

AI for Teachers

AI safety, Authenticity, and Redesigning Assessment with AI

Workshop | Christchurch | 16th April 2026



Benny Pan InspiraEd.com



Ko Taupiri te maunga

Ko Waikato te awa

No China ahau

Ko Benny toku ingoa

Ko Pan toku whaanau

Who am I?

Google TF

2017 Machine Learning
Google TensorFlow

International Project Coordinator

Asia Europe Foundation
International AI Training
Program across 18 countries

Science Teacher + AI Co-Lead

Design and roll out
school-wide AI curriculum
Rototuna Senior High
School

Project Lead

AI Forum NZ Education
Workstream


AI and Assessment
Project

academyEX

Master of Tech Future
AI and Critical Thinking in
Education

InspiraEd

Founder
AI Consulting and education



33%

Of the time Chatbot validate and encourage violent thoughts

7

Ongoing lawsuits related to AI-induced harm in teens

440,419

Reports on AI-generated child exploitations in 2025

Turn & Talk

Which stat surprised you most? What does it mean for your school and community?

33% of the time Chatbot validate and encourage violent thoughts

7 Ongoing lawsuits related to AI-induced harm in teens

440,419 reports on AI-generated child exploitations in 2025

You Will Leave With 3 Things

1. A visual framework for how students & teachers may use AI with use case
2. Custom GPTs and Gems ready to use to help AI policy drafting and prompting
3. Slash commands to save time

Your Concerns

What you told us — and what it means.

Your voice

78% Students misusing AI or over-relying on it

76% Academic integrity / cheating

24% Don't know effective AI prompts

The curriculum doesn't mention AI.

Students are already using it.



AI LITERACY

How Generative AI Actually Works

AI learns from and mimics human thinking

How AI Actually Works

Google [quick draw](#)
Teachable Machine [demo](#)

AI learns from human examples – trillions of them


It predicts the **next token**. A token can be a word, part-word or a pixel.

That's it. Trillions of examples, and it picks the most likely next token.

Pattern-matching mimic how human interprets the external world


That's why WHAT you give it changes EVERYTHING you get back.

Camel drawings from billions of users



THE AI KITCHEN: TRAINING, INFERENCE, & THE KNOWLEDGE CUT-OFF

AI LIFECYCLE TIMELINE



SUMMARY: The AI (Chef) only knows what happened in the **TRAINING PHASE** (Blue Zone). It uses that fixed 'Internal Cookbook' to work in the **INFERENCE PHASE** (Orange Zone) but it is blind to any new information or events that happened after the **KNOWLEDGE CUT-OFF DATE**.

⚠ This paper contains text that might be offensive. ⚠

PERSONA FEATURES CONTROL EMERGENT MISALIGNMENT

Miles Wang* Tom Dupré la Tour* Olivia Watkins* Alex Makielov* Ryan A. Chi*
Samuel Miserendino Johannes Heidecke Tejal Patwardhan Dan Mossing*

OpenAI

ABSTRACT

Understanding how language models generalize behaviors from their training to a broader deployment distribution is an important problem in AI safety. *Bealey et al. (2025b)* discovered that fine-tuning GPT-4o on intentionally insecure code causes “emergent misalignment,” where models give stereotypically malicious responses to unrelated prompts. We extend this work, demonstrating emergent misalignment across diverse conditions, including reinforcement learning on reasoning models, training on various synthetic datasets, and in models trained without safety training. To “fix” the mechanisms behind this generalized misalignment, we introduce a “fixing” approach using sparse autoencoders to compare internal model states before and after fine-tuning. This approach reveals several “mis-

BUSINESS INSIDER

NASDAQ ▲ +0.39% S&P 500 ▲ +0.52% AAPL ▲ +0.08% NVDA ▼ -0.06% MSFT ▼ -0.02% AMZN ▼ -0.34% META ▼

EXCLUSIVE

Scale AI exposed sensitive data about clients like Meta and xAI in public Google Docs, BI finds

By Charles Rollet, Effie Webb, Shubhangi Goyal and Hugh Langley



505.22954v1 [cs.AI] 29 May 2025

Darwin Gödel Machine: Open-Ended Evolution of Self-Improving Agents

Jenny Zhang^{1,2} Shengran Hu^{1,2,3} Cong Lu^{1,2,3} Robert Lange^{1,3} Jeff Clune^{1,2,4}
¹University of British Columbia ²Vector Institute ³Sakana AI ⁴Canada CIFAR AI Chair
{jennyzt, srhu, conglu}@cs.ubc.ca, robert@sakana.ai, jeff.clune@ubc.ca

