



IEEE Standard for Algorithmic Bias Considerations

IEEE Computer Society

Developed by the
Software Systems and Engineering Standards Committee

IEEE Std 7003™-2024

STANDARDS

IEEE Standard for Algorithmic Bias Considerations

Developed by the

Software & Systems Engineering Standards Committee
of the
IEEE Computer Society

Approved 11 December 2024

IEEE SA Standards Board

Abstract: The processes and methodologies to help users address issues of bias in the creation of algorithms are described in this standard. Elements include but are not limited to: criteria for the selection of validation data sets for bias quality control, guidelines on establishing and communicating the application boundaries for which the algorithm has been designed and validated to guard against unintended consequences arising from out-of-bound application of algorithms, and suggestions for user expectation management to help mitigate bias due to incorrect interpretation of systems outputs by users (e.g. correlation vs. causation)

Keywords: algorithm, algorithmic system, artificial intelligence, bias, discrimination, ethics, fairness, IEEE 7003™

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Copyright © 2025 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 24 January 2025. Printed in the United States of America.

IEEE is a registered trademark in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

PDF: ISBN 979-8-8557-1619-1 STDGT27576
Print: ISBN 979-8-8557-1620-7 STDPD27576

IEEE prohibits discrimination, harassment, and bullying.

For more information, visit <https://www.ieee.org/about/corporate/governance/p9-26.html>.

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.

Important Notices and Disclaimers Concerning IEEE Standards Documents

IEEE Standards documents are made available for use subject to important notices and legal disclaimers. These notices and disclaimers, or a reference to this page (<https://standards.ieee.org/ipr/disclaimers.html>), appear in all IEEE standards and may be found under the heading “Important Notices and Disclaimers Concerning IEEE Standards Documents.”

Notice and Disclaimer of Liability Concerning the Use of IEEE Standards Documents

IEEE Standards documents are developed within IEEE Societies and subcommittees of IEEE Standards Association (IEEE SA) Board of Governors. IEEE develops its standards through an accredited consensus development process, which brings together volunteers representing varied viewpoints and interests to achieve the final product. IEEE standards are documents developed by volunteers with scientific, academic, and industry-based expertise in technical working groups. Volunteers involved in technical working groups are not necessarily members of IEEE or IEEE SA and participate without compensation from IEEE. While IEEE administers the process and establishes rules to promote fairness in the consensus development process, IEEE does not independently evaluate, test, or verify the accuracy of any of the information or the soundness of any judgments contained in its standards.

IEEE makes no warranties or representations concerning its standards, and expressly disclaims all warranties, express or implied, concerning all standards, including but not limited to the warranties of merchantability, fitness for a particular purpose and non-infringement. IEEE Standards documents do not guarantee safety, security, health, or environmental protection, or compliance with law, or guarantee against interference with or from other devices or networks. In addition, IEEE does not warrant or represent that the use of the material contained in its standards is free from patent infringement. IEEE Standards documents are supplied “AS IS” and “WITH ALL FAULTS.”

Use of an IEEE standard is wholly voluntary. The existence of an IEEE standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard.

In publishing and making its standards available, IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity, nor is IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing any IEEE Standards document should rely upon their own independent judgment in the exercise of reasonable care in any given circumstances or, as appropriate, seek the advice of a competent professional in determining the appropriateness of a given IEEE standard.

IN NO EVENT SHALL IEEE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO: THE NEED TO PROCURE SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE PUBLICATION, USE OF, OR RELIANCE UPON ANY STANDARD, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE AND REGARDLESS OF WHETHER SUCH DAMAGE WAS FORESEEABLE.

Translations

The IEEE consensus balloting process involves the review of documents in English only. In the event that an IEEE standard is translated, only the English language version published by IEEE is the approved IEEE standard.

Use by artificial intelligence systems

In no event shall material in any IEEE Standards documents be used for the purpose of creating, training, enhancing, developing, maintaining, or contributing to any artificial intelligence systems without the express, written consent of IEEE SA in advance. “Artificial intelligence” refers to any software, application, or other system that uses artificial intelligence, machine learning, or similar technologies, to analyze, train, process, or generate content. Requests for consent can be submitted using the [Contact Us](#) form.

