ENVIRONMENTAL AI REGULATION: THE GLOBAL PUSH TOWARDS GREEN AI AND SUSTAINABLE ARTIFICIAL INTELLIGENCE

Detailed Summary

1. Title and Central Theme

The Global Landscape of Environmental AI Regulation: From the Cost of Reasoning to a Right to Green AI

The paper examines the rapidly increasing **environmental cost of Artificial Intelligence**, particularly the electricity, water and carbon emissions associated with modern AI systems.

Its central argument is that **AI is becoming more environmentally expensive at the same time that legally required transparency about those environmental costs remains inadequate—and in some respects has actually declined.**

The authors argue that existing environmental regulation is focused primarily on **data centres and physical infrastructure**, rather than on the individual AI models and applications responsible for increasingly intensive computing.

They therefore propose a new regulatory approach based on three pillars:

1. **Mandatory model-level environmental transparency**
2. **New user rights to choose or reject environmentally costly AI**
3. **International coordination to prevent regulatory arbitrage**

2. Why the Paper is Important

The paper begins from a major contradiction.

AI adoption is expanding rapidly across:

- business
- communication
- search engines
- advertising
- software
- image generation

- video generation
- coding
- enterprise applications

At the same time, the environmental consequences of this expansion remain poorly understood.

The authors particularly highlight the emergence of **reasoning models**, which became increasingly widespread in 2025. Their use in enterprise applications reportedly increased dramatically, while compute-intensive image and video-generation systems are also becoming increasingly integrated into advertising and other commercial activities.

The paper gives a striking example involving Coca-Cola: its Christmas 2025 advertising campaign reportedly generated approximately **70,000 trial videos using AI**, but there was no disclosure concerning the computing resources or energy required for those generations.

This illustrates the authors' broader concern:

AI is becoming embedded in everyday activities, but its environmental footprint remains largely invisible to users.

3. The Three Major Contributions of the Paper

The authors identify three principal contributions.

First: Measuring AI's growing environmental impact

The paper brings together empirical evidence concerning the environmental costs of:

- generative AI search
- reasoning models
- AI inference
- other increasingly compute-intensive applications.

It particularly contributes new evidence concerning the difference between reasoning and non-reasoning models.

Second: Comparative legal analysis

The paper compares environmental AI regulation across **10 countries plus the international framework**, examining whether governments regulate:

- voluntarily or mandatorily
- facilities or models
- with high or low environmental ambition.

Third: New legal proposals

The authors go beyond existing discussions and propose:

- a right to use digital infrastructure without unnecessary generative AI
- a right to choose a green AI model
- mandatory environmental disclosure
- international coordination.

4. Environmental Cost of AI

The paper explains that the environmental impact of AI is not limited to electricity consumed while training a model.

A complete environmental assessment can involve:

- electricity
- carbon emissions
- water consumption
- manufacturing of computing hardware
- data-centre infrastructure
- cooling systems
- idle computing capacity
- experimentation
- inference.

Earlier research concentrated heavily on **model training**.

The paper argues that this is no longer sufficient.

Increasingly, attention must shift toward **inference**—the computation required every time people actually use an AI model.

5. The Importance of Inference

This is one of the paper's most important points.

Training occurs when a model is created or substantially updated.

Inference occurs every time the model responds to a user.

With billions of AI interactions, inference can become a much larger cumulative environmental burden than the original training process.

The paper notes that the final model training run represented only **5.7% of OpenAI's reported 2024 compute budget**, while Research & Development represented **64.5%**. This illustrates why focusing exclusively on the final training run provides an incomplete picture of AI's environmental footprint.

6. The Jevons Paradox

The paper also discusses the **Jevons Paradox**.

The basic idea is:

When technology becomes more efficient and cheaper, people may use substantially more of it, causing total resource consumption to increase.

Thus, improvements in AI efficiency do not necessarily guarantee lower overall environmental impact.

For example:

- AI models become more efficient.
- Computing becomes cheaper.
- AI becomes available to more people.
- People use AI much more frequently.
- Total energy consumption may therefore rise.

The authors believe there is evidence suggesting this rebound effect may already be occurring, although the lack of comprehensive data makes its precise scale difficult to establish.

7. Case Study: Generative AI in Web Search

Traditional search engines generally relied on **extractive approaches**.

They searched existing documents and presented relevant information.

