studio and deployed within Microsoft Teams — the platform staff already use as their primary collaboration interface. Each agent references internal documents stored in dedicated SharePoint folders. All data remains within SUTD's Microsoft tenant, satisfying institutional governance requirements without requiring new infrastructure or vendor agreements.

Copilot Studio's architecture also enables future connectivity: individual bots can be linked to other Microsoft services (Outlook, Excel, Power Platform) when workflows require it, and can be connected to each other as cross-domain patterns emerge. The modular foundation makes growth natural rather than disruptive.

**THE STRATEGIC DECISION:
Portfolio Over Platform**

The most consequential decision in SUTD's AI journey was architectural. When the Business Excellence Office began exploring AI deployment in 2023, the tempting path was a single, general-purpose assistant — one interface, everything in one place. This option was rejected for reasons that are directly instructive for any institution facing the same choice.

A general assistant for a university must navigate procurement rules, safety regulations, HR policies, research ethics requirements, IT procedures, student services, and more. These domains have different vocabularies, different update cycles, different ownership structures, and different definitions of what constitutes a 'good' answer. A single bot would either be superficial across all domains or deeply accurate in one while unreliable in others. Neither outcome served SUTD well.

Design Principle	Rationale	Benefit
Clear Ownership	Each agent and chatbot is associated with a defined function and a named domain owner who maintains the knowledge base, handles escalations, and updates content as policies change.	Unambiguous accountability; faster iteration when edge cases or errors arise.
Domain Specificity	Procurement queries, initiative assessments, chemical compliance checks, campus wayfinding, and research ethics guidance require different vocabulary, logic, and forms of good output.	Each agent is calibrated to the realities of its domain — not to an average of all domains.
Ease of Maintenance	Smaller, focused agents and chatbots are easier to update as policies evolve. A redesign of one bot does not affect others. New bots can be piloted without disrupting existing ones.	Reduced fragility; incremental improvement without cross-contamination.
Modularity & Growth	Each agent or a chatbot is a component. Related bots can be connected in Copilot Studio over time as cross-domain work patterns are identified. The portfolio can grow organically.	No forced integration; new domains can be added as they are identified, without architectural rework.

A portfolio of narrow, purpose-scoped agents and chatbots does not represent a compromise on ambition. It represents a higher standard of institutional reliability.

The portfolio's breadth is itself significant. Every major administrative function of the university — from student services to EHS compliance to research governance — is now supported by at least one purpose-scoped AI tool. Staff no longer need to email functional teams for first-level policy questions in any of these domains. The agents and chatbots absorb the routine; the humans handle the consequential.

SPOTLIGHT

Illustrating the Portfolio in Practice

Spotlight — SUTD Wiki: Student Services & Co-Creation

The SUTD Wiki is distinctive in the portfolio — not only for what it does, but for how it was created. It was conceived as a conversational interface for students to access SUTD-specific institutional knowledge: policies, procedures, services, and administrative guidance. Students initially prototyped the Wiki using ChatGPT to explore what conversational access to institutional knowledge might look like in practice. Business Excellence then partnered with these students to deploy the Wiki on Microsoft Teams, making it available within an institutionally governed environment.

This student–staff co-creation model is itself one of the portfolio's most transferable innovations. Departments contribute real use cases and institutional knowledge. Students prototype, test, and refine solutions as part of their learning and contribution to the university. The result is rapid innovation that remains aligned with governance requirements. It is a model that SUTD intends to sustain as the portfolio continues to grow.

The screenshot displays a search result within the SUTD Wiki. At the top right, the search query "how to apply for hostel" is visible. Below it, the SUTD Wiki logo is shown. The main content area is titled "How to Apply for Hostel at SUTD" and includes a section for "Application Steps" with five bullet points: logging into the housing portal, clicking on the 'APPLICATIONS' tab, paying the application fee, completing the final step, and a note that incomplete applications will not be processed. Below this is an "Important Notes" section with five bullet points regarding address accuracy, application term sessions, personal details, session selection, and waitlisted students. An "Eligibility" section follows, stating that only full-time SUTD students are eligible. At the bottom, there is a "Message Copilot" input field with a plus sign on the left and a microphone icon on the right. A small disclaimer at the very bottom reads "AI-generated content may be incorrect".

Figure 2: SUTD Wiki — multi-turn conversation demonstrating natural language knowledge retrieval by a student user.

THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

Skills and Tasks	Before AI	Now (With Portfolio)
Design	Only a small number innovate.	All can innovate with new models of partnership (e.g. staff and students)
AI	No AI support for institutional knowledge retrieval. Staff relied on search, informal consultation, and personal memory, leading to delays and inconsistent results. Knowledge was fragmented across intranet pages, email threads, and individual staff.	15+ purpose-scoped agents and chatbots each own a discrete institutional workflow. Staff receive consistent, policy-aligned guidance on demand. Every chatbot has a named domain owner. Decision authority remains with humans. All operate in Tool Mode: they retrieve, surface, and structure institutional knowledge. None generate free-form recommendations or make consequential decisions without human review. Building the agents however often including working with AI in teammate mode.
Domain Knowledge	Expert knowledge held by Procurement, EHS, IT, HR, and other functional teams — inaccessible on demand, retrieved through email or informal consultation.	Domain experts encode their knowledge into agents training data. Students co-created the SUTD Wiki, establishing a reusable innovation model: student prototyping + institutional governance. Expert authority is preserved for all consequential decisions.

What is also striking is that the same three design commitments hold across the portfolio: each bot owns exactly one workflow domain, all operate in Tool Mode (but building them was often done with AI as a teammate), and domain experts retain accountable authority over consequential decisions.

SCALING IMPLICATIONS

Organisations building their own chatbot portfolios should assess readiness across five dimensions. The following framework provides implementation guidance and identifies best-fit organisational contexts.

Scaling Principle	Implementation Guidance	Best-Fit Context
Start narrow, scale wide	Begin with one high-frequency, clearly bounded workflow. Validate the model before expanding. SUTD’s earliest bots (ProcureBot, SUTD Wiki) provided the blueprint for those that followed.	Any institution building its first agent or chatbot
Assign ownership before building	Name the domain owner before development begins. Ownership determines knowledge quality, escalation paths, and maintenance discipline.	All deployment contexts
Platform first, portal second	Deploying on an existing platform (Teams, Slack, intranet) removes the need for users to adopt new tools. Adoption follows the workflow, not the technology.	Organisations with existing collaboration platforms
Co-creation as a model	Students prototyping and institutions governing is replicable. Identify technical communities — student groups, innovation labs, internal champions — who can prototype cheaply.	Universities, training institutions, innovation-rich corporates
Portfolio governance at scale	As portfolios grow beyond 10 bots, formalise a governance register: each bot’s owner, knowledge source, review cycle, and escalation path. Without this, the portfolio becomes unmanageable.	Maturing AI portfolios in any sector

New Role: The Forward Deployment Designer (FDD)

SSG x SGTECH: Ensuring Human Needs Are Met in the Age of AI

EXECUTIVE SNAPSHOT

<p>Days/ Weeks</p> <p>Stalled to Buildable AI initiatives parked for months reworked into clear proposals</p>	<p>~3 Months</p> <p>Prototype to Beta From project restart to working AI agents</p>	<p>Low Code, High Fidelity</p> <p>FDD Prototypes with AI Non-technical professionals driving AI adoption through prototypes</p>
---	---	---

CONTEXT

Imagine you work at a company that knows it should be using AI, but nobody quite knows where to start. The tools are out there. The hype is everywhere. But the gap between having AI tools and using them well feels enormous.

That's where a Forward Deployment Designer (FDD) comes in.

An FDD is someone who sits right at the crossroads of people and technology. They are domain experts first—people who deeply understand a particular industry, function, or workflow. They know how real people work, what frustrates them, where time gets wasted, and what good outcomes look like. And they use that understanding to figure out exactly where AI can step in to make things better.

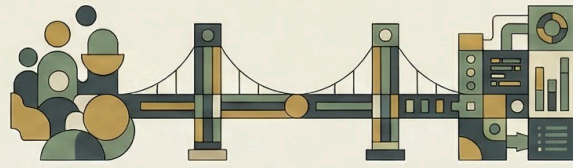
Critically, the FDD then uses AI tools to build a working prototype that demonstrates the solution, turning understanding into something people can test and react to.

Think of an FDD as a translator who also builds. They take messy, real-world business problems, turn them into clear plans, and then use AI tools to create working solutions. From a small app that solves one team's daily headache to a system-wide workflow that transforms how an entire organisation operates.

The FDD is distinct from the Forward Deployment Engineer (FDE) role popularised by Palantir. The FDE builds production-grade systems, while the FDD figures out what should be built and prototypes it.

The Forward Deployment Designer: Bridging Human Insight and AI Capability

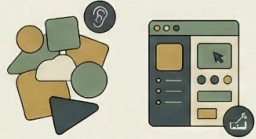
The Forward Deployment Designer (FDD) is a new role designed to bridge the gap between a company's business needs and AI's technical capabilities. By combining deep domain expertise with AI-assisted prototyping tools, the FDD identifies high-value problems and builds functional solutions that move projects from "Ideas" to "execution."



DEFINING THE ROLE & THE DISTINCTION

The Translator-Builder Hybrid

FDDs use domain expertise to translate messy business problems into functional, AI-powered prototypes.



FDD (What to Build) vs. FDE (How to Scale)



Forward Deployment Designer (FDD)

- Primary Focus: People, workflows, and problem definition
- Core Output: User journeys and AI built prototypes
- Analogy: The concept car designer (drivable prototype)

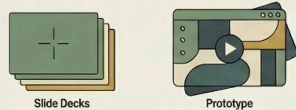


Forward Deployment Engineer (FDE)

- Primary Focus: Building, integrating, and deploying tech
- Core Output: Production code and data pipelines
- Analogy: The production engineer (highway safety at scale)

The Prototype is the Minimum Bar

FDDs are expected to deliver working, testable solutions, not just static slide decks or briefs.



EXPECTED IMPACT & WORKFLOW

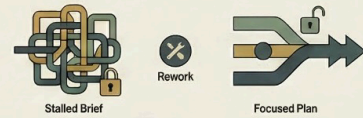
From Deep Listening to Rapid Prototyping

FDDs map human workflows to identify "automation sweet spots" where AI can provide immediate relief.



Rescuing Stalled AI Initiatives

FDDs unlock projects by reworking vague, overly ambitious briefs into focused, achievable plans.



Collapsing the Implementation Loop

By using AI to build, FDDs bypass traditional handoffs, moving from discovery to testing in weeks.

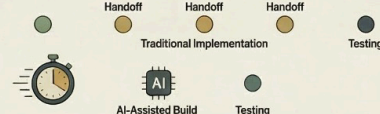


Figure 1: The FDD as a Bridge

ORIGIN

The Beginning of the Forward Deployment Designer

The role emerged from research and practice, and born from a pioneering collaboration between SGTech and SUTD.

First the research: the trainers and facilitators – from diverse domain backgrounds – of the AI Impact Series were all able to design and prototype medium to high fidelity prototypes with the workshop participants. The participating professionals and companies were thus all able to secure quick wins to see the value of AI within days and weeks.

Then the practice: a pioneering collaboration – building on the above – between SGTech and SUTD where an experienced communications professional was brought in to help with stalled AI initiatives and work on the missing links between AI aspiration and execution. In one instance, a project went from stalled to beta testing in three months.

“The idea of doing something new and different, with AI, was exciting to me. And indeed, it was. Somehow, once I started, something unexpected happened. Things fell into place nicely, and the work felt instinctive.”

THE STRUCTURAL PROBLEM

Why AI Projects Stall

Most organisations today face the same challenge: they know AI could help them, but they cannot connect the dots between their day-to-day work and what AI can realistically do. Three failure patterns recur consistently:

- **Projects stall.** Good AI ideas get written up but sit on the shelf because the briefs are too vague, too ambitious, or disconnected from what AI can deliver.

- **The wrong problems get prioritised.** Without someone who understands both the workflow and the technology, teams chase flashy solutions instead of addressing real pain points.
- **Engineers and business teams talk past each other.** Technical teams struggle to understand what users actually need, while business teams struggle to describe their problems in ways engineers can act on.

This is a widespread challenge. Many organisations find that good AI ideas sit on the back burner because project briefs do not translate well into buildable solutions. Opportunities are lost. Without someone who understands both the domain and the AI, and how to design buildable solutions around them, those opportunities stay invisible.

The FDD’s job is see what is possible, how to link everything and everyone, and turn potential into possibility into prototype. The FDD helps organisations:

- **Unlock stalled projects.** By reworking vague or overly ambitious AI proposals into focused, achievable plans.
- **Spot sweet spots.** By mapping out how people work and pinpointing where tasks can be handed to AI.
- **Prototype to make AI adoption realistic.** Through prototyping, a FDD helps organisations – especially those who cannot to hire a full AI team – to take the first practical steps using existing tools and resources.

THE DISTINCTION

FDD vs. Forward Deployment Engineer

If you have come across the term Forward Deployment Engineer (FDE), popularised by companies like Palantir, the distinction is important. The simplest way to think about it: the FDE builds production-grade solutions; the FDD figures out what the solution should be and delivers a working prototype using AI tools. The two roles are complementary, not competing.

Dimension	Forward Deployment Designer (FDD)	Forward Deployment Engineer (FDE)
Primary Focus	People, workflows, and problem framing	Building, integrating, and deploying code
Core Output	Problem statements, user journeys, project briefs, and working prototypes built with AI tools	Working systems, data pipelines, configured platforms
Background	Any domain expertise—communications, design, consulting, operations—paired with willingness to build with AI	Software or data engineering, computer science
Technical Depth	AI fluency — uses AI tools to build medium- to high-fidelity prototypes	Deep technical skills — writes and deploys code
When Involved	Before and during — discovery, framing, iteration	During and after — build, deploy, maintain
Analogy	The concept car designer who builds a drivable prototype	The production engineer who makes it safe for the highway at scale

In an ideal setup, the FDD and FDE work hand in hand. The FDD does the discovery, framing, and prototyping; the FDE takes that working prototype and builds it into a production-ready system. When both roles are present, AI projects move faster, with fewer misunderstandings and better alignment between what the business needs and what gets built.

ROLE CONVERGENCE

When the Two Become One

That said, the table above describes where the two roles *started*. Today's AI tools are rapidly redrawing the boundary between them.

With tools like ChatGPT, Claude, Cursor, and a growing ecosystem of no-code and low-code AI platforms, an FDD is now expected to build a medium- to high-fidelity prototype—a working chatbot, an automated workflow, a functional dashboard—without writing traditional code. The AI does the heavy technical lifting; the FDD provides the problem understanding and design intent that guides it. This is not an optional extra – it is the baseline.

This means the FDD's deliverable cannot be a document that says “here's what we should build.” It is a **working** design. Stakeholders can click through it, test with real data, and give concrete feedback. That is a fundamentally different conversation from reviewing a slide deck.

