

Governing AI by Autonomy and Consequence

Applying the ACTAM™ model to a real AI system

Guest lecturer	Ali Alasiri — Chairman, BlackBaz Holding; former CEO, Saudi e-Government Program (Yesser)
Lecture	"The Autonomy Dividend: Governing AI for Strategic Value in the Saudi Economy", 30 September 2026
Weighting	Counts toward your course grade as determined by your course instructor
Length	1,200 words ± 10%, excluding the cover block, tables, references and appendices
Due	14 October 2026, 23:59 AST
Submit to	ali@alasiri.net • file name: ACTAM_StudentID_LastName.pdf (PDF only)
Template	Use the ACTAM Answer Template. Submissions in any other structure are returned unmarked.

1. What this assignment asks of you

The lecture argued one proposition: an AI system's governance obligation should follow its position on two axes — how much autonomy it holds, and how much consequence it carries. This assignment asks you to apply that proposition to a real system and defend your reasoning in writing.

You are not asked to agree with the model. A submission that applies it carefully and then shows where it fails will score higher than one that applies it uncritically.

Learning outcomes

- Translate a governance framework into a repeatable assessment procedure and apply it to a real organisation.
- Score a system on defined dimensions and justify each score from evidence rather than impression.
- Convert an assessment into decision rights: who approves, who owns, what evidence is kept, and by when.
- Evaluate the limits of the framework you used, and cite primary sources to support your claims.

2. Choose one brief

Pick the one closest to your track. All three are marked against the same rubric and use the same template.

Brief	What you do	Best suited to
A Governance audit	Pick one AI system a real Saudi or international organisation has deployed. Score it on both axes with the five-step method, place it in a zone, name the accountable owner, and recommend three proportionate controls.	Management, strategy and operations students
B Return-on-Trust case	Build the business case for an AI governance programme at a named company: expected loss avoided, value delivered earlier, market access, minus programme cost. Source every input and give a sensitivity range.	Finance and accounting students

Brief	What you do	Best suited to
C Regulatory readiness	A Saudi firm sells an AI product into the European Union. Map its obligations on both the EU and Saudi timelines, state what it must be able to evidence, by when, and who inside the company owns each item.	Legal, compliance and public-policy students

3. What every submission must contain

1. The completed cover block: name, student ID, section, brief chosen, word count, date, and the signed declaration.
2. The six-row assessment table (decision, autonomy, consequence, zone, owner, evidence) — this is the spine of the work.
3. A named accountable human and a review date. An assessment without an owner and a date is incomplete by definition.
4. At least five sources, of which at least three are primary: the standard, the law, the company's own disclosure, the regulator's decision. Not this brief, not the lecture slides, not an AI chatbot.
5. A section stating what the framework failed to surface about your system.
6. The AI-use disclosure, completed honestly. Using an AI tool is permitted; not disclosing it is an integrity violation.

4. The method, in five steps

Reproduce this reasoning in your answer. Scores are averages of the sub-dimensions, each rated 1 to 5.

Step	What you do	What the marker looks for
1 · Describe the decision	What decision does the system touch, for whom, how often, at what value?	A specific decision, not a product description.
2 · Score autonomy	Speed (human-paced → real time), scale (one case → millions), depth (recommends → decides → acts). Average the three.	Each score justified in one sentence from observable facts.
3 · Score consequence	Severity, reversibility, breadth, exposure. Average the four. Score consequence from the affected person's side as well as the company's.	Reversibility reasoned about explicitly — it is the dimension most often ignored.
4 · Place it, name the owner	Read the zone off the matrix: Experiment, Product Control, Executive Approval, or Prohibited / Exceptional. Name one accountable human role.	The zone follows from the two scores, and the owner has the authority the zone implies.
5 · Set the evidence	What proof, at what cadence, reviewed by whom, by what date.	Evidence a third party could actually inspect.

The four zones

Zone	Autonomy × consequence	Governance demand
Experiment	Low × low	Sandbox, logged, fast iteration.
Product Control	High × low	Release gates, monitoring, a human escalation path.
Executive Approval	Low × high	Risk-committee sign-off, meaningful human review of each decision.
Prohibited / Exceptional	High × high	Board or ministerial approval, or do not deploy.

Where a law applies, the law governs. The matrix is a management heuristic. If your system falls under the EU AI Act, SDAIA's National AI Risk Management Framework or the PDPL, say so and cite the instrument.

5. How your work is marked

Weight	Criterion	What earns the marks
30%	Correct use of the method	Both axes scored and justified, sub-dimensions shown, zone read correctly from the scores.
25%	Evidence and sourcing	Primary sources cited, figures traceable, estimates labelled as estimates, no unsupported numbers.
25%	Business judgement	Controls proportionate to the risk; the recommendation survives a sceptical CFO or regulator.
20%	Accountability	A named owner, an escalation path, a review date, and an evidence plan that someone could audit.

Marks are for reasoning, not for conclusions. Two students can score reversibility 2 and 4 and both score well if each argues it from evidence. Marks are lost for an unnamed owner, a missing review date, an unsourced figure, or a word count outside the range.

6. Submission, integrity and what happens next

File format	PDF only, produced from the ACTAM Answer Template.
File name	ACTAM_StudentID_LastName.pdf — for example ACTAM_202312345_Alqahtani.pdf
Send to	ali@alasiri.net with the subject line: ACTAM assignment — StudentID — LastName
Deadline	14 October 2026, 23:59 AST. Late submissions are marked at the course instructor's discretion.
Feedback	Each submission is marked against the rubric above. A consolidated report is sent to the College of Business, which records the grade against your course.
AI use	Permitted for research, structuring and language. You must disclose which tool, for what, in the declaration. Every citation an AI tool gives you must be opened and verified — fabricated references are the single most common failure and are treated as a citation error, not a style issue.
Originality	Your own work, in your own words. Quoted material in quotation marks with a source.

A note on why the disclosure exists. In 2025 a global consulting firm delivered a report to an Australian government department containing references to academic papers that did not exist and a fabricated quotation from a court judgment; the work had been produced with a generative AI tool that was not disclosed, and the firm refunded part of the fee. The lesson is the assignment: a brand is not a control, and disclosure plus verification is.

Questions: ali@alasiri.net. Ask before the deadline rather than explaining afterwards.

ACTAM™ is a pending SAIP service-mark application (class 42, filed 13 September 2026). The framework names are descriptive and openly citable. Cite as: Alasiri, A. (2026) 'Digital Trust, Governance and Risk Management in the AI Era', ACTAM™ framework, Alfaisal University guest lecture, Riyadh.