CAN/CIOSC 101:2019

NATIONAL STANDARD OF CANADA

REAFFIRMED 2021-10

Ethical design and use of automated decision systems

03.100.02; 35.020; 35.240.01

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

- Page left intentionally blank -

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Table of Contents

Introduction vi

1 Scope **9**

2 Normative references **9**

3 Terms and definitions **9**

4 Management of ethical risk **12**

4.1 Risk management framework 12

4.2 Ethics by design 15

4.3 Deployment of automated decision system 16

4.4 Monitoring and maintenance of automated decision systems 16

4.5 Appeals and escalations of decisions rendered by automated decision systems 17

Annex A (informative) 18

Annex B (informative) 20

B.1 AI Principles 20

B.2 Ethical Impact Assessment Tools 20

Bibliography 22

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

- Page left intentionally blank -

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Foreword

CIO Strategy Council (CIOSC) is not-for-profit corporation providing a national forum for public and private sector members to transform, shape and influence the Canadian information and technology ecosystem.

CIOSC standards are developed in accordance with the *Requirements & Guidance – Accreditation of Standards Development Organizations*, 2017-04-05, established by the Standards Council of Canada (SCC).

Attention is drawn to the possibility that some of the elements of this Standard may be the subject of patent rights. CIOSC shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of this Standard are included in the Introduction.

For further information about CIOSC, please contact:

CIO Strategy Council
304-555 Legget Drive, Tower A
Ottawa, ON K2K 2X3
ciostrategycouncil.com

A National Standard of Canada is a standard developed by a Standards Council of Canada (SCC) accredited Standards Development Organization, in compliance with requirements and guidance set out by SCC. More information on National Standards of Canada can be found at www.scc.ca.

SCC is a Crown corporation within the portfolio of Innovation, Science and Economic Development (ISED) Canada. With the goal of enhancing Canada's economic competitiveness and social well-being, SCC leads and facilitates the development and use of national and international standards. SCC also coordinates Canadian participation in standards development, and identifies strategies to advance Canadian standardization efforts.

Accreditation services are provided by SCC to various customers, including product certifiers, testing laboratories, and standards development organizations. A list of SCC programs and accredited bodies is publicly available at www.scc.ca.

Please note there is an interpretation that accompanies this reaffirmed Standard. The reaffirmation can be found in the same location this Standard was downloaded.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

- Page left intentionally blank -

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Introduction

This is the First Edition of CAN/CIOSC 101:2019, Ethical design and use of automated decision systems.

CAN/CIOSC 101:2019 was prepared by the CIO Strategy Council Technical Committee 2 (TC 2) on the ethical design and use of automated decision systems, comprised of over 100 thought leaders and experts in artificial intelligence, ethics, and related subjects. This Standard was approved by a Technical Committee formed balloting group, comprised of 5 producers, 1 government / regulator / policymaker, 4 users, and 9 general interests.

All units of measurement expressed in this Standard are in SI units using the International system (SI).

This Standard is subject to technical committee review beginning no later than one year from the date of publication. The completion of the review may result in a new edition, revision, reaffirmation or withdrawal of the Standard.

The intended primary application of this Standard is stated in its scope. It is important to note that it remains the responsibility of the user of the Standard to judge its suitability for a particular application.

This Standard is intended to be used for conformity assessment.

ICS 03.100.02; 35.020; 35.240.01

CETTE NORME NATIONALE DU CANADA EST DISPONIBLE EN VERSIONS FRANÇAISE ET ANGLAISE

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

- Page left intentionally blank -

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Ethical design and use of automated decision systems

1 Scope

This Standard specifies minimum requirements in protecting human values and incorporating ethics in the design and use of *automated decision systems*.

This Standard is limited to *artificial intelligence (AI)* using *machine learning* for automated decisions.

This Standard applies to all organizations, including public and private companies, government entities, and not-for-profit organizations. It provides a framework and process to help organisations address *AI* ethics principles, such as those described by the OECD:

- *AI* should benefit people and the planet by driving inclusive growth, sustainable development and well-being.
- *AI* systems should be designed in a way that respects the rule of law, human rights, democratic values and diversity, and they should include appropriate safeguards – for example, enabling human intervention where necessary – to ensure a fair and just society.
- There should be transparency and responsible disclosure around *AI* systems to ensure that people understand *AI*-based *outcomes* and can challenge them.
- *AI* systems must function in a robust, secure and safe way throughout their life cycles and potential risks should be continually assessed and managed.
- *Organisations* and individuals developing, deploying or operating *AI* systems should be held accountable for their proper functioning in line with the above principles.

