

Assessment Brief: AI Portfolio

Individual portfolio, 100% of the module mark

Weight	100% of the module mark
Size	3,000 words equivalent: code and write-up 1,000; discussion 1,500; reflection 500
Deadline	3.30pm on Wednesday 6 January 2027
Submit	One Word or PDF document, plus a .zip file of your code, on the module page. The AI declaration (Section F) goes on your title page.
AI tools	Not permitted at any stage of the portfolio (see Section F)
Learning outcomes	All four: (a), (b), (c) and (d)

Module learning outcomes. By the end of the module you should be able to:

- (a) develop an understanding of the main areas of Artificial Intelligence, including symbolic and statistical approaches;
- (b) understand different methods for search, planning and reasoning;
- (c) design and discuss applications of Artificial Intelligence to a particular scenario;
- (d) demonstrate awareness of the uses and implications of Artificial Intelligence technologies within businesses and organisations.

Every portfolio is assessed against all four outcomes, whichever problem you choose for Part 1.

A The scenario

Kingsgate Couriers (a fictional company) delivers parcels around Winchester and Hampshire. It has 40 vans, a warehouse and a small customer service team. The managing director has heard a lot about AI and wants to know where it could help. She has listed seven problems:

#	Problem	What the company wants
1	Routes	Each van visits 30 to 60 addresses a day. Find good routes that save time and fuel.
2	Loading the vans	Parcels are stacked in each van so the first delivery is on top. Work out how to load and reorder them.
3	Delivery rules	Apply the company's rules automatically, e.g. parcels worth over £500 need a signature; some goods cannot go on electric vans.
4	Late deliveries	Estimate the chance that a delivery will be late, using weather, traffic and last year's records.
5	Driver rotas	Build weekly shift rotas for 40 drivers, respecting holidays, rest breaks and preferences.
6	Parcel labels	Read handwritten house numbers and postcodes from photos of parcel labels.
7	Customer messages	Sort incoming emails and reviews (complaint or praise, urgent or not) and answer common questions.

Your portfolio has three parts:

- **Part 1:** choose **one** of the seven problems in the table above and build a small working AI system for it. Section B gives the starting code and the minimum to do for each problem.
- **Part 2:** advise Kingsgate Couriers. In 2a you cover **all seven** problems briefly. In 2b and 2c you go into depth on **one** problem of your choice: it can be your Part 1 problem or a different one. In 2d you discuss the implications for the company as a whole.
- **Part 3:** reflect on your learning.

Part	Weight	Size
1. AI system implementation	30%	Code, plus a write-up of up to 500 words
2. Discussion and topic research	55%	1,500 words
3. Reflection	15%	500 words

B Part 1: AI system implementation (30%)

How to choose

1. Choose **one** of the seven problems.
2. Find that problem's row in the table below. The row tells you which practical code to start from and the minimum you must do.
3. Use only that row. You cannot mix rows: for example, if you choose problem 2 (Loading the vans), you start from `blocks.py` and write loading problems; you cannot use the classifier code or `deliveries.csv` from problem 4.

You can go beyond the minimum, as long as you stay with your chosen problem.

Every row starts from code we used in a weekly practical, so you are adapting something you have already run, not starting from nothing. All the practical code is in the module's code repository: github.com/sakibanwar/comp5001-ai-code.

Your problem	Start from	Minimum to do
1. Routes	Week 2: <code>maze.py</code>	Make a grid map of a small area with a depot and a customer. Find a route with BFS, DFS and A* (Manhattan distance). Compare the states explored and the route length.
2. Loading the vans	Week 3: <code>blocks.py</code>	Treat parcels as blocks. Write at least 3 loading problems where the goal order matches the delivery order. Show and explain the plans.
3. Delivery rules	Week 4: <code>logic.py</code>	Encode at least 6 company rules and facts. Use model checking to answer at least 3 questions about specific parcels.
4. Late deliveries	Week 7: <code>banknotes0.py</code>	Use <code>deliveries.csv</code> (provided). Train at least 3 classifiers and compare their accuracy with always guessing "on time".
5. Driver rotas	Week 6: <code>schedule0.py</code> or <code>schedule1.py</code>	Assign at least 6 drivers to at least 5 shifts with at least 4 different constraints. Show one valid rota.
6. Parcel labels	Week 8: <code>handwriting.py</code>	Train the digit recogniser, then change the network (layers, units or dropout) at least twice. Compare accuracy and test it with <code>recognition.py</code> .
7. Customer messages	Week 9: <code>sentiment.py</code>	Write at least 30 positive and 30 negative customer messages. Train the classifier and test it on 10 new messages. Explain its mistakes.

