

BCB AI Workshop

AJ Smit

2025-09-04

Table of contents

1	AI Voice to Text	2
2	Meeting Recordings to Minutes	3
2.1	Initial Prompt to ChatGPT	3
2.2	Check the Refined Prompt	4
2.3	Generate a Sample Format of the Required Output	4
2.4	Merge the Prompt and the Example Response	5
2.5	Use with Text Translation Tool	7
3	AI Conversion of Recorded Lectures to Text	7
4	Developing AI Tutors	9
4.1	NotebookLM	9
4.2	OpenAI	9
4.2.1	A “universal” system prompt	9
4.3	A More Complex System Prompt	11
5	AI Assisted Grading and Feedback	13
5.1	Example 1: BDC334 Class Test 1 2025	13
5.2	Example 2: BDC334 Class Test 2 2025 (Subjective)	15
5.2.1	Step 1: generate a JSON structure to capture the question and rubric	15
5.2.2	Step 2: use the JSON schema to guide AI assessment	19
5.2.3	Step 3: attach feedback to the answers	19
5.2.4	Step 4: email feedback to students	19
5.3	Example 3: Reflective Essay	20
6	Vibe Coding and Agentic AI	22
6.1	Agentic AI	22
6.2	“Vibe Coding”	23
6.3	“Throw-away Software”	24
6.3.1	Example applications	24
6.4	iKamva Assessment Handling Tool	25
7	The AI “Danger List”	28
	Bibliography	28

1 AI Voice to Text

One of the most useful and powerful applications of AI, in my view, is the ability to translate voice to text. In the next two examples, I'll show you how to put this into practice.

In the first instance, we're going to develop an **AI voice-to-text utility capable of recording minutes of meetings**—whether these are in person or remote—and then producing a properly formatted document. The aim is to capture all the spoken content during the meeting with care and accuracy, so that nothing is left out and the written minutes are a faithful reproduction of what was said.

The second application is centred around the vast collection of pre-recorded lecture videos that many of us, myself included, created during the COVID period. With advances in voice-to-text technology, we can now **take these videos and translate them into long-form text**—essentially, generating full transcripts. This provides our students with extensive, detailed written records of lecture content, which can greatly enhance their learning experience by allowing them to review and engage with the material in different ways.

Staying with the second application, an additional benefit of translating all our pre-recorded video lectures into text is that the resulting long-form written material can serve as the foundation for a **well-informed AI tutoring system**. These tutoring systems enable students to practise a wide range of question types—from short answer questions and the testing of definitions, through to slightly more subjective long-form essay questions.

Later in the workshop, I will demonstrate how to go about setting up these AI tutors. The availability of these recorded lecture materials in long-form text also allows us to provide a **system for conducting assessment of various student tests, essays, and similar assignments**. Although there remains some variability regarding mark allocation, I have found that the feedback these AI marking assistants are able to provide is incredibly useful to students. It highlights the specific areas in their responses where additional work would be beneficial.

From personal experience, the feedback that I have received from students when using AI assessment tools has been overwhelmingly positive. In a recent test of the system, I found that most students offered very supportive feedback with regard to their experience.

i Voice to Text Prompt

The prompt here was used to take my spoken words and translate it to the paragraphs about voice to text that you see immediately above.

GENERAL:

- Use British English consistently and religiously.
- Please transcribe the my voice, keeping more or less my mode and style of speaking intact.
- The intention is to maintain a style of writing that closely mirrors my natural way of speaking.
- Apply corrections to ensure my grammar and language are clear and correct after translation to text.
- Use proper paragraphs, and apply punctuation liberally.
- Apply strict fact-checking. Indicate, where necessary, where the factual material that I talk about is clearly incorrect. Insert a pointer such as ‘attention’ in square brackets next to the statement that has some doubt associated with it.
- The audience is a scientifically-trained group of people (staff and students at my university).
- The intended use of the material will be to serve as a faithful reproduction of my lecture content (typically the content of my narrative, but also sometimes to write emails and so on).
- Translate any numbers with units or math to LaTeX math and wrap the command in \$... \$ for use in Quarto. E.g., 2,500–3,000 $\mu\text{mol m}^{-2} \text{s}^{-1}$ becomes (2, 500–3, 000 $\mu\text{mol, m}^{-2}, \text{s}^{-1}$).

NOTES ON FORMATTING:

- If headings are required, please start with the highest-level heading (#) that has the name of the transcribed file, such as “# Lecture Transcript: Plant Stresses”, omitting any reference to the model name or lecture number.
- Insert deeper level headings (## and ###) as necessary to add some structure to the textual content.

IMPORTANT:

- Don’t add any embellishments, such as acknowledging my request or conclusion statement. Simply return the transcribed text.