Requirements in this Standard are principles-based and recognize an organization's governing practices may depend on its size; ownership structure; nature, scope and complexity of operations; strategy; and risk profile. Organisations are expected to take *reasonable and responsible measures* to adopt and implement the principles in this Standard.

This Standard is intended to be used in conjunction and integrated with the organisation's compliance programs, including but not limited to, existing privacy, cybersecurity, data governance, complaints and appeals, and legal programs.

2 Normative references

There are no normative references in this document.

3 Terms and definitions

For the purposes of this document, the following terms and definitions apply:

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

accountable
obligation to report, explain, or justify algorithmic decision-making as well as mitigate any negative social impacts or potential harms.

administrative decision
any decision that affects:

- the rights of individuals or communities;
- the health or well-being of individuals or communities; or
- the delivery of consequential services.

affected persons
an individual, entity, community or ecosystem subject to or impacted by *administrative decision* of an *automated decision system*.

algorithm
A set of rules for transforming the logical representation of data. In machine learning applications, algorithms may range from locked to continuously adaptive.

[SOURCE: ISO/IEC 22050:2002, Information technology — Data interchange on 12,7 mm, 384-track magnetic tape cartridges — Ultrium-1 format (modified)]

artificial intelligence (AI)
{in the context of this standard} information technology that leverages machine learning to perform tasks that would ordinarily require biological brainpower to accomplish, such as making sense of spoken language, learning behaviours, or solving problems.

assessor
person who leads and conducts *ethical impact assessment*.

[SOURCE: ISO/IEC 29134:2017 (modified)]

automated decision system
technology that provides information, recommendations, or makes administrative decisions using *artificial intelligence*.

competent personnel
suitably trained or qualified by knowledge and practical experience to enable the required task or tasks to be carried out properly.

consequential service
services that affect human well being such as finance, aid (or assistance), housing, education, recruiting, entitlements.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

ethical impact assessment
a framework to help *organisations* better understand and reduce the potential ethical risks associated with *automated decision systems* and to provide the appropriate governance, oversight and reporting/audit requirements that best match the type of application being designed.

ethics by design
protecting human values and incorporating ethics through technology, product and services design.

generalization
the ability for one machine learning model to be effective using different data sets towards a different application. For example: using a machine learning model originally intended for fraud detection towards cybersecurity.

human in the loop
automated decision system requires human input for some part of task execution, or a human is accountable for executing the *administrative decision*.

human on the loop
automated decision system can operate independently but under human oversight, with intervention as needed, i.e., a human oversees and may override the automated *administrative decision*.

human out of the loop
an administrative decision is made by an *automated decision system* without human intervention, i.e., human cannot override the automated *administrative decision*.

machine learning
technology that is self-learning/improving and can build predictive models from examples, data, and experience, rather than following pre-programmed rules (i.e., deep learning, unsupervised and supervised learning).

material negative impacts
defined by the organisation.

may
optional or possible.

organisation
an individual, entity, community or ecosystem incorporating ethics in the design and use of *automated decision systems* that use *machine learning*.

outcome
intended result as defined by the organisation.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

reasonable and responsible measures[1]
the method(s) used achieves the prescribed *outcome*.

should / is expected
expectation commensurate with each organization's nature, size, business and complexity as well as its structure, economic significance and risk profile.

4 Management of ethical risk

4.1 Risk management framework

4.1.1 The *organisation should* assess the suitability of automated decision making for the intended administrative purpose.

4.1.2 *AI* risk management *should* be part of an *organisation's* compliance programs, including but not limited to, existing privacy, cybersecurity, data governance, complaints and appeals, and legal programs.

NOTE 1: *AI* risk management is a continuous process from project conception, delivery, operations and maintenance through to disposal or decommissioning.

NOTE 2: An organisation should consider incorporating AI principles in its ethics policy and associated objectives – See Appendix B.