Abstract

Most of today's AI systems are constrained by human-designed, fixed architectures and cannot autonomously and continuously improve themselves. The scientific method, on the other hand, provides a cumulative and open-ended system, where each innovation builds upon previous artifacts, enabling future discoveries. There is growing hope that the current manual process of advancing AI could itself be automated. If done safely, such automation would accelerate AI development and allow us to reap its benefits much sooner. This prospect raises the question of how AI systems can endlessly improve themselves while getting better at solving relevant problems. Previous approaches, such as meta-learning, provide a toolset for automating the discovery of novel algorithms but are limited by the human design of a suitable search space and first-order improvements. The Gödel machine [116], on the other hand, introduced a theoretical approach to a self-improving AI, capable of modifying itself in a provably beneficial manner. Unfortunately, this original formulation is in practice impossible to create due to the inability to prove the impact of most self-modifications. To address this limitation, we propose the Darwin Gödel Machine (DGM), a novel self-improving system that iteratively modifies its own code (thereby also improving its ability to modify its own codebase) and empirically validates each change using coding benchmarks. In this paper, the DGM aims to optimize the design of coding agents, powered by frozen foundation models, which enable the ability to read, write, and execute code via tool use. Inspired by biological evolution and search, the DGM maintains an archive of generated code and iteratively refines it. In this search, the DGM tries to create a new self-improving agent.

Agentic Misalignment: How LLMs could be insider threats

21 Jun 2025

Alignment

Highlights

- We stress-tested 16 leading models from multiple developers in hypothetical corporate environments to identify potentially risky agentic behaviors before they cause real harm. In the

Judge backs AI firm over use of copyrighted books


A copyright judge has ruled in favor of an AI firm, stating that its use of copyrighted books for training is fair use.

Natalie Sherman and Lucy Hooker
BBC News



Netsafe reports 68% rise in sextortion in NZ as AI deepfake threats increase

By Eva de Jong
Multimedia journalist - NZ Herald · 5 May, 2025 05:55 AM · 5 mins to read



Satya Nadella wants AI to solve real problems after Microsoft cuts 6,000 jobs

News / Technology / News / Satya Nadella wants AI to solve real problems after Microsoft cuts 6,000 jobs

cuts 6,000 jobs, more layoffs likely in July

In a recent interview, Microsoft CEO Satya Nadella emphasised that AI's real test lies in fixing real-world issues, and not just demos.

Listen to Story



Issues around AI

Issues with generative AI in Education

	What is it?	How it affects education	What skills it affects negatively	Examples	Mitigation strategies
1.Hallucination and misinformation	LLM hallucinates due to the non-deterministic nature, AI can make mistakes, or unexpected outputs with a confident tone . This affects both accuracy and relevance .	⊖ Errors wrapped in elegance and confidence ⊖ Correct but irrelevant info ✓ Unexpected great ideas (for ideation)	•Analytical thinking •Evaluation	✓ Same prompt generates different ideas each time ⊖ Hard to detect errors in Deep Research ⊖ Irrelevancy can be very subtle	🔗 Upload document 🔗 Google Notebook LM (strictly grounded to your docs) 🔗 Turn on internet access and double check sources
2. Sycophancy and echo chamber effect	When AI knows your preferences and your history really well, it tends to produce results that please you, and align with your value system	⊖ Discourage thinking outside the box ⊖ Intensify echo chamber effect as it knows you more ⊖ Product-driven ✓ AI more aligned with your values	•Creativity •Divergent thinking	✓ Highly relevant to your situation, no need to explain ⊖ AI likely gives you solutions similar to what you tried in the past	🔗 Turn off memory if possible 🔗 Temporary mode 🔗 Thinking buddy custom AI 🔗 AI Studio compare mode to see more perspectives
3. Cognitive atrophy <u>ResearchTL & Dr</u>	Outsourcing thinking to AI could potentially make our brains less capable . . . to some extent	⊖ Less friction → shallow learning ⊖ Less developed higher order thinking ✓ Boost productivity	•Critical thinking •Transfer from concept to solution	⊖ Students submit high-quality essay but do not understand ⊖ Lower expectations on learning (too focused	🔗 Use AI to remove the noises 🔗 Use AI to increase learning friction - thinking buddy, dedicated AI-on and AI-off tasks, 🔗 Encourage human-led activities that involve higher order thinking and cognitive struggles

1. Hallucination and misinformation

Hallucination (dreaming or recall error) usually happens because:

- In a fresh new chat, it does not know your context
- Pattern matching error
- Context window (aka memory) almost full



1. Hallucination and misinformation



Three girls are playing in a play room near window. You can see a flying yellow bird about to hit the window

Why AI Gets Things Wrong

Hallucination

Hallucination rate is getting lower. But citation is the worst area where 56% of citations are fabricated.

