

**IEEE P7003 Working Group
Meeting Minutes
2nd August 2018 /3:00 P.M. – 4:30 P.M. UTC
Teleconference**

1. Call to Order

The meeting was called to order at 15:05 UTC
A quorum has been established and noted.

2. Roll call and Disclosure of [Affiliation](#)

Establishing a quorum, group membership, and iMeet roster: Liz informed people of the rules. It is very important that we record whether or not a quorum has been achieved in each meeting; a quorum is 1/5 of voting members, or 3 voting members, whichever is greater. Voting members are those that have attended 2 consecutive meetings, gaining voting rights from their third meeting. Members lose voting rights after missing 3 consecutive meetings but can still participate in meetings. The roster is now kept in a read-only form on iMeet and it is very important that affiliations and email addresses are kept up-to-date. Voting membership will become more important when it comes to finalizing documents for the standard.

The list of attendees present is attached.

3. Approval of August Agenda

Motion to approve the meeting agenda from 2nd August 2018. The agenda was approved as submitted without objection.

4. [IEEE Patent Policy](#) (Call for Patents)

The call for patents was raised; no one raised any concerns or any comments for consideration.

5. Approval of 5 July meeting minutes

Motion to approve the minutes from July 5th, 2018. It was noted that the comments under point II.vii and II.ix should be swapped. This has been done and the minutes approved.

6. Announcements

Timeline for preliminary document integration:

September: prepare section drafts for integration by end of month

All sections will have a preliminary version added into a single outline document. This document will give a better overview of how well the

different sections are progressing, what gaps there are, and incompatibilities between sections.

October: check overall consistency and gaps of preliminary document

Everyone can add comments to document, then in January have a face-to-face meeting to resolve comments. (Ansgar attended P7001 meeting earlier this week). *From chat: Hans Jessen queried how much funding we would be looking for.*

Action: Ansgar to check procedures for setting up meetings, funds etc.

Note: this is intended as interim check on document coherence

September integration does not need completed section, but useful to integrate current state to check connections.

7. Accenture tool presentation

Rumman will provide a 15 min presentation of the Accenture "fairness tool", followed by Q&A

The toolkit was inspired by the FAT* conference talk on 21 definitions of fairness. It attempts to create predictive parity in algorithms, and has three parts, one of which concerns data and the other two the model.

Mutual information analysis – user selected sensitive variables, and their level of influence over non sensitive variables. Currently working on different levels of user control

Disparate impact – looks at error rates and ways to repair clustering of variables ie bias. Trade-off between repair level and level of global predictive accuracy

Predictive parity – is the model acting differently to different subgroups?

Hoping for results from working with clients as benchmark analysis.

Discussion: Matthew S: How does the tool cope with new categorisations? For example, the definitions of gender are expanding. How would the tool deal with comparing historical data in which there are only 2 categories (eg male/female) with newer data with extra categories (male/female/other, male/female/non-binary etc)

Will the quantification through financial costs and benefits mean companies try to de-bias their systems only to a certain dollar amount?

Mark U: Perhaps use cases could be an answer to identify problem rather than a toolkit that tells you how to solve it. Felt that the framework is data science friendly, and this is something we should be considering in terms of the working group and what is easy to use by the intended community.

Pascal J: The explanation of the toolkit relates well to the taxonomy section of the standard. Pascal and Lyria have discussed several problem definitions that tie into what was shown in the toolkit. There is a need to talk a lot more about the categories that we are using – protected categories are easy but some are not protected everywhere, and then there are other variables that are not protected but might predict something else (eg shoe size). What about the granularity of the group? It might be beneficial to discriminate against a smaller group whilst benefiting larger group. They are trying to find specific very small use cases from cases where it may be ok to attach bias to an individual for example.

Ansgar: What is the richness of data needed to do these kinds of predictions? The example used seems very rich, a new system might not have this hugely rich data

set about predicted outcomes beyond a particular validation set. How much data makes it work reliably?

Also pointed out in chat "Mark U: There is an unfortunate coincidence with the use of FAIR in infosec / risk management. Something to consider as we build acronyms and references

https://en.wikipedia.org/wiki/Factor_analysis_of_information_risk" "Francien D: In the Netherlands there was also confusion with the acronym FAIR used in a responsible data science consortium. "Responsible" in this context in fact referred to aspects of data management: Findable, Accessible, Interoperable, Reusable. I have the impression though that they realized how this creates confusion (that may backfire, especially while speaking of "responsible data science" - which seems to point to more societal concerns than the more internal, functional concerns this FAIR referred to), They now speak about FACT, where the F stands for fairness, A for accuracy, C for confidentiality and T for transparency. See <http://www.responsible-datascience.org/>"

Rumman was knocked off the call but it is hoped that the discussion can continue over email or at another time.

Action: This discussion will continue via email and/or iMeet.

8. Updated Outline Discussion

I. Update on Use Cases

Adam: The outcomes evaluation group are using some mini use cases but not sure if they meet the criteria for main use cases. They looked for where there is evidence of bias having occurred at least once, based on news articles etc.

Action: Adam to send to group, to see if they could be developed into full use cases.

II. Topic updates – 5min summaries

i. Key concepts

No update.

ii. Taxonomy (Pascal)

Call last week. Coming close to point where the core is more or less complete, distinguishes basic types of bias, captures the Accenture tools basic idea about disparate impact and so on.

There are several key ideas that need to be taken into account, but the idea of justified bias will be outside scope of the section and should be talked about in the rest of the standard – where and how needs to be addressed.

