

Q1 2023 AI PTF Teleconference Notes

A set of five teleconferences were organized in the first quarter of 2023 to make progress, with the following suggested topics:

- AI taxonomy
 - A model of AI risk
 - A synthesis paper of what the AI PTF has done since its inception
-

16 January 2023 Meeting

This meeting was inadvertently convened on a U.S. holiday, preventing at least one person from attending. The participants were:

- Claude Baudoin (co-chair)
- Ademola Adejokun (Lockheed Martin)
- Nick Stavros (Jackrabbit Consulting)
- Karl Gosejacob (GOSEJACOB)
- Mike Abramson (present but did not intervene)

Taxonomy Content

The Dokuwiki implementation of the current draft taxonomy within the private pages of this AI PTF wiki was demonstrated to Ademola, who expressed interest in working on this. Claude also showed the current state of the IEEE P3123 working group's AI terminology and Ademola asked for access. Claude e-mailed the Secretary of the P3123 WG to request those privileges for Ademola.

Taxonomy Representation

We discussed (in her absence) Elisa's intent to use this taxonomy as a use case for the new Multiple Vocabulary Facility (MVF) specification. Nick said that he would like to know what the desired format is, so he might write the JavaScript code to export the wiki content as an MVF file.

We also discussed the need for users who are not ontology/taxonomy experts to edit the taxonomy. This was partially based on finding that moving a taxon from one place to another is not obvious, even in the wiki. Nick should create a simple explanation of how to do that in the "Administrivia" section of the taxonomy pages.

Claude reminded others that he had created a [graphical visualization](#) of an earlier version of the taxonomy (shown as a hyperbolic tree by the Unilexicon tool), and said he could revive and update it, but that tool only has limited capabilities. This led Nick to suggest that we should capture the requirements for a good vocabulary management tool. Claude said that this is an interesting idea, but it is not specific to AI so that it would be a deliverable (discussion paper?) for the Ontology PSIG.

Risk Model

Karl said that he is interested in working on an AI risk model. This work should take into account the classification contained in the draft EU AI Act, as well as the NIST AI Risk Management Framework.

(Note that IEEE had started a study group in Q3 2022 about “AI implementation risk tiers” to respond to the NIST framework, but that effort was shut down because the study group was too small, couldn't agree on the scope, or even on what “implementation risk tiers” really meant. While this should serve as a cautionary note, it also means that there is no overlap with or competition from IEEE if we do this.)

30 January 2023 Meeting

The participants were:

- Claude Baudoin (co-chair)
- Ademola Adejokun (Lockheed Martin)
- Nick Stavros (Jackrabbit Consulting)
- Karl Gosejacob (GOSEJACOB)
- Alan Johnston (MIMOSA)
- Arnaud Billion (IBM France)
- Elisa Kendall (Thematrix Partners)

We went through a round of introductions since we had a couple of new participants.

Arnaud Billion ([LinkedIn](#)) is a PhD in Intellectual Property Law who is an ethics advisor at IBM France. He also collaborates with [Responsible Computing](#), an initiative launched by IBM Germany which became a managed program of OMG in 2021. Some of his work bears on whether the productions of AI systems should be copyrighted, while another area of work is captured in his book “Governed by Machines” (this refers to our reliance on prescriptive rules of governance, as opposed to natural law, not to computing machines).

Discussion on Legal Issues

Arnaud's introduction led to a discussion with Karl, who commented on his work on the IP aspects of AI, including the difference between EU and US laws, or the fact that in most countries a copyrighted work can only be licensed, not sold: the author remains the owner of the work. An exception is Switzerland, where the copyright to a work can be sold.

Arnaud thinks that OMG can play a role in “explaining to lawyers and the outside world that AI is just another form of data transformation.”

Alan said that he is regularly meeting with academics who have projects in data science and analytics that “creep” into the AI area, and they increasingly ask what can be patented. The university patent attorneys, whose job is to constantly watch for inventions that can be protected and monetized, are struggling to find answers. This input, combined with Arnaud's and Karl's work, supports the idea of

developing an OMG discussion paper on “AI and Intellectual Property” or a similar title.

Taxonomy

Claude reported that in his work with the IEEE P3123 Working Group, he added a taxonomy dimension to the terminology table under development. This simply consisted of adding a column called “Kind of” and, for now, making all forms of machine learning (supervised, unsupervised, transfer, etc.) “kinds of” machine learning – a term that was missing and that he added, together with a definition from the author who apparently invented the term.

