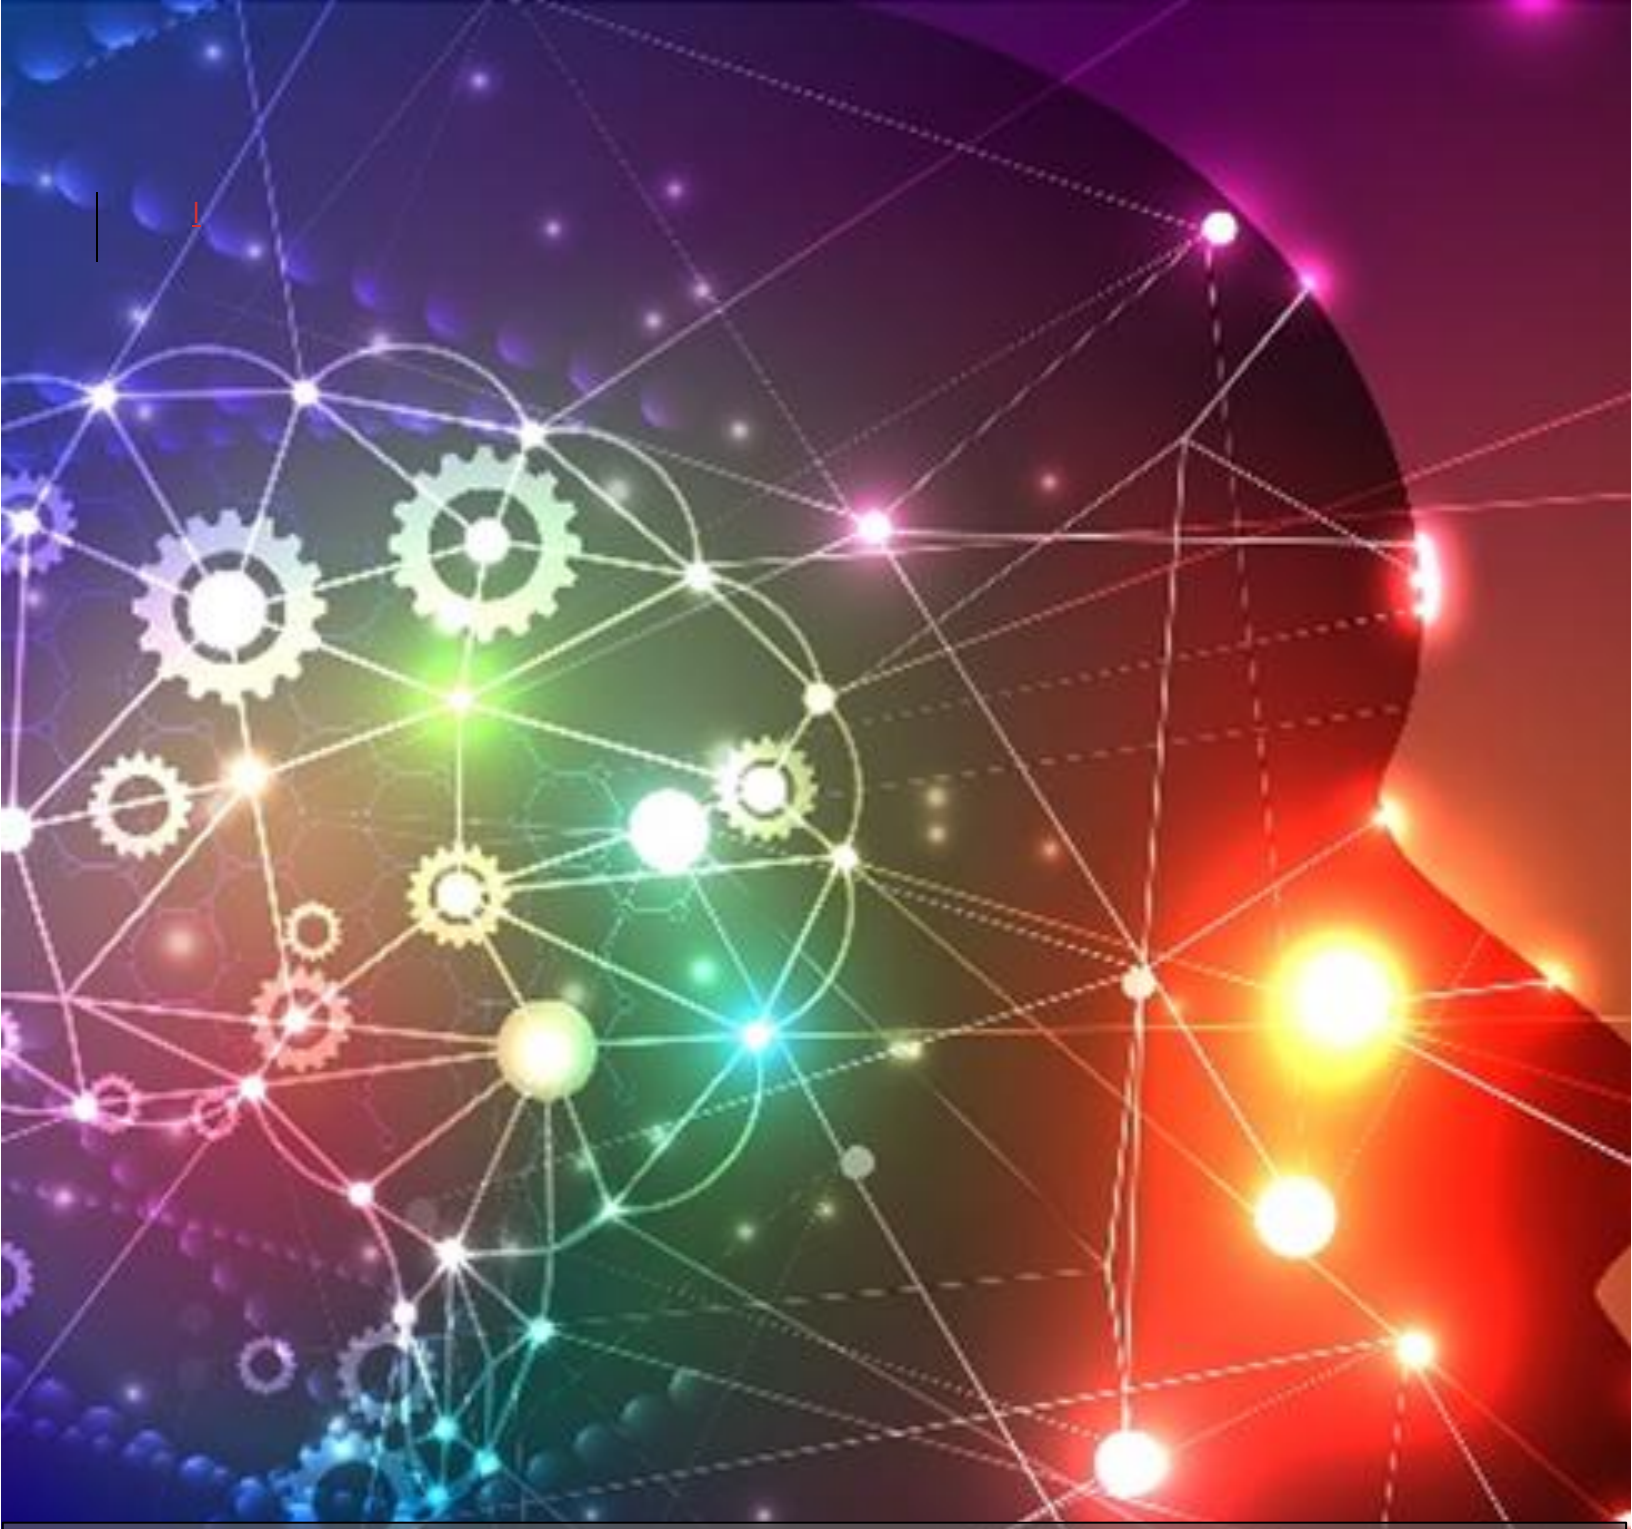

Artificial Intelligence and OMG Standards

Document omg/19-08-01
26 August 2019

Claude R. Baudoin, cébé IT & Knowledge Management
Larry L. Johnson, Technical Director, OMG
Pete Rivett, CTO, Adaptive Inc.

OBJECT MANAGEMENT GROUP®

Copyright Notice

Acknowledgements

The authors would like to thank the following OMG members and supporters for their expertise on the various standardization efforts discussed in this paper. In particular, we extend our sincerest gratitude to Sridhar Iyengar, IBM; Uwe Kaufmann, ModelAlchemy; Elisa F. Kendall, Thematix Partners LLC; Dr. Said Tabet, DELL; Bobbin Teegarden, OntoAge; and Ron Zahavi, Microsoft.

About OMG

The Object Management Group® (OMG®) is an international, open membership, not-for-profit technology standards consortium with representation from government, industry and academia. OMG Task Forces develop enterprise integration standards for a wide range of technologies and an even wider range of industries. OMG's modeling standards enable powerful visual design, execution and maintenance of software and other processes. Visit www.omg.org for more information.

Table of Contents

1 Introduction and Executive Summary

The Board of Directors of the Object Management Group believes that the time has come for defining standards in Artificial Intelligence (AI). Such standards will greatly accelerate and improve the creation of useful AI applications by reducing the amount of low-level interchange, integration, or interfacing work currently required.

The OMG is therefore establishing an initiative in AI standardization.

OMG is the best place for industry, academia, government and non-profit organizations to come together to define such standards, because of the best practices it has developed over the last thirty years and can easily extend to this new scope of work.

We initiated this paper to convey this message to both current and prospective members of OMG. The paper explains the following.

- An expanded scope of AI has emerged from the renaissance of the discipline since the beginning of the millennium. AI now reaches into many domains and integrates new technologies such as ontologies, vision recognition, analytics and machine learning for the industrial IoT, and more.
- When a technology area reaches a certain degree of maturity, standards can enable innovation—rather than impede it—by freeing organizations from having to constantly worry about the "plumbing" of systems or re-inventing platform techniques and tools.
- OMG is uniquely placed to lead this effort because it has a proven process, it has developed a series of foundational capabilities that can bootstrap the effort, and it covers many industry domains that can benefit from AI standards.

Based on these observations, we propose actions to facilitate the development by our members, in liaison with other organizations, of a roadmap for AI-enabling standards, and we invite all interested parties to join us in this effort.