At the highest level of skill, the FDD begins to resemble an FDE. Not because they've learnt software engineering, but because AI has lowered the technical barrier enough that someone with deep problem understanding and design thinking can now produce what used to require an engineering team. The FDD who can both define the problem *and* deliver a working prototype becomes extraordinarily valuable, because they collapse what used to be a handoff between two roles into a single faster loop.

THE PRACTICE

A Day in the Life of an FDD

“The work was about understanding people – their tasks, frustrations and goals. It was about asking the right questions on user experience, identifying friction points, and understanding enough of AI to know where it can be applied productively. It's design thinking and application.”

Listen and Learn. The day starts with conversations. You sit with a team—operations, marketing, membership coordinators—and ask them to walk you through how they work. What do they spend most of their time on? What's tedious? What falls through the cracks? You are not there to pitch AI. You are there to understand people.

Map and Analyze. With a clear picture of the workflow, you start mapping it out. You sketch user journeys, identify friction points, and flag the steps where AI could genuinely help as they are in practice. You review existing AI project briefs, looking for where objectives went off track or outcomes were unrealistic.

Redesign, Prototype, and Align. You rewrite the problem statement to be specific and clear. You define success in concrete terms. Then you build. Using AI-assisted tools—a large language model to generate code, a no-code platform to wire up a workflow, or an AI copilot to build an interface—you create a working prototype. Instead of walking the technical team through a document, you show them a working design and say: “This is what we need. How do we make it production-ready?”

Iterate and Refine. As the project moves forward, you stay involved. You put the prototype in front of real users. You gather feedback on the working design they can actually interact with. You use AI tools to iterate quickly, sometimes turning around a new version the same afternoon. The FDD is not a one-and-done consultant – they are embedded from first conversation to beta testing and beyond.

The FDD Workflow

Five phases from discovery to deployment



Figure 2: A typical day for an FDD — from listening and learning to mapping, redesigning, and iterating.

Throughout the above, where needed, the FDD will also work with the organisation on:

- introduction of AI usage and/or prototypes to enterprises that adhere to organisation’s internal policies, governance and compliance
- risk assessments and cost benefit analysis where applicable
- workflow redesign, process redesign and job roles redesign where applicable

THE SKILL SET

Skills an FDD Needs

The FDD role does not require a traditional technical background. What it does require is domain expertise—deep knowledge of a particular industry, function, or workflow—and a willingness to design and build solutions with AI. The foundational skills are ones that many professionals already have.

Skill	What It Looks Like in Practice
Active Listening & Empathy	Understanding how people actually work; the real picture, not what they say in meetings.
Problem Framing	Taking vague or complex challenges and defining them in clear, specific terms that others can act on.
Workflow Mapping	Documenting how tasks flow through an organisation, spotting bottlenecks, identifying where AI fits.
Design Thinking	Approaching problems from the user’s perspective, prototyping solutions, iterating on feedback.
Clear Communication	Writing briefs and outcome definitions that both business leaders and technical teams understand.
AI Literacy & Fluency	Understanding what AI can and cannot do, and being able to use AI tools hands-on; prompting, testing, and building working solutions are a core part of the role.
Stakeholder Management	Working across teams, keeping everyone aligned on goals, timelines, and expectations.

WHY IT MATTERS

The Age of AI Needs Translators

We are living through a period where AI tools are becoming more powerful and more accessible by the month. Large language models can draft documents, analyse data, and automate tasks that used to take hours. But the technology is only as useful as the problems you point it at.

Most organisations do not fail at AI because the technology is not good enough. They fail because nobody took the time to properly understand the problem, talk to the people doing the work, and translate that understanding into something an AI system can act on. That is a human problem, not a technical one. The FDD role matters because:

- **AI is democratising, but guidance is not.** The tools are available to everyone, but knowing how to use them well still requires someone who can think clearly about problems and solutions.
- **Organisations, especially SMEs, need practical help, not theory.** Small and medium enterprises lack resources for full AI teams. They need someone who can walk in, understand operations, and create actionable plans and prototypes.
- **It creates a career pathway for domain experts.** Anyone with deep knowledge of their field—whether healthcare, education, logistics, finance, communications, or operations—already has the most important ingredient: understanding how things work on the ground. AI tools handle the technical side; the domain expertise is what makes the solution worth building.
- **Understanding and translating remains the bottleneck.** As AI gets smarter, the limiting factor is not processing power. It is whether anyone has taken the time to understand what needs solving and for whom
- **The prototype is the minimum bar.** AI tools have collapsed the distance between “I understand the problem” and “I’ve built something you can try.” An FDD is expected to frame the problem, design the solution, and deliver a working prototype. This is what makes the role a force multiplier.

The Forward Deployment Designer is the person who figures out what to build and builds it. The tools have changed. The barrier has dropped. What has not changed is what matters most: understanding people, their problems, and what a good solution looks like. That is the foundation and the prototype is the proof.

DESIGNING A HUMAN-AI INNOVATION EDGE FOR ENTERPRISES

STRATEGY USE CASE

Leadership: Can AI Replace Leaders?

Testing C-Suite Digital Twins for Strategic Reasoning

EXECUTIVE SNAPSHOT

80–90%

Simple Decision Accuracy

Across 16 frontier models on binary preference questions

30–50%

Complex Decision Accuracy

Multi-variable trade-offs, value conflicts, ambiguous objectives

Culture

Bias Can Be Bi-directional

AI can have cultural bias, humans too; both risks must be governed

CONTEXT

Digital twins started as an engineering idea: replicate a physical asset in software, stress-test it, and optimize performance. This study asks a parallel question for organizations: can we build a cognitive twin of strategic decision-making? Using a Cognitive Digital Twin approach, we benchmarked 16 frontier AI models against the reasoning patterns of a senior strategist, surfacing a critical “Fidelity Gap” between simple and complex decision-making that has direct implications for enterprise AI governance.

WHY THIS MATTERS NOW

In July 2025, DBS Group’s board told its incoming CEO: “The CEO’s job can be replaced by AI.” Billionaire investor Ray Dalio unveiled “Digital Ray,” an AI clone trained on 50 years of Bridgewater data—performing at ~95% fidelity on his thinking. Deloitte’s Center for Board Effectiveness now highlights digital twins as a boardroom tool; the Conference Board’s 2026 C-Suite Outlook puts AI at the centre of executive strategy; and HBR’s survey of digital leaders confirms AI as a top investment priority. We see the same interest first-hand: whenever we share this research, the response from leaders is immediate and intense.

KEY FINDINGS

- Across 16 frontier models, accuracy is high (80–90%) on straightforward binary preference questions but drops sharply (30–50%) on complex multi-variable trade-offs
- We call this the “Fidelity Gap”
- Potential for further research: the role of culture – much is discussed about models having cultural bias. But what if the converse is a risk too – what if decision makers have cultural and institutional

‘TRILINGUAL’ SKILLS (DESIGN x AI x DOMAIN)

Design: Innovating C-suite cognitive digital twins with AI as tool, teammate, and/or neither. Designing new organization structures and workflows with AI cognitive digital twins.

AI: Tool mode automates simple decisions that are repeatable (instead of delegating to another employee). Teammate mode for complex decisions, using the “Fidelity Gap” into an advantage by

context or even biases that align with specific models' reasoning norms?

assisting as a devil's advocate, identifying potential blind spots, and even acting as a muse.

Domain Knowledge: Human decision remains central because as AI accelerate, decisions will become more complex.

When DBS's board raised the spectre of AI-replaceable CEOs, they were not speaking hypothetically. Nor was Ray Dalio when he built an AI clone of himself.

But how close are we really?

To find out, we ran a controlled experiment: we built a Cognitive Digital Twin of a senior strategist and tested 16 frontier AI models against it. The study was featured at the world's first agents4science 2025 workshop (organised by Stanford University), and was one of only 11 papers spotlighted from over 250 valid submissions (top 5%).

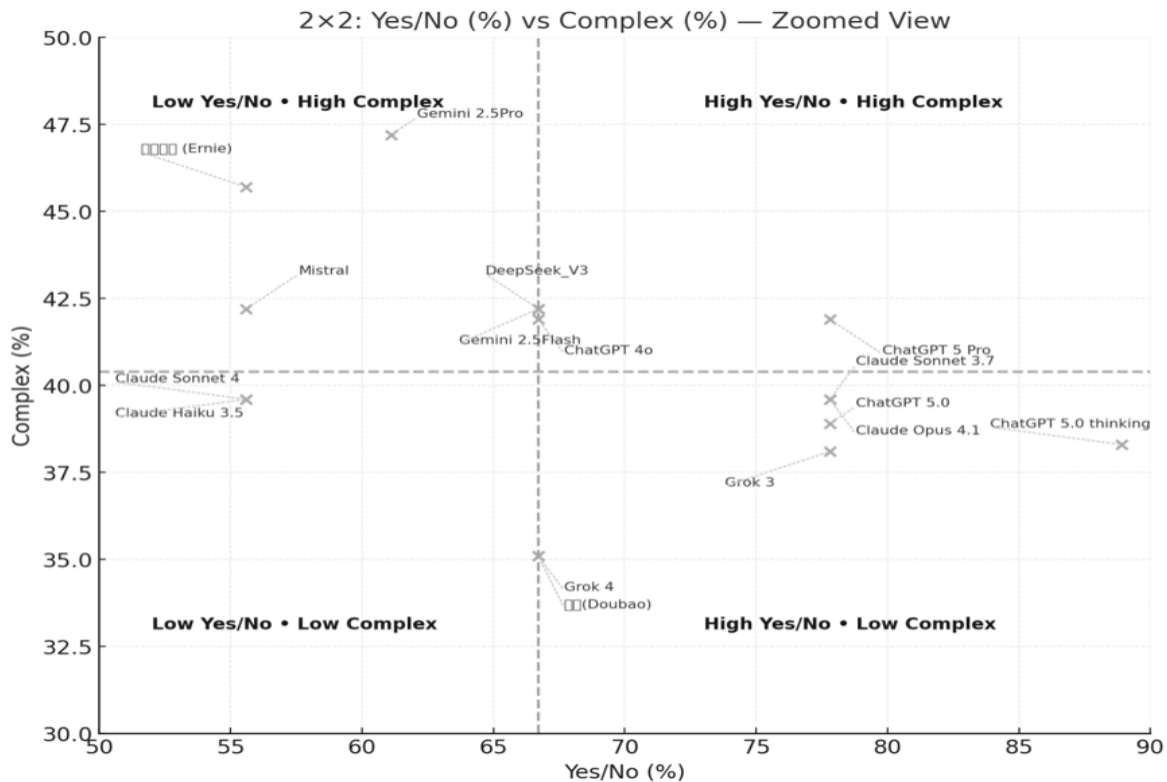


Figure 1: Accuracy of models for simple (yes/no) and complex questions — the Fidelity Gap visualised across 16 frontier AI models.

THE CORE FINDING

The “Fidelity Gap”: AI Sounds Like a Leader, but Doesn’t Think Like One

Here is the uncomfortable truth: when we gave these models straightforward binary preference questions, they got it right 80–90% of the time. But the moment decisions involved real-world complexity—conflicting values, ambiguous objectives, multi-variable trade-offs—accuracy collapsed to 30–50%. Even the strongest model in our benchmark could not break through 48%.

We call this the Fidelity Gap.

The models can reproduce the *voice* of decision-making. They do not yet reliably reproduce the *logic* for complex decision-making.

This matters because organisations risk deploying AI as if fidelity was a given. Our data says otherwise. And the gap is systematic and predictable under complexity (i.e. not random), which means it can be governed.

There is also a cultural dimension. Much is discussed about models having cultural bias. But what if the converse is a risk too – what if decision makers have cultural and institutional context or even biases that align with specific models’ reasoning norms? For obvious governance reasons, this is a question that demands deeper research.

THE TRILINGUAL FRAMEWORK

Turning the Gap into an Edge: Design × AI × Domain

The Fidelity Gap is a limitation that can be turned into a design opportunity.

Once you know where AI is reliable and where it is not, you can redesign decision workflows accordingly. Simple, repeatable decisions can be automated outright (replacing delegation to another employee). For complex decisions, the AI’s lower accuracy ‘bug’ can be exploited to become a ‘feature’: it can serve as devil’s advocate, blind-spot detector, and creative muse—precisely because it reasons differently from the human it shadows.

Skills and Tasks	Then (Past wave of AI)	Now (Current wave of AI)
Design	Not possible to design cognitive digital twins.	Innovating C-suite cognitive digital twins with AI. Designing new organisation structures and workflows.
AI	In the past wave, AI mainly supported decisions through analytics and prediction, producing static outputs like dashboards, reports, and alerts. It had low autonomy and limited integration into end-to-end workflows.	In Tool mode, the AI cognitive twin automates simple decisions that are repeatable (instead of delegating to another employee), absorbing the repeatable 80–90% of routine decision-support work. Leaders retain control of the critical 10–20%. In Teammate mode, AI plays a different role for complex decisions, turning the Fidelity Gap ‘bug’ of its lower accuracy into a “feature” by assisting as a devil’s advocate, blind spot detector, and even as a muse.
Domain Knowledge	Critical: leaders scanned information, framed problems, set decision criteria, drove option generation, and pressure-tested trade-offs. AI played a narrow support role.	Even more critical because leaders must know how to take advantage of cognitive twins in their decision processes. With accelerating advances in AI, decisions will become even more complex, requiring deeper domain knowledge.

Table 1: The Cognitive Digital Twin through the trilingual lens — Design × AI × Domain Knowledge. The infographic below provides one example of how these shifts reshape organisational decision-making.

At the same time, much remains open for the impact on AI cognitive twins in C-suite decision making. For example, the infographic below provides one example on how all these changes organisational decision making. There could be many more, and it is worthwhile researching and experimenting further.

From Tool to Teammate: The Evolution of AI in Organizational Decision-Making

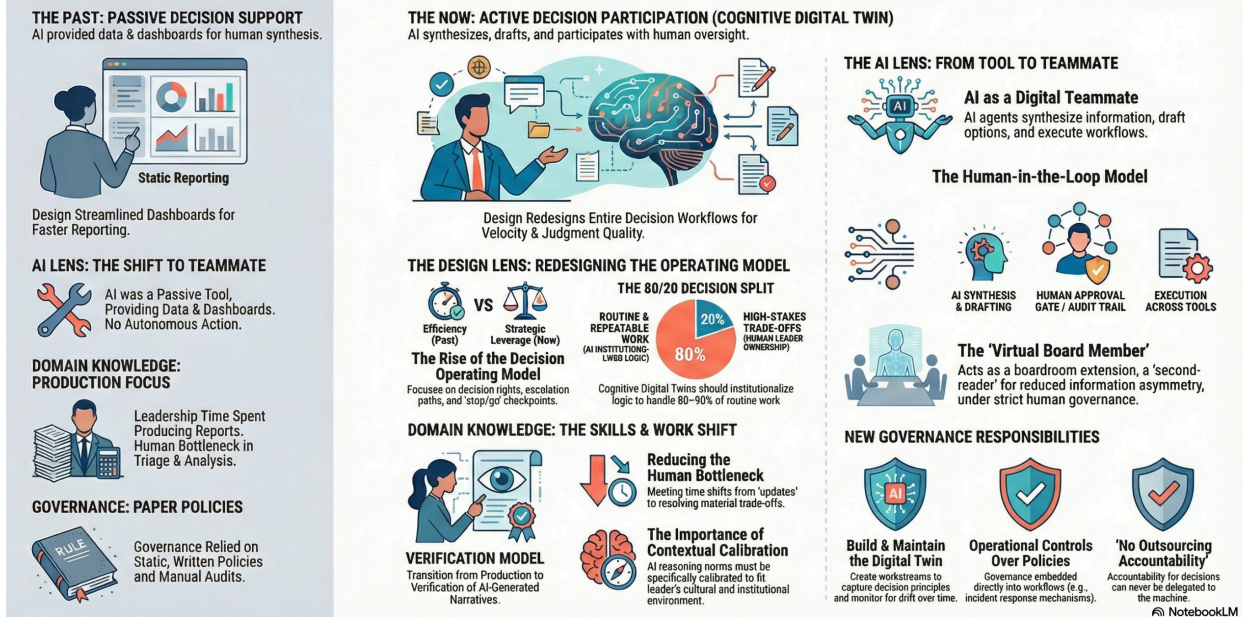


Figure 2: From Tool to Teammate: The Evolution of AI in Organisational Decision-Making.

