Council for International Organizations of Medical Sciences

11th meeting of the CIOMS Working Group WG XIV on Artificial Intelligence in Pharmacovigilance

25th and 26th of June 2025, virtual meeting

Meeting Minutes

Overview:

Summary

The CIOMS Working Group (WG) XIV on Artificial Intelligence (AI) in Pharmacovigilance (PV) held its 11th meeting in Geneva, Switzerland on the 25th and 26th of June 2025, with some WG members joining virtually in the afternoon of the second day. The purpose of the meeting was to discuss the Public Consultation comments on the draft report received over 1st May to 6th of June. The final WG meeting, which will be held virtually in early September 2025, will be used for discussing the final details of the draft report as we prepare for publication.

Minutes of discussion

Day 1

1. Welcome and opening remarks

- Hervé Le Louët, Immediate Past President of CIOMS, opened the meeting as Chairman for Day 1.
- He introduced (in absentia) the new CIOMS President, Stella Blackburn.
- The CIOMS WG XIV report does not cover just a timely topic, but a necessary one. Whether algorithmic, probabilistic or generative AI, it is already reshaping the field. From Individual Case Safety Report (ICSR) processing and signal detection to natural language processing and

predictive modelling, AI is moving from concept to practice. In a domain like PV, innovation must be matched with clear principles for validation, oversight, and governance.

- The CIOMS WG XIV report has put forward a clear set of foundational principles:
  - Risk-based approach aligned with the stakes of each AI Use Case;
  - Calibrated human oversight, whether in-the-loop, on-the-loop, or in-command;
  - Strong expectation for validity and robustness;
  - Clear and actionable transparency, not only on data and outputs, but also on model architecture and communications;
  - Firm stance on data protection, particularly in the context of Generative AI (GenAI) and Large Language Models (LLMs); and
  - Strong commitment to fairness and bias mitigation, especially in underserved or vulnerable populations.
- This report is not about prescribing rigid, technical specifications; instead, it promotes a robust, future-facing framework that balances scientific innovation with regulatory and ethical responsibility. It aligns with global thinking from the EU AI Act, EMA reflection papers, FDA guidance, WHO ethical frameworks, and OECD recommendations.
- The report was submitted for Public Consultation, and we received a wide range of comments, and this marks an important step in the WG XIV work's review.
- The WG is looking forward to optimising the use of the report going forward.
- Sue Le Roux provided logistics for the days and made available copies of the most recently printed CIOMS reports:
  - Severe cutaneous adverse reactions (SCAR);
  - Benefit-risk balance for medicinal products.
- Sanna Hill was rapporteur.

2. Public Consultation – Summary and work process going forward

- Through the Public Consultation, the WG XIV received 1,024 specific comments linked to a line, paragraph, section or chapter of the draft report; and these have been inserted into the report as comments. They can also be found in Tab 1 of the accompanying Excel spreadsheet. In addition, the draft report includes the WG's own unresolved comments from before the Public Consultation.
- Furthermore, the WG received 36 general comments linked to several chapters or the whole report. They can be found in Tab 2 of the accompanying Excel spreadsheet.
- The list of companies and organisations that provided comments was reviewed, as well as the countries they represented.
- Some 15+ commentators opted not to be listed in the final report.
- The accepted work process is to address comments during the WG meeting, for the Chapter and Section Leads to continue working on the comments over the summer, and for the Editorial Team to meet to discuss at the end of the summer.
- All work is to be done in track changes for the full WG's visibility.
- The Chapter and Section Leads are asked to log how the comments are addressed. The Chapter and Section Leads are requested to note in the report their initials at the end of each comment and to provide an outcome, e.g. "Done", "Declined", "Need to consult WG"; and especially in the case of declined comments, to give some explanation as to why a given decision was taken, e.g. "Addressed in other chapter", "Beyond scope", "Will not recommend one method above another". The phrasing does not have to be harmonised but it is preferable for it to be succinct. This information will enable the CIOMS Secretariat to answer questions afterwards regarding the comments if needed.