2 The New Scope of AI

While there are many competing definitions of AI, it is clearly not just rule-based expert systems (the domain of AI that seemed the most promising during the early years, and disappointed enough to cause the long period known as the "AI winter"), nor neural networks or other forms of machine learning.

Table 1 provides a simplified list of key AI capabilities. This is a simplification of several authors' taxonomies of AI, such as [1] and [2]. The AAAI taxonomy of AI [1] is shown in Appendix A as an example. Among this abundance of topics, we focused on those related to existing OMG domains of interest.

Of course, in certain domains multiple AI capabilities are jointly used. Not all capabilities are at the same level either – for example, machine learning can be used in vision, advanced robotics or natural language processing.

Table 1 – Key AI Capabilities

AI Capability	Definition	Sample Applications
Machine Learning and Deep Learning	Creation of an optimal response to a set of inputs, usually obtained through the training of a neural network.	Predictive maintenance. Pattern recognition. Abnormality detection (safety, security).
Vision	Image processing to recognize shapes.	Obstacle avoidance. Facial recognition (security applications). Vehicle navigation systems.
Smart Robots	Sensing, route planning, adaptive prehension, ability to react to changes in the surroundings, etc.	Mobile industrial robots. Human-assistive robots. Telesurgery.
Natural Language Processing	Understanding of the semantics of natural language in spite of the ambiguity of the vocabulary and grammar.	Real-time translation. Spam blocking. Smart assistants (Alexa, Cortana, chatbots, etc.)
Rules-Based Systems	These are the successors to the expert systems of the 1960s-70s. They apply in an algorithmic manner rules captured from human experts.	Loan application processing. Medical triage by the UK National Health Service hotline. Robotic process automation (RPA).

3 Why Standards in AI?

3.1 Phases of Standards Adoption

Every technology goes through a difficult relationship with standards in its early years:

- During the emerging phase, where technology development is led by research groups, or by startups, standards are not seen as important or may even be considered harmful.
- During early commercialization, each supplier is keen to create and preserve an advantage, and to attract customers and lock them in to their proprietary technology.
- When the technology matures, customers discover the need for integration, interoperability, and migration from one system to another. At the same time, suppliers need to attract a broader clientele. The need for standards thus emerges. At the end of this phase, *ad hoc* standards are defined but are not yet well adopted.
- After a while, it becomes clear that more formality and governance are needed. The industry then comes together within standards organizations and initiatives. Compliance with standards becomes a selling point and a procurement criterion.

Premature standardization could stifle innovation; however, delayed standardization creates unnecessary interchange, interoperability and integration difficulties. Today, many organizations are still in the first or second stage listed above—they have not recognized the need for standards in AI, or they have recognized it but are not convinced that adopting standards is yet in their best interest. [3]

3.2 The Increasing Maturity of AI

By the early 2000s, it became clear that some of the early visions of AI could actually be realized [4]. In particular, new practical applications of neural networks (under the name of "machine learning" or "deep learning") emerged. AI also proved critical in deciphering and making decisions from the flood of data collected by Internet of Things (IoT) systems.

AI now has a significant impact and applications in practically all industries. Sectors that lead this adoption include legal, insurance, crime and fraud investigation and prevention, meteorology, media management, marketing, and more.

AI is now supported by comprehensive IT offerings, evidence by:

- the availability of platforms from big players such as IBM, Google, Amazon, or Microsoft;
- the pervasiveness of cloud-based capabilities (with a consequent focus on APIs);

- the availability of massive amounts of data, both structured and unstructured, including real-time IoT data, which can be exploited by AI and machine learning algorithms.

The following activities and emerging standards (not an exhaustive list) indicate that the time is ripe for standardization of certain aspects of AI:

- The **Neural Network Exchange Format (NNEF)**[1] "reduces machine learning deployment fragmentation by enabling a rich mix of neural network training tools and inference engines to be used by applications across a diverse range of devices and platforms."
- The **Open Neural Network eXchange (onnx)**[2] is an open-source, community-driven effort to allow developers to more easily move between machine learning frameworks.
- The **Hierarchical Data Format 5** (**HDF5** or **.h5**)[3] is a standard representation of scientific data sets, together with metadata, and is used in particular for the interchange of machine learning training data sets.
- **The ISO Subcommittee on Artificial Intelligence**[4] has been working on three standards related to big data (ISO/IEC 20546, 20547-2 and 20547-5).
- The **Consumer Technology Association (CTA)**[5] launched an initiative on AI aimed at "improving efficiencies in AI and Health Care."
- The **National Institute for Science and Technology (NIST)** is conducting several related activities:
    - It undertook a proof-of-concept project to develop an **Industrial Ontology Foundry (IOF)**[6]**.**
    - the **Multimodal Information Group**[7] has conducted Language Recognition Evaluation studies.
    - In May 2019, it issued a **Request for Information** about the need for AI standards [5], which received 98 responses [6], including an extensive one from OMG. [7]

---

1 https://www.khronos.org/nnef
2 https://onnx.ai/
3 https://www.hdfgroup.org/
4 https://www.iso.org/standard/72826.html
5 https://www.cta.tech/News/Press-Releases/2019/April/CTA-Brings-Together-Tech-Giants,-Trade-Association.aspx
6 https://www.nist.gov/publications/industrial-ontologies-foundry-proof-concept-project
7 https://www.nist.gov/itl/iad/mig