Official statements

A statement, written or oral, that is not processed in accordance with the IEEE SA Standards Board Operations Manual is not, and shall not be considered or inferred to be, the official position of IEEE or any of its committees and shall not be considered to be, or be relied upon as, a formal position of IEEE or IEEE SA. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that the presenter’s views should be considered the personal views of that individual rather than the formal position of IEEE, IEEE SA, the Standards Committee, or the Working Group. Statements made by volunteers may not represent the formal position of their employer(s) or affiliation(s). News releases about IEEE standards issued by entities other than IEEE SA should be considered the view of the entity issuing the release rather than the formal position of IEEE or IEEE SA.

Comments on standards

Comments for revision of IEEE Standards documents are welcome from any interested party, regardless of membership affiliation with IEEE or IEEE SA. However, **IEEE does not provide interpretations, consulting information, or advice pertaining to IEEE Standards documents.**

Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Since IEEE standards represent a consensus of concerned interests, it is important that any responses to comments and questions also receive the concurrence of a balance of interests. For this reason, IEEE and the members of its Societies and subcommittees of the IEEE SA Board of Governors are not able to provide an instant response to comments or questions, except in those cases where the matter has previously been addressed. For the same reason, IEEE does not respond to interpretation requests. Any person who would like to participate in evaluating comments or revisions to an IEEE standard is welcome to join the relevant IEEE SA working group. You can indicate interest in a working group using the Interests tab in the Manage Profile & Interests area of the [IEEE SA myProject system](#).¹ An IEEE Account is needed to access the application.

Comments on standards should be submitted using the [Contact Us](#) form.²

¹ Available at: <https://development.standards.ieee.org/myproject-web/public/view.html#landing>.

² Available at: <https://standards.ieee.org/about/contact/>.

Laws and regulations

Users of IEEE Standards documents should consult all applicable laws and regulations. Compliance with the provisions of any IEEE Standards document does not constitute compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

Data privacy

Users of IEEE Standards documents should evaluate the standards for considerations of data privacy and data ownership in the context of assessing and using the standards in compliance with applicable laws and regulations.

Copyrights

IEEE draft and approved standards are copyrighted by IEEE under U.S. and international copyright laws. They are made available by IEEE and are adopted for a wide variety of both public and private uses. These include both use by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making these documents available for use and adoption by public authorities and private users, neither IEEE nor its licensors waive any rights in copyright to the documents.

Photocopies

Subject to payment of the appropriate licensing fees, IEEE will grant users a limited, non-exclusive license to photocopy portions of any individual standard for company or organizational internal use or individual, non-commercial use only. To arrange for payment of licensing fees, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; +1 978 750 8400; <https://www.copyright.com/>. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Updating of IEEE Standards documents

Users of IEEE Standards documents should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect.

Every IEEE standard is subjected to review at least every 10 years. When a document is more than 10 years old and has not undergone a revision process, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE standard.

In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit [IEEE Xplore](#) or [contact IEEE](#).³ For more information about the IEEE SA or IEEE's standards development process, visit the IEEE SA Website.

³Available at: <https://ieeexplore.ieee.org/browse/standards/collection/ieee>.

Errata

Errata, if any, for all IEEE standards can be accessed on the [IEEE SA Website](#).⁴ Search for standard number and year of approval to access the web page of the published standard. Errata links are located under the Additional Resources Details section. Errata are also available in [IEEE Xplore](#). Users are encouraged to periodically check for errata.

Patents

IEEE standards are developed in compliance with the [IEEE SA Patent Policy](#).⁵

Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken by the IEEE with respect to the existence or validity of any patent rights in connection therewith. If a patent holder or patent applicant has filed a statement of assurance via an Accepted Letter of Assurance, then the statement is listed on the IEEE SA Website at <https://standards.ieee.org/about/sasb/patcom/patents.html>. Letters of Assurance may indicate whether the Submitter is willing or unwilling to grant licenses under patent rights without compensation or under reasonable rates, with reasonable terms and conditions that are demonstrably free of any unfair discrimination to applicants desiring to obtain such licenses.

