



The National Council
for Artificial Intelligence



Egyptian Charter for Responsible AI

2023 v1.0

Introduction and Background

As the use of Artificial Intelligence (AI) systems becomes more prevalent, the need has arisen for all stakeholders to become more aware of their potential risks and limitations. Despite their undeniable benefits, AI systems, if designed, deployed, or used incorrectly, can pose significant risks. Among those are: biased or wrong results, data drift, lack of transparency, lack of legal responsibility, lack of fairness and equality, to name but a few.

Governments, international organizations, as well as large corporations are therefore waking up to the necessity of governing AI projects properly to ensure those risks are mitigated or minimized. As part of its efforts to build an AI industry, which started over two years with the launch of the National AI Strategy, Egypt has taken a leading role in drafting several of the ethics guidelines around AI in different international organizations such as the OECD, UNESCO, G20, and the Group of Governmental Experts (GGE) on Lethal Autonomous Weapons Systems (LAWS), in addition to other international organizations and initiatives. In addition, leading teams within the African Union and the League of Arab States (LAS) work to unify ethical recommendations for AI on the regional level, to ensure that the priorities, needs, and special circumstances of our societies are considered.

Egypt's efforts were met with worldwide recognition as it became the first Arab and African country to adhere to the OECD AI Principles, and an early adopter of the UNESCO Recommendation on the Ethics of Artificial Intelligence. As most of these recommendations are non-binding and quite generic, it is incumbent upon the different countries to develop their own local interpretations of these guidelines and translate them into actionable insights and policies for decision makers in government, academia, industry, and civil society.

Document Scope

This document serves as the first attempt at articulating Egypt’s interpretation of the various guidelines on ethical and responsible AI, adapted to the local context and combined with actionable insights to help ensure the responsible development, deployment, management, and use of AI systems in the country. It draws upon the guidelines developed by the OECD, UNESCO, WHO, IEEE, EU, as well as by many leading countries such as Singapore, UK, US, Australia, and others. Although most international guidelines are still non-binding, it is expected that a set of standards would soon be issued by international organizations for AI systems worldwide (the EU AI Act being the first such standard).

In that regard, this document serves two purposes:

1. To be a “soft launch” to empower citizens to expect and demand the best from the use of AI and for all stakeholders to be aware of ethical considerations related to AI and incorporate those considerations into their AI adoption plans.
2. To signal Egypt’s readiness to follow responsible AI practices, something many investors as well as AI ranking bodies measure when assessing a country’s readiness for AI investment and adoption. It would also help to communicate Egypt’s needs and priorities to foreign AI developers looking to develop or market their products in the country.

This document shall be reviewed on an annual basis to ensure its continuous currency and relevance. A public consultation shall also take place prior to each revision to ensure that the perspectives of all stakeholders are considered.

The document is divided into two parts:

- **General Guidelines**, which are overarching rules applicable to all members of the ecosystem,
- **Implementation Guidelines**, which are technical considerations, mainly applicable to any entity developing, deploying, or managing an AI system.

Each guideline is tagged with the most relevant key principle of responsible AI.