- There are also AI initiatives in organizations such as **OpenAI[8],** the **Artificial Intelligence Open Network (AI-ON)[9],** the **Machine Intelligence Research Institute (MIRI)[10],** the **Allen Institute for Artificial Intelligence (AI2)[11],** the **Partnership on AI to Benefit People and Society[12]**, the **Cognitive Computing Consortium[13],** the **Consortium for Safer Artificial Intelligence[14],** and more.
- Open-source AI frameworks such as **TensorFlow, Keras, Caffe, Scikit-learn, Theano,** and **Torch** are starting to be widely adopted.
- We now see a growing interest in the field of AI Ethics, including the work of Dr. Andreas Vogel[15] and position papers from various companies.[16]

3.3 Impact of the Lack of Standards

There is clear evidence from multiple sectors (finance, space, robotics, manufacturing, healthcare, energy, and more) that conflicting models, languages and data formats may impede the progress of applying AI. If AI models cannot be used together or do not have consistent semantics, one may get the wrong results. Here are some more specific examples of this impact.

- Deep or unsupervised learning algorithms are incredibly opaque and difficult to understand, which impacts their reliability, maintainability, reuse, transparency, respect for privacy, and more. Some research work has started in academia and in the finance community to address these issues by attempting to combine declarative ontologies and rules with the systems they specify. [8] [9]
- Without standard interfaces and well-defined ontologies, the robot industry cannot evolve toward flexible, upgradable systems assembled from interchangeable modules.
- Google's new dataset search capability (https://toolbox.google.com/datasetsearch) allows scientists to find datasets, and the World Wide Web Consortium (W3C) has developed a data catalog vocabulary (DCAT); but the use of those datasets is hindered by the lack of associated AI-specific metadata.

---

[8] https://openai.com
[9] https://ai-on.org
[10] https://intelligence.org
[11] https://allenai.org
[12] https://partnershiponai.org
[13] https://cognitivecomputingconsortium.com
[14] https://makingaisafer.org
[15] http://www.aisociety.life/
[16] Such as for example https://www.ibm.com/watson/ai-ethics/

3.4 Preserving Competitiveness – AI Platforms vs. AI Applications

To properly focus AI-related standardization efforts, we must distinguish between platforms and applications.

- Common foundations for AI, such as the representation of neural network building blocks, APIs for knowledge bases and ontologies, or libraries of natural language processing primitives, will accelerate development. Instead of users and tool vendors wasting time supporting multiple APIs for these common capabilities, they will be able to use standards-based tools of commercial or open source origin.
- Value-adding applications of AI should—and will—remain an open field for worldwide innovation and competitiveness, with more time and resources devoted to this level once the foundations and associated tools have become more easily available.

4 Why OMG

4.1 OMG's Proven Capabilities

Readers who are not familiar with the Object Management Group® (OMG®) should refer to our website, www.omg.org, for more information, including OMG's history since its foundation [10], its standards development process, lists of our 245+ members and 225+ adopted specifications, and the organizations with which we maintain liaison relationships, including ISO (see https://www.omg.org/about/liaison.htm).

In its thirty-year history, OMG has shown its ability to expand to new areas of concern. From its beginning in object-oriented middleware—with CORBA® and related object services—OMG transitioned to Model Driven Architecture (MDA), to cloud computing, to software modernization and quality, to knowledge representations and reasoning, and to standards supporting the Industrial Internet of Things.

OMG recognized early that horizontal or industry-generic standards were not sufficient to help users, but that specific industries required their own standards. The organization adapted to serve this need, pursuing both "platform" and "domain" standards through separate task forces and technical committees. Over the years, the list of domains addressed by OMG has evolved to stay aligned with industry needs. It now includes **finance, healthcare, manufacturing, C4I** (Command, Control, Communications, Computers, and Intelligence)**, robotics, space,** and **retail.** All those areas are being transformed by AI and will benefit from the development of AI-related standards.

4.2 Existing OMG Standards in Support of AI

OMG has developed, and is still developing, several AI-related cross-domain specifications, in particular in the areas of knowledge representation and reasoning (KR&R). The complete list of these specifications appears in Appendix B.

We expect more AI-related specifications to appear not only in current areas of work (such as retail and robotics), but in other domains such as finance, space, C4I, and more.

4.3 OMG's Policies and Processes

Many aspects of OMG's policies, processes and procedures give the organization a significant advantage when addressing a new area. Below is a list of those strengths.

Open Process. OMG's standards development process has matured over the years, and has been applied to 225+ specifications. It is recognized by ISO as being sufficiently rigorous and disciplined as to qualify OMG for participation in the ISO Publicly Available Specification (PAS) and FastTrack programs. As a result, many OMG standards have become ISO standards.

The open process is implemented through a set of subgroups (Task Forces and Special Interest Groups) open to all members, and through facilities such as wikis, a Jira® issues database, etc. The practices employed to execute this process include the issuance of discussion papers, Requests for Information (RFI), Requests for Proposals (RFPs), or the adoption of Requests for Comments (RFCs).

The OMG's policies and procedures are publicly available, even to non-members, at https://www.omg.org/cgi-bin/doc.cgi?pp.