Never trust AI citation. One of the highest rate of hallucination



Why AI Gets Things Wrong

Hallucination – stress test


AA-Omniscience Hallucination Rate

AA-Omniscience Hallucination Rate (lower is better) measures how often the model answers incorrectly when it should have refused or admitted to not knowing the answer. It is defined as the proportion of incorrect answers out of all non-correct responses, i.e. $\text{incorrect} / (\text{incorrect} + \text{partial answers} + \text{not attempted})$.


— Paid License

— Free License


Artificial Analysis



What is context window?



Low usage of context window after 2min → Lots of free context (free memory) -> **low** rate of hallucination



HIGH usage of context window after 15min → Very little free context (free memory) -> **HIGH** rate of hallucination

Context window comparison


	ChatGPT – 5 family	Claude 4.6	Gemini 3 Pro	Gemini 3 Flash
Max context window	400k	1 million	1 million	1 million

But the effective context window is 100K regardless of the max window


Best practice

- **One specific** task per chat
 - ❌ “Help me write a report template for my geography, history and social legal study classes.”
 - ✅ In one chat - “Help me write a report template for my geography class”
 - ✅ In a second new chat - “Help me write a report template for my history class”
- Only upload what the AI needs to know
 - If you have too many things to upload, dedicate one chat session just to reorganize and structure everything in one document 📄 feed this doc as upload in a new chat
- Use **thinking mode** when possible
 - Thinking mode makes less mistake, but it takes slightly longer.
- Use text instead of screenshot
 - Any visual will take up more tokens and AI will have to use up more memory.

Switch to thinking mode



Free and Paid




Paid only

2. Sycophancy


Agreeing with you even when you're wrong

Prompting: Tell it "challenge my thinking" or "what am I missing?"

MEDQUAD



AMPS



Are we getting dumber? ([MIT study](#))

Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task[△]

Nataliya Kosmyna¹
MIT Media Lab
Cambridge, MA

Eugene Hauptmann
MIT
Cambridge, MA

Ye Tong Yuan
Wellesley College
Wellesley, MA

Jessica Situ
MIT
Cambridge, MA

Xian-Hao Liao
Mass. College of Art
and Design (MassArt)
Boston, MA

Ashly Vivian Beresnitzky
MIT
Cambridge, MA

Iris Braunstein
MIT
Cambridge, MA

Pattie Maes
MIT Media Lab
Cambridge, MA

United States




Figure 1. The dynamic Direct Transfer Function (dDTF) EEG analysis of Alpha Band for groups: LLM, Search Engine, Brain-only, including p-values to show significance from moderately significant (*) to highly significant (***).

Most people don't read beyond the Abstract

and an AI judge (a specialty built AI agent).

We discovered a consistent homogeneity across the Named Entities Recognition (NERs), n-grams, ontology of topics within each group. EEG analysis presented robust evidence that LLM, Search Engine and Brain-only groups had significantly different neural connectivity patterns, reflecting divergent cognitive strategies. Brain connectivity systematically scaled down with the amount of external support: the Brain-only group exhibited the strongest, widest-ranging networks, Search Engine group showed intermediate engagement, and LLM assistance elicited the weakest overall coupling. In session 4, LLM-to-Brain participants showed **weaker neural connectivity and under-engagement** of alpha and beta networks; and the Brain-to-LLM participants demonstrated higher memory recall, and re-engagement of widespread occipito-parietal and prefrontal nodes, likely supporting the visual processing, similar to the one frequently perceived in the Search Engine group. The reported ownership of LLM group's essays in the interviews was low. The Search Engine group had strong ownership, but lesser than the Brain-only group. The LLM group also fell behind in their ability to quote from the essays they wrote just minutes prior.

As the educational impact of LLM use only begins to settle with the general population, in this study we demonstrate the pressing matter of a likely **decrease in learning skills** based on the results of our study. The use of LLM had a measurable impact on participants, and while the benefits were initially apparent, as we demonstrated over the course of 4 months, **the LLM group's participants performed worse** than their counterparts in the Brain-only group at all levels: neural, linguistic, scoring.

3. Are we getting dumber?

? How did they use LLM in the study: they must use ChatGPT as their **sole resource** for the essay writing task, with **no access to the internet or other applications**.