REFERENCES

Jayasekara, Dinithi, & Huang, Qian. (2025 Oct). Simulating strategic reasoning: A digital twin approach to AI advisors in decision-making. Presented at Agents4Science, organised by Stanford University. Retrieved from <https://openreview.net/forum?id=L4arZChBJD>

Strategy: Forecasting Fast vs Slow

Co-Forecasting With AI to Keep Pace With Accelerating AI Advances That Happen in Weeks

EXECUTIVE SNAPSHOT

Weeks vs Months

Complete foresight cycles in weeks (even days) instead of months/years

Fast vs Slow

Task- and interaction-level forecasts refreshed regularly instead of slow and irregular updates

Human x Agentic

Ensemble of AI models stress-tested by human domain experts

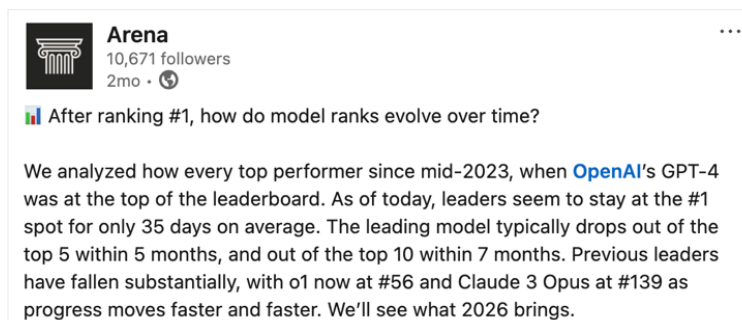
CONTEXT AND WHY THIS MATTERS NOW

Across workshops and meetings, we are often asked “How do we keep up with AI?”, “How do we know what will be disrupted?”, “How can we see into the future?”. Many forecasts from organisations such as the World Economic Forum and consulting companies (e.g. Deloitte, PwC etc) also point out the rapidly changing job landscape and the need to decide and act fast.

For years, to answer these questions, forecasting for strategy, processes improvements, and workforce planning relied on a stable formula: examine role descriptions, identify broad trends, and prepare for change through phased reskilling and job redesign. External experts were consulted over months, sometimes years. This model works when change is gradual.

The accelerating advances of AI today demolishes that assumption. Frontier models advance in weekly leaps (see right), while organisations adapt only incrementally.

At the same time, automation potential is highly uneven. Some tasks within a role may be readily augmented, while others remain stubbornly dependent on judgement, negotiation, and political sensitivity. A role is therefore not simply



Arena
10,671 followers
2mo · 🌐

📊 After ranking #1, how do model ranks evolve over time?

We analyzed how every top performer since mid-2023, when **OpenAI's GPT-4** was at the top of the leaderboard. As of today, leaders seem to stay at the #1 spot for only 35 days on average. The leading model typically drops out of the top 5 within 5 months, and out of the top 10 within 7 months. Previous leaders have fallen substantially, with o1 now at #56 and Claude 3 Opus at #139 as progress moves faster and faster. We'll see what 2026 brings.

Figure 1 Source: https://www.linkedin.com/posts/arenaai_after-ranking-1-how-do-model-ranks-evolve-activity-7415485814291329025-t5ub/

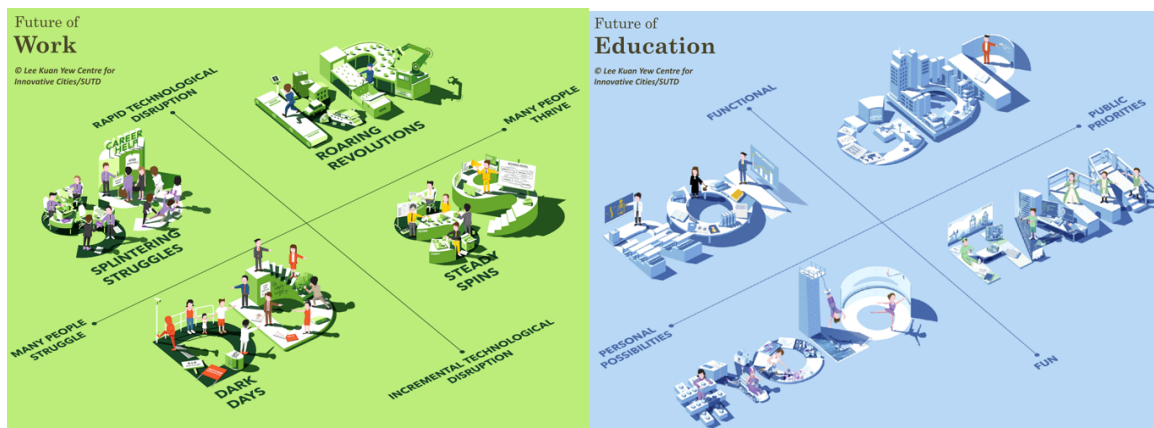
"at risk" or "safe." It is internally uneven, and generic forecasts are poor proxies.

In the Future of Innovation Lab at SUTD, we have developed a co-forecasting process that combines human expertise with AI to generate strategic foresight in weeks, instead of months and years, with fast refresh to keep pace with accelerating advances.

FORECASTING FAST Tried and Tested

This work builds on over a decade of applied forecasting and foresight research at the Lee Kuan Yew Centre for Innovative Cities (LKVIC), Singapore University of Technology and Design (SUTD).

It began with the National Research Foundation-funded **Living Digital 2040** project – a 25-year horizon study of work, education, and healthcare that developed scenario-based foresight for Singapore's future. The study produced scenarios for the Future of Work, Future of Education, and Future of Healthcare (see two examples below), each exploring how technology would reshape these domains through the interplay of different driving forces.



From there, the methods were refined through a series of real-world projects across all sectors:

- **With the unions:** Working with the National Trades Union Congress (NTUC) in 2019: Polarising of Job Opportunities: Charting New Pathways and Adopting New Technologies, we examined how job polarisation affects Singapore's workforce and proposing new pathways for workers.
- **With industry:** A 2020 study with the Institute of Singapore Chartered Accountants (ISCA) and Ernst & Young showing the share of accounting tasks done by machines would surge by 80% within three to five years. This was one of the first task-level analyses of automation impact in a professional services sector.
- **With government:** Producing the first-of-its-kind "A Guide to Job Redesign in the Age of AI" for Singapore's Infocomm Media Development Authority (IMDA) and the Personal Data Protection Commission (PDPC). The Guide – developed under the guidance of the Advisory Council of the Ethical Use of AI and Data – distilled the approach to a single word: task. It offered a sector-agnostic, step-by-step framework for businesses to rethink jobs at the task level when implementing AI. We have also done more recent work with other organisations (e.g. Public Service Division).
- **During the pandemic:** Working with the MICE (Meetings, Incentives, Conferences, Exhibitions) events sector and the unions, as it transformed from physical to hybrid models in under ten months.

The research documented how an entire industry redesigned its work, talent, and business models under extreme time pressure. The lessons map directly to how organisations must now adapt to AI disruption.

Each project sharpened the same core insight: forecasting must be fast, granular, and grounded in actual task structures instead of generic role descriptions. And most recently, we have accelerated the process further through agentic AI that combines different models for speed and rigour.

PROCESS INNOVATION Co-Forecasting With AI

We have since taken our forecasting methods and processes further to take full advantage of AI as a Tool, Teammate, and/or Neither.

Task-level forecasting. Roles are translated into task-level units and scored through an ensemble of AI models until the forecasts stabilise. These scores are then projected through multiple scenarios to develop robust final scenarios that guide strategy, planning, and decision-making. As a result, we can account for the fact that different tasks are disrupted to different degrees, different roles are disrupted to different degrees, and different individuals are disrupted to different degrees.

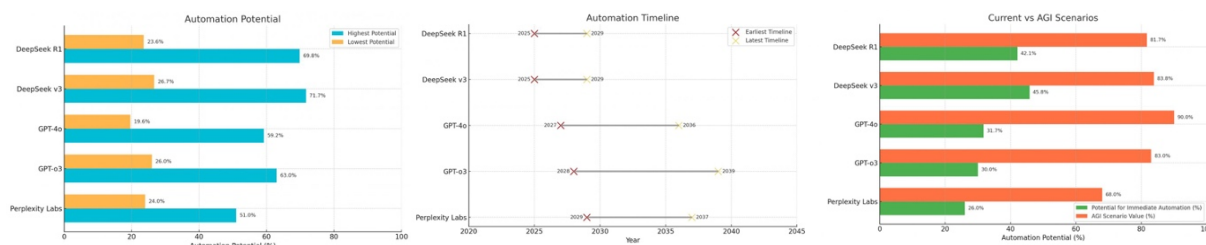


Figure 1: DeepSeek models exhibit the highest automation potential among the models evaluated. The variance between different models suggests the importance of considering both conservative and aggressive scenarios when forecasting AI-driven job disruption.

Figure 2: There is considerable divergence in forecast uncertainty across models. DeepSeek models project a relatively narrow window (2025–2029), implying more confidence in the timing of AI-driven automation. In contrast, models like GPT-4o and GPT-4.3 forecast much wider timelines, reflecting greater uncertainty about when AI technologies will reach sufficient maturity and adoption levels to significantly impact the job market.

Figure 3: While the potential for immediate automation is moderately low given the current technology, the high AGI Scenario Value shows that the job tasks can be highly automated in the event AGI is available in the near-term. This highlights an urgency for policy makers and organizational leaders alike to rethink jobs and make them more resilient.

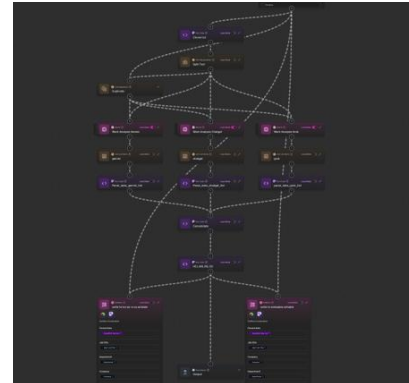
Figure 2 Source: <https://www.sutd.edu.sg/stories-listing/ai-job-disruption-forecasting-issue-1-conventional-job-forecasting-vs-llms-powered-models/>

Interaction mapping. On its own, task-level analysis would still leave a familiar problem: numbers can show where disruption pressure may rise, but not how work is socially organised. This is where interaction mapping becomes essential. Drawing on interviews, focus groups, and AI-assisted qualitative coding, we map who hands what to whom, which activities are still carried by unassisted human exchanges, where AI assistance is already present, and where more hybrid human–AI teaming is emerging. Interaction maps reveal persistent pockets of human centrality, together with the relational exchanges that drive performance but often go undocumented.

Work is not merely a bundle of tasks. It is also a bundle of interactions. Value is produced from applying skills to skills, as well as through coordination, stakeholder engagement, negotiation, ethical adjudication, and interpretive judgement.

Agentic workflows. When the above are built into an agentic workflow (see right), each co-forecast can be generated in weeks – sometimes days for simpler analysis – instead of months or years. Forecasts can also be refreshed (also in weeks instead of months or years) whenever there are new AI advances.

In our methods, humans remain in the loop throughout, and domain expertise matters even more than before to assess these forecasts, and to strategise, plan, and decide with them.



GROUNDED IN RESEARCH EVIDENCE

Teaming With AI for Forecasting Works

The power of AI-assisted forecasting is now backed by growing research evidence:

- **LLM ensembles rival human crowds.** An ensemble of 12 LLMs made probabilistic predictions on par with 925 human forecasters in a three-month tournament. The "wisdom of the silicon crowd" was statistically indistinguishable from the human crowd (Schoenegger, Tuminauskaite, Park, Valdece Sousa Bastos, & Tetlock, 2024, Science Advances).
- **LLM assistants boost human accuracy.** When 991 participants could consult LLM assistants while making forecasts, their accuracy improved by 24–28% compared to a control group. Even a "noisy" assistant improved performance, suggesting that access to any frontier LLM can meaningfully augment human judgement in cognitively demanding tasks (Schoenegger et al., 2025, ACM Transactions on Interactive Intelligent Systems).
- **LLMs are approaching superforecaster parity.** On the ForecastBench dynamic benchmark, state-of-the-art LLMs now outperform non-expert public participants, and linear extrapolation of current improvement trends projects LLM–superforecaster parity for late 2026 (Karger, Bastani, Chen, Jacobs, Halawi, Zhang, & Tetlock, 2025, ICLR).

Our co-forecasting approach builds on these capabilities while keeping human expertise central. For us, teaming with AI as a powerful tool takes place within a human-led process.

IMPLICATIONS FOR ORGANISATIONS

From Foresight to Redesign

The practical payoff of co-forecasting at this level of granularity is agile work redesign before disruption hardens into confusion or skill mismatch.

Tasks can be broken down further and each micro-task reclassified as human-led, AI-assisted, or co-creative, and these classifications traced across a workflow. Managers can ask: which activities must remain human because of their relational, ethical, or interpretive character? Which can be delegated to AI because they are repetitive and auditable? And where should genuinely collaborative human–AI work be cultivated?

Precision upskilling. Because each task sits at the intersection of an automation estimate and an interaction type, blanket training responses are unlikely to work. Work that remains strongly human-centred requires deeper investment in relational, interpretive, and strategic craft. Work in more exposed analytical roles needs stronger capabilities in prompting, verification, workflow orchestration, and oversight. With our process, we can customise to each person, and target specific tasks and interactions.

THE TRILINGUAL FRAMEWORK


Building ‘Trilingual’ Skills and Tasks (Design × AI × Domain)

Dimension	Then (Past wave of AI)	Now (Current wave of AI)
Design	Forecast and foresight cycles designed around presumptions that changes happen in years.	Co-forecasting with AI for strategic foresight accounts for accelerating AI advances that take place in weeks. Cycles must be rapidly generated, refreshed, and stress tested.
AI	Digital systems primarily supported officers through dashboards, retrieval, and narrow automation.	Generative AI participates more directly as tool, teammate, and/or neither across the entire spectrum of activities. Ensemble approaches and agentic workflows combine multiple models for speed and rigour.
Domain Knowledge	Domain knowledge given time to develop.	Domain knowledge needs to keep pace with accelerating advances; even deeper domain expertise is now needed to assess rapid changes in the external environment.