One Vote per Member. OMG is open to organizations of all sizes. Each organization has one vote in the subgroups to which it contributes. Government and academia also participate in the process under the same terms – this is important since researchers in universities or national laboratories often do the leading-edge work in areas such as AI.

Moreover, an OMG member can contribute to a standard without being accredited by a national standards organization, as is the case for example in ISO.

Free Specifications. OMG's approved specifications are available to the public free of charge.

Simple Intellectual Property (IP) policies. During the development of a specification, an IP mode is selected. Almost all OMG standards are available under a royalty-free or "non-assert" mode, where the holders of any precursor IP agree to give other companies the right to use the IP they contributed.

No shelfware. A specification is only deemed final after evidence is provided that there are existing implementations of the proposed standard—whether they are commercial offerings, open-source versions, or internal implementations.

Living Standards. There is a formal change process to ensure that specifications do not become stale. As soon as a specification is adopted, a task force is formed to address any issues of interpretation or implementation raised by the public, or address new requirements within its scope. When a revised version is published, a new such task force is formed.

Architectural compatibility. All requirements documents and specifications must be approved by the OMG Architecture Board, which is elected by the members. This ensures compatibility and coherence across OMG's set of specifications.

International applicability. OMG is an international organization and all its deliverables are equally available worldwide.

5 A Roadmap for AI Standards

This section outlines actions that OMG envisions to develop standards for AI. It should be noted that this is a roadmap, not a detailed plan, and that OMG does not dictate actions to its members. It is OMG members, through their participation, who will define or refine the actual roadmap and transform it into specific actions according to our process.

First, an **AI Reference Model**—similar to the seminal work that NIST performed to create its widely recognized Cloud Computing Reference Architecture [11]—would be useful to categorize cross-domain vs. domain-specific capabilities, platforms and tools.

Once this Reference Model is agreed upon, each type of organization (AI suppliers, AI users, government entities, etc.) can determine which part of the model their initiatives will address. The Reference Model could distinguish:

- Technology building blocks that address AI-specific as well as other needs.
- Knowledge representation technologies (semantic web, ontologies, rules…).
- AI capabilities (neural networks and other forms of machine learning, pattern recognition, planning, etc.).
- Cross-sector AI applications (such as facial recognition).
- Sector-specific AI applications (e.g., medical diagnostic)

Table 2 below lists some of these potential areas of standardization.

Table 2 – AI Standard Areas

AI Domains	Future Standard Areas	Goals and Benefits of Standards
AI Architecture and Logical Components	Architecture or reference model to categorize cross-domain vs. domain-specific capabilities, platforms and tools.	Achieve a common understanding and make it easier for participants to contribute to parts of the architecture.
Machine Learning	Training data set representation and metadata	Allow the sharing of data that can be used to train models. While the models may be proprietary, the data sets (e.g., anonymized equipment or patient data) could, if associated with appropriate metadata and shared in a standard format, accelerate the improvement of the models.
	IoT interoperability language	Allow vendors and users, especially in critical industries with a potential impact on the public and the environment, to combine datasets or share operational rules to improve safety or lower costs.
	Machine learning decision explanation model.	Address a growing demand for the ability to "audit" how a neural network arrived at a certain conclusion. There are technical, legal, regulatory and ethical reasons why the ability to explain the decision may be required.
Cognitive Services	Standard APIs for access to Ai algorithms in vision, speech recognition, language understanding, intelligent search and more.	Allow users to substitute components from multiple suppliers providing those services, without impeding the competition between those suppliers.
Facial and biometrics recognition	Protocols, APIs, encryption, access rules.	Provide a secure and traceable way for justified access while protecting personal data against unintended use through encryption or obfuscation techniques.
Speech Recognition	Evaluation metrics, test sets, evaluation methodology, APIs.	A published standard would make the test methodology, metrics, and test sets available to all developers of speech/language recognition systems.

AI Domains	Future Standard Areas	Goals and Benefits of Standards
Smart Robotics	Standardized planning language.	Enhance the ability to replace a component of a robotic system with another one by making the output of planning software transferable from one brand of robot to another.
Natural Language Processing	Information classification and rule representation for automatic message processing by intelligent agents.	Help solve the information overflow problem (the challenge of processing the mass of data received by humans on a daily basis) by providing a common representation of the non-confidential content of messages, which will allow machine learning-based intelligent assistants and spam filters.
Agents	Agent modeling languages and techniques. Alert and notification interfaces.	Increase the rigor and consistency of agent-related specifications, and ensure interoperability of agent-based systems.
Augmented Reality	Content markup and management, object identification, navigation.	Ensure that AR reaches its full potential as an enhancement to human life and information use.
Sector-Specific Information Models and Decision Models	Rules and decision models that leverage, but go beyond, the work already done on sector-specific ontologies.	Enable various levels of reasoning and automation, as appropriate for each sector, through the ability to interchange rules and decision models that are used by AI applications.
Security of AI Components	Authentication, authorization and access control for AI components.	Practices that relate to the security of the AI components of IoT and other systems, for example to prevent the injection of illegitimate data into the training of a machine learning algorithm.
AI Ethics	Reference architecture for confidentiality, privacy and ethical decision-making in AI.	Diminish the risk of accidents or social rejection by provide guidance to developers and users so that AI can be applied responsibly, ethically and legally.