Modern search engines increasingly use **generative AI** to create summaries or answers.

The paper identifies three approaches:

- traditional extraction
- generative AI responses
- hybrid systems such as Retrieval-Augmented Generation (RAG).

The environmental implications are significant because generating a response can require considerably more computation than simply retrieving existing information.

8. Growth of Search Volume

The paper uses Google Search as an illustration.

Google reported that a traditional search in 2009 consumed approximately **0.3 Wh**.

Google reported in 2025 that a median text prompt to Gemini consumed approximately **0.24 Wh**.

At first glance, the newer figure appears lower.

But the authors point out that the **number of searches has dramatically increased**.

Approximately:

- **792 billion searches per year** in 2009
- nearly **5 trillion searches per year** in 2025.

If every 2025 search were powered by Gemini, the overall energy requirement would therefore be approximately **six times higher**, despite the lower estimated energy consumption per individual query.

This is a key illustration of the paper's argument:

Greater efficiency per query does not necessarily mean lower total environmental impact.

9. Case Study: Reasoning Models

The second major case study concerns **reasoning models**.

These models perform additional computational steps before producing an answer.

They often generate:

- longer reasoning chains
- more tokens
- more computational operations.

The AI Energy Score project found a dramatic difference between ordinary and reasoning models.

According to the paper:

- reasoning models use **approximately 30 times more energy on average**
- some reasoning models consume **150–700 times more energy** than their base models.

The reason is partly that reasoning systems can generate vastly longer outputs, with some producing hundreds of times more tokens than their non-reasoning counterparts.

10. GPT-5 and ChatGPT Example

The paper cites third-party estimates suggesting that a GPT-5 reasoning-enabled system could consume around **18 Wh of electricity for a medium-length 1,000-token prompt**.

It also cites an estimate that ChatGPT handled approximately **2.5 billion requests per day** in 2025.

On those assumptions, the annual electricity requirement could be comparable to the electricity consumption of **more than one million average U.S. households**.

Importantly, the authors caution that these are **third-party estimates**, rather than official figures from the AI provider.

This limitation itself demonstrates why mandatory disclosure is necessary.

11. The Global Regulatory Problem

The paper then moves from technology to law.

The authors identify three regulatory dimensions.

Dimension 1: Voluntary vs Mandatory

Some governments rely on:

- voluntary commitments
- industry standards
- best practices
- incentives.

Others impose:

- binding reporting requirements
- regulatory obligations
- penalties.

Dimension 2: Facility-Level vs Model-Level

Facility-level regulation regulates:

- data centres
- electricity consumption
- infrastructure.

Model-level regulation examines:

- specific AI models
- energy per inference
- water consumption
- model efficiency.

The authors argue that model-level regulation is crucial because two models operating in the same data centre can have radically different environmental footprints.

Dimension 3: Environmental Ambition

Countries differ substantially in whether environmental protection is:

- a major regulatory objective
- balanced against innovation
- subordinated to competitiveness.

12. European Union

The **European Union** is identified as the most advanced jurisdiction in terms of AI-specific environmental transparency.

Relevant frameworks include:

- Energy Efficiency Directive
- Corporate Sustainability Reporting Directive
- Corporate Sustainability Due Diligence Directive
- EU AI Act.

The Energy Efficiency Directive requires data centres with capacity of **500 kW or more** to report sustainability indicators to an EU database.

However, the authors identify important shortcomings.

Under the AI Act:

- General-Purpose AI providers have certain obligations concerning training compute and energy.
- High-risk AI systems have different requirements.
- **Inference is not adequately addressed in the binding legal framework.**
- Much of the relevant information is available to authorities rather than the public.

The 2025 General-Purpose AI Code of Practice extends documentation toward inference compute, but compliance with the relevant Code provisions is **voluntary**.

Therefore, even the EU's relatively advanced framework has significant limitations.

13. United States

The United States follows a substantially different approach.

The paper characterizes the federal approach as **deregulatory federalism**.

Traditional environmental laws include:

- Clean Air Act
- Clean Water Act.

Individual states have taken some additional measures concerning AI and data centres.

However, model-level energy disclosure remains absent.

The paper highlights California AB 2013, effective from 1 January 2026, which requires certain generative AI data disclosures but does **not require disclosure of energy consumption or compute**.