Another interesting insight is the **difference in information flow directionality** between the groups. Brain-only group showed evidence of greater **bottom-up flows** (e.g. from temporal/parietal regions to frontal cortex) during essay writing. This bottom-up influence can be interpreted as the brain's semantic and sensory regions "feeding" novel ideas and linguistic content into the frontal executive system, essentially the brain generating content internally and the frontal lobe integrating and making decisions to express it [76]. In contrast, LLM group, with external input from the bot, likely **experienced more top-down directed connectivity** (frontal → posterior in high-beta). Their frontal cortex was often in the role of integrating and filtering the tool's contributions (an external source), then imposing it onto their overall narrative. This might be to an extent analogous to a "preparation" phase in creative tasks where external stimuli are interpreted by frontal regions sending information to posterior areas [76]. Our results support

Regarding executive function, the results show Brain-only group's prefrontal cortex was highly involved as a central hub (driving strong theta and beta connectivity to other regions), indicating substantial executive control over the writing process. LLM group's prefrontal engagement was comparatively lower, implying that some executive functions (like maintaining context, planning sentences) were most likely partially taken over by the LLM's automation. However, the LLM group still needed executive oversight to evaluate and integrate LLM suggestions, which is reflected in the top-down connectivity they exhibited. So, while the **quantity** of executive involvement was less for LLM users, the **nature** of executive tasks may have shifted, **from generating content to supervising the AI-generated content**.

4. Are we getting dumber?

Our cognitive load is not vanishing; it's being *redistributed*

It indicates a shift towards a greater level of **high-order thinking**, you are

- 🤔 **Critique** - ask the **good questions** to guide the AI
- 🤔 **Thinker** - prompt AI to criticize you, challenge your assumptions
- 💻 **Chief editor** - responsible for **assessing** the AI's output for accuracy, bias, and relevance (**Gemini Double Check** function, [EDIT framework](#) etc)
- 🏗️ **Architect** - **connect** AI-generated components with our own insights

| The future isn't about *IF* we use AI, but **HOW**

Turn & Talk

What issues have you seen among students AND teachers?

How can you identify the issues? What do they look like?


The Assessment Framework

*How students use AI — and how
you know the work is theirs.*

From Consumer to **Co-Creator**

How students can actually use AI — and how you know the work is theirs

The AI Assessment Scale Is a Compass, Not a Ladder



Adapted from Perkins, M., Furze, L., Roe, J., & MacVaugh, J. (2025). The AI Assessment Scale. aiassessmentscale.com

What Does Each AIAS Level Look Like in NCEA?



English



Maths




Science

Level 1: No AI	Handwritten close reading under exam conditions	Supervised in-class problem solving, no digital tools	Hands-on practical assessment, e.g. ion analysis
Level 2: AI Planning	AI brainstorms perspectives — student drafts independently	AI explores variables for stats investigation — student collects data	AI helps plan fair test — student does practical work
Level 3: AI-Assisted Task	AI editing tools refine text — student keeps voice and argument	AI helps with data visualisation — student interprets results	AI helps draft report after practical — student shows understanding
Level 4: Full AI	AI co-creates video essay — assessed on analysis and AI use	AI builds a mathematical model — assessed on design and AI use	AI runs simulations — assessed on reasoning and tool strategy
Level 5: AI Exploration	Co-design an AI-powered interactive narrative	Build a custom AI tool to teach a maths concept	Create an AI feedback system for lab technique

We are here

LEVEL	WHAT AI DOES	WHAT STUDENTS DO	HOW TEACHERS CAN SUPPORT	SCIENCE EXAMPLE
1 No AI	Nothing. AI is not used at any stage.	Work independently. Complete all tasks by hand — conduct experiments, record observations, write explanations. Authentic errors and personal language are visible.	Observe and probe in real time. Ask questions during hands-on work: “Why did you do that step?” Watch for gaps between what students say vs. write — a sudden jump in writing quality is a red flag.	<i>Students test unknown substances in a supervised lab, recording what they see by hand. Teacher walks around asking them to explain their reasoning as they work.</i>
2 AI Planning	Brainstorms and scaffolds. AI suggests ideas, variables, or approaches. It helps students think through the planning stage only — not writing or analysis.	Own all decisions and outputs. Students can explain what they asked the AI and why . The actual experiment design, data collection, and write-up are entirely theirs. They keep an AI use log.	Teach good vs. bad prompting. Model the difference: “Tell me what to do” (too much) vs. “I’ve thought of X and Y — what am I missing?” (genuine planning). Provide a prompt scaffold card. Require an AI use declaration.	<i>Student asks AI: “I’m testing how temperature affects plant growth. I’ve controlled light and water — what else should I consider?” Then designs and runs the experiment independently.</i>
3 AI-Assisted Task	Refines and polishes. AI improves drafts, structures reports, or helps create graphs. It works on content the student has already produced.	Do the core work first, then improve. Students generate their own data, observations, and initial interpretations before AI touches anything. They must be able to explain everything verbally.	Set clear boundaries and verify. Specify which parts AI can help with. Require a “before and after” draft. Use a 2-minute oral check: if they can’t explain their own conclusions, the AI did the thinking.	<i>Student runs an experiment, records results, and writes an initial report. Then uses AI to improve the structure and grammar. Teacher compares drafts and asks: “Explain what this finding means.”</i>
4 Full AI	Co-creates content. AI generates drafts, runs simulations, or builds models. The output is a joint human-AI product.	Direct, evaluate, and validate. Students choose what to ask the AI, check its output for errors, and explain their reasoning. Assessment focuses on how they used AI, not just the final product.	Shift rubric to assess AI literacy. Ask: “Where did the AI get it wrong, and how do you know?” Assess critical evaluation alongside subject knowledge. Teach students to cross-check AI output. Mark the process, not just the product.	<i>Student uses AI to simulate how changing conditions affects a process, then is assessed on: why they chose those inputs, whether the AI’s output makes sense, and what limitations exist.</i>
5 AI Exploration	Powers student-designed tools. AI is the engine inside something the student builds — a feedback tool, interactive model, or teaching app.	Design, build, and defend. Students create an original AI-powered tool, test it with peers, refine it, and justify their design choices using subject knowledge.	Mentor through iterative design. Guide students through testing and refining their tool. Assess via a design journal, peer testing, and a live demo where the student fields questions.	<i>Student builds an AI chatbot that coaches classmates through a practical technique — assessed on design decisions, accuracy of the feedback logic, and a live demo to the class.</i>