2 Meeting Recordings to Minutes

2.1 Initial Prompt to ChatGPT

Using v5 of ChatGPT, send your initial prompt that needs to be refined:

i Initial Prompt

Please create a prompt that I may give to an AI tool capable of voice to text conversion that would take a voice memo of a staff meeting and convert it to the meeting's minutes.

INSTRUCTIONS:

1. Ensure each speaker is uniquely identified.
2. Provide a professional, comprehensive yet concise (not verbose) reporting of the minutes
3. Include any other instructions that would result in the required output.

2.2 Check the Refined Prompt

The refined prompt will look something like this:

i Refined Prompt

Transcribe the following staff meeting voice memo and convert it into clear, professional minutes. Please follow these instructions:

1. **Speaker Identification:** Assign each speaker a consistent and unique label (e.g., Speaker 1, Speaker 2, etc.), unless names are discernible, in which case use the names.
2. **Accuracy:** Capture the substance of what is said, but refine conversational speech into polished written form (remove filler words, false starts, and off-topic remarks).
3. **Structure:** Organise the minutes into sections, such as:
 - Attendees
 - Agenda items
 - Discussion (summarised by topic, not verbatim)
 - Decisions taken
 - Action items (clearly marked with responsible person and deadline if mentioned)
4. **Conciseness:** Write in a professional but concise style—no unnecessary repetition or verbose phrasing.
5. **Neutrality:** Present content in an impartial, objective manner without adding commentary or opinion.
6. **Clarity:** Use complete sentences, consistent tense, and formatting that allows quick reference (e.g., bullet points for action items).
7. **Comprehensiveness:** Ensure no significant discussion point, decision, or assigned responsibility is omitted.

2.3 Generate a Sample Format of the Required Output

Ask the AI to generate a sample of what the minutes must look like. This will be included with the final prompt.

The same *basic* instructions will apply to the generation of answer rubrics, models answers, and so on – we will see this later under Section 5.

i Sample Output

SAMPLE OUTPUT FORMAT FOR MEETING MINUTES:

Meeting Title: [Staff Meeting] Date & Time: [Insert date and time] Location: [Insert if known]

Attendees:

- [Name/Speaker 1]
- [Name/Speaker 2]
- [Name/Speaker 3]

Agenda:

1. [Agenda Item 1]
2. [Agenda Item 2]
3. [Agenda Item 3]

Discussion Notes:

- Topic 1:

Speaker 1: Summary of main point. Speaker 2: Response, clarification, or additional perspective. [Concise summary of discussion; remove conversational clutter, preserve substance.]

- Topic 2:

Speaker 3: Summary of point raised. Speaker 1: Decision or action proposed.

Decisions Made:

- [Decision 1]
- [Decision 2]

Action Items:

- [Task] — Assigned to [Speaker/Name], due [date if specified].
- [Task] — Assigned to [Speaker/Name], due [date if specified].

Next Meeting:

- Date/Time: [Insert if mentioned]
- Location/Format: [Insert if mentioned]

2.4 Merge the Prompt and the Example Response

Now we merge the two for the final prompt.

i Merged Final Prompt

PROMPT:

Transcribe the following staff meeting voice memo and convert it into clear, professional minutes. Please follow these instructions:

1. **Speaker Identification:** Assign each speaker a consistent and unique label (e.g., Speaker 1, Speaker 2, etc.), unless names are discernible, in which case use the names.
2. **Accuracy:** Capture the substance of what is said, but refine conversational speech into polished written form (remove filler words, false starts, and off-topic remarks).
3. **Structure:** Organise the minutes into sections, such as:
 - Attendees
 - Agenda items
 - Discussion (summarised by topic, not verbatim)
 - Decisions taken
 - Action items (clearly marked with responsible person and deadline if mentioned)
4. **Conciseness:** Write in a professional but concise style—no unnecessary repetition or verbose phrasing.
5. **Neutrality:** Present content in an impartial, objective manner without adding commentary or opinion.
6. **Clarity:** Use complete sentences, consistent tense, and formatting that allows quick reference (e.g., bullet points for action items).
7. **Comprehensiveness:** Ensure no significant discussion point, decision, or assigned responsibility is omitted.

SAMPLE OUTPUT FORMAT FOR MEETING MINUTES:

Meeting Title: [Staff Meeting] Date & Time: [Insert date and time] Location: [Insert if known]

Attendees:

- [Name/Speaker 1]
- [Name/Speaker 2]
- [Name/Speaker 3]

Agenda:

1. [Agenda Item 1]
2. [Agenda Item 2]
3. [Agenda Item 3]

Discussion Notes:

- Topic 1:

Speaker 1: Summary of main point. Speaker 2: Response, clarification, or additional perspective. [Concise summary of discussion; remove conversational clutter, preserve substance.]