General Guidelines

1. The primary goal of using AI in Government is the well-being of citizens, including combating poverty, hunger, inequality, illiteracy, and corruption; achieving prosperity and inclusion; increasing fairness and transparency; augmenting human capabilities; protecting the environment; invigorating economic growth and opening new markets and job opportunities for Egyptians. **[Human-Centeredness]**
2. Any end-user using an AI system has the fundamental right to know when he or she is interacting with an AI system and not a human being, for example in the case of automated call centers. **[Transparency and Explainability]**
3. No individual should be harmed by the introduction of an AI system. Special considerations must be taken to protect vulnerable and marginalized groups such as children, PWDs, and those with lower economic or educational level. Sample considerations include checking potential data bias, tuning system parameters periodically, and preferring development team diversity. **[Fairness]**
4. Appropriate mechanisms should be in place to allow anyone adversely affected by an AI system to challenge its outcome based on plain and easy-to-understand information on the factors, and the logic that served as the basis for the prediction, recommendation, or decision. **[Fairness]**
5. Documented policies and processes should be in place to quickly respond to and resolve any adverse outcomes caused by the unauthorized use of AI systems. **[Fairness]**
6. AI systems should not be designed to primarily replace human labor except in cases that pose danger or risk to human wellbeing. If job losses are inevitable as a side effect of an otherwise beneficial AI system, measures should be taken by the system owner (Government, private sector, or other) to ensure a fair transition for workers as AI is deployed, such as through training programs along the working life, support for those affected by displacement, and access to new opportunities in the labor market. **[Human-Centeredness]**
7. All stages of the lifecycle of the AI system, including data collection, hosting, and engineering, and system development, testing, deployment, continuous operation, monitoring, and maintenance, are subject to the relevant laws of the Arab Republic of Egypt, including laws of consumer protection, personal data protection, and anti-cybercrimes. **[Accountability]**

8. Certification mechanisms for AI systems or similar forms of regulation are introduced by the appropriate regulatory bodies in the different domains to ensure the safety, transparency, robustness, and reliability of AI systems based on each domain's requirements. **[Accountability]**

9. International efforts should be pursued continuously to develop guidelines for the responsible use of AI in military applications. **[Human-Centeredness]**

10. Ultimate responsibility and accountability for the behavior and outcomes of an AI system must always lie with natural or legal persons. AI systems should not be given legal personality themselves. To ensure this, any regulatory framework should be consistent with the principle of human oversight and establish a comprehensive approach focused on the actors and the technological processes involved across the different stages of the AI system's lifecycle. **[Accountability]**

11. Final human determination is always in place. This means that ultimately, humans are in charge of making decisions, and are able to modify, stop, or retire the AI system if necessary. Individuals with that power must be decided upon by the owner of the system. **[Security and Safety]**

12. All members of the AI ecosystem, especially academic and educational institutions, should promote capacity building and public awareness programs about AI development, including various AI technologies such as supervised, unsupervised, and reinforcement machine learning, and the opportunities and challenges brought about by those technologies. Those programs should encourage multi-disciplinary collaboration and should be accessible to technical and non-technical groups alike. **[Transparency and Explainability]**

13. AI systems that support entrepreneurship through innovative start-ups and MSMEs should be encouraged and made a priority in order to achieve economic prosperity and society welfare. **[Human-Centeredness]**

Implementation Guidelines

1. AI systems should be robust, secure, and safe throughout their entire lifecycle so that, in conditions of normal use, foreseeable use, misuse, reward hacking, or other adverse conditions, they function appropriately and do not pose unreasonable safety risk. **[Security and Safety]**
2. Ideally, any AI project should be preceded by a pilot or proof of concept (PoC) to ensure the technical viability of the solution. Specific success criteria should be set and only if those are met can the pilot be deemed successful and ready for large-scale implementation. **[Accountability]**
3. Additional measures should be in place in case of sensitive or mission-critical AI applications, including additional measures to ensure data protection, beneficiary engagement, and avoidance of any harm resulting from applications. **[Security and Safety]**
4. AI projects that go into production must be developed by qualified entities with proven experience in product-grade AI solution development. The teams should be diverse enough to include system architects, AIOps and QA engineers, cybersecurity experts, software engineers (non-AI engineers that develop the application or platform hosting the AAI models), data scientists, AI engineers (specialty will depend on the nature of the project), at least one domain expert, and one project manager. **[Accountability]**
5. Domain experts are a crucial part of any AI team. They are the professionals who understand the business problem and can guide the team in terms of data availability and quality, as well as validating the relevance of the results to the problem at hand. **[Accountability]**
6. Government entities, private companies, academic and research organizations, and any other entities developing AI systems should work with a representative sample of the beneficiaries of their AI systems. **[Fairness]**
7. Developers of AI systems must adopt a systematic risk management approach as part of the system development lifecycle, which augments and complements the usual software development lifecycle (SDLC) to include risks specific to AI systems such as privacy, digital security, safety, and bias. **[Security and Safety]**