Essential Patent Claims may exist for which a Letter of Assurance has not been received. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patents Claims, or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

IMPORTANT NOTICE

Technologies, application of technologies, and recommended procedures in various industries evolve over time. The IEEE standards development process allows participants to review developments in industries, technologies, and practices, and to determine what, if any, updates should be made to the IEEE standard. During this evolution, the technologies and recommendations in IEEE standards may be implemented in ways not foreseen during the standard's development. IEEE standards development activities consider research and information presented to the standards development group in developing any safety recommendations. Other information about safety practices, changes in technology or technology implementation, or impact by peripheral systems also may be pertinent to safety considerations during implementation of the standard. Implementers and users of IEEE Standards documents are responsible for determining and complying with all appropriate safety, security, environmental, health, data privacy, and interference protection practices and all applicable laws and regulations.

⁴ Available at: <https://standards.ieee.org/standard/index.html>.

⁵ Available at: <https://standards.ieee.org/about/sasb/patcom/materials.html>

Participants

At the time this draft Standard was completed, the Algorithmic Bias Considerations Working Group had the following membership:

Gerlinde Weger, *Chair*
Ansgar Koene, *former Chair*
Chris Clifton, *Vice Chair*
Julian Padget, *Technical Editor*
Liz Dowthwaite, *Secretary*

Rahaf Albalkhi	Allison Gardner	Satyam Mohla
Frédéric Andres	Andreas Gratz	Zvikomborero Murahwi
Andrea Basso	Jon Hagar	Theodore S. Nowak
Lyria Bennett Moses	Yohko Hatada	Annette D. Reilly
Raymond Robert Bond	Maya Indira Ganesh	Abel Pena Fernandez
Abhik Chaudhuri	Clare James	M S Prasad
Patrick Courtney	Jussi Leppälä	R.K. Rannow
Francien Dechesne	Maroussia Lévesque	Patricia Shaw
Paola Di Maio	Róisín Loughran	Adam L. Smith
Andrey Fajardo	Mathana	Ewelina Szczekocka
Ahmad Firdaus	José Meza	Mark Underwood

The following members of the individual Standards Association balloting committee voted on this Standard. Balloters may have voted for approval, disapproval, or abstention.

Butch Anton	Werner Hoelzl	Edwin Puerta
Lee Barford	Ronald Jarrett	Sridhar Raghavan
Lyria Bennett Moses	Piotr Karocki	R. K. Rannow
Juris Borzovs	Ansgar Koene	Annette D. Reilly
Pieter Botman	Ruth Lewis	Jhony Sembiring
Chris Clifton	Cari Miller	Patricia Shaw
Andrey Fajardo	Roberto Moreno	Thomas Starai
Andrew Fieldsend	Joanna Olszewska	Ron Starman
Andres Frederic	Moradeke Olumogba	Walter Struppler
Madhukar Gaganam	Thomas Owens	Mark Underwood
Jon Hagar	Julian Padget	M. Woolf
Sheri Harshberger	George Percivall	Janusz Zalewski

When the IEEE SA Standards Board approved this Standard on 11 December 2024, it had the following membership:

David J. Law, Chair
Jon Walter Rosdahl, Vice Chair
Gary Hoffman, Past Chair
Alpesh Shah, Secretary

Sara R. Biyabani
Ted Burse
Stephen Dukes
Doug Edwards
J. Travis Griffith
Guido R. Hiertz
Ronald W. Hotchkiss

Hao Hu
Yousef Kimiagar
Joseph L. Koepfinger*
Howard Li
Xiaohui Liu
John Haiying Lu
Kevin W. Lu
Hiroshi Mano

Paul Nikolich
Robby Robson
Lei Wang
F. Keith Waters
Sha Wei
Philip B. Winston
Don Wright

*Member Emeritus

Introduction

This introduction is not part of IEEE Std 7003™-2024, IEEE Standard for Algorithmic Bias Considerations.

This standard has evolved along with the maturity of the concept of ethical bias. The term ‘ethical bias’ is used to refer to bias that is wanted, at the time it is being assessed, because it contributes to the correct functioning of the Autonomous Intelligent Systems (AIS).