Higher level requires:

1. More training in teacher & students
2. More teacher supervision
- More discipline
3. Well-designed policy and framework

CORE HUMAN SKILLS IN THE AI AGE



Critical Thinking

Catching when AI is confidently wrong



Creativity

Connecting ideas no algorithm has seen together



Communication

Making someone trust you in 30 seconds



Collaboration

Getting five people who disagree to build something



Curiosity

Learning the next tool before anyone tells you to

Discussion

Where are your students on the spectrum? Pick one example.

2 min

Redesigning Assessment

Three principles to design authenticity in

Three Assessment Redesign Principles

1. Design authenticity IN

2. Value the process, not just the product. Aka build a learning portfolio

3. Make AI use explicit, not hidden.

Assumptions - AI policy in place, sufficient staff training, AI curriculum for students

4. Try ONE task first, not the whole assessment

Before and After

BEFORE (1 piece of evidence)

A typical written report on a chemistry technological development over time

AFTER (multiple evidence)

The portfolio of evidence may include:

- A written report
- Physical drawing
- AI-edited content
- Voice transcript

Same pattern works for any subject. Science: lab notebook + prompt log + oral explanation.

Maths: worked solutions + AI prompt log + explain-your-thinking video.

English: Written + Debate

1. ESS + SS assessment example

I was assessing:

- Student's ability to explain different aspects of ocean acidification
- Student's ability to explain how ocean acidification occurs using scientific principles

1. ESS + SS assessment example

Name 

ESS OA marking

Strategic Objective	Indicator	Baseline	Target	Current Status	Progress	Notes
1.1.1.1	1.1.1.1.1	1.1.1.1.1	1.1.1.1.1	1.1.1.1.1	1.1.1.1.1	1.1.1.1.1
1.1.1.2	1.1.1.2.1	1.1.1.2.1	1.1.1.2.1	1.1.1.2.1	1.1.1.2.1	1.1.1.2.1
1.1.1.3	1.1.1.3.1	1.1.1.3.1	1.1.1.3.1	1.1.1.3.1	1.1.1.3.1	1.1.1.3.1
1.1.1.4	1.1.1.4.1	1.1.1.4.1	1.1.1.4.1	1.1.1.4.1	1.1.1.4.1	1.1.1.4.1
1.1.1.5	1.1.1.5.1	1.1.1.5.1	1.1.1.5.1	1.1.1.5.1	1.1.1.5.1	1.1.1.5.1
1.1.1.6	1.1.1.6.1	1.1.1.6.1	1.1.1.6.1	1.1.1.6.1	1.1.1.6.1	1.1.1.6.1
1.1.1.7	1.1.1.7.1	1.1.1.7.1	1.1.1.7.1	1.1.1.7.1	1.1.1.7.1	1.1.1.7.1
1.1.1.8	1.1.1.8.1	1.1.1.8.1	1.1.1.8.1	1.1.1.8.1	1.1.1.8.1	1.1.1.8.1
1.1.1.9	1.1.1.9.1	1.1.1.9.1	1.1.1.9.1	1.1.1.9.1	1.1.1.9.1	1.1.1.9.1
1.1.1.10	1.1.1.10.1	1.1.1.10.1	1.1.1.10.1	1.1.1.10.1	1.1.1.10.1	1.1.1.10.1

Management Strategies

Marine Management: Analysis Table

analysis table below for **THREE** strategies to highlight the strengths and weaknesses
use specific statistics and examples from throughout New Zealand's Economic Exclusive Zone

1. Management Strategy: Rahui	
Helpful in achieving the objective of a healthy marine environment	Harmful toward achieving the objective of a healthy marine environment
Strengths	Weaknesses

TASK FOUR explanation

This is a suggested format. You don't have to strictly follow the format but you must cover all the content.