Claude reported, before Nick Stavros joined the meeting, that work is going on to redesign the OMG wikis so that (a) there is a single installation of the Dokuwiki engine, (b) OMG members can sign in using their OMG credentials instead of us having to issue separately managed credentials by hand.

Elisa said that she wants examples of terms in order to apply the Multiple Vocabulary Facility (MVF) to them, after which she can demonstrate this. Claude said that in order to finalize enough terms, he needs to resume work on a long list of action items he has.

AI Risk and Trustworthiness

The discussion on taxonomy veered toward a discussion of risk, and this in turn led to AI trustworthiness. Claude showed a report which he just described a couple days earlier in an e-mail to the AI mailing list:

Thanks to Clayton Pummill, who briefly attended our meetings when he was with Torch.ai and alerted me to this, I just looked through a white paper written by Jessica Newman, from UC Berkeley’s Center for Long-Term Cybersecurity (CLTC), which adds an extra dimension to the NIST AI Risk Management Framework.

The report is entitled [A Taxonomy of Trustworthiness for Artificial Intelligence](#) and subtitled “Connecting Properties of Trustworthiness with Risk Management and the AI Lifecycle” (no paywall, no signing in – how refreshing!).

As the subtitle indicates, the report creates a mapping between the concepts of the NIST AI RMF, in particular the lifecycle stages it defines (Plan and Design, Collect and Process Data, Build and Use Model, Verify and Validate, Deploy and Use, Operate and Monitor, Use or Impacted By) and the “characteristics of trustworthiness” (valid and reliable, safe, fair, secure and resilient, explainable and interpretable, privacy-enhanced, accountable and transparent, responsible practice and use). If you can imagine the resulting matrix of 7 stages by 8 characteristics, the author then goes on to define a set of properties within each cell of this matrix – sometimes just one property, often two to four, in one case 26 of them – for a grand total of 150 distinct properties.

The report also lists (and used as inputs) a number of existing frameworks for AI trustworthiness, and specifically highlights these:

- * The “Ethics Guidelines for Trustworthy AI” from the High-Level Expert Group on Artificial Intelligence
- * The EU AI Act, which we’ve discussed several times in our OMG AI PTF meetings
- * The White House Blueprint for an AI Bill of Rights
- * The NIST AI Risk Management Framework

This is not for the faint of heart (78 pages, 2 appendices, 69 footnotes...) but seems to be a really important piece of work for people interested in AI ethics and responsible computing.

13 February 2023 Meeting

The participants were:

- Claude Baudoin (co-chair)
- Ademola Adejokun (Lockheed Martin)
- Jürgen Boldt (OMG)
- Karl Gosejacob (GOSEJACOB)
- Elisa Kendall (Thematrix)

This meeting was basically a review of ongoing action items.

Claude reported that he sent a message to the UC Berkeley Center for Long-Term Cybersecurity (CLTC) in order to invite Jessica Newman, the author of the report on a “taxonomy of AI trustworthiness” discussed last time, to speak at our March 21 meeting.

The wiki still needs work. There are four past meetings (March and September 2020 and 2021) that have not been documented, and this is still needed, even though this is in the relatively distant past, if we are to at any point reconstruct a full history of our deliberations over the years.

Regarding the AI terminology work being done in parallel by the IEEE P3123 working group, Claude reported that:

- Someone has been adding terms and relations in the IEEE spreadsheet, but some of the relations are not purely “is a” relations. These relations may be useful to capture, but they are slowly moving the project past a taxonomy and toward a more complex conceptual model, potentially a knowledge graph or an ontology.
- The next meeting of the IEEE group is tomorrow, February 14, at 7:30 a.m. PST. Ademola said he had not received the invitation, which Claude then forwarded to him.

Regarding the OMG taxonomy:

- Nick Stavros was not present to give us an update.
- We briefly discussed (mostly for Jürgen's benefit) the desire to give people access to private wiki areas via single sign-on, based on their OMG credentials, instead of manually administering separate credentials.
- Claude said that if Elisa wants an example to test MVF, she could perhaps use the Industry IoT Consortium's vocabulary instead of waiting for the AI terminology to be complete enough. Claude sent her the official vocabulary report (PDF file) as well as a SKOS file created by Erin Bournival of Dell, who chairs the IIC Vocabulary Task Group.

Regarding legal issues (especially intellectual property) related to AI, Karl said that he would draft an invitation, with attached documents, to be sent to Arnaud Billion and Alan Johnston in order to initiate the collaboration mentioned at the previous meeting. Claude will forward that message.

