Artificial Intelligence: Principles & Definitions

1.0 Overview

Artificial intelligence (AI), while a newly trending topic in the public discourse, has been at work helping New York City provide public services for many years. AI technologies offer a wide range of opportunities to make government run better for New Yorkers. Prudent use of AI can improve operational efficiency, social equity, environmental sustainability, and more. At the same time, use of AI tools can pose a range of risks for individuals and communities, whether due to lack of appropriate governance, misuse, flawed design, or other factors. Use of AI without appropriate oversight and governance poses a range of potential risks, from benign inaccuracies or the misallocation of government resources, to cybersecurity risks such as new avenues for cyber-attacks or data spills. On a broader environmental and social scale, AI usage can consume large amounts of energy, generate excess carbon emissions, create electronic waste, and perpetuate discriminatory biases and social inequality. Furthermore, the complexity and indiscernible mechanics of many AI applications pose unique transparency and accountability challenges, which are particularly pronounced for governments working in service to the public. Addressing these varied risks will require a collaborative approach from all sectors of both government and industry.

2.0 Purpose

New York City agencies have been exploring emerging technologies and have developed strategies and procedures for building, procuring, implementing, and monitoring AI solutions. As AI technologies mature and become more accessible to agency users, a need for greater coordination and alignment of AI management practices citywide has become clear. The purpose of the guidelines herein is to establish core principles to support agency use of AI, and to create accessible, standardized definitions for commonly used terms that may be adapted citywide. This guideline will be updated as needed to reflect the changing technology ecosystem.

3.0 Authority

The Office of Technology and Innovation ("OTI") was formed under Mayoral Executive Order 3 of 2022 ("EO 3") in order to unify technology teams across government and centralize coordination around existing and emerging technologies. OTI serves as the city's central technology agency, leading "the development, coordination and implementation of the city's information technology, information security, information privacy and telecommunication matters." This guideline is issued under EO 3.

4.0 AI Principles

The New York City AI Action Plan of 2023 commits the city to designing and implementing a robust governance framework that aligns the city's use of AI with a defined set of principles. While these principles have been developed with reference to other efforts to codify AI principles, notably the White House's Blueprint for an AI Bill of Rights,[1] the New York City AI Principles ("AI Principles") reflect the particular needs and priorities of New York City, its agencies, and its residents. OTI encourages agencies to consider these principles as they develop and use AI solutions. Subsequent guidance and policy

---

[1] White House Office of Science and Technology Policy (OSTP), Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People, 2022, https://www.whitehouse.gov/ostp/ai-bill-of-rights/.

published by OTI will incorporate these principles to help ensure that they are integrated into all facets of AI governance.

4.1 Validity and Reliability

AI has great potential to address many problems faced by the city. However, AI tools do not offer one-size-fits-all solutions, and, like traditional technology tools, should be used, developed, and implemented in ways that are responsive to existing problems. Performance of AI tools can also depend on their design or maturity, and that performance may change over time, based on factors such as modifications to models, data drift, or context of use.

The city should therefore assess whether AI solutions are valid for their tasks and that they rely on appropriate data, metrics, and assumptions to solve the problem at hand. The city should also carefully monitor AI tools—from ideation to deployment—to ensure that those tools are performing reliably, especially as conditions change over time or tools are adapted to address new use cases.

4.2 Social Responsibility

AI tools and systems have varying types and degrees of impact on the public. AI can be used to help automate simple administrative tasks or to inform policy decisions that have a material public impact. In some cases AI tools can support social responsibility efforts: for example, they can act more consistently than humans or make materials more accessible at scale. But absent appropriate controls, AI systems may also perpetuate or exacerbate bias due to the greater scale of deployment, which may amplify historical biases in the training data. Without appropriate consideration and governance, decision-making that is assisted or automated by AI can result in unfair or inequitable outcomes for individuals or groups. Moreover, the development and use of AI tools can have serious resource impacts, both financial and environmental, that require ongoing consideration.

When using, developing, and implementing AI solutions, city agencies should seek to ensure that tools operate with fairness and equity. There are a variety of methods for addressing these principles at various points in the AI lifecycle. City agencies should consider examining training data for pre-existing biases, conducting robust testing of systems in a variety of settings including through public engagement, performing appropriate risk assessments prior to deployment, continuously monitoring deployed tools, integrating a human-in-the-loop for AI decision making, incorporating human alternatives/fallbacks, and/or implementing other risk mitigation tactics. The city should also consider how to evaluate the impact of AI usage on the environment, in line with its sustainability goals. OTI intends to support agencies in establishing and carrying out best practices for such methods.