REFERENCES

- Poon, K. W., Lee, H., Lim, W. K., et al. (2018). *Living Digital 2040: Future of Work, Education, and Healthcare*. World Scientific.
- NTUC Learning Hub (OTCi), & SUTD. (2017–2021). *Job polarisation and new pathways for workers in Singapore. Collaborative foresight research with unions.*
- Institute of Singapore Chartered Accountants (ISCA), Ernst & Young Singapore, & SUTD. (2019). *Automation of Accounting Work. Landmark study showing 80% of accounting tasks susceptible to machine automation.*
- Infocomm Media Development Authority (IMDA) & Personal Data Protection Commission (PDPC). Singapore. (2020). *A Guide to Job Redesign in the Age of AI. First-of-its-kind guide for workforce transformation.*
- MICE Singapore & SUTD. (2020–2021). *Pandemic foresight study: Physical to hybrid transformation of Singapore's events sector. Task-level redesign during crisis acceleration.*
- Poon, K. W., Willems, T., & Liu, W. S. Y. (2023). *The future of expertise: From stepwise domain upskilling to multifaceted mastery*. In *International Handbook on Education Development in Asia-Pacific* (pp. 1-19). Springer Nature Singapore.
- Willems, T., Hotan, D. J., Tang, J. C., et al. (2025). *Assessing employment and labour issues implicated by using AI*. arXiv:2504.06322.
- Willems, T., Poon, K. W., & Huang, Q. (2025). *Strategic foresight and co-forecasting with generative AI* [Conference paper]. IPPA Conference.
- Schoenegger, P., Tuminauskaite, I., Park, S., Valdece Sousa Bastos, P., & Tetlock, P. E. (2024). *Wisdom of the silicon crowd: LLM ensemble prediction capabilities rival human crowd accuracy*. *Science Advances*, 10(45).
- Schoenegger, P., et al. (2025). *AI-Augmented Predictions: LLM assistants improve human forecasting accuracy*. *ACM Transactions on Interactive Intelligent Systems*.
- World Economic Forum. (2025). *The Future of Jobs Report 2025*. Geneva.
- Deloitte. (2026). *2026 Chief Strategy Officer Survey: Amid persistent volatility, chief strategy officers face rising expectations and limited bandwidth.*
- PwC. (2026). *29th Annual Global CEO Survey: Leading through uncertainty in the age of AI*.
- McKinsey & Company. (2024). *The critical role of strategic workforce planning in the age of AI*.
- Karger, E., Bastani, H., Chen, Y.-H., Jacobs, Z., Halawi, D., Zhang, F., & Tetlock, P. E. (2025). *ForecastBench: A dynamic benchmark of AI forecasting capabilities*. *International Conference on Learning Representations (ICLR)*.

THE FUTURE OF INNOVATION



CASE STUDIES IN THIS SECTION

- 13. WORLD FIRST DESIGN AI UNIVERSITY: THE FUTURE IS BUILT, NOT TAUGHT
- 14. EJC: WHEN TEACHERS AND STUDENTS LEARNT AND INNOVATED TOGETHER
- 15. INNOVATING WHAT'S NEXT FOR MATURE WORKFORCES
- 16. ALL CAN INNOVATE: VOICES FROM CITIZENS
- 17. FUTURE OF CITIES: CITIZEN PARTICIPATION TO CITIZEN CO-CREATION

Transforming Lives:
All Can Innovate

The Future Is Built, Not Taught

SUTD: How Students and Faculty Are Co-Designing and Redesigning What It Means to Learn, Build, and Assess with AI

EXECUTIVE SNAPSHOT

Faster

First prototypes can be built in hours instead of days

Higher

Both faculty and students can aim higher for what they wish to build

Stronger

Students' skills are strengthened as they acquire adjacent and deep skills

CONTEXT

AI is transforming Institutions of Higher Learning (IHLs) along two simultaneous axes: **students are building for industry earlier than ever before**, and learning faster than before; and **faculty are redesigning how they teach and assess**. Taken together, the two axes map a structural shift in how innovation happens within education.

KEY FINDINGS

- Year 1–2 students delivered working solutions to industry partners and internal stakeholders, that would conventionally have been built by older students.
- AI collapses the prototyping and execution timelines, shifting the constraint from technical skill to design judgement
- Students and faculty adopt three distinct roles with AI: Tool, Teammate, and/or Neither, demonstrating sophisticated metacognitive agency
- Calibrated AI assessment shows promise for creative grading when rubrics are precisely encoded with human oversight
- Students transition from AI users to AI creators, deepening critical and creative engagement

“TRILINGUAL” SKILLS (DESIGN x AI x DOMAIN)

Design: Students and faculty designed solutions faster and could choose to aim higher. They also must design their AI tools and teammates carefully (see below).

AI: Working with AI as Tool, Teammate, and/or Neither, is a choice that depends on judgement and context.

Domain Knowledge: Both students and faculty developed and deepened new, adjacent, and domain expertise.

The future of Institutions of Higher Learning (IHLs) is being reshaped by a structural change that operates along two axes simultaneously. On one axis, students equipped with AI as a building partner are delivering working solutions to industry partners not in their final year, but in their first and second. They are learning faster as a result. On the other axis, faculty are redesigning how they teach, assess, and plan curricula, using

AI as calibrated co-intelligence that enhances pedagogical quality while preserving human judgement, instead of as a replacement.

Together, they reveal that the future of education lies in investing in the future of innovation. It is about designing the architecture that allows humans and AI to operate at their best, instead of the conventional assumption that one has to choose between human and artificial intelligences.

Axis 1: Students as Innovators — Building for Industry

Industry projects have traditionally been a final-year undertaking. Students spend three or four years accumulating the technical depth, domain fluency, and implementation skills needed to build something an industry partner can use. Working with AI as a Tool, Teammate, and/or Neither has broken that assumption. In SUTD’s Building 0 programme, numerous projects undertaken by Year 1 and Year 2 undergraduates are delivering working solutions to industry partners and institutional stakeholders (see table below).

Table 1: Student Projects from SUTD Building 0 — Working Solutions for Industry Partners

Project	Industry Partner	What Was Built	Student Year
Manpower Management System	Hospitality company	Working prototype in two meetings	Year 2
Conversational CAD	Manufacturing company	Tool running on consumer hardware with local LLM	Year 1
Privacy-First Transcription	Manufacturing company	Secure on-device transcription	Year 1
AI-Assisted Design	Urban sector company	Spatial concept produced and modelled	Team
AI-Native LMS	Internal stakeholder	Ongoing evaluation for adoption	Year 1
Guided EHS Assessment	Internal stakeholder	Prototyped in one session and submitted for review	Year 1
SUTD Course Comparisons	Internal stakeholder	Ongoing evaluation for adoption	Year 2

The students often had to pick up adjacent skills on their own (e.g. one of the students above had an architecture background and picked up coding on his own to build the solution). As they built, they also developed and deepened domain expertise. Their experiences and the trajectory of these innovations suggest that *“The high fidelity of today is the low fidelity of tomorrow. What counts as ‘good enough to show’ keeps advancing.”*

Axis 2: Faculty as Innovators — Redesigning How We Teach and Assess

While students are building for industry, faculty face a parallel transformation: how to teach, assess, and design curricula in a world where AI is no longer a course or subject to be taught and studied, but a force reshaping every aspect of pedagogical practice. Ongoing research and practice suggest the following insights and opportunities:

1. **AI can mark creative work accurately — if humans set it up properly first.**
AI can grade student design projects almost the same as a real teacher would. The catch is it took a lot of careful setup and testing to get there.
2. **AI grading is contextual – it works better in some classrooms than others.**
The same AI tool that works in one course/school might not work as well in another course/school. It is likely a lot of contextual tweaks are needed before it works somewhere else.
3. **Students already have good instincts about when to use AI and when to skip it.**
Students are able to make thoughtful choices about how they work with AI e.g. as a tool (e.g. a helper for quick tasks); as a teammate (e.g. as a creative partner to bounce ideas off); and/or neither (e.g. they deliberately avoided it so their own ideas stayed original).
4. **Students learn way more when they *build* AI tools, rather than just use them.**
Building AI solutions gave students a much better understanding of what AI is good at and where it falls short. It also made them feel more confident and connected to their classmates because they had to work as a team to build these solutions.
5. **AI can help teachers design better lessons, even in subjects they are less familiar with.**
AI can help teachers put together projects that link different subjects. Even teachers who were not experts in every subject involved were able to create well-organised, high-quality plans with the help of AI.

Taken together, the big picture that emerges is that AI works best in education when educators put real effort into setting and training it, and students are given the opportunities to actively shape how AI is used, including creating and building their own AI innovations.

THE TRILINGUAL FRAMEWORK

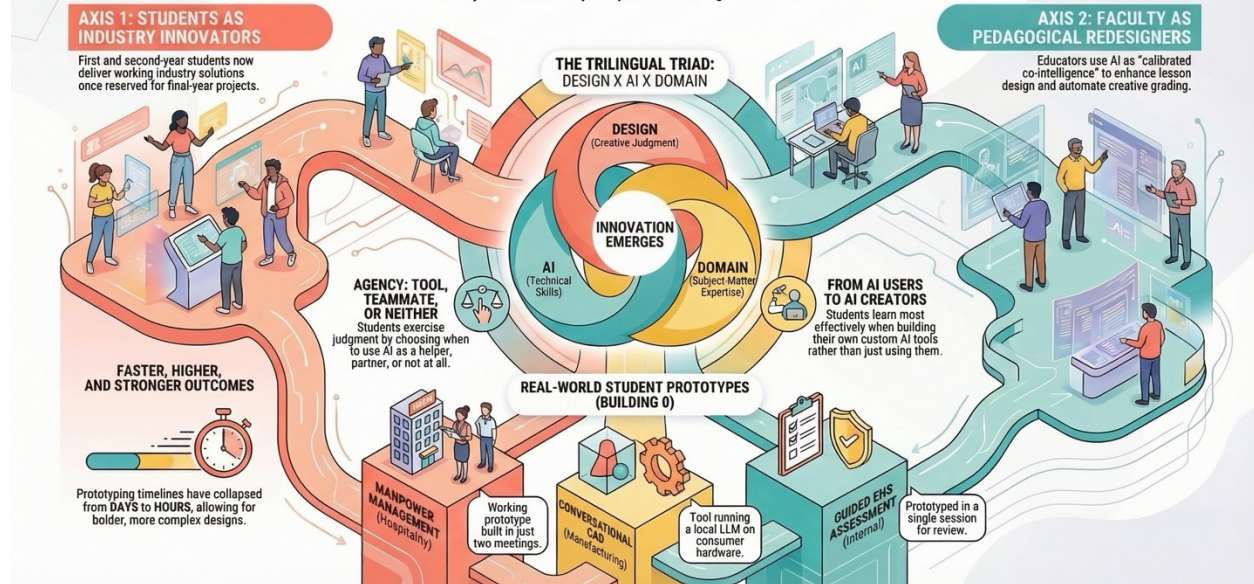
Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

THE TRILINGUAL LENS		
<p>Design</p> <p>Students and faculty designed faster and could choose to aim higher and bolder in what they design.</p> <p>Students and faculty also must design their AI tools and teammates carefully.</p>	<p>AI</p> <p>Working with AI as Tool, Teammate, and/or Neither, is a choice that depends on judgement and context.</p>	<p>Domain Knowledge</p> <p>Both students and faculty developed and deepened new, adjacent, and domain expertise.</p>

Institutes of Higher Learning now have an option: use AI to educate, or use AI to empower students and faculty to create as well. It opens up the opportunity that the future is less about what is taught by faculty to students, and more about what is built by them together.

THE FUTURE IS BUILT, NOT TAUGHT: SUTD'S HUMAN-AI INNOVATION EDGE

SUTD integrates AI, Domain Knowledge, and Design—the “Trilingual” framework—enabling students and faculty to deliver industry-ready solutions through active co-innovation.



REFERENCES

- Willems, T., Khan, S., Huang, Q., Camburn, B., Sockalingam, N., & Poon, K. W. (2025). To use or to refuse? Re-centering student agency with generative AI in engineering design education. arXiv:2510.19405. <https://arxiv.org/abs/2510.19405>
- Huang, Q., Sockalingam, N., Willems, T., & Poon, K. W. (2025). Designing knowledge tools: How students transition from using to creating generative AI in STEAM classroom. arXiv:2510.19342. <https://arxiv.org/abs/2510.19342>
- Khan, S., Liow, W. T., & Ang, L. K. (2025). Human or AI? Comparing design thinking assessments by teaching assistants and bots. arXiv:2510.16069. <https://arxiv.org/abs/2510.16069>
- Liow, W. T., Khan, S., & Ang, L. K. (2025). Co-designing interdisciplinary design projects with AI. arXiv:2510.16068. <https://arxiv.org/abs/2510.16068>
- Huang, Q., Willems, T., & Poon, K. W. (2025). The application of GPT-4 in grading design university students' assignment: An exploratory study. *International Journal of Technology and Design Education*, 1–25.
- Poon, K. W., Willems, T., & Liu, W. (2023). The future of expertise: From stepwise domain upskilling to multifaceted mastery. In: Lee, W. O., Brown, P., Goodwin, A. L., & Green, A. (Eds.), *International Handbook on Education Development in Asia-Pacific*. Springer, Singapore.
- SUTD Building 0: <https://www.sutd.edu.sg/future-of-innovation/x-intelligence/>
- Sockalingam, N., Huang, Q., Willems, T., & Poon, K. W. (2025). Innovative “tool-to-teammate” instructional strategy for using ChatGPT in higher education. In *EDULEARN25 Proceedings* (pp. 7891–7900). IATED.
- Huang, Q., & Willems, T. (2025). Empowering educators in the age of AI: An empirical study on creating custom GPTs in qualitative research method education. arXiv:2507.21074. <https://arxiv.org/abs/2507.21074>
- Huang, Q., & Poon, K. W. (2026). The Trilingual Triad Framework: Integrating Design, AI, and Domain Knowledge in No-code AI Smart City Course. arXiv:2603.05036. <https://arxiv.org/abs/2603.05036>

DESIGNING A HUMAN-AI INNOVATION EDGE FOR ENTERPRISES

STRATEGY USE CASE

EJC: When Teachers and Students Learnt and Innovated Together

Eunoia Junior College x SUTD: How We — Schools, Universities, Teachers, and Students — Can Change The Future of Education Together

EXECUTIVE SNAPSHOT

94%

Like Hands-On Activities

They gave the breakout sessions at least 5 out of 6

> 50

Prototypes Built by Teachers and Students

Working prototypes across subjects

1 Month

From Seed to Scale

From first training to institution-wide delivery by EJC's own staff

Something Different Happened in the Room

On 27 January 2026, nearly 200 participants at Eunoia Junior College (EJC) sat down for a day devoted entirely to AI. But this was not a typical professional development session where an outside expert talks and the audience listens. SUTD had designed something different: teachers and students learning side by side, with no assumption about who would take to it faster.