The paper further discusses the second Trump administration's emphasis on:

- AI competitiveness
- deregulation
- fossil energy
- reducing regulatory barriers
- faster data-centre development.

The authors characterize the resulting framework as:

voluntary + facility-level + general environmental law + declining environmental ambition.

14. China

China represents a different model.

China has pursued:

- green infrastructure
- energy efficiency
- renewable energy
- strategic data-centre relocation.

The "**Eastern Data, Western Computing**" initiative is particularly important.

It seeks to move energy-intensive computing toward western regions with greater energy and resource availability.

China has incorporated green development into several major policy frameworks, including:

- Made in China 2025
- New Generation AI Development Plan
- 14th Five-Year Plan
- Green and Low-Carbon Development of Data Centres initiative.

However, the major weakness is **transparency**.

China has introduced efficiency standards but does not provide the public with extensive AI-specific environmental information.

Thus, China's approach can be summarized as:

mandatory efficiency + facility regulation + limited public transparency.

15. Japan

Japan follows a **hybrid model**.

Its Energy Conservation Act requires certain energy-intensive businesses, including qualifying data centres, to report:

- energy use
- efficiency measures
- improvement plans.

Japan has also established a Power Usage Effectiveness target of **1.4 by 2030** for data centres.

However, Japan does not have comprehensive AI model-level energy disclosure.

Its approach combines:

- mandatory facility-level requirements
- voluntary AI-sector cooperation
- renewable energy planning
- innovation-oriented policies.

16. Major Global Regulatory Gap

The comparative analysis produces an important conclusion.

Facility-level regulation dominates worldwide.

Governments generally regulate the physical data centre rather than the individual AI model.

This is inadequate because a data centre may simultaneously run:

- efficient traditional software
- ordinary AI
- energy-intensive reasoning models
- image-generation systems
- video-generation systems.

Aggregate facility statistics therefore cannot show which AI systems are responsible for resource consumption.

17. India's Position

The paper specifically discusses **India**.

India's Energy Conservation Act 2001 permits the Central Government, in consultation with the Bureau of Energy Efficiency, to classify certain entities as energy-intensive and establish efficiency requirements.

India also updated its **Renewable Consumption Obligation rules in September 2025**.

However, data-centre operators had not yet been specifically designated under those rules at the time examined by the paper.

Various Indian states have introduced policies encouraging renewable-energy procurement for data centres.

Nevertheless, the paper concludes that India had **not yet developed a unified policy response** to the environmental impact of AI and data-centre operations.

This is particularly relevant because India is becoming an increasingly important digital and AI market.

18. Brazil, UAE and Other Countries

The paper also examines several other jurisdictions.

Brazil

Brazil is attempting to become an AI and data-centre hub while taking advantage of its approximately **87% renewable electricity grid**.

Its policies increasingly connect data-centre incentives with:

- renewable electricity
- water efficiency
- environmental reporting.

UAE

The UAE has introduced mandatory emissions reporting for large emitters and voluntary sustainable digital-service guidelines.

Its rapid AI infrastructure development has nevertheless created concerns about increasing electricity demand.

The paper's comparative table shows that **only the EU has model-level transparency**, while other countries largely remain at facility or corporate levels.

19. Regulatory Arbitrage

Another major concern is **regulatory arbitrage**.

If one country establishes strict environmental disclosure requirements while another does not, AI companies may locate computing infrastructure in the less regulated jurisdiction.

The result could be:

1. Computing moves to jurisdictions with weaker transparency.
2. AI services continue to be supplied globally.
3. Environmental impacts remain difficult to measure.
4. Companies avoid stronger disclosure obligations.

Because AI infrastructure is globally distributed, national regulation alone cannot fully solve the problem.

20. First Policy Proposal: Model-Level Transparency

The authors' first major proposal is to move from:

Data-centre reporting → AI model reporting

AI providers should disclose:

- energy consumption
- water consumption
- carbon emissions
- inference consumption
- compute locations.

The EU AI Act could be amended to incorporate these requirements, particularly concerning General-Purpose AI models.

21. Inference Consumption Disclosure

Providers should disclose the environmental impact of **inference**, not just training.

This should include:

Energy

How much electricity is required?

Water

How much water is consumed?

Carbon

How much CO_2-equivalent emissions are associated with the computation?

This information would allow:

- consumers

- businesses
- researchers
- governments
- NGOs

to compare models and understand their environmental consequences.