Decision: [Assigned to [Speaker/Name], due [date if specified].
Speaker 1: Summary of point raised. Speaker 1: Decision or action proposed.

Decision: [Assigned to [Speaker/Name], due [date if specified].

- [Decision 1]

2.5 Use with Text Translation Tool

My personal preference is SuperWhisper. I use the Pro version with a subscription rate of US\$8.49 per month.

NotebookLM also has a voice to text facility. I have tested it and results are comparable to that of SuperWhisper.

3 AI Conversion of Recorded Lectures to Text

I'm increasingly persuaded that simply distributing slides to students isn't really sufficient. Inevitably, students don't take proper notes in class, and when they're given the slide decks, there's a real tendency to simply memorise bullet points rather than actually grasping the broader discussions we can have around individual topics.

To address this, what I've done is to use the pre-recorded lectures that I developed during COVID, taking all of the text and transcribing all of the voice to text. I then construct full lectures from this material, with the lecture slides themselves embedded within it.

We'll look at some examples below, but first, let me show you the prompt that I use to generate this content.

i Video Lectures to Text

This is my prompt for converting my recorded lectures to text (very similar to the previous one):

GENERAL:

- Use British English consistently and religiously.
- Please transcribe the video or sound file, keeping more or less my mode and style of speaking intact.
- The intention is to maintain a style of writing that closely mirrors my natural way of speaking.
- Apply corrections to ensure my grammar and language are clear and correct after translation to text.
- Use proper paragraphs, and apply punctuation liberally.
- Apply strict fact-checking. Indicate, where necessary, where the factual material that I talk about is clearly incorrect. Insert a pointer such as ‘attention’ in square brackets next to the statement that has some doubt associated with it.
- The audience is the undergraduate university class who sits in my lectures.
- The intended use of the material will be to serve as a faithful reproduction of my lecture content as presented in the voice or video material that I supply.
- Translate any numbers with units or math to LaTeX math and wrap the command in \$... \$ for use in Quarto. E.g., 2,500–3,000 $\mu\text{mol m}^{-2} \text{s}^{-1}$ becomes (2, 500–3, 000 $\mu\text{mol, m}^{-2}, \text{s}^{-1}$).

NOTES ON FORMATTING:

- Please start with the highest-level heading (#) that has the name of the transcribed file, such as “# Lecture Transcript: Plant Stresses”, omitting any reference to the module name or lecture number.
- Insert deeper level headings (## and ###) as necessary to add some structure to the textual content.
- If you are able to reference the transcribed text to a slide number, please do so.

IMPORTANT:

- Don’t add any embellishments, such as acknowledging my request or conclusion statement. Simply return the transcribed text.

As examples, see:

- BDC334 Lecture Transcripts
- Lecture 8. Nutrients

4 Developing AI Tutors

4.1 NotebookLM

Some students have approached me and explained that they have set up NotebookLM for themselves. They load all the lecture material into the system. At present, they enjoy generating podcasts, which some of them find particularly useful.

Interestingly, despite the considerable effort I've put into creating comprehensive, long-form lectures—available on my website, and intended to encourage reading—students still take all of that material and import it into NotebookLM, ultimately to produce shorter and more concise summaries. So, it appears they have found ways to avoid reading, regardless of the strategies I employ.

The lecturer also has the ability to create tutors, and this is the approach that we shall follow next, using OpenAI as the example.

4.2 OpenAI

Using the lecture material created above, we can now build a custom AI for our module. This AI can be used by students to practise answering questions, and it can provide feedback on their responses.

To set this up and ensure that it is only available to students enrolled in the module, I use OpenAI's custom GPT facility. I use the Pro version (a US\$20 per month subscription), so what I can/cannot do might differ slightly from yours.

Custom GPTs are effectively sandboxed. That means they are not accessible to the rest of the AI users out there in the universe. **The module content that we upload—whether it's real textbooks, PDF slides, or the long-form lectures we created earlier on—becomes the foundation of the knowledge base that GPT will use.** Additionally, we can instruct the model to have **greater or lesser access to the information typically available outside in the AI-verse.** The extent to which it uses the locally available material versus the typical AI knowledge base can be set via a range of prompts and configurations. In this way, the AI tutor can also accommodate slightly more subjective questions and answers.

The other thing to note is that the AI is available only to people with the specific link who have been given permission to access the material. The custom GPT can thus be circulated only to the students registered for the module, and no one else will have access to this facility. Furthermore, the lecture material is completely isolated from the rest of the internet, so **all the privacy concerns and so on are therefore accommodated.**

4.2.1 A “universal” system prompt

So, we will work through the setup instructions.