The expectation was that students would be the more natural technology users, and the main dynamic would be younger learners helping older staff. That is not what happened. What emerged was genuine two-way collaboration. Students pushed creative boundaries: more exploratory, more willing to fail interestingly, more likely to find unexpected uses. Teachers brought something students could not: deep subject knowledge and the ability to see immediately how a tool connected to what they teach and what their students need.

“When the technology is genuinely new to everyone in the room, teachers and students learn better together than apart.”

The proof is in what people built.

By the end of the day, teachers had produced 49 working prototypes across three tracks — Active Learning, Cognitive Load, and Productivity. An Economics teacher built a policymaker simulation. A Chemistry teacher created an interactive bonding visualisation. A GP teacher designed an international relations role-playing game. A PE teacher gamified fitness training. One teacher built a CLT Design Coach — a tool that reviews lecture slides and suggests tweaks to reduce cognitive overload for students. Another built a CCA attendance dashboard to replace a manual process that had frustrated student leaders for years. And yes, one teacher built a drawing canvas controlled entirely by nose movement, for no reason other than to see if it could be done (It could.).

Students built too. One group created an assignment checker that reads uploaded documents and gives feedback based on the rubric. Another built a timetable optimiser to help students plan their free periods. A third built a push-up counter using AI vision (which has since inspired Project Work groups to explore how AI can give real-time feedback on sports form). These were not hypothetical. They were working applications, built in a day.



The feedback reflected the energy in the room. Staff described it as “fruitful” and “well planned.” Multiple people said having students present added an important perspective. And one comment captured what many felt: they wanted more protected time like this to keep experimenting.

“It will be great if we had more protected time like this to work on our projects!” — Survey respondent

What the Numbers Tell Us

93 staff members responded to a post-event survey. Of the 93 who completed the full set of questions, the results were striking. They were positive and revealed about what mattered most.

The hands-on breakout sessions were the clear standout, averaging 5.45 out of 6. Over half the respondents gave them a perfect score. People valued the format itself: hands-on, collaborative, and built around real experimentation rather than passive listening.

The keynote and panel discussion rated well too (5.06 out of 6), but several respondents said explicitly that they would have preferred less lecture time and more time to experiment.

Table 1: Post-event survey results (N = 93, on a scale of 1–6).

Survey Question	Average	% Rated 5–6	% Top Score
Breakout sessions were useful	5.45	94%	52%
Gained ideas for AI use cases	5.22	85%	38%
Keynote and panel Q&A was useful	5.06	78%	32%
Feel more equipped to apply AI	5.06	80%	29%

There is a meaningful gap in the data worth paying attention to. Staff came away with plenty of ideas (5.22) but felt slightly less confident about being equipped to act on them (5.06). It is a small difference, but it tells a story: people were inspired, and now they need more runway to turn that inspiration into practice. Some of the open feedback echoed this, with staff mentioning information overload and wanting more time to process what they had learnt.

Perhaps the strongest signal of all, almost all respondents said they wanted a school-provisioned AI account, showing that the baseline had moved up and there was demand.

Overall, the enthusiasm was consistent and the reach was broad. Science, Mathematics, English, Economics, Mother Tongue Languages, Arts, PE, and others all responded positively.

How It Worked

The model behind the day was straightforward. SUTD did not try to train every teacher at EJC directly – that approach does not scale and requires a permanent external presence. Instead, they trained a small core group at EJC first. Three weeks later, that core group turned around and led the Learning Day sessions for the rest of the school, organised by subject area: PE, English, Mathematics, Economics, Physics, and more. The entire arc – from first training to institution-wide delivery – took under a month.

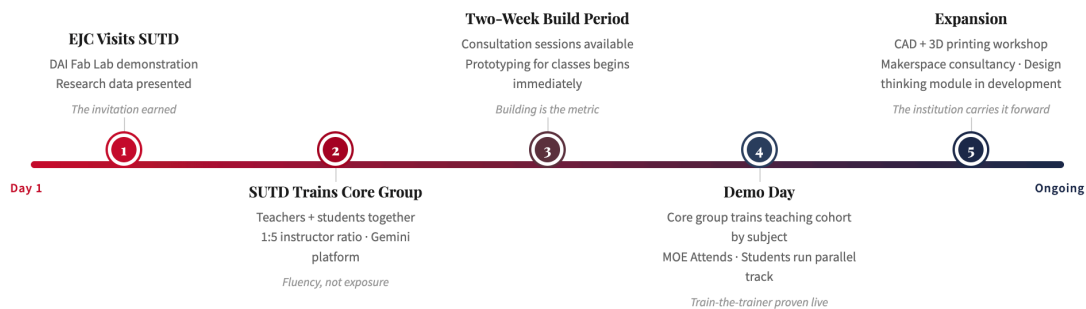
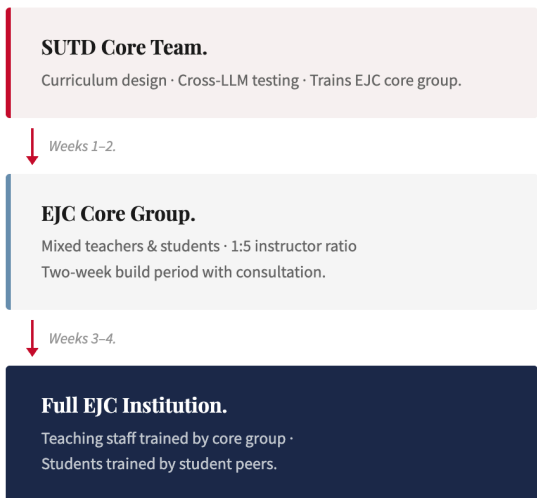


Figure 1: The one-month arc from first training session to institution-wide deployment.

A few design choices made it work. The curriculum was built around how to think with AI, not around any one platform, so it did not matter that EJC used Gemini while SUTD had trained on ChatGPT. Sessions ran at a 1:5 instructor-to-learner ratio, which sounds intensive but turned out to be necessary: when people are new to AI, they need someone nearby who can help them make sense of unexpected results in real time. And teachers and students were placed in the same sessions deliberately, which is where the two-way learning described earlier came from.

THE ARCHITECTURE.



WHAT HAPPENS IN THE ROOM.

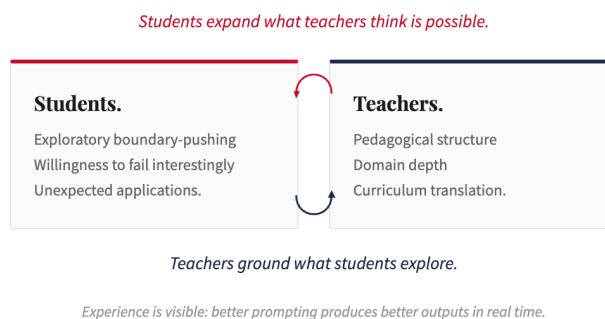
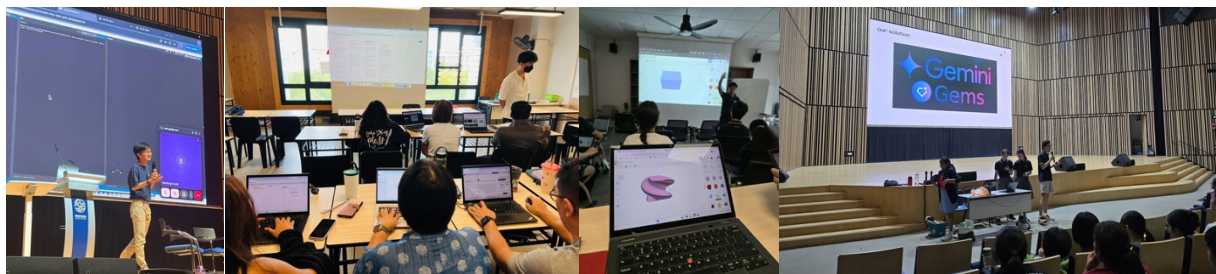


Figure 2: The train-the-trainer architecture (left) and the bidirectional learning dynamic observed within mixed cohorts (right).

The proof point was the Learning Day itself. When our friends from MOE observed the event, what they saw was not an outside team delivering a programme. They saw a school’s own people teaching their colleagues — people who had been learners themselves just weeks earlier. That is the difference between a workshop and a capability.



What Students Could Build Next

The teachers who pushed hardest for this programme were frustrated by a pattern they kept seeing in A-level project work: student outputs that defaulted to social media campaigns and posters. Not because the students lacked ambition, but because building something real — a working application, a deployable tool, a functional prototype — simply took too long given everything else competing for attention in an intensive curriculum.

AI has changed that equation. A working application that would have taken weeks of technical effort can now be co-developed in a single session. This is not a small efficiency gain. It is a shift in what is possible for students working within the real constraints of a JC schedule.

“Once an institution begins to imagine itself as a place that builds capability rather than merely delivers content, it becomes harder to go back.”

This is no longer a hypothesis as we saw earlier with what happened during Learning Day. To recap, EJC students built a push-up counter using AI vision that analyses form in real time. That project has since inspired Project Work groups to explore how AI can coach movement across different sports. Another student group built a timetable optimiser that helps peers make the most of free periods — a tool born from a problem every JC student knows. A third built an assignment checker that reads uploaded documents and gives rubric-based feedback. These were not prompted by teachers. Students identified the problems themselves and built the solutions in a day.

For other schools considering something similar, the message from this experience is simple: the barrier to building has dropped. The question is no longer whether students can create something real. It is whether the school is set up to let them try.

Just as importantly, the larger message is that when we choose to work with each other — schools, universities, teachers, students — we can change the future of education.



WSG x SUTD: Innovating What's Next For Mature Workforces

From Knowing to Doing: A Career Guidance Approach for Mature Professionals Using Embodied Learning, Design, and AI

EXECUTIVE SNAPSHOT

~380

professionals over 16 sessions

98%

say "Can now identify meaningful and achievable career goals"

98%

say "Gen AI useful and important for career transition"

CONTEXT

Singapore's mature workforce is the fastest-growing labour market segment. For the over-50 cohort, this transition arrives with particular complexity: participants carry decades of identity investment in specific roles, face uncertainty about how their skills transfer, and often hold unexamined assumptions about what is realistic at their age. These cannot be answered by skills matrices or job-matching algorithms alone. The answers must be discovered through embodied tools, AI that amplifies self-knowledge, and peers at the same life stage.

"This is one of the best courses I've attended because it is practical, participative, motivational and engaging. It helped us to scope up a better understanding of ourselves visually with all the physical tools such as career mapping, skill cards, job cards. The AI tool presented is easy to pick up and useful. It also helped you to want to begin working on actionable stuff when we go back to our daily lives."

— Participant

THE STRUCTURAL PROBLEM

The Upstream Challenge

Most mature professionals navigating career transitions face questions that skills matrix find hard to answer: who am I now, what has my accumulated experience actually made me capable of, and what is genuinely possible for someone like me? The inputs to this question — self-knowledge, values, narrative identity — are not yet available in structured form. They must be discovered.

This is the Self-Discovery Gap. Closing it cannot be accomplished through cognitive instruction alone. It requires an approach that engages body alongside mind, surfaces self-knowledge through doing rather than

telling, and uses AI not as a recommendation engine but as a reflective mirror that amplifies what the person has already discovered about themselves.

Pre-workshop baselines confirm the gap is real: participants arrive with strong growth mindsets: 4.5+ out of 5 on openness, curiosity, and learning orientation. The problem is that openness has not been converted into agency. That conversion is what the What's Next programme is designed to produce and what the data shows it achieves.

THE WHAT'S NEXT WORKSHOP APPROACH **The Open Loop Career Suite**

Career Navigation Through Embodied Tools and Design•AI

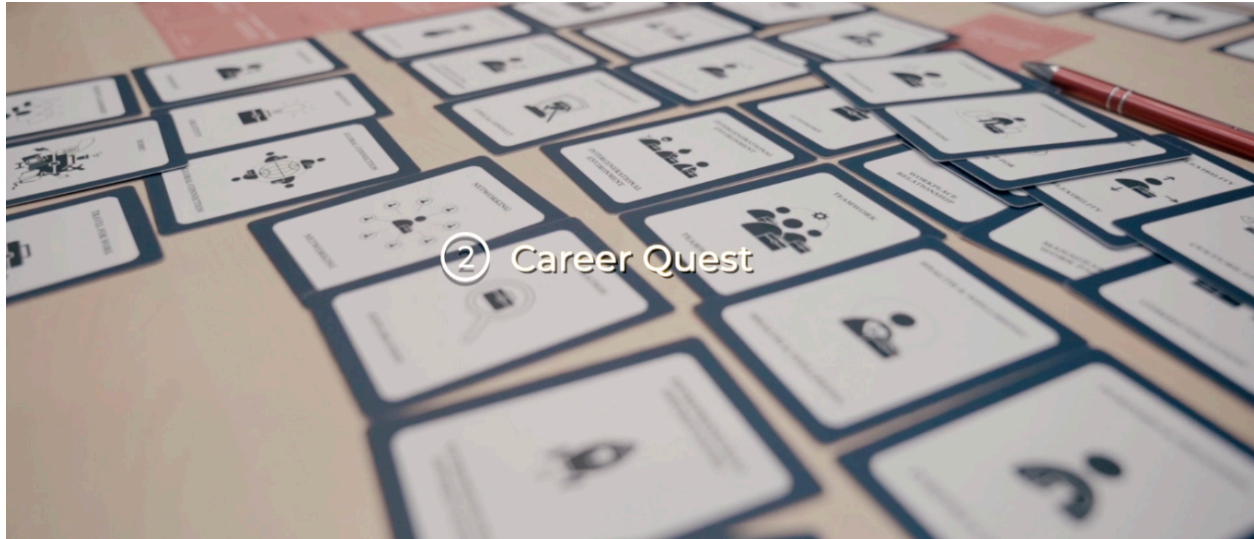
The What's Next: Reimagine Your Career Using Design•AI programme was developed by the Future of Innovation Lab and the Lee Kuan Yew Centre for Innovative Cities at SUTD in partnership with Workforce Singapore (WSG). Its theoretical foundations draw on Experiential Learning (Kolb, 1984), Transformative Learning (Mezirow, 1991), and Embodied Cognition research (Marshall, 2016). Tools were developed and tested – partly under a SkillsFuture (SSG)/Institute for Adult Learning (IAL) applied research grant to examine Open Loop Careers – with over 200 Singapore workers across ages, sectors, and career stages before launch, ensuring they reflect Singapore's labour market realities and the lived texture of mature professional experience.

The programme's tools are designed as a sequence to progressively close the Self-Discovery Gap.

