Augmented RenAIssance Manifesto (v.1.2)

Eng. PhD. Luigi Lella (https://www.linkedin.com/in/luigi-lella/), Dr. Chiara Lella (chiaralella00@gmail.com)

The Vision of "Augmented RenAIssance"

The Renaissance placed human reason, dignity, and responsibility at the center of progress. The Augmented RenAIssance reaffirms these values in the age of Artificial Intelligence. Artificial Intelligence is neither a sovereign nor an oracle.
It is a tool for amplification: of knowledge, of care, of administrative capacity, of collective intelligence. In the Augmented RenAIssance:

- Humans decide, institutions remain accountable.
- AI supports, informs, and amplifies human reason and responsible creativity.
- Technology serves the dignity and the freedom of the person, not replace it.
- Innovation is oriented towards the common good.
- Transparency and comprehensibility are fundamental democratic rights.

Foundational Principles (Non-Negotiable)

2.1 Human Dignity and the Centrality of the Person (*Dignitas Hominis*)
No AI system can replace human responsibility. Decisions with legal, social, or economic impact must remain under human guidance:

- Every high-impact algorithmic decision must be supervised, understandable, and overridden by a competent human.
- Citizens have the right to know when they are interacting with AI system, to understand its capabilities and limitations, and to request the intervention of a qualified human interlocutor at any time.
- Human dignity prevails over any consideration of computational efficiency.

2.2 Accountability and Justice (Iustitia et Responsabilitas)
The use of AI must respect the primacy of justice and responsibility for the common good:

- Every AI system must have an identifiable person responsible for its outcomes.
- AI-assisted decisions must be publicly justifiable according to criteria of fairness and proportionality, ensuring the algorithmic non-determinism on humans.
- Procedural justice (how the decision is reached) is as important as substantive justice (the content of the decision), considering the epistemic uncertainty of AI.
- Accountability requires organisational and infrastructural control as its precondition, including the ability to audit, reconfigure, substitute components and decommission.
- Outsourcing infrastructure does not mean outsourcing responsibility.
- Reparation when AI causes unjust harm.

2.3 Regulatory Compliance as an Ethical Imperative (Lex et Fides Publica)
All uses of AI must comply with all related regulations:

- Compliance with the EU AI Act, GDPR, NIS2, Data Governance Act, Digital Services Act, Cyber Resilience Act, Open Data and ePrivacy Directives, Fundamental Rights Charter.
- Compliance is a by-design requirement, not an ex-post constraint.
- Regulations are seen as protections of human dignity, not as obstacles to innovation.
- Violation of regulations is a violation of ethics, regardless of the legal consequences.

2.4 *Transparency and Eloquence of the Algorithm* (*Claritas et Eloquentia*)
Renaissance humanists sought to create citizens capable of speaking and writing with eloquence and clarity, capable of participating in civic life. In the Augmented RenAIssance, algorithms must also "speak" comprehensibly:

- AI outputs must be understandable by their users, not just by technicians.
- Traceability of data, prompts, and outputs is mandatory.
- Technical language must be translated accessibly to decision makers and citizens.

2.5 *Continuing Education and Specialization* (*Educatio Perpetua*)
Humanists saw education as essential for moral and intellectual development. Augmented RenAIssance requires ongoing training for all stakeholders:

- Lifelong learning programs for decision makers, technicians, and citizens.
- Sharing lessons learned, including failures.
- Creating communities of practice that develop collective wisdom in the use of AI.
- Recognition that technology evolves in time and requires continuous updating of technical, ethical, civic and critical skills and values.

2.6 Critical Reasoning and Practical Wisdom (Ratio et Prudentia)
The use of AI must be guided by critical reasoning and practical wisdom, not by uncritical technological enthusiasm or irresponsible delegation:

- Every AI implementation must be preceded by a critical analysis of necessity, proportionality, risks, and alternatives. Not every problem requires an AI solution.
- Decision makers must understand the blind spots of AI: what it can and cannot know.
- Consultation of various experts (technical, legal, ethical) before important decisions.

