

CSC 163: Computing, Ethics, and Society

Chapter 8

Governance Exercises Solution

Department of Computer Science,
Hofstra University

Exercise: Levels of governance

Place each measure at the level that implements it: corporate governance, national governance, international governance, or compute security (two measures per level).

Measure	Level
Incorporating as a public benefit corporation instead of a C-corp	
A board creates a special committee for AI safety risks	
Legal liability for harms caused by AI systems	
A value-added tax on GPUs	
A treaty with inspections by an agency (like the IAEA)	
A summit such as the UK AI Safety Summit	
Location verification on chips, enabling geofencing	
Offline licensing of AI chips	

Exercise Solution: Levels of governance

Corporate = legal and ownership structure inside a firm. National = laws, liability, taxes. International = treaties, summits, agencies. Compute security = security features built into chips.

Measure	Level
Incorporating as a public benefit corporation instead of a C-corp	Corporate governance
A board creates a special committee for AI safety risks	Corporate governance
Legal liability for harms caused by AI systems	National governance
A value-added tax on GPUs	National governance
A treaty with inspections by an agency (like the IAEA)	International governance
A summit such as the UK AI Safety Summit	International governance
Location verification on chips, enabling geofencing	Compute security
Offline licensing of AI chips	Compute security

Exercise: Compute threshold

A regulation requires a report for any training run above 10^{25} floating point operations (FLOP). A lab trains a model on 10,000 GPUs, each delivering 10^{15} FLOP/s at peak, with 40 % average utilization, for 30 days (1 day = 86,400 s).

- (a) Compute the total training compute in FLOP. Is the threshold exceeded?
- (b) The lab uses 8,000 GPUs instead (same utilization and duration). Exceeded?
- (c) What is the minimum number of GPUs needed to reach 10^{25} FLOP in 30 days?
- (d) Name three properties of compute that make it more governable than data or algorithms.

Exercise Solution: Compute threshold

- (a) Compute = GPUs $\times$ FLOP/s $\times$ utilization $\times$ seconds = $10^4 \times 10^{15} \times 0.4 \times 2.592 \times 10^6$ s $\approx$ **1.04×10^{25} FLOP**: the threshold is **exceeded** (by about 4 %).
- (b) 8,000 GPUs: $\approx$ **8.3×10^{24} FLOP**, **below** the threshold.
- (c) $N = 10^{25} / (10^{15} \times 0.4 \times 2.592 \times 10^6) \approx$ **9,650 GPUs**.
- (d) **Physical** (rivalrous, trackable), **quantifiable** (FLOP/s allows thresholds), **excludable** (few chip firms to regulate).

Exercise: Explosive growth

Some researchers argue AI-driven automation could raise growth up to ten times the current growth rate. Suppose the current growth rate is 3 % per year, so AI-driven growth would be 30 % per year. Doubling time = $\ln 2 / \ln(1 + g)$.

- (a) Doubling time at 3 % and at 30 %?
- (b) By what factor does the output grow in 10 years in each case?
- (c) Bottleneck illustration (Amdahl's law): automation makes 90 % of the work infinitely fast, the other 10 % stays at human speed. What is the maximum speedup of the whole process?
- (d) Name the bottleneck concept in which non-automated jobs get more expensive.

Exercise Solution: Explosive growth

- (a) 3 %: $\ln 2 / \ln 1.03 \approx \mathbf{23.4 \text{ years}}$. 30 %: $\ln 2 / \ln 1.3 \approx \mathbf{2.6 \text{ years}}$.
- (b) $1.03^{10} \approx \mathbf{1.34\times}$; $1.3^{10} \approx \mathbf{13.8\times}$.
- (c) Speedup = $1 / 0.10 = \mathbf{10\times}$: the non-automated tasks cap the gain.
- (d) **Baumol's cost disease**: non-automated jobs may see rising costs, offsetting the overall impact.

Exercise: True or false

For each statement write T or F. If it is false, correct it with a few words.

Statement	T / F	Correction (if false)
Releasing model weights is a reversible decision: they can be taken back if misuse appears		
An AI singleton reduces competitive pressures but is a single point of failure		
Shareholder theory says corporations should weigh employees, customers and society equally		
Legal liability alone is enough: a company can pay for the damages of an AI-enabled bioweapon		
Compute is physical, quantifiable and excludable, so it is a promising target for governance		
AI seems defense dominant, since good AIs can always neutralize bad AIs		

Exercise Solution: True or false

Each statement: 1 point for T/F, 1 point for the correction.

Statement	T / F	Correction (if false)
Releasing model weights is a reversible decision: they can be taken back if misuse appears	F	It is irreversible proliferation; released weights cannot be taken back
An AI singleton reduces competitive pressures but is a single point of failure	T	—
Shareholder theory says corporations should weigh employees, customers and society equally	F	That is stakeholder theory; shareholder theory: maximize value for shareholders
Legal liability alone is enough: a company can pay for the damages of an AI-enabled bioweapon	F	Liability is a limited tool; damages can exceed what a company can pay
Compute is physical, quantifiable and excludable, so it is a promising target for governance	T	—
AI seems defense dominant, since good AIs can always neutralize bad AIs	F	AI seems offense dominant (e.g., easier to engineer pandemics than to counteract them)