4.3 Information Privacy

Data is foundational in the development and use of AI, from training models to achieve desired objectives, to acting as inputs to generate new predictions, estimates, or other content. New York City recognizes, accordingly, that use of AI solutions may require agency collection, use, or

disclosure of identifying information, which is governed by New York City's Identifying Information Law.[2]

Where AI development and use involves collection, use, or disclosure of identifying information, New York City agencies must continue to comply with the Identifying Information Law and its associated citywide privacy policies.[3] This includes consideration of New York City's privacy principles,[4] which should be honored in all aspects of agency decision-making and operations impacting information privacy.

In order to promote confidence in the city's adoption of AI and its lawful and responsible data stewardship, and protect New Yorkers from potential unlawful and unethical uses of their data, the city must commit to extending existing privacy protections in agency AI development and use, and modify those protections as necessary to keep pace with technological advances.

4.4 Cybersecurity

The city can collectively harness the benefits of AI through secure enablement, with a continued mission to make New York City the most resilient city in the world. Security is central to the trustworthiness of AI technologies and its integration into the city's digital infrastructure. As technology continues to evolve, so will the cyber threat landscape. This demands further development of protections to thwart both traditional cyber threats, and those unique to AI such as novel machine-learning attacks. Accordingly, it is important that the city builds and leverages AI systems with security and resilience in mind, which includes the foundational elements of secure by design, and thoughtful investment in the expansion of citywide cybersecurity defenses.

The city is committed to advancing cybersecurity strategies that align with the global rise of AI, and bolstering preparedness for both current and future threats. This includes using public-private partnerships to strengthen secure-by-default practices, ensuring solutions are in alignment with citywide cybersecurity policies, and updating those policies to be reflective of the shifting technology ecosystem.

4.5 Trust and Transparency

New Yorkers reasonably expect to know how their government is working for them, and to understand the tools that agencies use to serve the public. It can be challenging to engender trust and understanding about public sector technology generally, as tools may be complex or require technical subject matter expertise to understand. These challenges are often greater with AI due to the complexity or opacity of the tools themselves, the varied impacts of their use on individuals and communities, broader discomfort with using technological tools in contexts

---

[2] Administrative Code of the City of New York, §§ 23-1201 to 23-1205.

[3] Chief Privacy Officer of the City of New York, Citywide Privacy Protection Policies and Protocols, 2023, https://www.nyc.gov/assets/oti/downloads/pdf/citywide-privacy-protection-policies-and-protocols.pdf.

[4] *Id.*; Transparency; public trust; accountability; data minimization; use limitation; responsible governance and stewardship; data quality, integrity, and accuracy; security safeguards; and equity.

and processes historically managed by humans, potential negative past experience with the use of such tools, or other factors.

In order to encourage greater public trust in the city's use of AI, agencies should carefully consider whether public engagement, disclosure or explanation of AI use, or human alternatives/fallback mechanisms may be beneficial. To this end, it is critical to understand how a given system may impact residents' safety, rights, or access to services, which may require interdisciplinary collaboration or direct engagement with impacted communities. Additionally, the city will continue to report on the use of "algorithmic tools" (including many AI tools) with a public impact on decision-making.[5]

5.0 Definitions

5.1 Objective and Use

The foundation of a strong policy framework is the common understanding of frequently used terms. Thus, the standardization of AI terminology will enable better support for agency implementation of AI projects and improved public accountability. OTI strongly encourages all agencies to integrate these definitions into their own policies or business governance practices. These terms will be updated or expanded periodically to reflect technological advances.

5.2 Approach

Definitions related to AI vary in terms of their historical use, the presence of widespread agreement across other organizations, and the precise boundaries of concepts that they represent. To acknowledge this variability, and the fact that other organizations including governments have begun efforts to standardize terminology, OTI has developed its definitions by referencing existing sources where available. Section 5.3 provides the list of defined terms, the definition for each, and additional notes or references to further explain the development of the definition.

5.3 Defined Terms

Term	Definition
Algorithm	A set of instructions that can be programmed and followed by a computer. Computer algorithms can range in complexity, from sending an email if a condition is true, to counting pedestrians in images or videos.[6]
Algorithmic Tool	Any technology or computerized process that is derived from machine learning, artificial intelligence, predictive analytics, or other similar methods of data analysis, that is used to make or assist in making decisions about and implementing policies that materially impact the rights, liberties, benefits, safety or interests of the public, including their access to available city services and

---

[5] Local Law 35 of 2022, Administrative Code of the City of New York, § 3-119.5, https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=4265421&GUID=FBA29B34-9266-4B52-B438-%20A772D81B1CB5.