They would also like to include experts from different fields in how far aspects of the taxonomy does not generalize esp. medical, policing, insurance, politics, finance, robotics. No medical/finance at moment. Please reach out to Pascal or Matthew S.

The basic idea of the taxonomy section has progressed far enough to put into publication at some point. Any suggestions/contributions are welcome.

Other sections should also be tied into the taxonomy section, and this is a checking step that needs to be included in the preparation of the final document. Some sections, eg legal, more obviously necessary but should be in all other sections. Adam pointed out that the outcomes

evaluation sections probably needs some merging, especially with definitions etc. Pascal and Adam will discuss.
Next call in next two or three weeks, no date yet, but will organize soon.

Action: Pascal to send details of next meeting to group mailing list

- iii. Legal frameworks
No update. It is important that we have someone from this team in the next meeting as we are behind in updates.
Action: Ansgar to contact team
- iv. Psychology (BC)
Looking to set up a call in two weeks
- v. Cultural aspects (Maya)
There are some new people in the group and it has been quite active in the last few days. There are a lot of questions and they need to get a sense of how to pitch it, so they are going to continue adding content and then try to have a call to resolve issues, although there is a bit of a timezone problem. Hopefully this will take place this month. Notes they are not dealing with culture as related to geography or nations etc, more on algorithmic culture. There are however different cultural norms to consider overall and the constraints and focus need to be confirmed. There is a lot of social science material in the doc which needs a lot of editing. Ansgar suggested to check with the legal frameworks team since they are dealing with national difference in laws.
- vi. System design stages (Chris C)
Not a huge amount progress since last time but new member which makes it a bit more manageable. Hoping to have a meeting soon. Current idea to use existing design process – CRISP-DM – which captures a wide scope of areas in system design. They need to look at alternatives and other standards groups, but Mark U commented in chat: *In NIST Big Data security we could not find anything better than CRISP-DM*
- vii. Person categorization (Srivathsan)
First version in next few weeks. Looking at psychology part of categorization, could do with some help, asked whether anyone from the psychology section might be interested.
- viii. Representativeness of data (Adam)
Group hasn't met. Some work in document, most of the people in the section are doing the outcomes evaluation at the moment. Anyone else who wants to take in on in the meantime get in touch.
- ix. Outcomes evaluation (Adam)

Use cases as above. Great session a few weeks ago, split into two groups: risk and impact assessments, and metrics. The IA part have a discussion paper to take back to subgroup. Metrics – Michael but not on call. Adopted definition of three types of techniques for OE, and have a set of draft discussion papers to turn into draft procedural standard.

- x. Evaluation of processing (Adam)
Adam and Michael had a meeting. Most of the issues crossover with OE except how to decompose system that contains multiple algorithm and look at each part. Need some people to drive that issue.
- xi. Resilience against manipulation
No update.
- xii. Documenting & transparency (Maya)
Some good questions being raised in the doc. Documenting new initiatives around transparency etc. Question of how general or specific does it need to be? Ansgar suggested referencing other existing standards around documenting of software. On good track to finish by September.
Adam – OE overlap with explainability and impact assessment documents. Maya to look at OE doc and speak to Adam.

Some sections are lagging due to people working on multiple sections and focusing on one currently. If anyone is not contributing, please find an area you are interested in and let us know.

9. Conferences and Whitepapers

- I. ITU Data for Good – journal paper
Progress made but maybe different journal due to submission deadlines
Matthew – how new technologies are enabling new data streams, and how things can be linked to an individual. Data breaches allow new threat modeling in terms of fraud.
- II. FAT* conference paper
No paper underway yet, deadline is 16th august for registering on intention to submit, paper on 23rd august. If want to do as part of p7003 need to move on it quickly.

Francien – special issue on social and cultural biases in information algorithms and systems – seems highly related to the work of the subgroup but no idea about journal. Francien can't contribute at this time.
Pascal – European symposium on societal challenges in computational social science, December this year, on bias and discrimination. Deadline for papers has passed but interesting track on comparative analysis of the same datasets which might be interesting.

<http://symposium.computationalsocialscience.eu/2018/>

10. Any Other Business

No other business was raised

11. Future Meetings

Thursday September 6th, 3pm UTC

Thursday October 11th, 3pm UTC

Thursday November 8th, 3pm UTC

Thursday December 6th, 3pm UTC

12. Adjourn

The meeting was adjourned at 16:29 UTC

Attendees:

Last Name	First Name	Employer/Affiliation	Voting
Biermann	BC	Heavy Projects & CAVAD	
Carrier	Ryan	Self	
Chowdhury	Rumman	Accenture	
Clifton	Chris	Purdue University	X
Costley	Jennifer	New York Academy of Sciences	
Courtney	Patrick	Tec-connection	X
Dechesne	Francien	Leiden University	X
Dowthwaite	Liz	University of Nottingham	X
Dwarakanarth	Anurg	Accenture	X
Fefegha	Alexander	Self	
Ganesh	Maya	Leuphana University	
Gardner	Allison	Keele University	
Hatada	Yohko	EMLS_RI	X
Jessen	Hans	EY	
Jurgens	Pascal	U of Mainz, Germany	X
KM	Srivathsan	TATA Consultancy Services	X
Koene (Chair)	Ansgar	University of Nottingham	X
Mandal	Sukanya	Self	X
McIntosh	Suzanne	Self	
Nadel	Larry	NIST	X
Ngounou	Charlie M	AfroLeadership	
Ruggeri	Salvatore	University of Pisa	
Smith	Adam L	Piccadilly Group	X
Stender	Matthew	Self	X
Underwood	Mark	Synchrony Financia	X
Weger	Gerlinde	Independent	X