22. Cumulative Environmental Reporting

The authors say that reporting only the energy required for one prompt is insufficient.

A model could have low energy consumption per query but millions or billions of queries.

Therefore, providers should also disclose **cumulative annual consumption**.

They consider two approaches.

Option 1: Granular Operational Metrics

Providers would report:

- total query volume
- median energy consumption
- mode
- standard deviation
- cache utilization.

This would allow researchers to understand variation between:

- simple questions
- complex reasoning
- image generation
- multimodal tasks.

Option 2: Aggregate Reporting

Providers would simply publish annual:

- electricity consumption
- carbon emissions
- water consumption.

For example, an AI provider could report the total annual environmental footprint of an entire model system.

This is easier to understand and compare but more difficult for researchers to independently verify.

23. Authors' Preferred Solution: Hybrid Reporting

The authors recommend combining both systems.

Providers should publish:

Headline figures

- annual electricity
- annual carbon emissions
- annual water consumption.

Supporting data

- query volume
- distributional energy data
- caching information.

This would combine:

simplicity + comparability + auditability.

24. Disclosure of Compute Locations

The authors also propose that providers disclose **where AI computation takes place**.

For example:

- country

- state
- computing region
- data centre region.

They should also disclose:

- whether local electricity grids are used
- whether dedicated energy sources are used
- the renewable/fossil composition of the electricity supply.

This matters because the environmental impact of one unit of electricity differs depending on whether it comes from:

- renewable energy
- nuclear energy
- natural gas
- coal.

Location disclosure would therefore help determine the actual carbon intensity of AI operations.

25. AI Energy Score

One of the most practical proposals is the introduction of an **AI Energy Score**.

The idea is similar to energy labels found on:

- refrigerators
- washing machines
- other appliances.

AI models could receive a simple:

1–5 star rating

based on their environmental efficiency.

This would allow ordinary users to understand complex environmental information without studying technical data.

26. Why Energy Labels Matter

The authors believe energy labels could influence consumer behaviour.

A user might see:

Model A – ★★★★★

versus

Model B – ★★

and choose the more efficient model where both perform adequately.

This could create a market incentive for AI companies to develop:

- smaller models
- more efficient architectures
- better inference systems
- lower-energy services.

Thus, regulation could influence not only disclosure but also **competition and innovation**.

27. Mandatory Energy Scores

The authors recommend that AI Energy Scores become mandatory.

They suggest that:

- the EU AI Office
- CEN
- CENELEC

could develop or endorse standardized AI energy-rating methodologies.

The rating should appear:

- on AI provider websites
- during registration
- when users select models

- on platforms hosting AI models.

Cloud providers should also disclose these ratings for business users.

28. Monetary Environmental Information

The authors also propose displaying standardized monetary information where possible.

For example:

- cost per token
- cost per unit of CO_2-equivalent.

This would give businesses and users another way to compare models.

The combination of:

energy score + monetary cost + environmental data

could encourage more rational technology choices.

29. Trade Secret Objection

AI companies may argue that environmental disclosures could reveal confidential information concerning:

- hardware
- software
- architecture
- infrastructure.

The authors reject this objection.

They argue that inference energy consumption depends on many variables, including:

- model architecture
- optimization
- quantization
- pruning
- batching
- scheduling

- caching
- cooling
- data-centre design.

Therefore, environmental consumption figures would not necessarily reveal enough information to reverse-engineer proprietary systems.

30. Standardized Benchmarks

Another major recommendation is the creation of **standardized AI environmental benchmarks**.

Without standardized testing, companies could:

- choose favourable test conditions
- use different methodologies
- report incomparable numbers.

The authors therefore propose independent benchmarking using standardized:

- hardware
- datasets
- prompts
- tasks
- measurement methods.

They specifically point to the **AI Energy Score** and open-source benchmarking approaches as potential foundations.

31. Independent Verification

Transparency should not depend solely on companies reporting their own numbers.

The authors propose giving qualified researchers access to relevant information so they can independently verify environmental claims.

They also propose stronger roles for:

- digital regulators
- researchers
- NGOs

- consumer organizations.

The EU AI Office, for example, could conduct:

- audits
- investigations
- compliance reviews.

Researchers could receive controlled access to information comparable to existing access rights under digital regulation.