4.1.3 The *organisation is expected* to appoint a person or persons (e.g., ethics board or equivalent advisory body) accountable for overseeing the process of including ethical risk as a consideration in the design and use of *automated decision systems*.

4.1.4 The *organisation should* define and implement processes and activities to address the potential risks associated with each stage of deployment of automated decision systems. Figure 1 provides a representative illustration of an *automated decision system* lifecycle.

[1] Anticipatory, reflective, deliberative and responsive to the *outcomes* of the use of the technology.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

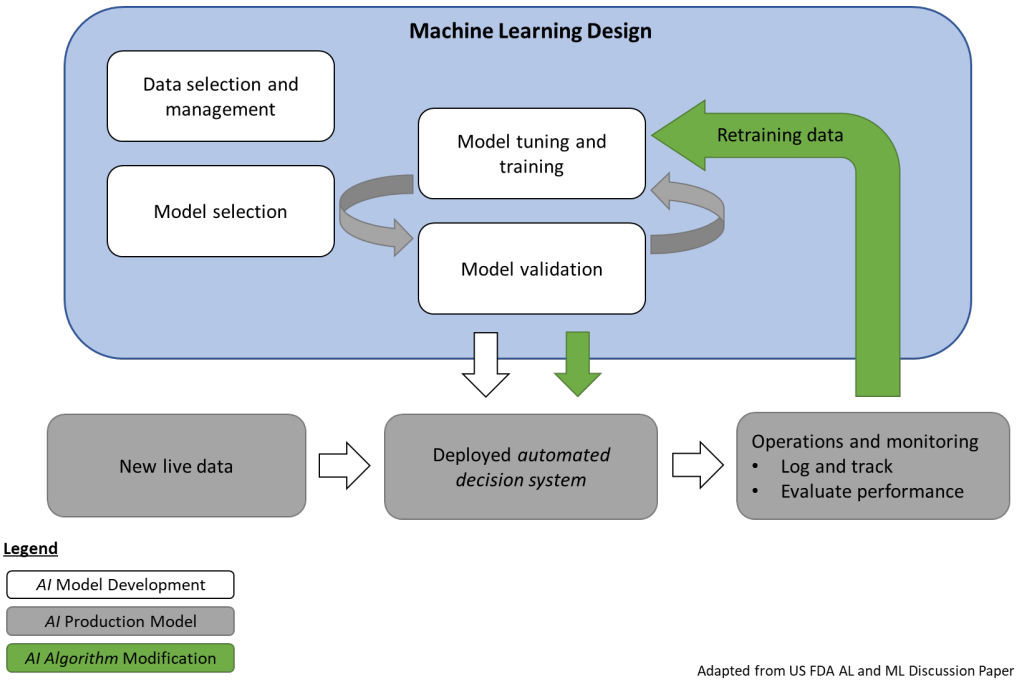


Figure 1 – Automated decision system lifecycle

4.1.5 The *organisation should* consider national and international guidance or standards in developing a risk management framework to address ethical considerations in the design and use of *automated decision systems*.

NOTE: Ethical considerations *may* include those described in Annex A.

4.1.6 The *organisation should* have a measurement system in place to evaluate the severity of the negative *outcomes* identified and to have in place risk mitigation and remediation processes that are commensurate for the severity identified.

4.1.7 The risk of a negative impact resulting from the design and use of *automated decision systems should* accompany the risks assessed by the *organisation* on new initiatives and business-as-usual processes. The *organisation should* establish and apply appropriate safeguards to address the risks identified.

4.1.8 Risks accepted by the *organisation should* be documented and signed off by a senior official of the *organisation* prior to production and testing.

NOTE: This may include review by an ethics board.

4.1.9 The *organisation should* monitor and track occurrences of unintended and unforeseen *outcomes* (including mis-use). The *organisation should* endeavour to address these risks in future iterations.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

4.1.10 An *ethical impact assessment should* form part of the framework to manage ethical risk in the design and use of *automated decision systems*. The *organisation* is required to define the process and applicability of the assessment.

4.1.11 The design and operational effectiveness of the risk management framework as well as its applicability *may* be subject to review by an independent third party.

NOTE: The independent third party *may* be internal or external to the *organisation*.

4.1.12 The *organisation should* determine triggers and thresholds to conduct a new or update an existing *ethical impact assessment*. These triggers *should* be appropriate for the nature and potential impact of the technology being designed and/or modified.