What to submit for Part 1

- Your code in a .zip file, with comments explaining what you changed. Your system must not call a generative AI service (for example an LLM).
- A write-up of up to 500 words in your document, covering: which problem you chose; what your system does and how the AI technique works; what you changed from the practical code; your test runs, with outputs or screenshots; and the limitations of your system.

C Part 2: Discussion and topic research (55%)

Write this as advice to Kingsgate Couriers' managing director, in four sections: 2a covers all seven problems; 2b and 2c go into depth on one problem; 2d covers the company as a whole. The seminar activities in Weeks 4 to 9 practise each section (see Section G).

2a. An AI map for Kingsgate Couriers (about 350 words)

Make a table with one row for **each of the seven problems**. For each problem, give: the specific AI method you would suggest (for example A* search, STRIPS-style planning, logical rules with model checking, a Bayesian network, constraint satisfaction, a classifier, a convolutional neural network, a Naive Bayes text classifier), not just "AI" or "machine learning"; whether it is symbolic or statistical; and one or two sentences on why it fits. For problems 1 to 4, also say in one sentence how the method works on that problem (for example, what the states, actions and goal are for routes).

2b. Comparing two approaches (about 450 words)

Pick **one** of the seven problems. It can be the problem you chose for Part 1 or a different one, and you use the same problem in 2c. Describe two different AI approaches that could solve it, ideally from different weeks (for example, A* search against hill climbing for routes, or rules against a trained classifier for customer messages). **At least one of the two must be a search, planning or reasoning method:** search, planning, logical reasoning, probabilistic reasoning or constraint satisfaction (Weeks 2 to 6). Say whether each approach is symbolic or statistical, and what difference that makes. Compare them on: the data or knowledge each needs; how well it would work; whether it can explain its decisions; and the cost to build and maintain. Say which one you recommend and why.

2c. Design of the chosen approach (about 350 words)

For the problem you chose in 2b, describe how your recommended approach would work at Kingsgate Couriers: its inputs and outputs; how the problem is represented (for example states and actions, rules, features, or training data); how you would test that it works; and what happens when it gets something wrong, including who checks its output. A simple diagram is welcome.

2d. Implications for the company (about 350 words)

Discuss the benefits and risks of using AI at Kingsgate Couriers. Cover at least: one ethical risk (for example bias, privacy or effects on staff); UK GDPR; and which EU AI Act risk level applies to **two** of the seven problems. Give at least one real example of an organisation that uses AI for a similar problem, with a source. Finish with a short recommendation: what should the company do first (for example, whether to build or buy)?

Research

Support your points with at least 5 sources, cited in the text and listed at the end using Harvard referencing. Use the module reading (Russell and Norvig; Poole and Mackworth) and at least one source about AI in real organisations.

D Part 3: Reflection (15%)

Write about 500 words in the first person. Use these questions as a guide:

1. What did you learn about AI that you did not know before this module?
2. What went well in your implementation, and what was hard? How did you deal with it?
3. What would you do differently if you started the portfolio again?
4. Which learning outcome do you feel strongest in, and which weakest? Why?
5. How might you use what you learned in future study or work?

Tip: a few lines of notes after each week's seminar and practical make this part much easier to write.

E Marking criteria

Each criterion is marked on the University's scale. The table shows what the main bands look like. Weights add up to 100%: Part 1 = 30%, Part 2 = 55%, Part 3 = 15%.

Criterion (weight)	70+ Excellent	50–69 Good	40–49 Pass	Below 40 Fail
Part 1 Working system (12%)	Runs correctly; well adapted to the problem; clear, commented code	Runs with minor issues; sensible changes to the practical code	Runs, but changes are small or only partly work	Does not run, or is the practical code with no real changes
Part 1 Testing and write-up (18%)	Thorough tests with clear results; technique explained accurately; limitations well judged	Several tests; technique explained with small gaps	Some testing; explanation is basic	Little or no testing or explanation
2a AI map (12%)	A specific, suitable method for all seven problems with clear reasons; symbolic and statistical used correctly; methods for problems 1–4 explained accurately	Most problems matched to specific methods; some reasons or explanations thin	Some suitable matches; methods vague ("use AI") or explanations missing	Few or no suitable matches
2b Comparison (14%)	Accurate comparison on all four points, including a search, planning or reasoning method; symbolic and statistical difference well explained; well-justified choice	Sound comparison with some gaps	Mostly describes the approaches rather than comparing them	Missing, inaccurate, or no search, planning or reasoning method