2.7 *Orientation to Common Good and Civic Virtue* (*Bonum Commune et Virtus Civica*)
AI should be used only where it can be concretely demonstrated to improve public or organizational value, serving the common good:

- No speculative or excessive automation: AI should be used only where it adds measurable value to public service, evaluating it with interdisciplinary methodologies.
- Success metrics include not only efficiency but also equity, accessibility, digital inclusion and citizen trust (including through their involvement in the project).
- Improving citizen service takes precedence over technological innovation.

2.8 Data Quality and Truth (*Veritas et Integritas Fontium*)
Artificial intelligence systems can amplify human capabilities only if they are based on high-quality, accurate, representative, and reliable data:

- Mandatory Data Quality Assessment: Before using datasets for training or inference, a structured assessment of data completeness, accuracy, representativeness, timeliness, and provenance must be conducted. Datasets must be periodically reviewed.
- Data lineage and traceability: It must be possible to trace the data from the original collection to the final output, identifying every intermediate transformation.
- Principle of Representativeness: Training data must fairly represent all populations served, avoiding underrepresentation of vulnerable or minority groups.
- Source verification for RAG systems: In Retrieval-Augmented Generation systems, documentary sources must be validated, authoritative, and up-to-date.
- Rejection of "toxic" data: Datasets containing discriminatory information, harmful stereotypes, misinformation, or privacy violations must be excluded.
- Human-in-the-loop validation: Significant data samples must be validated by human experts before use, particularly in critical domains (health, justice, welfare).
- Preserve data assets: Carefully consider whether to share the organization's data (e.g. population's data in open format).
- Public ownership: Data generated or used by AI systems in the organization remain public assets. Contracts may not transfer secondary use rights to the supplier without explicit authorization.
- Portability and return: Each contract must provide for the return of used and generated data in open formats at the end of the relationship.
- Location: Sensitive data must reside on resilient infrastructures and systems under the control of the public administration (on-premises or EU private cloud). The architecture (cloud/hybrid/on-prem) must be proportionate to the data classification.
- Data cybersecurity: encryption at rest and in transit as minimum requirement.

2.9 Sustainability and Computational Temperance (Temperantia et Custodia Naturae)
The use of artificial intelligence must be environmentally and energy sustainable:

- Environmental Impact Assessment for AI: Every energy-intensive AI system (large language models, intensive training) must include an environmental impact assessment in terms of energy consumption, CO2 emissions, and water use for data center cooling.
- Principle of Energy Proportionality: The computational power employed must be proportionate to the public value generated. Next-generation generative models should not be used when lighter models or traditional methods adequately solve the problem. Balance between innovation/risk, efficiency/rights, automation/control.
- Preference for efficient architectures: e.g., edge computing when possible to reduce data transfers; quantized or compressed models when appropriate; Retrieval-Augmented Generation (RAG) instead of full LLM fine-tuning; Intelligent caching to avoid repetitive computations.
- Use of renewable energy: Data centers hosting AI systems must prioritize using energy from renewable sources. In cloud supply contracts, preference should be given to providers with certified carbon-neutral policies.
- Lifecycle thinking: Consider the environmental impact of the entire lifecycle of the AI system, from the training phase (which can require weeks of GPU time) to the continuous inference phase.
- Batch processing vs. real-time: When possible, prefer nighttime batch processing (when the power grid has surplus renewable energy) over continuous real-time processing.
- Responsible decommissioning: When an AI system is decommissioned, document the lessons learned to prevent other organizations from repeating the same computationally expensive experiments.

2.10 *Fortitude and Courage in Responsible Innovation (Fortitudo)*
The digital transformation of the public sector requires courage:

- Be early adopters when innovation genuinely serves the public good.
- Create regulatory sandboxes where innovative solutions can be safely tested.
- Publish honest reports on the successes and failures of AI projects.
- Document lessons learned even when they reflect errors of judgment.
- Resist political or economic pressure to implement AI systems that would violate fundamental rights.
- Say "no" to vendors when their solutions do not meet the principles of the manifesto.
- Do not be paralyzed by technophobia or excessive fear of change.

Continuously updated open licensed tools and materials

1. Ensuring AI awareness and AI literacy
2. Involving citizens in choosing use cases
3. Support the drafting of public-private partnership projects
4. Use Cases Repositories

License: Creative Commons Attribution 4.0 International (CC-BY 4.0) You are free to share and adapt this manifesto, provided appropriate credit is given.