Bias is inherent and in many cases wanted. For example, a good result from a search engine should be biased to match the interests of the user as expressed by the search-term, and possibly refined based on personalization data.

When ‘no bias’ is wanted it generally means that unwanted bias should be minimized, as defined by the context within which the AIS is being used.

In the absence of malicious intent, unwanted bias in an AIS is generally caused by:

- Insufficient understanding of the context of which the system is a part. This includes an incomplete understanding about who could be affected by the algorithmic decision outcomes, resulting in a failure to test how the system performs for specific groups. This can be mitigated when those responsible for the AIS have a diversity of perspectives. Diversity includes individuals with different cultural backgrounds, skin colour, appearance, gender/sex, age, experience, disabilities and education. The relevant clauses in the standard are those concerning requirements (Clause 4), stakeholder identification (Clause 6) and risk and impact analysis (Clause 8).
- Failure to rigorously map decision criteria. When algorithmic decisions are considered more objectively trustworthy than human decisions, more often than not this assumes that algorithmic systems follow a clearly defined set of criteria with no hidden agenda. However, the complexity of system development may embed hidden decision criteria derived from the data used by the system. The failure to correctly map decision criteria can be mitigated by ensuring a balance of influence from across stakeholder groups throughout the life cycle of the AIS. The relevant clauses in the standard are those concerning data representation (Clause 7) and again stakeholder identification (Clause 6).
- Inadequate monitoring in the operations stage. AIS are complex systems, which means they have emergent properties, some of which may only become apparent long after development. Amongst many issues, the most notable is drift, arising either from changing alignment between the original data from which the model was built and ‘live’ data, or from the model’s outputs no longer aligning with changing expectations. Thus, bias that was ethical in the past may no longer be so and indeed bias that was unethical may become so. Consequently, an AIS needs to be monitored throughout its operation and the associated mitigations can propagate back through to the initial steps in its creation. The relevant clauses in the standard are all of those identified so far with the addition of evaluation (Clause 9).

This standard sets out an approach for the mindful consideration of bias throughout the AIS life cycle from inception through to decommissioning. This document focuses on what to do and why, but not how to do it, because the precise actions are inevitably context specific, and should be determined and justified by those responsible. Rather than providing a checklist for algorithmic bias consideration, the aim is to support those using the standard in working out the appropriate considerations for their context of use. This bias profile, central to the process, contains the various initial, intermediate and final versions of the documents. These documents are associated with each activity that capture information for considering what constitutes an acceptable level of algorithmic bias risk for the circumstances that hold at the time. Each activity feeds forwards and feeds backwards to facilitate the bias consideration process. Some biases may be caught before they do harm through foresight and experience, others may only manifest later, from which hindsight can inform the next iteration. The iterative process ensures that risk and the stakeholder impact are assessed at all

milestones of the AIS life cycle. Conducting iterative risk assessments at appropriate life cycle stages allows for fulsome and targeted mitigation strategies. This standard provides organizations, that develop, implement and use AIS, an approach to mindfully consider and then optimize for ethical bias.

Acknowledgments

Permissions have been granted as follows:

- Material excerpted from ISO standards, as cited in this standard, is used with permission of the American National Standards Institute (ANSI) on behalf of the International Organization for Standardization. All rights reserved.
- The definition of stakeholder (Clause 3) is reproduced with modification from IEEE Std 7010-2020 with permission from IEEE SA.
- Figure 1 “Bias-considerate Development” reproduced with permission from Julian Padget, © 2024.