Introduction

- ☐ The issue of ocean acidification
- ☐ The impacts of ocean acidification
- ☐ You can copy the information from Task 3

Societal response to the issue

- ☐ Choose an societal group - this can be a group of people, a company, a council, a government body etc
- ☐ Describe who they are
- ☐ Describe what they do
- ☐ How are they bringing about positive impacts in the **economic** and **social** aspects?
- ☐ Evaluate **ONE** initiative of the society group:
 - ☐ The initiative could be technological, legal or cultural
 - ☐ Discuss 3 advantages
 - ☐ Discuss 3 disadvantages

Personal response (personal opinion)

TASK FOUR Letter to the...

Task Four: Write a letter to the Environment Minister of New Zealand

Maintaining marine biodiversity towards a sustainable future
It is very important that you use succinct data and research information to back up your own ideas

Include:

- An introduction describing what the scientific issue is and how it impacts on individuals and society (focus on **ocean acidification**)
- An evaluation of a **societal** response to the issue of **ocean acidification**. This may be a special interest group, i.e., government agency, law-making body, company or population of a town, city, country or planet. They must have consequences for a sustainable future for the marine environment in terms of the **economic** and **social** aspects also.
- Evaluate** means - is this an opinion? Is there evidence to support it? Is the evidence valid? Is there bias? Is there political influence? Are there commercial pressures?
- A justification of your **personal** response to your chosen issue. This can be your own response, that of your family or whānau, or a small group of which you are part.
- Justify** means - Means including evidence to support the stance
- A future projection in your own opinion on:
If more of these strategies were put in place in the future, what do you think would happen in terms of **Social**, **Environmental** and **Economic** aspects. Do you think more

TASK FOUR Management...

Marine Management: Analysis Table

analysis table below for **THREE** strategies to highlight the strengths and weaknesses
use specific statistics and examples from throughout New Zealand's Economic Exclusive Zone

1. Management Strategy:	
Helpful in achieving the objective of a healthy marine environment	Harmful toward achieving the objective of a healthy marine environment
Strengths	Weaknesses
Moving forward into the future - what technology and scientific research could influence this management strategy?	
Threats - what EXTERNAL threats can have an impact on the management strategy "Things that are not part of the management strategy"	

TASK THREE: Impacts 2023

Impacts: Human induced impacts on the marine environment

TASK: You will complete the table below that explains **three human-induced impacts** in New Zealand and will also write an in-depth evaluation of the issue and the impact on individuals and society in the four aspects TWO. Write about why this is an important issue; and show that you understand the links between the science and the impact.

During your research, you must keep a digital logbook that contains processed information. Processed information includes:

- Paraphrased information
- Highlighted or underlined information from the original texts
- Information that has been sorted or rearranged

An example of a digital logbook can be found [here](#).

For the ESS standard, you must go in depth in the **Science of Ocean Acidification** issue.

If you are using AI assistance, revise the **AI lesson** and you must show evidence of evaluating AI output. For each question you answer, attach a completed AI evaluation form.

Impact	Explanation of the impact
1. Global warming	This must include statistics, species and examples of locations throughout the world. Maps, images, diagrams and graphs are encouraged. A. How do greenhouse gases contribute to global warming?

TASK TWO Aspects of...

Sustainability and Human Relationships with the Marine Environment

Sustainability involves the development of ways of thinking and acting to meet the needs of the present without compromising the ability of future generations (of all living things) to meet their own needs

Environmental	Social
This aspect acknowledges the need to enhance and maintain the systems that sustain all life on Earth. It includes the structure of natural ecosystems and the interactions between them and calls for guardianship/kaitiakitanga of our environment.	This aspect acknowledges the need for equity within generations, and within and between ethnic and social inclusive of people's mental and physical well-being and their communities based on a fair distribution of resources.
Cultural	Economic
This aspect acknowledges the need to nourish and share attitudes and values that represent diverse worldviews, and the political need for all to express their views freely and to participate in decision-making. Addressing these needs can build resilience for the future.	This aspect acknowledges the interactions of humans with the environment in using resources to create goods and services for their lives. It acknowledges the resource use must occur within the capacity of our planet. It encourages sustainable use of resources.