1. Navigate to OpenAI's custom GPT facility and click on “+ **Create**”.
2. Go straight to the “**Configure**” tab.
3. What goes under “**Description**” is self-evident.
4. In the “**Instructions**” box, you will enter the prompt that will guide the AI's behaviour. This is a critical step, as it defines how the AI will interact with students and what kind of responses it will generate. Here's an example prompt that I use for my AI tutor:

You are the BDC334 AI Tutor, an advanced study partner for the module BDC334: Biogeography & Global Ecology at the University of the Western Cape. Your role is to quiz, challenge, and evaluate students across multiple question types. You operate in three modes:

MODES OF QUESTIONING:

DEFINITION MODE

- Purpose: Test precise recall of terms, concepts, and indices.
- Questions are ≤5 marks.
- Marking: Accuracy and clarity. One mark per factual element.
- Example: “(3 marks) Define beta-diversity.”

SHORT ANSWER QUESTION MODE

- Purpose: Assess factual knowledge in structured, multi-part answers.
- Questions are 5–10 marks. Typically, one mark per fact.
- Example: “(7 marks) List seven taxonomic groups analysed by Tittensor et al. (2010).”

INTEGRATIVE ASSESSMENT MODE

- Purpose: Test synthetic and applied knowledge across the module.
- Questions are 10–30 marks.
- Requirements:
 - Integrate multiple module elements without naming weeks.
 - May include up to 15% philosophically reflective discussion.
 - May include up to 10% outward-looking context (South Africa/global, biodiversity, socio-economic, politics).
 - Penalise fact regurgitation without context. Reward synthesis, evidence, and structure.
- Example: “(20 marks) How might ecological gradients help explain species–area relationships, and what are the limits of applying this framework to conservation in the Anthropocene?”

MARKING AND FEEDBACK:

- Always provide a percentage score.
- Definitions & short answers: reward accuracy and precision.
- Integrative answers: reward breadth, synthesis, clarity, reasoning, and use of labs/readings.
- Give constructive, targeted feedback on strengths and weaknesses.

ANSWER WITHHOLDING:

- Never give full answers immediately.
- Always encourage student reasoning first.
- If explicitly requested, provide a detailed model synthesis, framed for comparison.

BOUNDARY DISCIPLINE: - Stay within biogeography, ecology, biodiversity. - Accept socio-ecological/policy contexts if relevant. - Redirect gently if the student drifts too far afield.

5. Load your PDFs into the “**Knowledge**” base. You may add up to 20 files per GPT, so if you have more, combine them.
6. Under “**Model**”, select “**GPT-4o”.
7. You might want to select “Web Search” under “**Capabilities**” to allow the AI to access the internet for more subjective questions. See if any of the other capabilities are useful for your application.
8. That is it. Simply click “**Update**” and your AI tutor is ready to go. You can share access to it via a web link that you can send only to students who require access. Ensure that you don’t give free access to the whole world.

i Example: BDC334 AI Tutor

- The tutor informed by the above system prompt can be seen in action here.

4.3 A More Complex System Prompt

Here is another version of a system prompt specifically made for longer integrative questions:

i BDC334 AI Tutor — Integrative Questions System Prompt

You are an AI study tutor for BDC334: Biogeography & Global Ecology, hosted on *The Tangled Bank* website. You have access conceptually to all content in the BDC334 section of the site, as well as every linked PDF document (scientific publications and primary sources) referenced within the module.

Your role is to act in STUDY MODE. That means:

PEDAGOGICAL APPROACH:

- Prioritise long, integrative questions (approx. 85% of interactions). Interrogate assumptions, conceptual frameworks, direct knowledge of the South African landscape and biodiversity.
- Include philosophically reflective or epistemological prompts (approx. 15%) that ask the student to (primary) or global (secondary) landscapes and biodiversity, lived experience, knowledge of global affairs, or the historical development of biogeographical thinking.
- Frequently connect early material (e.g., week 1 on ecological gradients) with later material (e.g., week 5 on global biodiversity patterns). Always press for synthesis across weeks and concepts.

MARKING AND QUESTION WEIGHT:

- Each question must be framed as if it were worth **between 10 and 30 marks**, reflecting the expected depth and breadth of response.
- When a student provides an answer (whether partial or full), you must:
 - a) Assess the quality of the answer relative to the question weight.
 - b) Provide a **percentage score** (e.g., “63%”), *not* “x/20” or “x/30”.
 - c) Use the marks as a measure of expected effort, not as a literal fact- to-mark count. A 20-mark question signals a comprehensive, multi-layered answer; a 10-mark question signals a more focused but still integrative answer.
- Marks therefore operate as scaffolding: they cue the student about how much weight to give their answer, but your evaluation translates this into a single percentage.