Contents

1	Overview	14
1.1	Scope	14
1.2	Purpose	14
1.3	Word Usage	15
1.4	Field of application	15
1.5	Limitations	15
1.6	Organization of the standard	16
1.7	Audience	16
1.8	Conformance	16
1.9	Disclaimer	17
2	Normative references	17
3	Definitions, acronyms, and abbreviations	17
3.1	Definitions	18
3.2	Acronyms and abbreviations	21
4	Requirements for Bias Consideration	22
4.1	Purpose	22
4.2	Inputs	22
4.3	Outputs	22
4.4	Actions	22
4.5	Outcomes	23
5	The Bias Profile	23
5.1	Purpose	23
5.2	The bias profile and the system life cycle	24
5.2.1	Requirements for Bias Consideration	24
5.2.2	Stakeholder Identification	25
5.2.3	Data Representation	25
5.2.4	Risk and Impact Assessment	26
5.2.5	Evaluation	26
6	Stakeholder Identification	26
6.1	Purpose	26
6.2	Inputs	26
6.3	Outputs	27
6.4	Actions	27
6.5	Outcomes	27
7	Data Representation	28
7.1	Purpose	28
7.2	Inputs	28
7.3	Outputs	28
7.4	Actions	28
7.4.1	Action purposes	28
7.4.2	Metadata	29
7.5	Outcomes	30
8	Risk and Impact Assessment	30
8.1	Purpose	30
8.2	Inputs	30
8.3	Outputs	31

8.4	Actions	31
8.5	Outcomes	31
9	Evaluation	32
9.1	Purpose	32
9.2	Evaluation of the AIS Design and Outputs	32
9.2.1	Purpose	32
9.2.2	Inputs	32
9.2.3	Outputs	32
9.2.4	Actions	33
9.2.5	Outcomes	34
9.3	Ongoing evaluation	34
9.3.1	Purpose	34
9.3.2	Inputs	35
9.3.3	Outputs	35
9.3.4	Actions	35
9.3.5	Outcomes	36
Annex A	(informative) Conceptualizing Bias	37
A.1	Introduction	37
A.2	Understanding Algorithmic Bias and Aspects of Human Bias	37
A.3	Categorizing Biases	38
A.4	Bias profile within the life cycle	38
A.5	Considerations	39
A.6	Implications and Recommendations	39
Annex B	(informative) Stakeholder Identification	40
B.1	Physical Attributes	40
B.2	Impacted Stakeholders	40
B.2.1	Categories of Impacted Stakeholders	40
B.2.2	Attributes of Impacted Stakeholders	41
B.3	Process to Identify attributes Impacted Stakeholders	41
B.4	Influencing Stakeholders	42
B.4.1	Identifying Influencing Stakeholders	42
B.4.2	Assessing Influence	43
B.4.3	Process of Determining the Influence of Influencing Stakeholders	43
B.4.4	Attributes of Influencing Stakeholders	43
B.4.5	Metrics for Influencing Stakeholders	44
B.4.6	Outcomes from Influencing Stakeholders	44
B.5	Diversity	44
B.6	Protected attributes	44
B.7	Defining Fairness	44
Annex C	(informative) Guidance for Risk and Impact Assessment	45
C.1	Tasks for Assessment	45
C.2	Additional Resources	45
Annex D	(informative) Measuring Bias	47
D.1	Metrics	47
D.2	Reference values	47
D.3	Granularity of testing	48
D.4	Additional resources	48
Annex E	(informative) Cultural Aspects	49
E.1	General	49

E.2	Impact of cultural bias on AIS	49
E.3	Impact of AIS bias on culture	50
E.3.1	Algorithmic culture	50
E.3.2	Problems with (mis)classification	51
E.3.3	Algorithms and polarisation	51
E.3.4	Dimensions of cultural bias	51
E.4	Accounting for diversity	52
E.5	Recommendations	53
Annex F	(informative) Bibliography	54

IEEE Standard for Algorithmic Bias Considerations

1 Overview

1.1 Scope

Computer algorithms and analytics are playing an increasingly influential role in government, business and society. They underpin information services and autonomous intelligent systems (AIS) including but not limited to artificial intelligence applications that involve symbolic and subsymbolic technologies and their hybridization. These technologies are having a direct and significant impact on human lives across a broad socioeconomic, political and cultural spectrum. Algorithms enable the exploitation of vast and varied data sources from public and private spheres to support human decision-making and actions that serve the diverse interests of the societies and economies in which they operate. However, alongside the benefits, their use is not without maleficent risk.