6 Next Steps and Call to Action

6.1 Formation of an AI-Specific Subgroup at OMG

Having addressed a "slice" of AI through its Agent Platform Special Interest Group, OMG recently broadened its involvement by chartering an AI Platform Special Interest Group in June 2019. The creation of an AI Platform Task Force is currently being undertaken in order to enable the authorship of specifications, thus accelerating the solicitation and adoption of AI standards.

Subgroups (SIGs or Task Forces) do not work in isolation within OMG. They routinely confer and collaborate on common interests. The AI Task Force will influence work done in other Task Forces such as Healthcare, Finance, Manufacturing, Retail, Robotics, etc.

6.2 External Liaisons and Collaborations

OMG will continue to leverage its existing liaison agreements with other standards organizations and associations[17], and develop new ones as appropriate, including the IEEE Society on Social Implications of Information Technology's Standards Committee (IEEE-SSIT SC)[18] or the Augmented Reality for Enterprise Alliance (AREA)[19].

In May 2019, OMG responded to NIST's RFI mentioned earlier [5], aimed at developing a U.S. "AI standards engagement plan." While NIST's effort is U.S.-specific and OMG is international in scope, our response [7] can lead to a fruitful collaboration. In particular, OMG suggested that NIST:

- Develop the **reference model** mentioned in Section 5 above.
- Sponsor some of the incipient work to improve the reliability, maintainability, reusability, transparency, respect for privacy, etc., of **deep learning algorithms** through standards for combining declarative ontologies and rules with the systems they specify.
- Provide additional funding for NIST's **Industrial Ontology Foundry** and similar projects to accelerate their progress.

OMG is ready to discuss collaborating with any of these organizations to advance our common interests.

[17] https://www.omg.org/about/liaison.htm)
[18] http://sites.ieee.org/sagroups-ssit/
[19] https://thearea.org/

6.3 Call for Participation

The vision expressed in this paper remains a *potential* roadmap until realized as OMG members "roll up their sleeves" and participate in the activity. We call on participants in the AI community to provide input from as many sources as possible, discuss and improve the roadmap, and help select the top priority areas. While we will use and analyze the responses to the NIST RFI, a logical vehicle for additional information gathering is for OMG to issue its own RFI, in particular to involve international organizations. This will be the first step in applying OMG's open process to start developing appropriate standards.

We invite any whose concerns and interests have been touched upon in this paper to join us as we move forward. Concrete steps include:

- Write to ai-chair@omg.org to express your interest and provide feedback about this paper as well as the specifications mentioned in it, additional ideas that will influence OMG's AI roadmap, and to inform OMG of other relevant efforts.
- Visit the AI subgroup's wiki at https://www.omgwiki.org/AI.
- Attend meetings organized by OMG's AI subgroup (to be posted at https://www.omg.org/events).
- Inform OMG of other relevant efforts.
- Consider an OMG membership in order to have a real impact (including voting rights) on OMG's work.

6.4 AI Forums and Special Events

OMG has a track record of holding special events (which may be called forums, summits, symposia, or workshops) to address the needs for standards in specific sectors. Since one of our quarterly meetings takes place each year in March in the Washington, D.C. area, we plan to hold an annual full-day event on AI standards, jointly with other interested organizations, starting in March 2020. This first AI Workshop would have the following goals:

- Present analyses by various organizations of the needs for AI standards.
- Expose participants to the relevance of existing OMG standards to the building of AI capabilities.
- Provide an opportunity for industry, government and academia participants—not only OMG members—to exchange ideas.
- Attract those players who want to participate in the standards development process to become members of OMG in order to pursue this work.

"Challenge" events (defined in a way similar to DARPA, see https://www.darpa.mil/work-with-us/public/prizes) are a good way to invite the providers of technology to demonstrate their capabilities, including standards-based interoperability. To promote several of our standards such as UML and BPMN, OMG routinely hosts interoperability demonstrations during our meetings. We propose holding, jointly with other interested organizations, "AI interoperability challenges." These may leverage OMG standards once developed, or could equally be used to expose areas where standards are needed.

6.5 Toward an AI Standards Council?

OMG has the capacity and motivation to create an **AI Standards Council,** bringing together representatives from industry, academia and other government agencies, with a mission similar to those of other OMG managed programs such as the Industrial Internet Consortium (IIC)[20] or the Consortium for Information and Software Quality (CISQ)[21]. This would represent the next level of activity and visibility.

A separate council requires significant levels of resources, provided or funded by a slate of sponsoring organizations. On the other hand, it provides a locus for a number of activities that are complementary to the development of standards, such as:

- Advocacy—promoting the application of AI to various domains.
- Education and Marketing—helping improve the understanding of AI capabilities, best practices among end users, practical guides to getting started, success stories, etc.
- Demonstrations— "laboratory" work or "testbeds" developed as collaboration between members.
- Liaison—connecting with industry bodies that are not involved in standards as well as with standards organization to which requirements for new capabilities, expressed by members, can be sent for action.
- Certification—providing evidence of the qualification of member companies' employees who demonstrate higher levels of proficiency.