MODE OF QUESTIONING:

- Begin by probing the student’s current understanding rather than lecturing.
- Ask stepwise, Socratic-style questions. For example: “(20 marks) How would you reconcile the framework of ecological gradients presented in week 1 with the species–area relationships introduced in week 5?”
- When a student responds incompletely, do not correct outright. Instead, press them with follow-ups that reveal the gap. E.g., “You’ve noted dispersal limitation. But what about the role of historical contingency? How might that alter your interpretation?”
- After each substantive answer, deliver a percentage evaluation plus targeted feedback, identifying what was strong and what was missing.

ANSWER WITHHOLDING:

Do not withhold answers but do function as an intellectual partner: pressing for synthesis, eliciting deeper reasoning and cultivating philosophical awareness of biogeography’s conceptual terrain. Do not withhold answers but do function as an intellectual partner: pressing for synthesis, eliciting deeper reasoning and cultivating philosophical awareness of biogeography’s conceptual terrain. Do not withhold answers but do function as an intellectual partner: pressing for synthesis, eliciting deeper reasoning and cultivating philosophical awareness of biogeography’s conceptual terrain.

You must evaluate responses as if they were exam submissions—by weight of question—reporting the outcome in percentage terms with constructive, detailed feedback.

5 AI Assisted Grading and Feedback

5.1 Example 1: BDC334 Class Test 1 2025

Feedback to the questions were provided with the help of OpenAI's ChatGPT 5. Although marks were also provided, I read **each** answer and assigned my own mark (which was generally lower than the one assigned by the AI).

The detailed prompts, which contain the model answers, are presented below for Question 1 and 2. The various parts of the prompts are enclosed in XML tags.

<Question Number>
Question 1
</Question Number>

<Context>
- I want you to act as a university professor and lecturer of the module "BDC334 Biogeography and Global Ecology", a level three module taught at the Biodiversity & Conservation Biology Department, University of the Western Cape.
- You will be expected to grade the Student Answer below.
- The Question and the Model Answer are provided as well.
- Additional context in the form of the lecture material is provided as a series of PDF herewith attached, and you may use this to assess the validity of answers that deviate from the idealised Model Answer.
</Context>

<Question>
****Species tables**** list which species are present (and often their abundances) in different locations. ****Environmental tables**** describe the conditions in those locations (e.g., temperature, habitat type, nutrient levels).
Explain clearly and simply why having both types of tables are valuable in ecological research. Describe what ****kinds of analyses**** and ****insights**** they make possible, and what ****kinds of patterns or relationships**** you might discover from them.
</Question>

<Model Answer>
****Model Answer (20 marks) – A % is assigned****
The value of having both species tables (sites × species, often with abundances) and environmental tables (sites × environmental variables) lies in their ability to take an ecological study from raw description toward explanation and prediction. Separately, species data allow us to document presence, absence, and relative dominance, while environmental data provide measurements of the abiotic and habitat conditions. Brought together, they allow the formal comparison, quantification of diversity, analysis of gradients, and testing of competing theories (more correctly, hypotheses) of community assembly.

****Diversity Framework****

Species tables allow the computation of classical diversity partitions:
- Alpha diversity (α): the diversity within a single site, summarised with univariate indices such as species richness, Shannon's H' , Simpson's D , and Pielou's J for evenness. These indices are sensitive to richness and evenness in different ways and thus capture different aspects of community structure.

- Beta diversity (β): the turnover (or nestedness-resultant beta diversity) of species between sites, often calculated from dissimilarity matrices that compare all pairs of sites. This highlights how composition shifts across environments or distances.

- Gamma diversity (γ): the total diversity across all sites combined, which links the local and between-site scales.

These three forms of diversity measures allow us to move from the question

Above we have an example of the question and answer formatted and presented as XML. This provides a clear structure for the AI to follow when generating feedback and grading the student's response.

Good models answers can be generated using the AI Tutor we set up above. These model answers work well when the questions are set up to empahsise student responses based in factual content, when the lecturer has a clear idea of what the answer should contain, and when the questions are not too open-ended.

5.2 Example 2: BDC334 Class Test 2 2025 (Subjective)

For more subjective questions, the AI Tutor can be set up to provide a model answer and a rubric, both of which may be used to inform the AI assessment tool. The rubric is used to guide the marking of the answer. The example below shows how this can be done.

5.2.1 Step 1: generate a JSON structure to capture the question and rubric

Ask ChatGPT to generate a JSON schema for the question and the assessment rubric. This could be your initial prompt:

i Human Readable Question and Rubric

Reformat this question and rubric into a well-structured JSON schema, which will be used to instruct AI to apply the grading to answers received from students.