4.1.13 In evaluating ethical risk, the *assessor should* consult *competent personnel* to review and provide input into the *ethical impact assessment* and in the selection of treatment options to safeguard against ethical risk identified and assessed. *Competent personnel may* include:

- Scientific or technical expertise, preferably both related to the field being reviewed and outside that field;
- Lay persons: lay persons *should* only be permitted to serve for a limited time so that they continue to provide an 'outside' perspective on projects;
- Members from outside the *organisation* (e.g. stakeholders or civil society organisation (CSO) representatives) nominated by their organisations in a transparent way and selected because of their competence and experience;
- End-user, or representative of the end-user group or organisation, for example, patients or senior citizens;
- Privacy expertise;
- Security expertise
- Ethical expertise; and
- Legal expertise**.**

4.1.14 The *organisation should* update the *ethical impact assessment* when alerted to *material negative impacts* (as defined by the *organisation*) on *affected persons* by a credible source.

4.1.15 The *organisation should* develop and implement a process to iteratively identify and address ethical and legal challenges both before and after production to adaptatively refine processes as the capabilities of the technology and the legal and regulatory framework evolve to adequately identify and prevent negative impacts to *affected persons*.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

4.2 Ethics by design

4.2.1 The *organisation should* include ethical considerations and impacts in the design and deployment of the *automated decision system(s)* on affected stakeholders.

4.2.2 The *organisation should* consider the *algorithm* lifecycle prior to deployment to ensure that appropriate controls are put in place to maintain confidence and trust in the system. This may include controls for initial training, tuning, deployment side training, update management and continuous adaptation.

4.2.3 The *organisation should* remain compliant with regulations and guidelines associated with the design and use of *automated decision systems*. The *organisation should* consult with legal services to ensure the impact from the use of the *automated decision system* are compliant with applicable laws, and human rights and privacy requirements.

4.2.4 The *organisation should* validate the *outcomes* of the *automated decision system* prior to deployment.

4.2.5 The *organisation should* consider specific points of human intervention in the design of the *automated decision system*.

4.2.6 The *organisation should* consider and implement user experience requirements consistent with *human in the loop*, *human over the loop* and *human out of the loop automated decision systems*.

4.2.7 The *organisation should* establish contingency systems and processes relating to the use of *automated decision systems* that are consistent with business continuity and reporting requirements of the *organisation*.

4.2.8 The *organisation should* develop processes to adequately describe the data used to initially train *machine learning* models. These descriptions should describe characteristics, motivations and skews that the data may represent.

4.2.9 The *organisation* should analyse data for unintended biases and other factors that may unfairly impact the *outcomes* and any persons affected by the decisions rendered.

4.2.10 The *organisation should* select *machine learning* models appropriate for the intended administrative decision process.

4.2.11 The *organisation should* consider limitations such as the *generalizability* of their dataset and *automated decision system* for the administrative decision under consideration.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

4.3 Deployment of automated decision system

4.3.1 The *organisation should* train the applicable personnel in the function, and implementation of the *automated decision system* to be able to review, explain and oversee automated decision-making.

4.3.2 The *organisation's* training program should consider specialized training requirements for *human in the loop*, *human on the loop* and *human out of the loop* decision making systems.

4.3.3 The *organisation should* consider the *algorithm* lifecycle when deploying *automated decision systems*.

4.3.4 For non-locked *algorithms*, *organisations should* review tuning data for suitability to avoid unintended degradation of the expected *outcomes*.

4.3.5 For non-locked *algorithms*, *organisations should* define a tuning period where additional human oversight may be required.

4.3.6 The *organisation should* establish reliability metrics which determine the point at which an *automated decision system* is deemed to have reached production readiness.

4.4 Monitoring and maintenance of automated decision systems

4.4.1 The *organisation should* create a monitoring process to review *outcomes* appropriate to the *algorithm* lifecycle. Static and tuned *algorithms* may require less frequent review than continuously adaptive *algorithms*.

4.4.2 The *organisation should* periodically review established performance, quality and reliability metrics.

4.4.3 Where changes to *automated decision systems* may degrade performance until appropriately tuned, *organisations should* develop processes to safeguard against negative *outcomes*.

