



WEEK 1 CLASS ACTIVITY — “SPOT THE AI”

Foundations of AI in Business · LO1 · ~60 minutes in the tutorial

Goal. Put the lecture’s *technology stack* to work. By the end you’ll be able to look at any real business system and say which layer of AI it uses, what job it does, and whether it’s likely to create real value or fizzle as a pilot.

You’ll work in teams of 3–4 (cross-campus teams pair up via the Google Meet breakout). You need this worksheet, the lecture’s stack diagram, and a willingness to argue.

THE TECHNOLOGY STACK (YOUR REFERENCE)

Layer	One-line test: “ <i>This system...</i> ”	Typical business use
Rule-based / expert	...follows explicit <i>if-then</i> rules a human wrote.	Eligibility checks, phone menus, tax logic
Classical machine learning	...predicts from patterns it learned in lots of past data.	Credit scoring, fraud detection, recommendations, churn
Deep learning	...perceives images, audio, or video.	Medical imaging, face unlock, voice transcription, defect cameras
Foundation model (generative)	...understands or generates free-form language/images.	Chatbots, copywriting, document Q&A, coding assistants
Agentic	...takes a <i>sequence of actions</i> toward a goal, using tools.	Auto-reordering stock, research agents, build-and-test agents
None	...is just ordinary software (no learning, no judgement).	A calculator, a sorting filter, a fixed dashboard

Remember from the lecture: higher is *not* better. The right answer is the cheapest, most reliable, most explainable layer that solves the actual problem.

PART A — CLASSIFY SIX SYSTEMS (≈30 MIN)

For **each** scenario: (a) pick the **most likely layer** (sometimes a system blends two — name both and say which dominates); (b) name the **business function** it serves; (c) judge **likely high-value or likely low-value pilot**, and say **why** in one line. There isn’t always one perfect answer — your *reasoning* is what matters.



Scenario 1 — “ShopEasy”

An online grocer shows each shopper a personalised “you might also need” list at checkout. It’s built from three years of purchase history; the more you shop, the better it gets. Nobody wrote rules about which products go together — the system worked it out.

- Layer: _____ Function: _____ Value (H/L) + why: _____

Scenario 2 — “QuickLoan”

A micro-lender’s website instantly tells applicants if they pre-qualify. Behind it: a checklist a compliance officer wrote — *if age ≥ 21 AND income ≥ X AND no default in 24 months → pre-qualify*. The logic is printed in the staff manual.

- Layer: _____ Function: _____ Value (H/L) + why: _____

Scenario 3 — “CareScan”

A private hospital uses a system that reads chest X-rays and highlights areas a radiologist should check. It was trained on hundreds of thousands of labelled scans. It can’t explain *why* it flagged a spot — only that similar-looking spots were often significant.

- Layer: _____ Function: _____ Value (H/L) + why: _____

Scenario 4 — “BrandVoice”

A marketing agency types a one-line request — *“write 5 cheeky Instagram captions for a durian-flavoured latte, Gen-Z tone”* — and gets back usable drafts in seconds, in plain English, no special setup.

- Layer: _____ Function: _____ Value (H/L) + why: _____

Scenario 5 — “StockMind”

A convenience-store chain runs a system that monitors sales and supplier prices, *decides on its own* when to reorder each product, places the orders, and reroutes deliveries if a supplier is late — escalating to a human only for unusual cases.

- Layer: _____ Function: _____ Value (H/L) + why: _____

Scenario 6 — “DashBoard Pro”

A logistics firm’s manager opens a screen showing today’s deliveries, total fuel cost, and a bar chart of late parcels. The numbers are pulled straight from the database and added up. The vendor’s brochure calls it “AI-powered.”

- Layer: _____ Function: _____ Value (H/L) + why: _____
-

PART B — DEFEND AND DISAGREE (≈15 MIN)

Pair with another team. Compare your answers to **Scenarios 3, 5, and 6** — the ones that spark the most disagreement. For each clash:

- State both views in one sentence each.
 - Decide: is this a genuine “it depends,” or did one team misread the technology?
 - **Scenario 6 in particular:** is the vendor’s “AI-powered” label honest? This is your first taste of a theme that runs all semester — *AI-washing*, where ordinary software gets an “AI” sticker to sell better.
-



PART C — BRIDGE TO YOUR PROJECT (≈10 MIN)

Individually, then share in your team:

Think of **one real annoyance** — something slow, repetitive, or frustrating in your own life, a job you've had, or a business you use. **(1)** Write it as a precise problem ("answer the simple 40% of questions that clog the helpline," not "improve service"). **(2)** Which stack layer would you *try first*, and why? **(3)** What's the cheapest version that might work?

Keep this. It's a candidate for your **Assessment 1** business problem, and you'll develop it next week with the Business Model Canvas.

FACILITATOR ANSWER KEY (DO NOT DISTRIBUTE)

Indicative answers — accept any well-reasoned alternative; the reasoning is the assessment target.



#	Layer	Function	Value	Teaching point
1 ShopEasy	Classical ML (recommendation)	Sales / merchandising	High — proven, improves with data, low stakes if wrong	The quiet workhorse from Era II; most enterprise AI value still looks like this
2 QuickLoan	Rule-based	Risk / lending	High for what it is — transparent, auditable, cheap	Higher isn't better: rules are the <i>right</i> call where you must explain a decision (ties to Week 3– 4 governance)
3 CareScan	Deep learning (vision)	Clinical decision support	High but high-stakes — needs human-in- the-loop; black-box explainability is a real risk	Perception layer; the explainability gap previews the governance block
4 BrandVoice	Foundation model (generative)	Marketing / content	Medium–High — fast and cheap, but “confidently wrong”; needs human edit	This is the layer your Week 6 work and MVP lean on
5 StockMind	Agentic (with ML underneath)	Operations / supply chain	Potentially high, genuinely risky — acting-not-just- answering raises failure stakes	The frontier; tie to Gartner's “40% of agentic projects cancelled by 2027”
6 DashBoard Pro	None — ordinary software	Reporting	N/A — it's AI-washing	The most important answer of the day: not everything labelled “AI” is AI

Debrief script (last 5 min, runs into the tutorial debrief): “Two things to take away. One — you can now *name the layer*, which is more than most managers can do. Two — notice that your value judgements depended less on how fancy the AI was and more on the *problem* and the *stakes*. Hold that. It's the 6% lesson, and it's the spine of Assessment 1.”