[6] National Institute of Standards and Technology (NIST), *The Language of Trustworthy AI: An In-Depth Glossary of Terms,* March 29, 2023, https://doi.org/10.6028/NIST.AI.100-3.

	programming as they build a model based on training data to complete a task with minimal human intervention.[15]
Materially Impact	To have, through a system's outputs or outcomes, a discrete, discernible, or otherwise identifiable effect that: • Limits or suspends rights or liberties; OR • Determines the eligibility for, the type of, or quantity or magnitude of benefits, city services, and/or resources; OR • Changes the risk of harm to a person or group of people.[16]
Natural Language Processing	The ability of a machine to process, analyze, and mimic human language, either spoken or written.[17]
Nominal Data Type	A Data Type that describes two or more categories which cannot be compared to determine which one is greater, or whose ordering is not relevant to the use case.[18]
Numerical Data Type	A Data Type that can be manipulated with arithmetic operations.[19]
Optimization	The computation of a solution to maximizing or minimizing an objective function, often subject to constraints on the variables.[20]
Ordinal Data Type	A Data Type that describes two or more categories, such that every pair of categorical labels can be compared to determine which one is greater.[21]
Predictive Model	A model used for forecasting outcomes based on anticipated future values of input variables.[22]
Robotic Process Automation	A preconfigured software instance that uses business rules and predefined activity choreography to complete the autonomous execution of a combination of processes, activities, transactions, and tasks in one or more unrelated software systems to deliver a result or service with human exception management.[23]
Sequence	A Data Type describing multiple elements, whose number may change each time a function is executed, and whose order of appearance is relevant.[24]
Tree	A Data Type describing multiple elements where some elements, but not all, can be ordered.[25]

---

[15] NIST supra note 6.

[16] Definition developed by OTI to support agency compliance with Local Law 35 of 2022.

[17] NIST supra note 6.

[18] Alan Agresti, *Analysis of Ordinal Categorical Data*, 2010.

[19] *Id*.

[20] Phillip E Gill, Walter Murray and Margaret H Wright, *Practical Optimization*, 1997.

[21] Alan Agresti supra note 18.

[22] Leo Breiman, *Statistical Modeling: The Two Cultures, Statistical Science* 16(3): 199-231, 2001.

[23] NIST supra note 6.

[24] Benjamin C. Pierce supra note 13.

[25] Donald E. Knuth supra note 10.

Tuple	A Data Type describing a fixed number of multiple elements, whose order of appearance is relevant.[26]

6.0 Ownership

This guidance is provided by OTI's Strategic Initiatives Division. For questions related to this document, please reach out to your Agency Relations Manager.

7.0 Related Laws, Policies, Requirements and Processes

7.1 Laws

- New York State Freedom of Information Law
- New York City Open Data Law
- New York City Identifying Information Law
- Local Law 35 of 2022

7.2 Citywide Policies

- New York City AI Action Plan of 2023
- Citywide Cybersecurity Policies and Standards
- Citywide Privacy Protection Policies and Protocols of the Chief Privacy Officer, City of New York

7.3 Processes

- Cloud Review
- Procurement
- Software Security Assurance Process (SSAP)

7.4 Agency-Specific Policies

- Various internal business processes and use policies
- Agency cybersecurity policies
- Agency privacy policies
- Agency Acceptable Use Policies

8.0 References

Alan Agresti, *Analysis of Ordinal Categorical Data*, 2010.

Benjamin C. Pierce, *Types and Programming Languages*, 2002.

Center for Security and Emerging Technology, *What are Generative AI, Large Language Models, and Foundation Models?*

Donald E. Knuth, *The Art of Computer Programming, Vol. 2: Seminumerical Algorithms,* 3rd ed., 1997.

---

[26] Benjamin C. Pierce supra note 13.

Helen Toner, *What Are Generative AI, Large Language Models, and Foundation Models?,* Center for Security and Emerging Technology, May 12, 2023, https://cset.georgetown.edu/article/what-are-generative-ai-large-language-models-and-foundation-models/.

IEEE Guide for Terms and Concepts in Intelligent Process Automation, IEEE Std 2755-2017, Sept. 28, 2017

Leo Breiman, *Statistical Modeling: The Two Cultures, Statistical Science* 16(3): 199-231, 2001.

Martin Abadi and Luca Cardelli, *A Theory of Objects*, 1996.

Matthew Hutson, *AI Glossary: Artificial Intelligence, In So Many Words,* Science, July 7, 2017.

National Institute of Standards and Technology (NIST), *The Language of Trustworthy AI: An In-Depth Glossary of Terms,* March 29, 2023, https://doi.org/10.6028/NIST.AI.100-3.

National Security Commission on Artificial Intelligence, *National Security Council on Artificial Intelligence: The Final Report*, *2021,* https://www.nscai.gov/2021-final-report/

Phillip E Gill, Walter Murray and Margaret H Wright, *Practical Optimization*, 1997.

White House Office of Science and Technology Policy (OSTP), Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People, 2022.

9.0 History and Authorship

Version	Change Description	Author(s)	Date
1.0	Inaugural version	Alex Foard, Jiahao Chen and Renata Gerecke, Strategic Initiatives Division	03/04/2024