This standard describes processes and methodologies to help users address issues of bias in the creation of algorithms and models. Elements include but are not limited to: criteria for the selection of data sets; guidelines on establishing and communicating the application boundaries for which the AIS has been designed and validated to guard against unwanted consequences arising from out-of-bounds application of an AIS; suggestions for user expectation management to help mitigate unwanted bias due to incorrect interpretation of systems outputs by users (e.g., correlation vs. causation).

1.2 Purpose

This standard is designed to provide individuals or organizations creating an AIS, certification-oriented processes and methodologies to produce clearly articulated accountability and clarity around how an AIS targets, assesses and influences the stakeholders of said AIS. This standard enables AIS creators to define, measure and communicate how bias is used in the AIS to users and regulatory authorities, and show that best practices were used in the design, development, testing and evaluation of the AIS to avoid unwanted differential impact on stakeholders.

1.3 Word Usage

The word *shall* indicates mandatory requirements strictly to be followed in order to conform to the standard and from which no deviation is permitted (*shall equals is required to*).^{6,7}

The word *should* indicates that among several possibilities one is recommended as particularly suitable, without mentioning or excluding others; or that a certain course of action is preferred but not necessarily required (*should equals is recommended that*).

The word *may* is used to indicate a course of action permissible within the limits of the standard (*may equals is permitted to*).

The word *can* is used for statements of possibility and capability, whether material, physical, or causal (*can equals is able to*).

1.4 Field of application

This standard is agnostic to the types of computational approaches used in the algorithmic systems it applies to, be they rule-based, statistical, machine learning or otherwise.

This standard applies to all algorithmic systems that are involved in selection, allocation, ranking, decision-making or any other processes in which some parties could receive different outcomes to other parties, making it possible for the system to exhibit unwanted bias. The standard is applicable when new algorithmic systems are designed, existing systems are used with data different to their training data, when systems are deployed in new contexts or the systems are updated and when systems are decommissioned.

This standard does not deny the role of operationally-justified bias in algorithmic processing as a fundamental element in information classification and decision making. This standard seeks to help with distinguishing and communicating the difference between wanted and unwanted bias, and thereby clarify the limits for appropriate use of that algorithmic system.

This standard acknowledges that the defining attributes of unwanted bias may depend on the social and societal context within which an algorithmic system is used. This standard therefore does not seek to enumerate specific algorithmic decision criteria or differential impacts to define unwanted bias, but rather to provide methods which can be tuned in accordance with individual contexts of use.

It is recognized that teams with diverse experience, age, education, gender and cultural backgrounds would contribute to a comprehensive consideration of bias and help minimize unwanted bias within the AIS throughout its life cycle.

1.5 Limitations

The scope of this standard includes methods and processes to help designers and developers in the creation of AIS in which bias, when required, is wanted. It does not include systems or processes to detect and evaluate problematic algorithms in actual use. Tracking and certification are out of scope for this standard.

⁶The use of the word *must* is deprecated and cannot be used when stating mandatory requirements; *must* is used only to describe unavoidable situations.

⁷The use of *will* is deprecated and cannot be used when stating mandatory requirements; *will* is only used in statements of fact.

1.6 Organization of the standard

The algorithmic bias identification and mitigation processes described in this standard are constructed to be consistent with the IEEE 70xx series of standards developed as part of the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems.

This standard is organized into normative clauses addressing specific stages in the AIS life cycle [B34] where actions shall be taken to consider the bias intent and to mitigate unwanted bias. The normative clauses are accompanied by several annexes of a) informative clauses intended to support the contextualization and cultural context within which algorithmic bias is situated and b) examples of process to support the normative clauses.⁸

The clauses are ordered around the central concept of the bias profile (Clause 5). The bias profile is the repository that holds the outputs and outcomes that support bias consideration. The ordering of the clauses in the standard follows the sequence in which the clauses are most likely to be used for the first time during a development process, but each shall be iteratively re-visited throughout the AIS life cycle, appropriate to the development, operation and changing circumstances of the AIS. The clauses are:

- a) Requirements for Bias Consideration (Clause 4): Establishing the essential documents and processes and methods to consider bias for a specific AIS within its context of use.
- b) Bias Profile (Clause 5): The record of the process of bias consideration throughout the AIS life cycle within a specific context of use.
- c) Stakeholder Identification (Clause 6): The attributes of all stakeholders that influence and are impacted by the AIS. It is to be noted that it is an identification process not an evaluation process.
- d) Data Representation (Clause 7): Review of the data to assess whether the data are representative of the stakeholders identified. This includes data provenance and the representative mapping of data against the stakeholders.
- e) Risk and Impact Assessment (Clause 8): An assessment of all potential risks and impacts from intended and unintended uses or applications. It is to be noted that the team assessing the risks should be representative of the stakeholders identified.
- f) Evaluation (Clause 9): Evaluating the AIS as a whole for bias through two processes: 1) assess for bias in the design and outputs of the AIS; 2) set up a process to perform the ongoing evaluation of the AIS.

1.7 Audience

The intended audience for this standard includes system, software, and hardware suppliers, procurers, developers, maintainers, system auditors, operators, users, and managers in both the supplier and procuring organizations or within their own organization in case of in-house development. The standard is also aimed at data scientists, artificial intelligence researchers and their research organizations.

1.8 Conformance

This standard can be used as a conformance document for projects and organizations claiming conformance to IEEE 7003™ standard on Algorithmic Bias Considerations.

⁸The numbers in brackets correspond to those of the bibliography in Annex F.

Use of the nomenclature of this standard for the parts of information for users (for example, chapters, topics, pages, screens, windows) is not required to claim conformance.

This standard may be included or referenced in contracts or similar agreements when the parties (called the procurer and the supplier) agree that the supplier shall deliver services and systems in accordance with the standard.

This standard may also be adopted as an in-house standard by a project or organization that decides to acquire information for users from another part of the organization in accordance with the standard.

1.9 Disclaimer

This standard establishes minimum criteria for helping ensure that AIS do not exhibit unwanted bias. However, implementing these criteria does not automatically ensure conformance to system or mission objectives, or prevent adverse consequences (e.g., loss of life, mission failure, loss of system safety or security, or financial or social loss). Conformance to this standard does not absolve any party from any social, ethical, moral, financial, or legal obligations.

2 Normative references

This standard has no normative references.

3 Definitions, acronyms, and abbreviations

For the purposes of this document, the following terms and definitions apply. The *IEEE Standards Dictionary Online* should be consulted for terms not defined in this clause.⁹

NOTE 1—For additional terms and definitions in the field of systems and software engineering, see ISO/IEC/IEEE 24765:2017 [B39], which is published periodically as a “snapshot” of the SEVOCAB (Systems and Software Engineering Vocabulary) database and is publicly accessible at computer.org/sevocab.¹⁰

NOTE 2—The following ISO/IEC and ISO/IEC/IEEE standards: ISO/IEC/IEEE 2382:2015 *Information technology – Vocabulary*, ISO/IEC 22989:2022 *Information technology — Artificial intelligence — Artificial intelligence concepts and terminology*, ISO/IEC/IEEE 24765:2017 *Systems and software engineering — Vocabulary* and ISO/IEC/IEEE 29119-1:2022 *Software and systems engineering – Software testing – Part 1: General concepts* are publicly available from <https://standards.iso.org/ittf/PubliclyAvailableStandards/index.html>.

NOTE 3—The following IEEE standards: ISO/IEC/IEEE 24748-7000-2022 *Standard for Systems and software engineering—Life cycle management—Part 7000: Standard model process for addressing ethical concerns during system design*, IEEE Std 7005-2021 *IEEE Standard for Transparent Employer Data Governance* and IEEE Std 7010-2020 *IEEE Recommended Practice for Assessing the Impact of Autonomous and Intelligent Systems on Human Well-Being* are publicly available through the IEEE GET program from <https://ieeexplore.ieee.org/browse/standards/get-program/page/series?id=93>.

⁹IEEE Standards Dictionary Online is available at: <http://dictionary.ieee.org>. An IEEE Account is required for access to the dictionary, and one can be created at no charge on the dictionary sign-in page.

¹⁰Notes in text, tables, and figures of a standard are given for information only and do not contain requirements needed to implement this standard.