Untitled document

Zechu interview questions

Task three impact

Multiple choice questions:

- What is the main greenhouse gas that causes ocean acidification?
 - A) Methane
 - B) Nitrous oxide
 - C) Carbon dioxide
 - D) Water vapourAnswer: C
- What is the chemical formula of carbonic acid, which forms when carbon dioxide dissolves in seawater?
 - A) CO₂
 - B) H₂CO₃
 - C) HCO₃⁻
 - D) CO₃²⁻Answer: B
- What is the name of the thing that can act as both an acid and a base, and is formed when carbonic acid dissociates in seawater?
 - A) Hydrogen ion
 - B) Hydroxide ion
 - C) Bicarbonate ion
 - D) Carbonate ionAnswer: A

Try just 1 task

1. ESS + SS assessment example

EDIT framework (student + AI to co-work)

EDIT example

2. Physics assessment example

I was assessing:

- Student's ability to explain how radiation impact the community in Chernobyl
(they don't need to memorize equations)
- Student's ability to explain why the place remains dangerous to these day

The evidence:

- Chernobyl physical map with explanation
- Digital whiteboard with video recording
- Some written evidence

Physics assessment example - Digital whiteboard snorkl

Feedback

Visible only to you

Answer: Correct Explanation: 3/4 - Strong

🕒 0:13

Great job explaining what half-life is and how it relates to the decay of radioactive substances! 🍌

🕒 0:30

I like how you used the blue circles to demonstrate the concept of half-life. It's clever!

🕒 1:45

Your analysis of the decay of Cesium and Tellurium in the Chernobyl evacuation zone is great but need to improve the accuracy. Cs-137 has a half life of 30 years, how much would decay away in 20 years?

Publish to Student

Benny Pan Private

🕒 1:59

Your feedback

Physics assessment example – explaining the hand annotated map

1. Students annotated the map and performed calculation first.
2. They took the annotated map and explain how Chernobyl had stayed dangerous for a long time.
3. AI was used to construct a draft feedback
4. I edited the feedback before releasing to students.

Assessment Redesign Sprint

15 minutes

Sprint Instructions

1. Pick ONE assessment you're running right now. Pick ONE task to allow AI to assist students.
2. Where are your students on the framework? Level 1 to 5
3. Redesign it:
 - Build authenticity IN
 - Value the process
 - Make AI use visible
 - Pick one task to try with AI
4. Then ask each other:
"Could a student still fake this?" "how can you tell?"

Share-Back

What assessment would you redesign?

What task(s) changed?

Could a student fake it?

Your Productivity Toolkit

*What to automate, which tools to use,
and three workflows that save hours.*

3%

of teachers say AI has greatly reduced their workload

Weekly AI users save 5.9 hours per week — six weeks of your year.

[RSC Teaching Survey, 2024 | Gallup/Walton Family Foundation, 2025](#)

What to Automate, What to Keep Human

GREEN — Automate freely

Admin emails & communications

Differentiated reading lists

Practice problem generation

Rubric drafting

Parent communication templates

Scheduling & admin tasks

YELLOW — Use with judgment

Lesson planning (review + adapt)

Feedback drafts (personalise first)

Resource creation (verify accuracy)

Student report comments

Curriculum mapping

Analysing student data trends

RED — Keep human

Final assessment judgments

Student wellbeing conversations

Building classroom culture

Parent-teacher conferences

Professional judgment calls


Anything involving student data

BAD PROMPT

Help me draft an email to parents about this event - to celebrate the end of term one, we are hosting a small gathering for students and parents to gather and socialize this Friday evening, 7pm at school.

Thought for a couple of seconds

Email



Subject End of Term One Gathering – This Friday at 7pm

Dear Parents and Caregivers,

To celebrate the end of Term One, we are hosting a small gathering for students and parents to come together and socialise.

The gathering will be held this Friday at 7.00pm at school. It will be a lovely opportunity for families to connect, celebrate the term, and enjoy time together as a school community.

We would love to see you there and hope many of our families are able to join us.

Kind regards,

[Your Name]

[Your Role/School Name]

BETTER PROMPT

“Help me draw off an email to parents about this event – to celebrate the end of term one, we are hosting a small gathering for students and parents to gather and socialize this Friday evening, 7pm at school. Keep the tone warm as parents are familiar with this recurring event. Keep the word count under 100 words. Keep short paragraphs as most parents view emails on their phone. Use emojis when appropriate.” (adding 2Cs – Context and Constraints)
Dear parents and whānau,


To celebrate the end of Term 1, we're hosting a small gathering for students and parents this Friday evening at 7pm at school. 🎉

It will be a lovely chance to come together, catch up, and enjoy some time as a school community.

We'd love to see you there and hope you can join us. 😊

Kind regards,
[Your Name]

Stop Retyping: Custom Commands



Stop Retyping: Custom Commands

The problem: You write the same prompt every time.

/feedback — “Read this student work. Give 2 strengths, 1 area for growth, and a next step. Be specific and honest. Use encouraging language.”