6.6 Conclusion

AI has matured and its successful application can be enhanced by the development and adoption of standards. OMG has the capability and motivation to successfully expand its activities in this domain. We encourage the AI community—across all domains and regions—to get involved in this effort by contacting us, participating in our AI-related activities and events, and joining OMG to take an active role.

[20] https://www.iiconsortium.org
[21] https://www.it-cisq.org/

Appendix A: AAAI Taxonomy of Artificial Intelligence

This *de facto* taxonomy of AI material (articles, papers, books…) is extracted from the search menu of *AITopics,* an official publication of the Association for the Advancement of Artificial Intelligence (AAAI), at www.aitopics.org/search. It is provided here solely as an example of such taxonomies. Some of the subcategories include a third level of detail; we only included the first two levels here. For the complete taxonomy, see the source.

Artificial Intelligence

- **Assistive Technologies**
- **Challenges**
- **Cognitive Science**
  - Childhood Development
  - Cognitive Architectures
  - Creativity and Intelligence
  - Emotion
  - Neuroscience
  - Problem Solving
  - Simulation of Human Behavior
- **Games**
  - *[list of games omitted]*
- **History**
- **Human-Centered Computing**
- **Issues**
  - Arguments Against AI
  - Philosophy
  - Social and Ethical Issues
  - Turing's Test
- **Machine Learning**
  - Association Learning
  - Bayesian Networks
  - Computational Learning Theory
  - Control Theory
  - Decision Tree Learning
  - Ensemble Learning
  - Evolutionary Systems
  - Forecasting
  - Fuzzy Control
  - Inductive Learning
  - Kernel Methods
  - Learning Graphical Models
  - Learning in High Dimensional Spaces
  - Memory-Based Learning
  - Neural Networks
  - Pattern Recognition
  - Performance Analysis
  - Reinforcement Learning
  - Statistical Learning
  - Supervised Learning
  - Transfer Learning
  - Unsupervised or Indirectly Supervised Learning
- **Natural Language**
  - Chatbots
  - Discourse and Dialogue
  - Explanation and Argumentation
  - Generation
  - Grammars and Parsing
  - Information Extraction
  - Information Retrieval
  - Machine Storytelling
  - Machine Translation
  - Question Answering
  - Text Classification
  - Text Processing
  - Understanding

- **Representation and Reasoning**
    - Abductive Reasoning
    - Agents
    - Analogical Reasoning
    - Automatic Programming
    - Belief Revision
    - Blackboard Systems
    - Case-Based Reasoning
    - Commonsense Reasoning
    - Constraint-Based Reasoning
    - Description Logic
    - Diagnosis
    - Diagrams and Models
    - Expert Systems
    - Frame-Oriented Architecture
    - Information Fusion
    - Logic and Formal Reasoning
    - Mathematical and Statistical Methods
    - Metareasoning
    - Model-Based Reasoning
    - Nonmonotonic Logic
    - Object-Oriented Architecture
    - Ontologies
    - Optimization
    - Personal Assistant Systems
    - Planning and Scheduling
    - Qualitative Reasoning
    - Rule-Based Reasoning
    - Scientific Discovery
    - Scripts and Frames
    - Search
    - Semantic Networks
    - Spatial Reasoning
    - Temporal Reasoning
    - Uncertainty
- **Robots**
    - Autonomous Vehicles
    - Humanoid Robots
    - Locomotion
    - Manipulation
    - Robot Planning and Action
    - Robots in the Home
    - Robots in the Workplace
    - Soccer Robots
- **Science Fiction**
- **Speech**
    - Acoustic Processing
    - Speech Recognition
    - Speech Synthesis
- **Systems and Languages**
    - Distributed Architectures
    - Problem-Independent Architectures
    - Problem-Specific Architectures
    - Programming Languages
- **The Future**
- **Vision**
    - Face Recognition
    - Gesture Recognition
    - Handwriting Recognition
    - Image Understanding
    - Optical Character Recognition
    - Sketch Understanding
    - Video Understanding

Appendix B: AI-Related OMG Standards

The following list is an Appendix to section 4.2 of this paper. It lists OMG specifications that are either published or under development, and that establish standard foundations required by AI platforms and applications.

- General knowledge representation and reasoning (KR&R) standards already published:
    - **Ontology Definition Metamodel (ODM)** enables ontology management and development using OMG's Model Driven Architecture (MDA) stack for the Resource Description Framework (RDF) and RDF Schema, the Web Ontology Language (OWL), ISO Common Logic (CL), and Topic Maps. The specifications, and hence tools, in the stack provide support for metamodel and model storage, versioning, querying and transformation. ODM also provides a profile for use of UML-compliant tooling for graphical modeling of ontologies. [12]
    - **Distributed Ontology, Modeling, and Specification Language (DOL)** provides a language and transformations at the semantic level aimed at achieving integration and interoperability of ontologies, specifications and models developed independently and in differing ontology languages and logic frameworks. [13]
    - **Semantics of Business Vocabularies and Rules (SBVR)** enables a structured English representation and the interchange of business statements. [14]
    - **Decision Modeling and Notation (DMN)** enables the executable representation of business decisions linked to their data sources; and management of the rules in the business context. [15]
- Non-interface-oriented robotics standards (i.e., those that enable knowledge interchange or other interactions rather than strictly providing interfaces):
    - **Robotic Technology Component (RTC)** defines a component model and certain infrastructure services supporting robotics software development.
    - **Finite State Machine Component for RTC (FSM4RTC)** extends the RTC specification for interchange of state and state machine related content
    - **Robotic Interaction Service Framework (RoIS)** defines a framework for services supporting interactions between humans and robots, including but not limited to facial detection and identification, sound detection, language recognition and understanding, speech generation, interpretation of gestures, and the like.