Criterion (weight)	70+ Excellent	50–69 Good	40–49 Pass	Below 40 Fail
2c Design (10%)	Clear, specific design that could be built; testing, failure cases and human checks considered	Reasonable design; some parts vague	General outline only	Missing, or does not fit the problem
2d Implications (14%)	Insightful discussion of benefits and risks; UK GDPR and EU AI Act applied correctly; real organisation example well used; practical recommendation	Relevant points, applied to the company	Generic points, weakly applied to the company	Missing, or not relevant
2 Research and referencing (5%)	At least 5 good sources, well used; accurate Harvard referencing	Adequate sources; minor referencing errors	Few sources; several referencing errors	Little or no research or referencing
Part 3 Reflection (15%)	Honest, specific and critical; linked to the learning outcomes; clear lessons and next steps	Some reflection on specific experiences	Mostly description of what was done	Missing, or no reflection

F Rules and policies

Use of AI tools: not permitted

You must not use generative AI tools at any stage of this portfolio. This includes:

- chatbots and assistants such as ChatGPT, Claude, Gemini, Microsoft Copilot, DeepSeek and Perplexity;
- agentic AI tools that write, edit or run code for you, such as Claude Code, Google Antigravity, Cursor, OpenAI Codex and Windsurf;
- AI features built into other software, including GitHub Copilot, AI code completion in VS Code or other editors, Gemini in Google Colab and Copilot in Word;
- AI translation and paraphrasing tools, such as DeepL or QuillBot, used to produce or rewrite your text.

The ban covers every part of the work: developing ideas, understanding the readings, choosing methods, writing or fixing code, testing, making tables and diagrams, and drafting or rewriting your text.

The one exception: you may use Grammarly or a similar tool to correct spelling, grammar and punctuation, and to improve the clarity of individual sentences you have written yourself. You must not use it to generate new content or to rewrite whole paragraphs or sections.

Studying AI is not the same as using it to do the work. In some practicals and seminars you will try AI tools because the activity asks you to (for example, Week 9 compares an LLM chatbot with a Naive Bayes classifier). That is fine, but nothing produced by an AI tool may go into your portfolio, and your Part 1 system must not call a generative AI service. Writing *about* AI tools in Part 2, with

sources, is of course allowed.

Before you start, switch off AI code completion in your editor (for example GitHub Copilot in VS Code) and the AI features in Google Colab (Tools > Settings) if you use it. Some are switched on by default.

What you can use: the module slides, seminar sheets and textbooks; software documentation, such as the Python, scikit-learn, TensorFlow and NLTK guides; ordinary web searches, but not the AI-generated answers or summaries that some search engines show; spell-checkers and grammar checkers, within the exception above; help from the module team; and discussion of ideas with classmates, though your code and writing must be your own.

Why: this portfolio tests whether *you* can build a small AI system, explain how it works, and judge where AI fits in an organisation. Those are the skills in the module's learning outcomes. An AI tool would be doing that work for you.

How this is checked:

- You may be asked to explain your code, your choices and your results in a short conversation with the module team. Being unable to explain your own work may be treated as suspected academic misconduct.
- Keep the version history of your code and your document (for example Git commits, or OneDrive version history). You may be asked to show how your work developed.
- Suspected use of AI will be handled under the University's academic misconduct procedures.

Declaration. Put this statement, with your name and the date, on the title page of your portfolio:

I confirm that I have not used generative AI tools at any stage in producing this portfolio or its code, other than to correct spelling, grammar and punctuation or to improve the clarity of individual sentences I wrote myself.

Word count and academic integrity

- Follow the University's word count policy. Tables in 2a count towards the word count; code, references and appendices do not.
- The portfolio must be your own work. Cite every source, including any code you did not write yourself beyond the practical code.

G Support

- **Weekly practicals** (Weeks 2 to 9) give you the starting code for every Part 1 problem. The code is at github.com/sakibanwar/comp5001-ai-code.
- **Seminar activities** practise Part 2 with other organisations: the design checklist (Week 4), comparing two approaches (Week 5), build or buy (Week 6), UK GDPR and the EU AI Act (Week 7), human oversight (Week 8) and pitching a recommendation (Week 9).
- **Week 10** brings the module together and includes a portfolio clinic, where you can get feedback on a draft.
- **Office hours:** Wednesdays 4 to 5pm, 30-minute slots, booked at calendly.com/sakib-anwar. You can also ask questions in any practical.

Suggested structure for your document

1. Title page: name, student number, word count, and the AI declaration
2. Part 1: Implementation write-up
3. Part 2a: AI map for Kingsgate Couriers
4. Part 2b: Comparing two approaches
5. Part 2c: Design of the chosen approach
6. Part 2d: Implications for the company
7. Part 3: Reflection
8. References

About deliveries.csv (Part 1, problem 4)

600 past deliveries, one per row. Columns: `distance_km`, `stops_before` (stops the van made before this one), `weight_kg`, `rain` (1 = raining), `rush_hour` (1 = yes), `driver_experience_years`, and `late` (1 = late, 0 = on time). The data is made up for this assessment. About one delivery in three is late.