8. Developers of AI systems should always strive to provide transparent and explainable AI solutions. The degree of explainability required will vary according to the application domain and project requirements, but project sponsors must be clear on the potential tradeoff between the accuracy/quality and explainability of any given model. When in doubt, developers should opt for simpler models with higher levels of explainability, without compromising the minimum desired quality and accuracy. **[Transparency and Explainability]**
9. Developers of AI systems are encouraged to examine and address the cultural impact of AI systems, especially Natural Language Processing applications such as automated translation and voice assistants impacted by the nuances of human language and expression. Such addressing should provide input for the design and implementation of strategies that maximize the benefits from these systems by bridging cultural gaps and increasing human understanding, as well as minimizing negative implications such as the reduction of use, which could lead to the disappearance of endangered languages, local dialects, and tonal and cultural variations associated with human language and expression. **[Fairness]**
10. All members of the AI ecosystem, including government agencies, academic and educational institutions, and private sector companies, should facilitate access by the scientific community to their data for research purposes, provided that such access does not come at the expense of privacy. **[Accountability]**
11. The use of any data must be pre-authorized by the data owner except for data available in the public domain. Personally identifiable data must be anonymized and/or encrypted depending on the domain, and express written consent from the data owner must be obtained according to applicable laws. Data inputs should be comprehensive, and as much as possible, disaggregated, with corrections of distortions like invisibility of minorities. **[Security and Safety]**
12. AI systems, especially data-driven models, must be monitored regularly during development to detect and address data drift. In those cases, the quality of the data must be reviewed and if needed, the underlying models need to be changed to accommodate changes in data. **[Fairness]**
13. Foreign companies looking to roll out their AI products in Egypt must adhere to these guidelines, and must also ensure that their models have been trained using local data, relevant to the Egyptian market and availed through law-abiding mechanisms, and that they adhere to local customs and religious and social traditions and norms. Proper testing of these systems must be performed to ensure their quality and accuracy before they are introduced to the Egyptian market. **[Fairness]**

14. All government AI projects must be preceded by a thorough impact assessment to ensure maximum benefit from the technology, while respecting the guidelines of responsible and ethical AI development. Specifically, the following questions should be asked:

- a. What is the problem to be solved, and is AI the best way to solve it or are there other ways that could be cheaper, faster, or more reliable?
- b. Is the data required for the project ready and of sufficient volume and quality to ensure the desired output?
- c. Are the underlying processes properly engineered? AI is not a solution for broken processes but a technique to optimize certain variables. If the underlying process is broken or inefficient, this problem will only be amplified by the use of an AI system.
- d. What is the financial impact of the solution, both direct (project cost) and indirect, including potential loss of jobs?
- e. What is the social impact, if any?
- f. What is the environmental impact, if any?
- g. Is the data available diverse enough to cover all potential use cases of the solution, in order to minimize bias? For example, in the case of healthcare solutions, is data available from different ethnicities, genders, age groups, and medical conditions, in addition to any other factors that might impact the outcome?

All of the above points must be weighed against the impact and expected result of implementing the solution using non-AI technologies. Only if the benefits (including positive impacts) outweigh the costs (including any negative impact), can the project be approved. **[Accountability]**

15. Government AI projects should be implemented using components from the National AI Platform once completed. Until then, any project should be implemented in a modular, service-oriented way, while using open source and white box/non-proprietary technologies to ensure transparency and maintainability. **[Accountability]**

16. Government AI projects, similar to digital transformation projects, should be commissioned and supervised by the Ministry of Communications and Information Technology (MCIT) in order to ensure compliance with these guidelines and the credibility and quality of data and developers involved in the development of AI systems. MCIT presents periodic updates on the status of those projects to the National Council for Artificial Intelligence (NCAI). **[Accountability]**