4.4.4 The *organisation should* include *automated decision systems* in their cybersecurity plans and activities to guard against manipulation, adversarial attack, fraud, theft, data leakage, or denial of services.

4.4.5 The *organisation should* safeguard both the data and the *algorithms* associated with *automated decision systems*, since it may be difficult, time consuming or impossible to retrain models.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

4.4.6 The *organisation should* create a sunset plan for the *automated decision system* to manage data risks after the product/service is no longer active.

4.5 Appeals and escalations of decisions rendered by automated decision systems

4.5.1 *Organisations should* develop and maintain an appeals and escalation process for *affected persons* to challenge decisions made by the *automated decision process* having negatively impacted them.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Annex A
(informative)

A Ethical considerations

- Is the *automated decision system* used in or could its potential results be used in …

    - a health care or care-giver context with high sensitivity data that is governed by privacy legislation and requires specific accountability, oversight and data governance?

      Data sets can come from electronic health records and health insurance claims, but some systems can use purchasing history, income data, criminal records and social media for information about an individual's health. Precautions and safeguards need to be in place where privacy, discrimination, psychological harm, compromise of the doctor-patient relationship and the circle of care, can raise ethical and legal concerns.

    - a military context (dual use) where the potential benefits (in logistics, detecting and responding to cybersecurity incidents and analysing data from the battlefield, or safeguarding ground forces) needs to be balanced against potential harms (legal, human rights, international convention violations) to the society and other countries?

    - a social benefit context where the use of algorithms that requires public oversight, such as in education, finance, and transportation impact the welfare of individuals. Algorithms used for budgeting, allocation of government resources, welfare, social support programs and other social/economic interventions, need to safeguard against bias and discrimination, as the social impact on human well-being will scrutinized for ethical and legal impacts?

- Could the *automated decision system* or the use of its potential results …

    - have a negative impact on public health, public safety or society in general?

    - have a negative impact on the rights (human rights, charter rights, legislated rights) and liberties of individuals or groups for example, the elderly, racialized groups, persons with disabilities, immigrants, children, women, First Nations, Inuit, Métis, and LGBTQ2S persons.?

      NOTE: Consider effects on the well-being and interests of individuals and groups in society, including the quality of work, and effects on social institutions and structures, democracy and important aspects of culture and cultural diversity.

    - have a negative impact on the environment, animals and/or plants.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

NOTE: Consider, amongst others, the direct and long-term effects on the environment (air, land, sea, space), animals and plants.

- raise concerns in terms of sustainable development (development that does not deplete natural resources)?
- raise concerns in terms of the law (local, regional, national, international)?
- become subject to misuse? Misuse (used in the wrong manner or for the wrong purpose) of algorithms can happen if ethical and legal implications are not built into the decision-making process and *outcomes* are not scrutinized though impact assessments.

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Annex B
(informative)

B Examples (not intended to be an exhaustive list)

B.1 AI Principles

[1] ACM US Public Policy Council. Principles for Algorithmic Transparency and Accountability. https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf

[2] IEEE. Ethically Aligned Design – A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems: First Edition Overview. https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e-overview.pdf

[3] IT Professional. Moral Responsibility for Computing Artifacts: The Rules. https://www.researchgate.net/publication/224239391_Moral_Responsibility_for_Computing_Artifacts_The_Rules

[4] Future of Life Institute. Asilomar AI Principles. https://futureoflife.org/ai-principles/

[5] OECD. Recommendation of the Council on Artificial Intelligence. OECD/LEGAL/0449 http://www.oecd.org/going-digital/ai/principles/

[6] Montréal Declaration for a Responsible Development of Artificial Intelligence. Montréal Declaration. 2017. https://www.montrealdeclaration-responsibleai.com/the-declaration

[7] The Future Society, Law & Society Initiative, Principles for the Governance of AI. https://thefuturesociety.org/2017/07/15/add-law-and-society-initiative-project/

[8] Toronto Declaration: Protecting the right to equality and non-discrimination in machine learning systems. 2018. https://www.accessnow.org/cms/assets/uploads/2018/08/The-Toronto-Declaration_ENG_08-2018.pdf

[9] UNI Global Union. Top Ten Principles for Ethical AI. www.thefutureworldofwork.org/media/35420/uni_ethical_ai.pdf