32. Second Policy Pillar: Right to Green Digital Infrastructure

The second major policy pillar goes beyond transparency.

The authors argue:

Knowing the environmental cost of AI is not enough if users cannot choose whether or how they use AI.

They therefore propose new **user rights**.

33. Right to Search Without Generative AI

Users should have a clear and persistent option to disable AI-generated search results.

Instead of requiring users to:

- type special commands
- disable AI repeatedly
- search for hidden settings,

the interface should provide a simple permanent switch.

The choice should persist across sessions.

34. Right to Use Digital Services Without Unnecessary GenAI

The proposed right should extend beyond search.

Increasingly, generative AI is being inserted into:

- messaging applications
- office software
- writing tools
- coding tools
- social platforms
- other digital services.

The authors are concerned that users may be nudged toward AI through **dark patterns**, such as prominent AI buttons or interfaces.

They therefore propose a general right to use digital infrastructure **without unnecessary generative AI components**.

35. Three Reasons for This Right

The authors identify three reasons.

Environmental reason

Avoiding unnecessary AI can reduce energy consumption and carbon emissions.

Cognitive reason

Constantly outsourcing:

- writing
- thinking
- researching
- problem-solving

to AI could contribute to cognitive offloading and potentially weaken critical thinking.

Consumer-protection reason

Users often do not possess the same information as technology companies.

Therefore, the better-informed party should disclose relevant information and allow meaningful user choice.

36. The "Penalty Default" Concept

The authors draw on economic theory concerning **penalty defaults**.

The party with more information should bear a greater disclosure burden.

Applied to AI:

- AI companies know the environmental cost.
- Users generally do not.
- Therefore, companies should disclose information before users are pushed toward AI use.

The authors suggest that AI should not automatically be the default where a traditional, less energy-intensive alternative is available.

37. Third Policy Proposal: Right to a Green Model

The paper does not argue that users should always avoid AI.

Instead, users who want AI should have the right to select an **environmentally optimized model**.

A "green model" would be the most energy-efficient model within the provider's available models that can perform the required task.

For example, if several models can answer a simple question:

- Model A uses very little energy.
- Model B uses significantly more.
- Model C is a large reasoning model.

The user should be able to choose Model A as the default green option.

38. Reasoning Should Be "Off" by Default

The paper makes a particularly strong recommendation concerning reasoning models.

Because reasoning can dramatically increase energy consumption, platforms should:

Default:

Reasoning OFF

User action:

User actively turns reasoning ON.

Interface:

The system should clearly indicate that reasoning requires additional computing resources.

This would preserve user choice while discouraging unnecessary high-intensity computation.

39. International AI Environmental Governance

The third major pillar is **international cooperation**.

AI's environmental impact crosses national borders.

For example:

- a user may be in India
- the AI company may be based in the US
- computation may take place in Europe
- electricity may come from a regional grid
- water consumption may occur in a drought-prone area.

Therefore, national regulation alone cannot adequately govern AI's environmental footprint.

The authors argue for a **Global AI Compact** incorporating common:

- disclosure standards
- environmental rights
- measurement methodologies.

40. Role of the United Nations

The authors identify the UN as a potential coordinating institution.

Possible participants include:

ITU

For technical and digital governance expertise.

UNEP

For environmental expertise.

UNFCCC

For climate governance.

A global framework could establish minimum standards applicable regardless of where AI infrastructure is located.

41. Bottom-Up Governance

The authors recognize that an international treaty may be difficult in the current geopolitical environment.

Therefore, they also recommend a **bottom-up approach**.

Participants could include:

- universities
- researchers
- AI companies
- industry associations
- NGOs
- civil society.

They could develop common:

- measurement methods
- benchmarks

- standards
- reporting protocols.

These standards could eventually become the foundation for formal international law.

42. Overall Comparative Picture

The paper's comparative analysis can be summarized as follows:

Jurisdiction	Main Approach	Model-Level Transparency	Environmental Direction
EU	Mandatory disclosure	Yes, but limited	Relatively ambitious
USA	Deregulation/general environmental law	No	Declining ambition
China	State-directed efficiency	No	Increasing ambition
Japan	Hybrid	No	Moderate
UK	Hands-off/facilitation	No	Lower ambition
Canada	Hands-off	No	Limited
India	Developing framework	No	Not yet unified
Brazil	Facility regulation + incentives	No	Relatively ambitious
UAE	Facility reporting + voluntary guidance	No	Moderate
Singapore	Facility efficiency + guidance	No	Moderate

The most important finding is that **model-level environmental transparency remains extremely rare**.