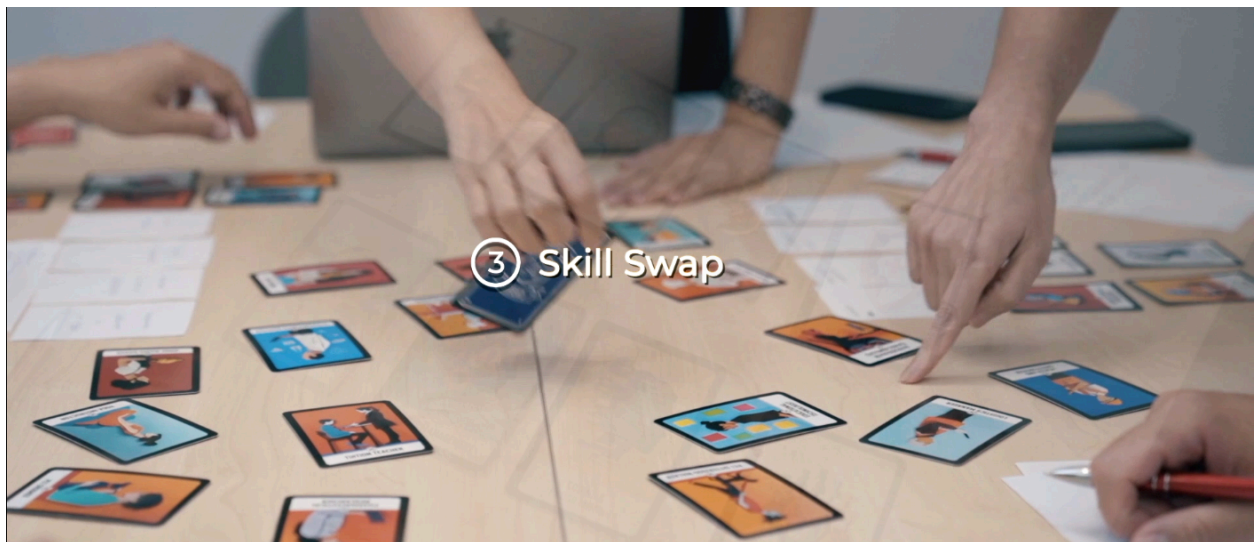
Career Journey Maps: A visual, active reflection tool drawing on reminiscence therapy principles. It activates episodic memory to surface patterns and transferable strengths across a career history, often revealing coherence that participants did not know was there; fragmented in retrospect does not mean incoherent in capability.



Career Quest Cards: A gamified, tactile card-sorting exercise. Participants physically handle and arrange cards representing work values, preferences, and trade-offs. The physical act of sorting is the intervention. It produces clarity about priorities that abstract questionnaires cannot reach, because the body knows things the conscious mind has not yet articulated.



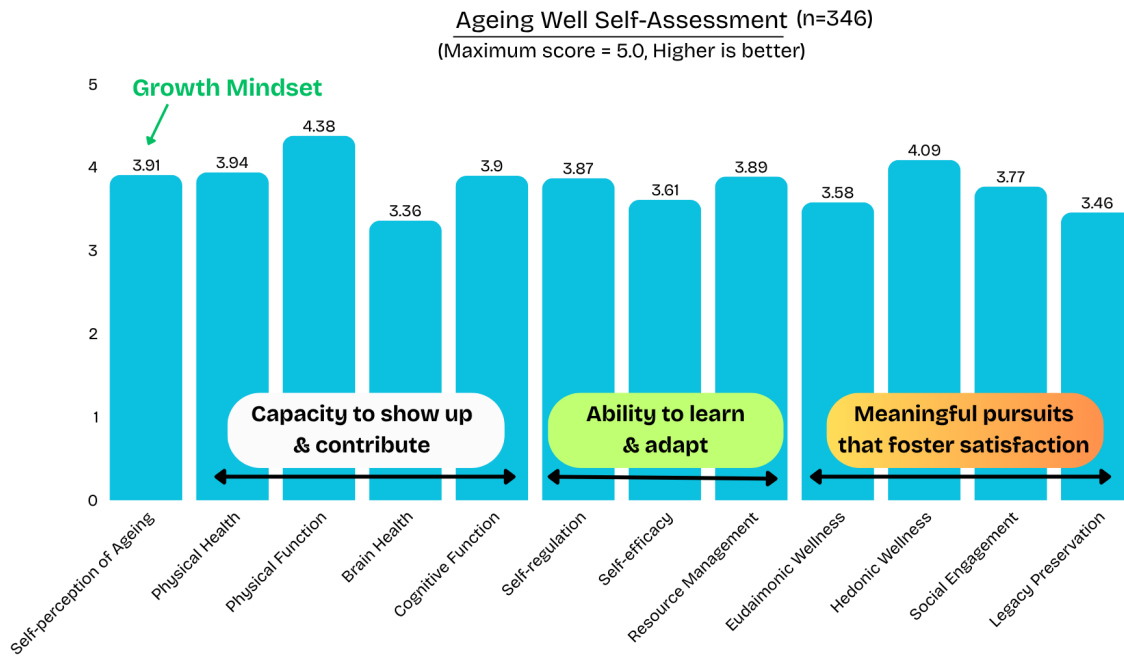
Skill Swap: A structured role-playing game that makes unfamiliar career roles tangible by placing participants inside them. Participants expand their sense of what is plausible, correct misconceptions about entire career categories, and discover transferable capabilities they had written off as irrelevant.



What's Next GPT: A custom AI tool that ingests the outputs of the preceding exercises — Journey Map patterns, Career Quest priorities, Skill Swap discoveries — and synthesises them with one's employment history into a personalised, explorable action space. The result is possibilities that feel earned rather than assigned — built from what the person has already discovered, not from generic market data.

Design-AI — Build Your Own Bot: Every participant builds a Custom GPT encoding their own domain knowledge or career application. The act of building — not watching, not using, but creating — is the final closure of the Self-Discovery Gap. A person who has built an AI tool from their own expertise has concretely discovered something new about what they know and can do.

The programme also addresses what most career programmes ignore: that physical health, cognitive vitality, and employability are inseparable at this life stage. The Ageing Well for Employability module treats these not as background context but as active levers: plannable factors that participants can work on alongside their career direction. The impact of this module is borne out in the data.



Self-assessment profile (copyright: Kang Kok Hua)

The combination of the above gives the programme its most distinctive structural feature: weaving AI capability with human connection and peer community – because self-discovery is not a solitary act, and career transitions are not navigated alone. Open-ended responses reveal what participants actually valued: feeling respected rather than patronised, being in a room of peers at the same life stage, and having someone at the table who ensured the AI was accessible to them personally.

EVIDENCE AT SCALE

What the Data Shows Across 16 Sessions

Across 16 sessions (October 2025–March 2026) and 373 participants, the findings tell a consistent story.

Career Clarity, Openness, and Agency

Participants leave with measurably more clarity and confidence.

Satisfied with <i>What's Next</i> programme	Can identify meaningful, achievable career goals	Clearer, more confident navigating career journey
98%	91%	88%

Of those who entered undecided about their next career direction, 80% left with a defined direction — most choosing to continue working in some form.

Entered undecided 69	Left with a direction 55 (80%)	Still undecided 14 (20%)
--------------------------------	--	------------------------------------

Mindset and Life-Stage Readiness

On the item "I trust myself to take charge and move forward, even when things are uncertain", the proportion scoring Strongly Agree rose from 58% to 66% — a +8.4 percentage point shift in the form of self-trust that determines whether broadened horizons translate into actual steps.

Across the programme, 93.1% of participants found the Ageing Well topic relevant to their career planning, and 93.8% said it helped them see ageing pillars as plannable factors and active levers rather than fixed conditions.

Insight Turned into Action Through Design•AI

AI-related items score highest across both workshop days because every participant uses it to build something personal: a Custom GPT encoded with their own career history, values, and priorities. Agreement that GenAI is important for career transition and useful for career exploration both exceeded 98%. Participants leave convinced that Design•AI belongs in their career toolkit.

“One key takeaway for me was witnessing the power of an AI-driven career navigation tool — almost like having a personalised pocket career coach. Trained on research and layered over our own life Journey Maps, Career Quest insights, and Skill Swap exercises, this tool can generate future and emerging career transition options that were both fascinating and thought-provoking. Some of them genuinely blew my mind! Another highlight was learning how to create a custom GPT. Seeing the sheer number of creative custom GPTs dreamt up and designed by the class was amazing — a real testament to what happens when people, design and AI come together.”

— Participant

THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

The distinctive structure of What’s Next is the integration of Design x AI x Domain, closing the Self-Discovery Gap that precedes the skills gap. This is clear from the quote below and elaborated in the table that follows.

“What stood out wasn’t the content. It was the learning design. What’s Next took a very different approach from conventional career programmes. Instead of expert talks or prescriptive advice, the experience was built around sense-making, experimentation and guided reflection. As a learning designer, it was refreshing to see how intentionally the workshop was structured for senior professionals... Career transitions, especially later in life, need more than advice. They need well-designed experiences that respect lived experience and translate insight into next steps.”

— Participant

Table 1: What's Next through the trilingual lens — Design × AI × Domain Knowledge

Skills and Tasks	Then (Past Wave of AI)	Now (Current Wave of AI)
Design	<ul style="list-style-type: none"> • Career guidance designed as information delivery: workshops, slide decks, coaching sessions that explain options • Participant positioned as recipient of expert advice; burden of synthesis falls on the individual • Scale achieved through standardisation: same content delivered to everyone 	<ul style="list-style-type: none"> • Career guidance redesigned as an embodied system: tools sequence physical, cognitive, and AI-mediated experiences to produce clarity • Participant positioned as builder and co-creator; AI synthesises the person's own inputs • Scale through structural design: high-contact delivery with facilitators and peers speeds up the pace at which participants take action
AI	<ul style="list-style-type: none"> • AI as information source: participants shown what AI can do but do not interact as builders • Interaction transactional and one-directional: AI provides outputs; participants receive them • AI literacy treated as conceptual knowledge: understanding what AI is, not how to use it purposefully 	<ul style="list-style-type: none"> • AI as teammate and personal asset: every participant builds a Custom GPT encoding their own experience and domain knowledge • What's Next GPT synthesises personal inputs into a personalised action space with specific, testable suggestions • AI literacy produced through doing: building a Custom GPT creates durable understanding and confidence
Domain Knowledge	<ul style="list-style-type: none"> • Mature professionals' decades of domain expertise treated as a credential, not a living asset 	<ul style="list-style-type: none"> • Participant expertise is the primary input to AI: tools are most powerful when participants bring rich domain knowledge, turning their expertise into a living asset

MORE READINGS

What's Next Programme: <https://www.sutd.edu.sg/academy-course/whats-next-reimagine-your-career-using-design-ai/>

SUTD Future of Innovation Lab: <https://www.sutd.edu.sg/future-of-innovation/x-intelligence/>

REFERENCES

Kolb, D. A. (1984). *Experiential Learning: Experience as the Source of Learning and Development*. Prentice-Hall.

Marshall, J. (2016). Embodied cognition and design research. In P. Lloyd & E. Bohemia (Eds.), *Proceedings of DRS 2016*.

Mezirow, J. (1991). *Transformative Dimensions of Adult Learning*. Jossey-Bass.

Merriam, S. B. (2004). The role of cognitive development in Mezirow's transformational learning theory. *Adult Education Quarterly*, 55(1), 60–68.

Poon, K. W., Willems, T., & Liu, W. (2023). The future of expertise. In Lee, W. O. et al. (Eds.), *International Handbook on Education Development in Asia-Pacific*. Springer.

World Economic Forum. (2025). *The Future of Jobs Report 2025*. World Economic Forum.

Ministry of Manpower Singapore. (2024). *Re-employment of older employees*. <https://www.mom.gov.sg>

SkillsFuture Singapore. (2025). *SkillsFuture Movement*. <https://www.skillsfuture.gov.sg>

MyCareersFuture. (2025). *Jobs to consider in the second half of your career for mature workers*. <https://www.mycareersfuture.gov.sg>

GovTech Singapore. (2025). *Considering a mid-career switch in Singapore? Meet Career Kaki*. <https://www.tech.gov.sg>

Citizens: All Can Innovate

Voices from Citizens

EXECUTIVE SNAPSHOT

The most powerful results came from the interplay between designing AI capabilities to meet deep lived experiences. Across Fitness Education, Healthcare & Rehabilitation, and Parenting & Home Learning, each of the four citizens below discovered the power of AI through a moment of genuine need. The result was amplification of skills, agency, and human connections.

CONTEXT

When Ordinary People Do the Extraordinary With AI

What happens when a fitness instructor uploads a student’s posture photo to GPT-4o? When a father points his phone at his daughter’s homework and asks ChatGPT to explain grammar? When a patient recovering from a rare blood disorder uses AI to decode a 40-page insurance policy? Or when a mother wants to help her son with his mother tongue subject?

These are not stories about technologists or data scientists. They are stories about citizens like you and me. Each discovered AI through a moment of genuine need. And each found that the most powerful results came from the interplay between designing AI capabilities to meet their own deep lived experiences. In doing so, they also amplified their skills and connection to other humans.

Contributor	Domain	Description
The Fitness Instructor	Health & Wellness	Nearly two decades of Pilates and fitness teaching. Uses AI as a brainstorming buddy and diagnostic second pair of eyes, while bringing embodied teaching and emotional sensing.
Patient	Healthcare & Insurance	Diagnosed with rare blood disorder (POEMS Syndrome). Explored how AI could have accelerated medical journey. Used ChatGPT for insurance policies, test results, and rehabilitation.
The Working Father	Home Education	Father of three (ages 10, 7, 5) who practises “guided co-use” of ChatGPT, adapting supervision by each child’s age and personality. Developed a layered protocol for family AI use.
The Working Mother	Home Education	Initially anti-device, she became a Custom GPT creator to support her son’s Chinese language learning. Built “Ninja Comics”—a personalised learning tool combining interests with structured practice.

THE TRIGGER MOMENT

Domain Needs and the Need for Design: How AI Entered Everyday Life

Across all four of them, AI adoption began with a moment that was often spontaneous, and yet always rooted in genuine need.

The father’s daughter was stuck on a grammar question. He wanted to do his best for her. On impulse he pointed his phone at the worksheet. ChatGPT delivered a clear explanation, and the children were immediately hooked. From that single moment, AI became woven into the family’s daily learning routine.

The fitness instructor’s journey followed a similar arc. When she asked GPT-4o to explain ribcage breathing, it delivered a structured explanation with metaphors that was faster and more accessible than any textbook. And the big shift came when she began uploading photos to get a second opinion on movement form.

“It’s like GPT was my second pair of eyes. But I had the hands, the heart, and the judgement.”
 — Fitness Instructor

The patient’s health journey offered a more poignant entry point. His path to diagnosis was marked by fragmented care. Each symptom – headaches, difficulty walking, high platelets, protein in urine – managed in isolation.

No one had the full picture.

It was only after advanced tests revealed POEMS Syndrome that the pieces fell into place. Looking back, the patient recognised that AI could have served as an early “system thinker,” clustering symptoms and suggesting questions to ask specialists.