Karl also mentioned that Oliver Klein, who talked about the EU AI Act at a couple of past AI PTF

meetings, is now with United Internet AG, and that he would probably have some news to share at the next meeting. Karl will lay the groundwork for an invitation to speak again on March 21.

As a result, Claude updated the wiki page about potential future speakers, and in doing so he noticed that we had not followed up last year with Mary Armijo (FACE Consortium) or Simon Mettrick (BAE Systems). So he sent them messages to ask about their interest to speak next month.

Elisa said that Evan Wallace (NIST) has proposed one talk that might fit better on the AI PTF agenda than on the ontology PSIG's. She will dig up the e-mail in question and forward it to Claude, who has not seen it.

27 February 2023 Meeting

The participants were:

- Claude Baudoin (co-chair)
- Ademola Adejokun (Lockheed Martin)
- Karl Gosejacob (GOSEJACOB)
- Elisa Kendall (Thematrix)
- Bobbin Teegarden (OntoAge)

Alan Johnston (MIMOSA) sent apologies due to family reasons.

Claude reviewed the notes from Feb. 13 and gave updates on the progress of several items:

- Elisa and he talked to Jessica Newman (UC Berkeley CLTC) and she is now a confirmed speaker on March 21
- He talked to the NIST people (Milos and Serm) and Milos will talk about their ML Ontology for Industrial Applications on March 21
- IEEE P3123 continues to add to its terminology. Angie Qarry has suggested adding some relations between terms besides "kind of" relations, so the work is slowly creeping toward being more than a taxonomy. In ANSI Z39.19 terms, it would be called a thesaurus; in our terms it might be called an ontology.
- Claude sent Elisa the IIC Vocabulary in PDF form, as well as a SKOS file created by Erin Bournival (Dell) using a program she wrote. Another person at IIC is looking into how to do the forward generation in the future: SKOS -> XML -> Word -> PDF (which is what should have been done from the beginning).
- Claude asked Oliver Klein if he would speak about the EU AI Act on March 21, but Oliver's work at his new organization is related to the cloud, not to AI. Oliver recommended Christian Rudelt.
- The tentative [agenda of the March 21 meeting](#) was prepared accordingly (with typos that have since been fixed).

When Bobbin Teegarden joined the meeting, we discussed the co-chair situation. Bobbin will send Mike Bennett (OMG Technical Director) a note to relinquish her position. Claude will ask Davide Sottara (Mayo Clinic) whether he has thought more about the possibility of taking on this role, or whether someone else at the Mayo Clinic might be interested.

Regarding the agenda of the March 21 meeting, Claude asked what is the presentation that Evan Wallace (NIST) had mentioned as possibly fitting more in the AI PTF agenda than the Ontology PSIG's.

Elisa believes that it is precisely about the ML Ontology that Milos Drobnjakovic is planning to present. Claude will ask Evan if we want to make this presentation a joint session (or if we just informally invite the OntoPSIG members to attend this talk).

Claude will also find out who is in charge of the Responsible Computing consortium; even though they don't seem to be meeting during OMG TC week, their members should be interested in Jessica Newman's talk on AI Trustworthiness.

Karl Gosejacob offered to contact Christian Rudelt at BDI to ask him to present an update on the EU AI Act.

Recap of Action Items

Several action items are mentioned in the above text.

- Karl: contact Christian Rudelt about giving us another update on the EU AI Act
 - Karl: propose collaboration with Alan and Arnaud on a discussion paper on “AI and Intellectual Property”
 - Claude: document the meetings (this page, done)
 - Claude: add definitions in the OMG AI taxonomy (see table of action items generated by Nick)
 - Claude: finalize the notes from the March 2020, Sept. 2020, March 2021, and Sep. 2021 meetings (Google Docs)
 - Claude: update the Unilexicon visual taxonomy to reflect the current state of our taxonomy
 - Claude *and others*: brainstorm and capture requirements for a vocabulary management tool, submit to Ontology PSIG
 - Nick: document how to move a taxon from one part of the tree to another
 - Karl: start an AI Risk [meta?]model
 - Claude: structure the wiki pages better, as the current page is starting to be too big. Create subpages for the reference architecture, the taxonomy, the risk model when it exists, and other future topics.
-

The next meeting is the plenary AI PTF meeting on March 21 (see agenda at the above link).

From:
<https://www.omgwiki.org/ai/> - **Artificial Intelligence PTF**

Permanent link:
https://www.omgwiki.org/ai/doku.php?id=q1_2023_ai_ptf_teleconference_notes

Last update: **2023/02/27 20:13**