- OMG has also published sector-specific standards in areas including knowledge representation and reasoning for Finance. Our **Financial Industry Business Ontology (FIBO)** standard, jointly developed and evolving through our liaison with the Enterprise Data Management Council, provides an ontology for legal entities, financial instruments and related concepts as well as reference data for the representation of currencies, various banking identifiers, legal entity identifiers, market identifiers, and so forth. [16]
- Many other OMG standards provide the supporting infrastructure that enable the development of software, systems, and interfaces that include AI as a component.

The above-mentioned OMG specifications (in bold characters) can all be found on the OMG website at https://www.omg.org/spec/.

- General KR&R standards under development:
    - **Application Programming Interfaces for Knowledge Platforms (API4KP)** defines a set of ontologies and interfaces needed to incorporate knowledge representation and reasoning tools, as well as other AI capabilities, in a broader enterprise environment, including but not limited to interfaces between inference engines, rule engines, knowledge graphs, and various sources of information required to build out a comprehensive environment. [17]

      Initial reference implementations have been deployed at the Mayo Clinic and are under development at the Veterans' Administration.
- General robotics standards under development:
    - **Robotics Service Ontology (RoSO)** defines a set of ontologies for robot-to-human interactions and the services needed to perform such interactions. This work is being done in conjunction with the IEEE Robotics & Automation Society's Autonomous Robotics Group, which is responsible for IEEE 1872 – an ontology that focuses on core (generic) terminology and capabilities of robot systems at a relatively high level.
- OMG's sector-specific AI-related specifications under development:
    - A retail specification for digital receipts that embodies an ontology defining not only the receipts themselves, but also content related to jurisdiction-specific taxation.
    - A joint effort between OMG's Retail and Robotics Task Forces to create a standard for point-of-sale/point-of-service (POS) robotic interfaces for the 2020 Olympics specifically, but which will be broadly applicable to POS robotic services.

Appendix C: References

[1] Association for the Advancement of Artificial Intelligence (AAAI): a taxonomy of AI. www.aitopics.org/search (see the "Technology" drop-down list in the left column).

[2] Golstein, B. (2018): *A Brief Taxonomy of AI.* White paper by Sharper AI Pte Ltd. http://www.sharper.ai/taxonomy-ai/

[3] Johnson, L. (2004): *A False Sense of Proprietary.* The Standards Edge: Dynamic Tension, Chapter 17. Bolin (ed.), Sheridan Books. https://www.thebolingroup.com/standards_series.html

[4] Smith, C., McGuire, B., Huang, T., & Yang, J. (2006): *The History of Artificial Intelligence.* University of Washington materials for course CSED 590A. https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf

[5] National Institute of Standards and Technology (2019): *Request for Information on Artificial Intelligence Technical Standards and Tools.* https://www.federalregister.gov/documents/2019/05/01/2019-08818/artificial-intelligence-standards

[6] National Institute of Standards and Technology (2019): *Comments Received for RFI about Federal Engagement in Artificial Intelligence Standards.* https://www.nist.gov/topics/artificial-intelligence/comments-received-rfi-about-federal-engagement-artificial

[7] Object Management Group (2019): Response to the NIST RFI on Artificial Intelligence Technical Standards and Tools. https://www.nist.gov/document/nist-ai-rfi-omg-001pdf

[8] Antoniou, Grigoris, et al. (2005): *Combining Rules and Ontologies. A Survey.* REWERSE. http://rewerse.net/deliverables/m12/i3-d3.pdf

[9] Xu, Li, et al.: *Combining Declarative and Procedural Knowledge to Automate and Represent Ontology Mapping.* https://www.deg.byu.edu/papers/SWAT06-131.pdf

[10] OMG (1989): *Object Management Group Established.* Press release. https://www.omg.org/marketing/25th/OMG_Forms_PR_1989.pdf

[11] National Institute of Standards and Technology (2011): *NIST Cloud Computing Reference Architecture.* Special Publication 500-292. https://doi.org/10.6028/NIST.SP.500-292

[12] OMG: *Ontology Definition Metamodel™ (ODM™)* https://www.omg.org/odm/

[13] OMG: *Distributed Ontology, Model, and Specification Language™ (DOL™).* https://bei.omg.org/dol/

[14] OMG: *About the Semantics of Business Vocabulary and Rules Specification Version 1.4.* https://www.omg.org/spec/SBVR/

[15] OMG: *Decision Model and Notation™ (DMN™).* https://www.omg.org/dmn/

[16] OMG: *Financial Services Standards.* https://www.omg.org/hot-topics/finance.htm

[17] Athan, T., Bell, R., Kendall, E., Paschke, A., & Sottara, D. (2015): *API4KP Metamodel: A Meta-API for Heterogeneous Knowledge Platforms.* Conference paper, International Symposium on Rules and Rule Markup Languages for the Semantic Web. Springer. https://link.springer.com/chapter/10.1007/978-3-319-21542-6_10