CLINICAL LABORATORY				
SPECIAL PROTEIN STUDIES		Collection Date: 19 July, 2024	Time: 10:05:00 AM	
	RESULT	REFERENCE	UNITS	RESULT
Serum Kappa FLC	16.60	(6.70 - 22.40	mg/L	
Serum Lambda FLC	150.00	(8.30 - 27.00	mg/L	
Serum Kappa/Lambda Ratio	0.11	(0.31 - 1.56		

HAEMATOLOGY				
		Collection Date: 19 July, 2024	Time: 10:05:00 AM	
	RESULT	REFERENCE	UNITS	RESULT
RBC	红细胞计数	4.95	(4.01 - 5.36	mB/L
HGB	血红蛋白	15.2	(13.0 - 16.9	g/L
HCT	红细胞比积	46.6	(35.6 - 50.6	%
MCV	红细胞平均体积	94.1	(79.1 - 97.4	fL
MCH	红细胞平均血红蛋白	30.7	(25.5 - 33.1	pg
MCHC	红细胞分布血红蛋白度	32.7	(30.9 - 34.8	g/dL
RDW	红细胞体积分布宽度	15.0	(12.2 - 15.4	%
MPV	血小板平均体积	8.0	(6.8 - 10.4	fL
PLT	血小板计数	496	(140 - 460	thB/L
WBC	白细胞计数	7.58	(4.20 - 11.00	thB/L
Neutrophils	中性粒细胞	77.5	(42.5 - 74.0	%
Lymphocytes	淋巴细胞	12.5	(21.1 - 43.2	%
Monocytes	单核细胞	2.8	(4.0 - 10.0	%
Eosinophils	嗜酸性粒细胞	6.0	(0.0 - 5.6	%
Basophils	嗜碱性粒细胞	1.2	(0.0 - 1.0	%
Neutrophils #		5.87	(1.90 - 7.20	thB/L
Lymphocytes #		0.95	(1.20 - 3.40	thB/L
Monocytes #		0.21	(0.20 - 0.70	thB/L
Eosinophils #		0.45	(0.04 - 0.56	thB/L
Basophils #		0.09	(0.00 - 0.09	thB/L

Looking at this report on its own, without knowing your history, here's what it shows and could suggest to a clinician:

Special Protein Studies

- Serum Kappa FLC: Normal (16.6, ref 6.7-22.4)
- Serum Lambda FLC: Markedly elevated (150, ref 8.3-27.0)
- Kappa/Lambda Ratio: Very low (0.11, ref 0.31-1.56)

Interpretation:

There is monoclonal excess of lambda free light chains. This kind of abnormality often raises concern for plasma cell or lymphoid disorders (e.g. multiple myeloma, AL amyloidosis, MGUS, or related conditions). The ratio imbalance is the key red flag here.

Haematology

- Red cell parameters (RBC, Hb, HCT, MCV, MCH, MCHC, RDW): All within normal range → no anemia.
- Platelets (406, ref 140-460): Normal.
- Total WBC (7.58, ref 4.20-11.00): Normal overall count.

Differential white cells:

- Neutrophils: Slightly high (77.5% vs 42.5-74.0).
- Lymphocytes: Low (12.5% and absolute 0.95 vs 12.0-3.40).
- Monocytes: Low (2.8% vs 4-10).
- Eosinophils: Mildly elevated (6.0% vs 0-5.6).
- Basophils: Upper end of normal (1.2%).

Interpretation:

The white cell distribution shows a relative neutrophilia and lymphopenia, with mild eosinophilia. This can reflect:

- recent infection, inflammation, or stress,
- effect of medications (e.g. steroids),
- immune system imbalance.

Overall Picture (without prior knowledge)

- Strong abnormality in protein studies → suggests an underlying plasma cell disorder producing excess lambda light chains. This is the most significant finding.
- Stable red cells and platelets → no evidence of anemia or marrow suppression at present.
- White cell shifts (neutrophil up, lymphocytes down, eosinophils up) → could point to immune dysregulation, stress, or treatment effect.

If a doctor were reviewing this in isolation, they would likely:

- Flag the lambda light chain overproduction as the key abnormality.
- Consider ruling out plasma cell disorders with further work-up (e.g. bone marrow studies, imaging, SPEP/UPEP, immunofixation).
- Note that blood counts are otherwise preserved, which is reassuring.

Would you like me to now reframe this in simple, patient-friendly language (like how a doctor might explain it in clinic)?

Figure 1: The patient’s clinical laboratory report—AI could help patients interpret complex test results in plain language.


In every case, the pattern is the same: AI entered through a door that domain knowledge had already opened.

AI ROLES

AI as Tool, Teammate, and/or Neither. One of the most illuminating themes is the fluid, shifting role that AI plays. The father articulated a framework that resonated across every project: AI operates as a tool when it simply provides information in one direction; it becomes a teammate when it co-drives learning bidirectionally; and it can slide into being a mere toy when interactions go off-track.

AI as Tool: Speed, Structure, and Retrieval. At its most basic, AI served as a powerful information retrieval and structuring tool. The fitness instructor used it to plan 60-minute classes and recall anatomy knowledge. The patient used AI to upload insurance policies and ask targeted questions. The mother used it to conduct spelling tests and grade writing in a subject she had little knowledge in.

AI as Teammate: Co-Creation and Partnership. The more significant shift occurred when AI moved to co-creating. For the fitness instructor, GPT-4o analysed a student’s photo and independently identified the same postural issue she had noticed, and together shortlisted the solution. The collaboration worked precisely because both human and AI brought distinct strengths.



Evaluation of Your Toe Tuck:

- Ball of the Foot:** It appears that the ball of your foot is not fully making contact with the mat. This could suggest:
 - Limited ankle or toe mobility.
 - Tightness in the arches of your foot or plantar fascia.
- Toes Curled or Over-Tucked:** It seems your toes might be over-curved, which can create discomfort. Ideally, the toes should be tucked under with the toes extending from the base, not tightly curled.

Suggestions to Improve Toe Tuck:

- Stretch and Mobilize Toes:**
 - Gently flex and extend the toes daily.
 - Use your hands to stretch the toes backward (into an extended position) to increase range of motion.
- Calf and Plantar Fascia Release:**
 - Roll the bottom of your foot over a massage ball or a tennis ball to release tension in the plantar fascia.
 - Perform calf stretches to improve ankle flexibility.
- Use a Prop for Support:**
 - Place a folded towel under the balls of your feet to reduce the range and avoid discomfort while you build flexibility.
- Distribute Weight Evenly:**
 - Focus on shifting some of your weight toward your hands and lift slightly out of the toes to avoid overloading them during practice.

Figure 2: GPT-4o’s analysis of a “tucking the toe” movement matched the instructor’s own diagnosis.

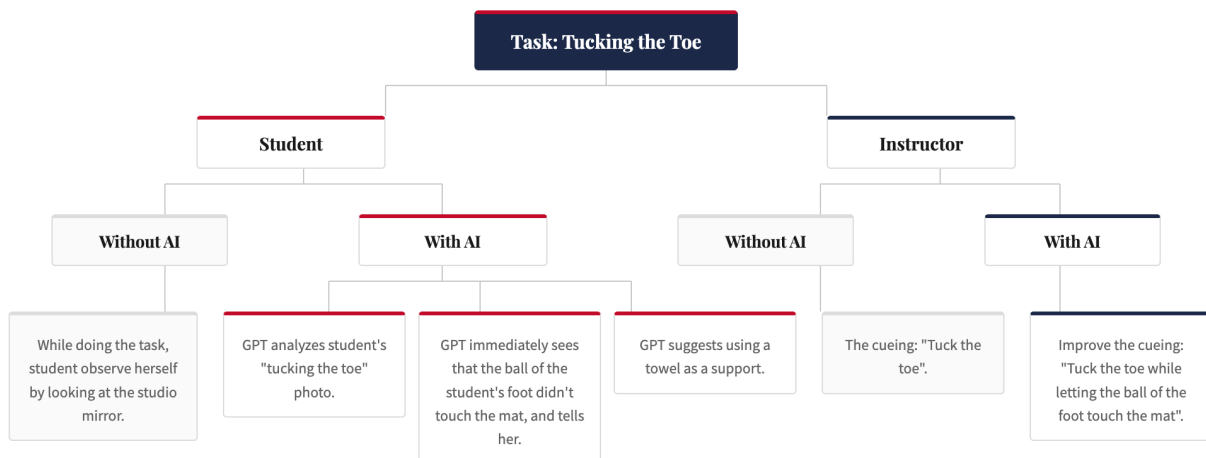


Figure 3: The Human x AI diagram—showing how student, instructor, and AI interact for the “tucking the toe” task.

For the mother, GPT became a teammate in a different sense: it took on the emotionally draining parts of practice—asking spelling questions, grading responses, giving patient feedback—while she sat beside her son, observing effort, body language, and engagement.

The patient experienced AI as a companion during rehabilitation. ChatGPT provided form corrections and flagged posture issues, lightening the mental load during long recovery nights.

“AI didn’t replace doctors, but it lightened the load and reduced the loneliness.”

— The Patient, Health Case Study

When AI Is Neither: The Irreducibly Human. Every contributor was equally clear about where AI fell short. The fitness instructor emphasised that AI cannot feel – it cannot sense when a student is struggling, crack a joke, or adapt a movement in real time. The father pointed out that GPT cannot read his children’s emotions.

“GPT can teach the what. But only a human can teach the how it feels.”

— Fitness Instructor

RESPONSIBLE INTEGRATION

Human-AI Working Together

A striking consensus emerged: AI outputs must be filtered through human judgement, and the quality of that filter depends entirely on domain knowledge. The fitness instructor ignored a GPT-generated cue that felt “too generic, maybe even risky.” The patient found that ChatGPT sometimes gave “overly generic or even misleading cues.” The mother knew when her son was guessing answers even though GPT could not tell.

AI is also a mirror. If you ask good questions, you get good insights. But you need domain knowledge to ask the right questions and to assess the answers, demonstrating that AI design is epistemically dependent on domain knowledge. The quality of an AI system is fundamentally constrained by the depth of human expertise embedded in its design.

“AI makes you more human, if you use it wisely.”

— Fitness Instructor

The father introduced a new framework: “guided co-use.” He described a layered protocol shaped by each child’s age and maturity. His eldest (10) could use GPT relatively independently. His middle child (7) needed active redirection, what the father called “needle moments” when human judgement stitches learning back on track. His youngest (5) simply observed, absorbing how the family interacted with AI.

“WHEN GPT JOINED OUR HOME” - Parent 1

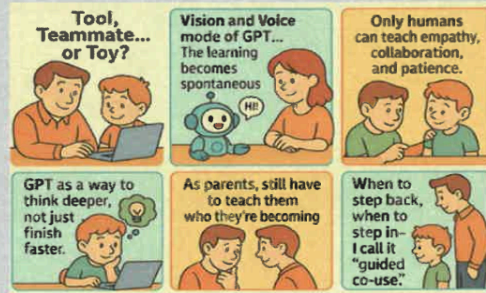
Characters:



Researcher: Curious, methodical, and passionate about the intersection of education and AI.



Working Father: Intentional, creative, grounded. A working father of three, he guides his kids in personalized learning using ChatGPT.



A reconstructed conversation between a Researcher and a Father of three children—ages 10, 7, and 5—on how ChatGPT reshapes daily learning, parenting, and ethical decision-making in the home:

RESEARCHER

How did GPT first become part of your home learning routine?

FATHER

I had a GPT Plus subscription and saw online that the Vision Mode had been released. Around that time, my daughter was doing revision and got stuck on the difference between *who* and *whom*. My wife and I debated it, but we weren't confident. So I pointed my phone at the worksheet and asked GPT. It gave a very clear explanation, with examples, and the kids were immediately hooked.

2. To _____ does that wallet belong?
 (1) who (2) whom
 (3) which (4) whose

RESEARCHER

So that one moment changed things?

FATHER

Yes. It wasn't just the explanation. It was how quickly GPT responded, how accurate it sounded, and how it kept offering more context when we asked follow-up questions. What surprised me was how curious the kids became. They didn't just want answers. They wanted to understand.



My son loves the explanations & the feedback. especially to listen to them in voice mode with fun tone.

My kids love the voice mode too!

for parents supporting their children in learning a language that isn't their own language/MT can be a real struggle. I can relate to this with my son learning Chinese. LLMs can play a big role supporting parents overcome these struggles.

Thank You! It helps to know I am not alone on this journey

Figure 5: “When GPT Joined Our Home”—a comic strip illustrating family AI interactions.

Both parents raised ethical concerns directly transferable to other domains. Over-reliance was the primary worry: if children stop thinking and let GPT do the work, they lose critical thinking development. Privacy was another concern, particularly around Vision Mode with young children. These parenting protocols mirror the professional practices in other case studies. Such as the fitness instructor’s judgement about when to accept or reject AI suggestions, and the patient’s verification against his doctors’ guidance.

“GPT can teach the what. But we, as parents, still have to teach the why, and the who they’re becoming.”

— Working Father

The lived experiences of all four reveal an additional critical insight: as **AI** knowledge grows, **design** knowledge evolves in response, and **domain knowledge** remains the constant anchor. The three forms of knowledge co-evolve. Working well with AI this way creates space for deeper human engagement, enables faster pattern recognition, and amplifies creative output.

REFERENCES

FOI Explainer Issue 7 – Flexible Intelligence: Stretching Skills, Creativity and Imagination

<https://www.sutd.edu.sg/wp-content/uploads/2025/06/X-Intelligence-Issue-7.pdf>

FOI Explainer Issue 10 – Parent Intelligence: Learning with AI at home

<https://www.sutd.edu.sg/wp-content/uploads/2025/11/X-Intelligence-Issue-10.pdf>

FOI Explainer Issue 11 – Flexible Intelligence: A What-if Journey Through Medical Treatment, Insurance & Rehabilitation

<https://www.sutd.edu.sg/wp-content/uploads/2025/11/X-Intelligence-Issue-11.pdf>

Huang, Q., & Poon, K. W. (2025). Human and AI collaboration in fitness education: A longitudinal study with a Pilates instructor. arXiv:2506.06383.

<https://arxiv.org/abs/2506.06383>

Future of Cities: From Citizen Participation to Citizen Co-creation

SUTD x Community: Prototyping Citizen-Created Cities of Tomorrow

EXECUTIVE SNAPSHOT

~40 Student Participants From three local educational institutions	5 Custom AI Agents As co-facilitators alongside human facilitators	>8 Master Plan Proposals Conceptual proposals produced	3 hrs Workshop Duration Single-session participatory studio
--	--	---	---

CONTEXT

In the past two years, students from SUTD's Master of Science in Urban Science, Policy and Planning (MUSPP) programme have been experimenting with how Design AI can empower citizens to co-create the city of tomorrow. This goes beyond the citizen participation of today, to designing and creating the future of the city together. In 2025, the focus was on reimagining the future of education. In 2026, students reimagined the future of master planning in the age of AI.