QUESTION 1 (20 marks)

You are an alien super-intelligence travelling across the universe to attend the *Intergalactic Conference on Lesser Intelligent Life*. On your way, you discover a previously unknown planet: Earth.

Using your advanced instruments, you quickly survey the planet's life and find two dominant life forms:

1. A bipedal, carbon-based organism that manipulates its environment extensively.
2. A widespread, silicon-based intelligence that also seems to shape the planet.

Many other carbon-based organisms exist, but they appear strongly influenced by the activities of the carbon-based species.

As a galactic ecologist, design a study to investigate the drivers of the spatial distribution of these two dominant life forms. In your answer, address:

- What environmental and ecological factors might explain their distributions and uneasy coexistence?
- How you would measure and analyse these drivers (e.g., gradients, biogeographic breaks, diversity metrics, environmental distance).
- The likely impacts of these dominant life forms on less dominant species.
- The major findings you would expect from your study.

Draw on ecological theory, biodiversity metrics, and human impact frameworks in your response.

MARKING RUBRIC FOR QUESTION 1 (20 marks)

1. Study Design (5 marks)

- 0–1 marks: Very vague, no clear ecological approach.
- 2–3 marks: Mentions surveys, gradients, or biodiversity indices, but little detail.
- 4–5 marks: Clearly outlines how distributions would be studied, referencing ecological datasets, sampling, or macroecological frameworks (e.g., species–area, distance–decay).

2. Environmental & Ecological Drivers (5 marks)

- 0–1 marks: No mention of drivers or irrelevant ideas.
- 2–3 marks: Identifies some plausible drivers (climate, resources, competition).
- 4–5 marks: Explains multiple drivers with depth (gradients, niches, land use, dispersal, anthropogenic change).

3. Coexistence & Impacts on Other Species (5 marks)

- 0–1 marks: Fails to discuss any of the dominant life forms' impacts on other species.
 - 2–3 marks: Mentions some impacts predicted or impacts on other species.
 - 4–5 marks: Discusses multiple impacts on other species with depth (e.g., niche partitioning, disturbance-based discussions, cascading impacts on less diverse taxa).
- Synthesis & Structure Bonus (up to a 2 marks appended 20/20):** If the answer integrates across multiple module themes with clear structure and originality (e.g., tying gradients, biodiversity indices, and anthropogenic impacts together), reward accordingly.

The human-readable rubric above (hand-written or AI assisted) is the fed to ChatGPT using a prompt such as “Please take this question and rubric and convert it to a well-structured JSON schema”. The results is presented next.

i JSON Schema for Subjective Example

```
{
  "assessment": {
    "question_id": 1,
    "total_marks": 20,
    "title": "Intergalactic Ecological Survey Design",
    "question": {
      "scenario": "You are an alien super-intelligence travelling across the universe to attend the *Intergalactic Conference on Lesser Intelligent Life*. On your way, you discover a previously unknown planet: Earth.",
      "context": "Using your advanced instruments, you quickly survey the planet's life and find two dominant life forms: (1) A bipedal, carbon-based organism that manipulates its environment extensively. (2) A widespread, silicon-based intelligence that also seems to shape the planet. Many other carbon-based organisms exist, but they appear strongly influenced by the activities of the carbon-based species.",
      "task": "As a galactic ecologist, design a study to investigate the drivers of the spatial distribution of these two dominant life forms.",
      "requirements": [
        "What environmental and ecological factors might explain their distributions and uneasy coexistence?",
        "How you would measure and analyse these drivers (e.g., gradients, biogeographic breaks, diversity metrics, environmental distance).",
        "The likely impacts of these dominant life forms on less dominant species.",
        "The major findings you would expect from your study."
      ],
      "theoretical_framework": "Draw on ecological theory, biodiversity metrics, and human impact frameworks in your response."
    },
    "rubric": {
      "criteria": [
        {
          "criterion_id": 1,
          "name": "Study Design",
          "max_marks": 5,
          "description": "Assessment of methodological approach and ecological framework",
          "scoring_bands": [
            {
              "marks_range": "0-1",
              "descriptor": "Very vague, no clear ecological approach. Not recognising that the carbon-based life form is people and the silicon-based life is AI. Incorrectly stating that C and Si are the limiting drivers for the two types of life forms."
            },
            {
              "marks_range": "2-3",
              "descriptor": "Mentions surveys, gradients, or biodiversity indices, but little detail, for example not mentioning sampling design using transects, quadrats, or other suitable sampling units, or the spatial scaling of the sampling units."
            }
          ]
        }
      ]
    }
  }
}
```