Instead of typing 30 words, you type one. Your best prompts become reusable tools.

Custom Commands examples (copy and edit it)

Whenever I type a slash prompt starting with "/", treat it as a command.

/sum - summarize in 5 bullet points. Each bullet point must be backed up by evidence and justification.

/chatsum - produce a detailed, well-structured summary of this conversation. Make sure the flow is logical so that I can copy this summary to start a new chat.

/starter [type] [topic] [year level] - the first parameter is the activity type — one of: "hook", "recall" (retrieval practice drawing on prior learning), "predict" (students make a prediction), or "discuss" (a think-pair-share prompt). The second parameter is the topic. The third parameter is the NZ year level in the format Y1–Y13. If no type is given, default to "recall". The activity must be completable in under 5 minutes.

/tidy [format] - the first parameter is the desired output format — one of: "bullets" (clean bullet points), "paragraph" (flowing prose), "table" (structured into columns), or "checklist" (actionable to-do items with checkboxes). The user will then paste messy, rough, or hastily typed notes (from a meeting, PD session, discussion). Clean up the notes into the specified format: fix spelling and grammar, organise logically, remove duplicates, and clarify vague points. Do not add information that was not in the original notes.

/question - Ask me challenging questions to help reveal my blind spots and assumptions. The goal less sycophancy.

AI That Works While You Don't

Scheduled Tasks in Gemini Studio

“Please summarize all emails in the last 72 hours, including those where I have replied.

Identify:

- the actionable points that require my attention.
- Who I should reply to in each actionable point”

You get sent a summary every day. You know exactly who to reply to, if you have replied to and the summary.

Try one new skill

15 minutes

Your Turn: Hands-On Practice

Build your library of prompts using slash commands

Try the Prompt Builder Use the Custom GPT or Gemini Gem to build better prompts for something you do regularly — report comments, lesson starters, feedback on student work. Let it interview you, push you to think.

Try the Policy Drafter Use the Custom GPT or Gemini Gem to start drafting an AI policy for your school. Answer its questions. In 10 minutes you'll have a framework you can take to your board—not a template, **your school's version. It has Te Tiri o Waitangi as core value.**

Note – always turn to “**Pro**” in Gemini Gem

WIFI

BHS_Guest

grass21830dog



Prompt Maker

It will help you make better prompts, and suggest the right documents to go with it

This Gem is shared. Instructions may change.

Ask Gemini



Tools

Gemini 3

Fast

Answers quickly



Thinking

Solves complex problems

Pro

Advanced maths and code with 3.1 Pro

Google AI Plus

Get more access to models and features

[Upgrade](#)

Fast ▾



Prize Draw

Thank you to our sponsor:

LearnCoach

Internal formative feedback coming up

Free access for 2026

Your voice will help us t know
what matters (survey)

5 minutes

References

[Gallup/Walton Family Foundation \(2025\). Three in Four Teachers Using AI Weekly Saving Six Weeks a Year. \[gallup.com\]\(#\)](#)

[McKinsey \(2025\). How Artificial Intelligence Will Impact K-12 Teachers. \[mckinsey.com\]\(#\)](#)

[OECD \(2026\). Digital Education Outlook 2026. \[oecd.org\]\(#\)](#)

[Twinkl \(2025\). AI in Education Survey: What UK and US Educators Think. \[twinkl.com\]\(#\)](#)

[RAND Corporation \(2025\). AI Integrity Policies in Schools. \[rand.org\]\(#\)](#)

[MDPI Education Sciences \(2025\). Brief Prompt-Engineering Clinic: +0.70 Literacy, -0.58 Anxiety.](#)

[Gibreel et al. \(2026\). Prompt Engineering Competence Predicts Knowledge Acquisition. \[Nature Scientific Reports\]\(#\).](#)

[Qian \(2025\). Goal-oriented, Iterative, Feedback-driven Framework. \[J. Educational Computing Research\]\(#\).](#)

[NZQA \(2025\). 876 Assessment Breaches; 59 AI-specific. \[NZ Herald\]\(#\).](#)

[Ministry of Education NZ \(2026\). Generative AI Guidance for Schools. \[education.govt.nz\]\(#\)](#)

[NZ Privacy Commissioner \(2025\). Annual Survey: 67% Concerned About Children's Privacy.](#)

[Bastani et al. \(2025\). Generative AI Can Harm Learning. \[PNAS\]\(#\).](#)

[Anthropic \(2025\). Education Report: How University Students Use Claude.](#)

[Deakin University/StudyFinds \(2025\). 56% of AI Citations Fabricated or Contain Errors.](#)

[PPTA \(2024\). Workload Survey: 86% Say Workload Increasing.](#)

[AI and cognitive Stunt podcast interview](#)