KEY FINDINGS

- In 2025, three custom AI agents reimagined learning
- In 2026, eight coherent, strategically distinct master plan proposals were produced in a single three-hour session by participants with widely varying planning experience
- Five custom AI agents were embedded as co-facilitators throughout the process
- AI-supported rapid prototyping enabled teams to visualise, test against stakeholder personas, receive critique, and iterate, all within the same session
- Post-workshop Reimaginator app transforms eight proposals into integrated options for planning scenarios

'TRILINGUAL' SKILLS

(DESIGN x AI x DOMAIN)

Design: Compressed the planning and prototyping cycles regardless of prior training.

AI: Teammate, not tool e.g. five purpose-built agents act as specialist co-facilitators covering site intelligence, expert perspectives, stakeholder simulation, design method guidance, and critical review.

Domain Knowledge: Shifts from prerequisite to evaluative lens; the critical skill is knowing when an AI output is right or wrong, and articulating why.

Reimagining the Future of Education. In 2025, MUSPP students experimented on how they could recreate a research methods course using Design-AI. As part of a broader exploration into citizen co-creation, students were challenged to go beyond simply using AI, to designing the future with it.

Without writing a single line of code, student teams built three custom GPT tools: the Interview Companion, a multimodal simulator for practising qualitative research interviews; the Urban Observer, a Socratic field guide for structured urban observation; and Buddy Buddy, a personalised pre-class study companion. The results were promising (you can read about what happened with the Interview Companion in one of the preceding Strategy Use Cases in this same publication).

The experiment and the promising results raised a provocative question for the future of cities: if citizens can be empowered to build their own AI tools, what else might they co-create?

Co-creating the Future of Master Planning. What would it take for communities to be genuine co-designers of the cities they live in, rather than respondents to plans already drawn? In 2026, MUSPP students designed and facilitated a workshop to explore that question.

Rather than treating AI as a finishing tool, the workshop embedded Design-AI into the facilitation process itself. Five custom AI agents worked alongside human facilitators across all eight teams simultaneously.

The result: eight coherent, strategically distinct conceptual master plan proposals were produced within a three-hour window, by people who brought widely varying levels of planning experience, AI fluency, and design backgrounds.

AI AS CO-FACILITATOR

The Custom Agent Suite

Table 1: The five AI co-facilitators and their roles in the workshop.

#	AGENT	ROLE IN THE STUDIO
1	Site Expert	A resident expert pre-loaded with current site data — land use maps, transport connections, tenant mix, government plans, and comparable global precedents. Teams could query it at any point like consulting a knowledgeable colleague.
2	The Great Master	Draws on the wisdom of selected architects and urban planners to generate ideas, debate trade-offs, and bring global design precedents into the room.
3	Personas' Perspectives	Simulates eight distinct stakeholder voices — from local students and researchers to elderly residents, investors, event managers, and tourists — giving teams live feedback on how different groups would experience their proposals.
4	Design Methods Guru	A design thinking advisor built on SUTD's Design Thinking Methods, guiding teams to the right tool for their current challenge — from affinity mapping and 'How Might We' framing to storyboarding and prioritisation matrices.

5	Guest Crit	Critiques proposals against SUTD’s design thinking rubric (Define, Discover, Develop, Deliver), providing pointed feedback from engineering and architectural perspectives.
---	-------------------	---

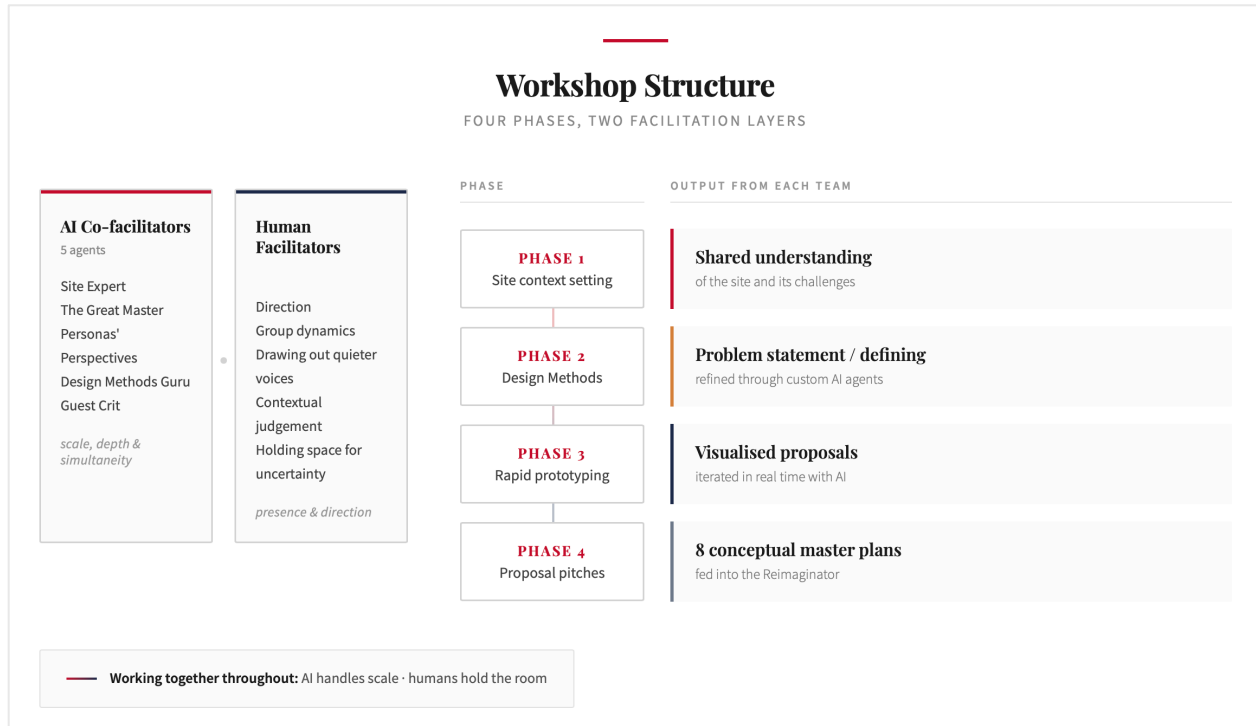


Figure 1: Workshop structure — four phases and two facilitation layers.

What AI Co-Facilitation Made Possible. Using AI for rapid prototyping produced two effects that reinforced each other. The first was speed: rather than spending time describing an idea in the abstract, teams could put a spatial concept on screen within minutes so that everyone could see, react to, and build on. The second effect was depth: because teams could generate an idea, visualise it, test it against stakeholder personas, receive feedback, refine it, and visualise again — all within the same session — discussions went further than they would have through conversation alone.

“They work together as a continuous, powerful loop — master planning, design thinking, and AI tools.”
— Studio participant

After the workshop, four themes came through consistently: planning felt different at this scale — participants described thinking about urban problems more holistically; AI was experienced as a creative partner, not a shortcut, with around two-thirds identifying real-time spatial visualisation as their team’s primary use; some teams’ problem statements shifted mid-session, triggered by AI agent challenges rather than facilitator instruction; and the design thinking framework provided a shared language and sequence that teams could orient around, particularly when working through disagreement.

From Eight Proposals to Hundreds of Possibilities

In most planning workshops, the proposals produced on the day represent the end of the road. Here, they were the beginning. A synthesis application called the Reimaginator was developed using Google AI Studio. It takes the ideas from all eight proposals and allows users to mix and match elements, adjust priorities, and generate new combinations. Each combination is rendered as a visual and an animated flythrough with an analysis of the planning trade-offs involved. Rather than community participation ending when the session ends, the Reimaginator keeps those ideas alive and explorable, in forms that anyone i.e. not just planning professionals, can engage with to explore options for the future.

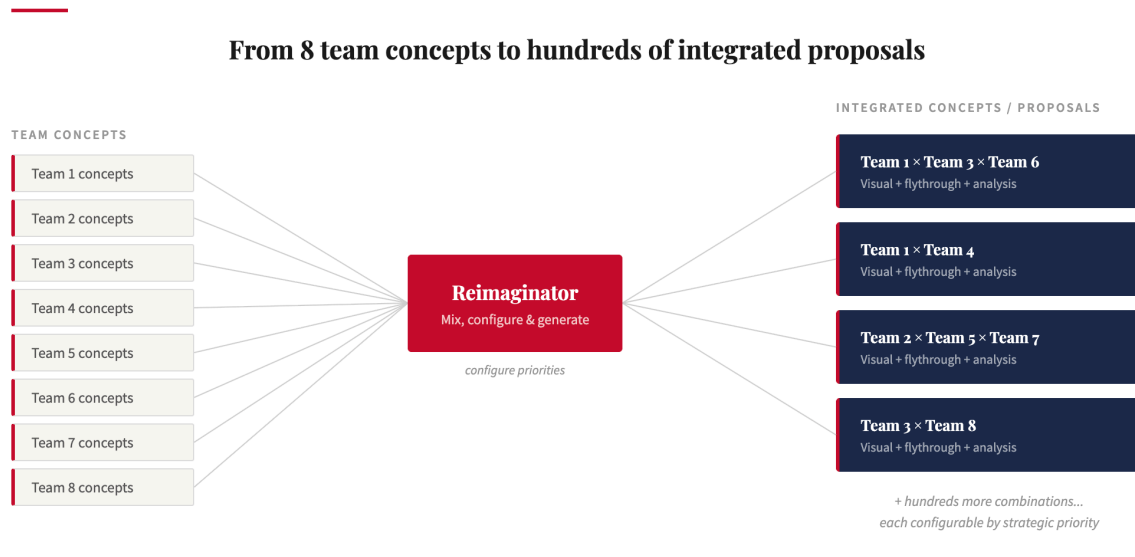


Figure 2: From eight team concepts to hundreds of integrated proposals via the Reimaginator.

AI as tool vs AI as co-facilitator

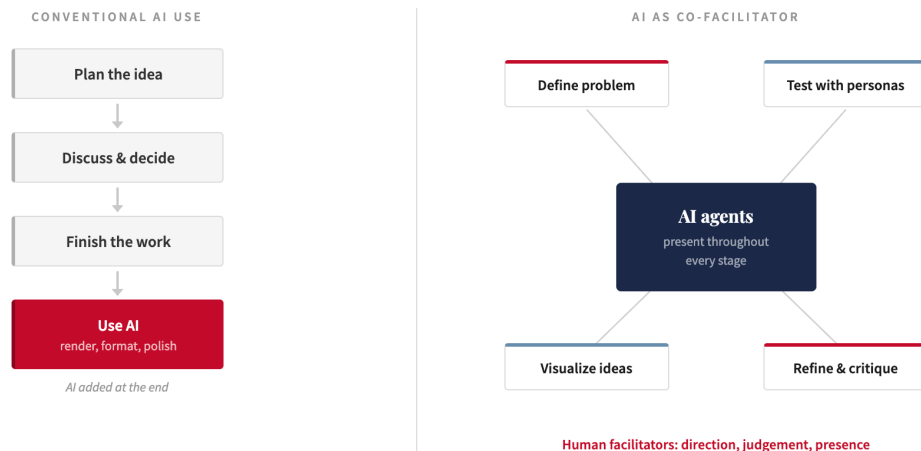


Figure 3: AI as tool (conventional) vs AI as co-facilitator (this workshop).

THE TRILINGUAL FRAMEWORK

Building ‘Trilingual’ Skills and Tasks (Design x AI x Domain)

Table 2: The AI-Augmented Master Planning Studio through the trilingual lens.

SKILLS & TASKS	THEN (CONVENTIONAL STUDIO)	NOW (AI-AUGMENTED STUDIO)
Design	Extended timelines gate-kept participation; meaningful contribution required prior technical knowledge. Scaling meant adding more human facilitators.	Compressed cycle: research, analysis, visualisation, and iteration within a single session — accessible regardless of prior training. AI agents co-facilitate all teams in parallel; human facilitators shift to holding space for judgement and creative risk.
AI	AI as production tool: one-directional, used after decisions are made — rendering, formatting, summarising. Visualisation required specialist software.	Five purpose-built agents as specialist co-facilitators covering site intelligence, expert perspectives, stakeholder simulation, design method guidance, and critical review. AI is a teammate, not a tool. Post-workshop synthesis app extends AI beyond the session.
Domain Knowledge	Domain expertise was prerequisite; weaker teams produced thinner outputs regardless of effort. Tasks front-loaded before spatial exploration could begin.	Domain knowledge shifts from prerequisite to evaluative lens — the critical skill is knowing when an AI output is spatially or contextually wrong. New tasks: encoding planning context into agents; curating synthesis outputs against real planning standards.

The same pattern runs through all three dimensions.

Good design is about the decisions that are easy to overlook. How the AI agents are set up, how the facilitation is structured, how the outputs are used after the session. AI works best not as an occasional tool but as an active partner built into the process from the start.

And planning expertise does not disappear. It shifts from being a barrier to entry into becoming the critical filter through which AI outputs are evaluated and refined.

Planning with communities — rather than simply for them — remains an ongoing aspiration in urban practice. The cities of the future will face planning challenges of increasing complexity, from climate adaptation to rapid demographic change, that no single profession can navigate alone. What this workshop offers is one grounded example of how co-creation with AI as Tool, Teammate, and/or Neither might help bring more voices meaningfully into that process.

METHODOLOGY

This report is a first draft of the future. It captures the first wave of case studies from ongoing research, experimentation, and deployment across Singapore's organisations, schools, and workforce development institutions. Additional case studies will be added as new partnerships produce evidence and new patterns emerge. The intention is less to present a finished body of work, and more to share what we have learnt so far. And to do so quickly enough for it to be useful.

Research Methods

The evidence presented across these case studies draws on multiple methods of insight collection. These include oral histories and structured interviews with programme participants and organisational leaders; pre- and post-workshop survey data with validated instruments; analysis of publicly available documents, annual reports, and institutional records; direct observation of workshop delivery and facilitation; and review of participant-generated artefacts including AI prototypes, custom GPTs, and workflow analyses. Each case study identifies its primary evidence base.

On Working with AI as a Tool, Teammate, and/or Neither

This set of Strategy Use Cases is itself an experiment. To practise what we preach, AI was used as Tool, Teammate, and/or Neither throughout the research, writing, and publication process. All findings were verified by human researchers and validated by the participating organisations.

Most importantly, we found that who did what was the least interesting question.

Sometimes we did better than the AI.

Sometimes the AI did better than us.

Oftentimes, we did better because we pushed each other and the AI hard.

Other times, the AI pushed us to be better.

In the end, it mattered less to us who or what was better.

What mattered more was which outcome was better for you, the reader.

And that is precisely the high bar we set for ourselves (and the AI).

We hope what we wrote is worthy of that high bar for you.



DESIGNING A HUMAN-AI INNOVATION EDGE FOR ENTERPRISES

Strategy Use Cases on Building "Trilingual Skills"
(Design x AI x Domain) for the Future of Innovation

Discussions about AI often focus on what AI can do. And what AI can do to people. The Strategy Use Cases in this briefing focus instead on what AI can do for people. And what people can do with AI. They show leaders ten decisions they can make that put humans at the heart of our age of AI.

April 2026

This publication was developed by the Singapore University of Technology and Design (SUTD), and supported by SGTech and SkillsFuture Singapore (SSG).