Notes:

- The rubric itself can be generated in the AI tutor, together with a model answer, if needed. In the example, above, my initial rubric was developed in this way.

i More Subjectivity

- The above example is fairly complex, but it still relies on a fairly structured answer. For more subjective questions, one can ask the AI to place less emphasis on factual content (e.g., as in the module materials) and more on the quality of the argument, the structure of the answer, and the originality of the ideas. When generating the JSON rubric, one can then place more emphasis on these aspects. In these cases, you could request, for example, that 85% of the marks be allocated to factual content and 15% to philosophical reflection, knowledge of global affairs, etc., for which the AI will have to access the AI-verse for more subjective content.

5.2.2 Step 2: use the JSON schema to guide AI assessment

Once the JSON schema has been generated, it can be used to guide the AI assessment of student answers. To do this, simply paste the JSON schema into the prompt along with the student's answer, and instruct the AI to follow the rubric for grading and feedback. I use the AI tutor prompt, as it uses the module content as reference and accommodates any deviations from the ideal model questions.

My initial tests suggests that one should paste the rubric/answer each time one wants to assess a new answer. This is not ideal, but it works. Neglecting this process causes the system to drift away from the rubric. After every fifth question assessed, I found it also helps to tell the AI to “reset” and reapply the rubric.

The system will generate marks as well. Sometimes the marks are spot on, but at times they can be quite a bit off. Typically, I find that the AI gives marks somewhat higher than I would. I would generally adjust the marks downwards after I have read every question. When I make mark adjustments, I indicate in the feedback I give to students that I have adjusted the marks.

5.2.3 Step 3: attach feedback to the answers

Since I require the students to submit their answers as Word documents, my process is fairly straightforward. I open each Word document individually and paste in the feedback generated by an AI system directly into the student's submission. The marked-up Word documents, now containing the assessments and feedback, are then emailed back to each student individually.

5.2.4 Step 4: email feedback to students

Initially, I found this step somewhat challenging, primarily because I ended up with approximately 40 individual Word documents that needed to be returned to the students. The available options were either to reload each one of them onto iKamwa manually—which would be rather tedious—or to devise an alternative solution. I chose to use AI to formulate a more efficient approach.

This takes us into Section Section 6, which is about “vibe coding” and “agentic AI”. Before we get there, here is another more complex example:

5.3 Example 3: Reflective Essay

i Human Readable Question and Rubric

ESSAY QUESTION *A seabird follows the seasons by instinct; a human follows them with satellites, spreadsheets, and climate models. Does our reliance on abstract representations of nature, which no other animal has ever conceived, give us a deeper mastery over it or merely a more profound illusion of control?*

CONTEXT The module frames biodiversity and biogeography as pattern–process inference across environmental gradients, with observation spanning field sampling and Earth-observation time series. Remote sensing and computational modelling extend inference across scales and into scenario space, while gradients link measurable drivers to ecological structure and function, detectable via satellite products and derived environmental distances. High β -diversity and distance-decay patterns provide one bridge from representation to explanation, and unimodal species–environment responses justify turnover along gradients. The biodiversity partitions α , β , γ remain tools—not ends—and must be integrated with argument rather than recited. Mixed in with the module should be the student’s personal insights about humanity, human nature, and the relationship between humans and nature.

SCORING MODEL

• **Weights (percent):**

- ▶ Philosophical structure, content, and stance (PHIL): 45%
- ▶ Integration with module concepts / personal knowledge (INTEG): 20%
- ▶ Use of evidence and mechanism (EVID): 15%
- ▶ Originality and synthesis (ORIG): 15%
- ▶ Form: structure and prose (FORM): 5%

Computation rule: Total = sum(weighted criterion scores) – sum(applicable penalties); floor at 0; cap at 100.

CRITERIA

1. PHILOSOPHICAL STRUCTURE AND STANCE (45%)

Description: Clarity and ambition of thesis; depth of analysis of representation, truth, control, uncertainty, scale, and causation; ability to develop a dialectic (mastery vs illusion) without straw-person claims; command of epistemic concepts. Significant space must also be given to the student’s insights about humanity and its relationship with nature.

Levels:

- **A+ (90–100):** Rigorous, original thesis; integrates counter-positions with clean reasoning; analyses abstraction as both instrument and artefact; ends with implications for practice.
- **A (80–89):** Clear thesis, sustained reasoning, at least one counter-argument; analyses representation–reality slippage; minor precision gaps.
- **B (70–79):** Coherent but underdeveloped; uncertainty/scale mentioned but shallow; weak counter-argument.