43. Core Legal Problem Identified by the Authors

The fundamental problem can be expressed as:

Environmental costs are increasing faster than legal transparency.

AI models are becoming:

- larger
- more capable
- more widely used
- more computationally intensive.

But legal systems largely continue to ask:

"How much energy does the data centre consume?"

rather than:

"How much energy does this particular AI model consume when answering this particular type of query?"

The authors believe this distinction is crucial.

44. Why Facility-Level Regulation is Insufficient

Imagine a data centre contains:

- Model A: low-energy classification
- Model B: normal language model
- Model C: image generator
- Model D: reasoning model
- Model E: video generator.

A facility-level report might simply state:

Total annual electricity = X GWh.

That figure does not reveal which model caused the consumption.

Consequently:

- regulators cannot effectively target high-impact models
- businesses cannot choose efficient models
- consumers cannot make informed choices

- researchers cannot accurately calculate AI's environmental footprint.

This is why the authors advocate model-level reporting.

45. Main Conclusion of the Paper

The paper concludes with three principal findings.

Finding 1: New AI modalities dramatically increase energy consumption

Reasoning models and generative AI search can increase energy consumption substantially.

Reasoning models may consume **150–700 times more energy** than their base models.

Finding 2: Transparency is inadequate

Paradoxically, environmental pressure from AI is increasing while legally required transparency remains limited.

The EU is the only jurisdiction identified as having specific model-level disclosure obligations, but these remain incomplete and largely inaccessible to the public.

Finding 3: Transparency must be combined with user rights and international cooperation

The authors therefore advocate:

- model-level disclosure
- inference reporting
- compute-location reporting
- standardized energy scores
- independent benchmarking
- right to avoid unnecessary GenAI
- right to choose green models
- reasoning-off defaults
- international coordination.

46. The Paper's Central Message in Simple Words

The paper essentially says:

AI should not be treated as environmentally neutral simply because its environmental costs are hidden inside data centres.

If AI is becoming a major consumer of electricity and water, society needs to know:

How much does each AI model consume?

Where is the computation taking place?

What is the carbon intensity of the electricity?

How much water is being used?

How does one model compare with another?

And, importantly:

Users should have the right to choose.

They should be able to say:

- "I don't want AI for this task."
- "I want the traditional version."
- "I want the environmentally efficient model."
- "I don't need reasoning for this simple question."

47. Key Recommendations at a Glance

A. Mandatory Transparency

AI companies should disclose:

- energy
- water
- carbon
- inference consumption
- query volumes
- compute locations.

B. Standardization

Governments should establish common:

- measurement protocols
- benchmarks
- environmental metrics.

C. AI Energy Score

Models should receive standardized **1–5 star environmental ratings**.

D. Right to Opt Out

Users should be able to permanently disable unnecessary AI features.

E. Right to a Green Model

Users should be able to select the most environmentally efficient model.

F. Reasoning-Off Default

High-energy reasoning should require affirmative user consent.

G. Independent Auditing

Researchers and regulators should have access to information necessary to verify corporate claims.

H. International Framework

UN bodies and governments should work toward common global standards.

48. Final Assessment

This paper makes an important shift in the discussion surrounding AI regulation.

Traditional AI governance generally asks questions such as:

- Is AI safe?
- Is AI fair?
- Does AI protect privacy?

- Is AI transparent?
- Is AI discriminatory?

The authors add another fundamental question:

Is AI environmentally sustainable?

Their answer is that current governance is **not sufficient**.

The paper's most significant conceptual contribution is the idea that environmental sustainability should become not merely an obligation imposed on AI companies but also a **consumer right**.

The proposed "**Right to Green AI**" combines environmental regulation with consumer protection and digital rights.

In this framework, a sustainable AI ecosystem would have:

Transparency → Measurement → Comparison → User Choice → Efficient Models → International Accountability

The authors ultimately argue that AI innovation and environmental sustainability do not have to be opposing goals. Instead, **proper regulation can make environmental efficiency itself a dimension of AI competition and innovation**.

-XXXX-

Mudita Learning and Knowledge Private Limited
Email: info@muditabooks.com
Date of post: 28.08.2026