[10] Women Leading in AI. Principles for Responsible AI. https://womenleadinginai.org/report2019

B.2 Ethical Impact Assessment Tools

[1] Government of Canada. Algorithmic Impact Assessment. https://open.canada.ca/aia-eia-js/?lang=en

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

[2] Norton, Rose, Fulbright. Artificial Intelligence. Ethics Risk Toolkit. https://www.insidetechlaw.com/publications/ethics-risk-toolkit

[3] NuEnergy.ai™. Machine Trust Index™ (MTI). http://www.nuenergy.ai/machine-trust-index/

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

Bibliography

[1] 07-58681:2008 United Nations Declaration on the Rights of Indigenous Peoples

[2] ACM US Public Policy Council. Principles for Algorithmic Transparency and Accountability. https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf

[3] CEN CWA 17145-1:2017, Ethics assessment for research and innovation - Part 1: Ethics committee

[4] CEN CWA 17145-2:2017, Ethics assessment for research and innovation - Part 2: Ethical impact assessment framework

[5] European Commission. Independent High-Level Expert Group on Artificial Intelligence: Ethics Guidelines for Trustworthy AI. 2018

[6] Food and Drug Administration (FDA). Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML) – Based Software as a Medical Device (SaMD) – Discussion Paper and Request for Feedback. 2019.

[7] Future of Life Institute. Asilomar Principles. https://futureoflife.org/ai-principles/

[8] Government of Canada. Algorithmic Impact Assessment. https://open.canada.ca/aia-eia-js/?lang=en

[9] IEEE. Ethically Aligned Design – A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems: First Edition Overview. https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e-overview.pdf

[10] integrate.ai. Responsible AI in Consumer Enterprise. 2018

[11] IT Professional. Moral Responsibility for Computing Artifacts: The Rules. https://www.researchgate.net/publication/224239391_Moral_Responsibility_for_Computing_Artifacts_The_Rules

[12] Norton, Rose, FullBright. Artificial Intelligence: Ethics Risks Toolkit. https://www.aitech.law/publications/2017/q2/ethics-risk-toolkit/

[13] NuEnergy.ai™. Machine Trust Index™ (MTI). http://www.nuenergy.ai/machine-trust-index/

[14] OECD. Recommendation of the Council on Artificial Intelligence. OECD/LEGAL/0449 http://www.oecd.org/going-digital/ai/principles/

[15] Open Data Charter. https://opendatacharter.net/principles/

[16] IAP2 Core Values for the Practice of Public Participation. https://www.iap2.org/page/pillars

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.

[17] IAP2 Spectrum of Public Participation. https://www.iap2.org/page/pillars

[18] ISO 10218-1:2011, Robots and robotic devices -- Safety requirements for industrial robots -- Part 1: Robots

[19] ISO 26000:2010, Guidance on social responsibility

[20] ISO/IEC 29134:2017, Information technology – Security techniques – Guidelines for privacy impact assessment

[21] Montréal Declaration for a Responsible Development of Artificial Intelligence. Montréal Declaration. 2017. https://www.montrealdeclaration-responsibleai.com/the-declaration

[22] The Future Society, Law & Society Initiative, Principles for the Governance of AI. https://thefuturesociety.org/2017/07/15/add-law-and-society-initiative-project/

[23] The Information Accountability Foundation (in collaboration with the Office of the Privacy Commissioner of Canada). Report for the Big Data Assessment for Canadian Private Sector Organizations Project. 2017

[24] Toronto Declaration: Protecting the right to equality and non-discrimination in machine learning systems. 2018. https://www.accessnow.org/cms/assets/uploads/2018/08/The-Toronto-Declaration_ENG_08-2018.pdf

[25] Treasury Board Secretariat of Canada. Standard on Automated Decision-Making. 2018

[26] UNI Global Union. Top Ten Principles for Ethical AI. www.thefutureworldofwork.org/media/35420/uni_ethical_ai.pdf

[27] Unfairness By Algorithm: Distilling The Harms Of Automated Decision-Making. December 2017, Future of Privacy Forum.

[28] Women Leading in AI. Principles for Responsible AI. https://womenleadinginai.org/report2019

© CIOSC 2019 – All rights reserved. Unauthorized reproduction is strictly prohibited.