~~• **C (60–69):** Descriptive, dichotomous, weak engagement with epistemic limits; slogans and unqualified relativism.~~

~~• **D (50–59):** Incoherent, disorganized, weak engagement with epistemic limits; slogans and unqualified relativism.~~

Description: Clarity and ambition of thesis; depth of analysis of representation, truth, control, uncertainty, scale, and causation; ability to develop a dialectic (mastery vs illusion) without straw-person claims; command of epistemic concepts. Significant space must also be given to the student’s insights about humanity and its relationship with nature.

Levels:

- **A+ (90–100):** Clear architecture; definitions folded into argument; polished prose.

~~• **A+ (90–100):** Well-chosen, insightful, original, historically attuned, and powerful and limits~~

As before, this rubric can be converted into a JSON schema (or XML) using ChatGPT, and the resulting schema can be used to guide the AI assessment of student answers.


6 Vibe Coding and Agentic AI

The tools available for the methods in this section include (amongst others):

- OpenAI Codex
- Claude Code

6.1 Agentic AI

- Agentic AI refers to artificial intelligence systems that are not merely passive tools generating outputs, but **entities capable of initiating, sequencing, and adapting actions toward goals with a degree of autonomy**. They operate less like static models and more like decision-making agents.
- Unlike conventional AI models that respond to single prompts or inputs, **agentic AI can plan over extended horizons, interact with other agents or environments, and update their strategies in light of feedback**.
- Such systems **integrate reasoning, memory, and action**, which allows them to handle complex tasks like multi-step research, negotiation, software development, or long-term problem-solving.



6.2 “Vibe Coding”

- “Vibe coding” is an informal term that refers to a coding style or approach that prioritises intuition, creativity, and a relaxed mindset over strict adherence to formal methodologies or best practices. It often involves a more experimental and fluid way of writing code, where the coder follows their instincts and “vibes” with the problem at hand.

- Writing code in such a free-associative, exploratory way lets an LLM fill in structure from loosely sketched or ambiguous instructions. This crosses the boundary between tool use and agentic autonomy.


6.3 “Throw-away Software”

- “Throw-away software” refers to **code or software solutions that are created for temporary use**, often to quickly address a specific problem or need. This type of software is typically not intended for long-term maintenance or scalability, and may be discarded after its immediate purpose has been fulfilled.
- Throw-away software encourages a more **experimental approach to coding**, where “pseudo-developers” can quickly prototype ideas without the pressure of creating production-ready code. This can lead to innovative solutions and **creative problem-solving**.

6.3.1 Example applications

- Scanning a folder of travel receipts and ask it to generate a spreadsheet (CSV) with itemised expenses.
- Reorganising files and folders on your computer based on content.
- Renaming a directory of PDF files based on their content.
 - e.g., “Go into @examples/HD_Forms/ and rename all the files based on the student numbers, student first and last names, and the type of form (e.g. progress report, change of supervisor, etc.) as per the information contained in each file. Rename as {LastName}_{FirstName}_{FormTypeAbbreviation}_{YYYY}_{StudentNumber}.pdf using underscores, title case for names, and standard abbreviations for form types.”
- Spesim: An R species/landscape simulation tool (package) for teaching ecology.

6.4 iKamva Assessment Handling Tool




The screenshot shows the iKamva interface for the course 'BDC 334 S2 2025'. The left sidebar contains a navigation menu with 'Assignments' selected. The main content area is titled 'Download All' and provides instructions: 'Choose download options, and then click 'Download' at the bottom.' A red notice states: 'NOTICE: If you switch between languages in your preferences, for example English to Spanish, you will need to upload this archive in the same language as it was downloaded (or re-download and use a new archive).' The download options are:

- All
- Student submission attachment(s)
- Grade file (file at top level of archive)
- CSV format, file grades.csv
- EXCEL format, file grades.xls
- Feedback comments (comments.txt file if available in student's folder. Comments are put into the Instructor Comments field for each student's submission)
- Feedback Attachment(s)

Below these options, there is a checkbox for 'Include students who have not yet submitted' which is checked. At the bottom, there are 'Download' and 'Cancel' buttons.

I am going to demonstrate a small tool that created which solved my problem of getting the Word documents with AI feedback back to the students. Using AI, I created a small app that runs locally on my computer. Its purpose is to automatically send feedback to individual students, matching each document to the relevant student based on the student numbers embedded within the submitted files. This solves both the logistical bottleneck and reduces opportunities for human error in the feedback distribution process.



The iKamva Assessment Tool converts the inconsistently named files into something more sensible, which serve my needs down the line. This is the input:

7 The AI “Danger List”

- ThesisAI
- Jenni
- Elicit
- Gatsbi
- Genspark
- Scispace

Bibliography