- The Chapter and Section Leads are asked to not delete comments. The comments can be marked as resolved. Sanna will eventually delete the comments once she has recorded them in the Excel spreadsheet.
- Around September, once the Chapter and Section Leads send their chapters and sections to Sanna, she can then move the information regarding how the comments were dealt with into the Excel spreadsheet. The Chapter and Section Leads do not need to complete the Excel spreadsheet.
- The Chapter and Section Leads are asked to only work in their own chapters/sections for the sake of version control.
- All WG members will have visibility of the draft report with the track changes and the accompanying Excel spreadsheet showing how the comments were addressed before publication.
- The following WG members will be looking after their chapters and sections with regard to addressing the Public Consultation comments:

  Preface: Walter
  Executive Summary: Manfred
  Ch 1: Introduction: Vijay (and Andrew)
  Ch 2: Landscape: Niklas and Benny
  Ch 3: Risk-based approach: Julie
  Ch 4: Human oversight: Julie
  Ch 5: Validity & Robustness: Niklas
  Ch 6: Transparency: Niklas
  Ch 7: Data privacy: Walter
  Ch 8: Fairness & Equity: Beth
  Ch 9: Governance & Accountability: Hua
  Ch 10: Future considerations: Tom
  Appendix 1: Glossary: Douglas
  Appendix 2: Comparison table: Benny
  Appendix 3: Use cases: Vijay at 11th WG meeting (later Taxiarchis)
  Appendix 4: Content related to:
  Explainability: Niklas
  Fairness and Equity: Beth
- As we edit the draft report, i.e. the Word document, the line numbering will become more and more different from what we started with. If the comments make reference to line numbers, please check the original PDF doc that was used for the Public Consultation.
- If some Chapter Teams prefer to work on SharePoint, please ask Beth or Sanna for help with uploading a chapter/section.
- Some edits will require follow through within the report, e.g. at the Glossary and the Use Cases, and perhaps also with hyperlinks and the list of acronyms.
- Regarding the length of the report, although some comments recommend shortening the document, and whereas quick solutions for shortening the report are encouraged, we are not proposing investing time into reducing the length of the report. Many readers will read the report electronically and not from cover-to-cover.
- The WG decided on some ground rules:
  - It is acceptable to decline comments that are not actionable or not clear;
  - We can reach out to commentators for clarity if necessary;
  - If no reference has been provided, we can decline proposed content;
  - Where vendors put forward citations of their own systems, we will not include proposed content - the exception would be if it is put forward in a citable form.

3. Public Consultation - General Comments

- **Placement of Chapter 2.** In line with one of the comments (below in blue), we discussed the possibility of moving Chapter 2 on Landscape analysis to the appendices; or alternatively, reducing Chapter 2 and moving part of it to the appendices. The advantage of moving Chapter 2 would be to reach the Principles in the report sooner; however, Chapter 2 helps to set the scene, and many readers dip in and out of a book rather than reading through from start to finish, and so maybe the order of items with the report is satisfactory. It was decided to keep Chapter 2 where it is and to add a note at the start of the chapter to say that the Principles begin at Chapter 3.

  General comment
  PIC/S Good Pharmacovigilance Practices Expert Circle Working Group on AI and Machine Learning
  Strongly recommend that Chapter 2 on Landscape analysis is moved to an appendix. It is very useful information; however, you want the readers to go straight from the introduction to scope to core principles. I believe you may lose some readers along the way who might not make it to the core principles if Chapter 2 on landscape analysis is not moved to the appendix.

- **Link section on Robotic Process Automation back to AI definition.** The comment on Robotic Process Automation (RPA) (below in blue) suggests that RPA does not fall under AI but is simply about rules-based automation. The definition of AI in the report comes from OECD, which is a very broad concept of AI and includes almost any use of computers to do a task that humans would otherwise do. At this place in the text, perhaps we need to point back to the OECD definition in the report.

  702
  Christine Prendergast, PIC/S Good Pharmacovigilance Practices WG on AI and ML, HPRA

  Concern or text in question: Several organisations who process large numbers of case reports have also automated repetitive, labour-intensive tasks using so-called robotic process automation (RPA) technologies.

  Comment or suggestion for solution: It was my understanding that RPA was not AI, but rather rules-based automation. Could this be explained in the text here. It will be important for all stakeholders to know the difference between something like a rules-based RPA vs AI solution.

- **Elaborate on the implications of the AI definition.** The comment on disproportionality analysis (below in blue) helps to elucidate a point about how we have a very broad definition of AI, but the extended question is "what kind of AI is this?". In risk-based approach, we do things differently depending on the type of AI we are dealing with. We risk confusing readers because they arrive at e.g. the Landscape analysis chapter or the Use Cases, and they will think "this is not AI". So even though we start the Introduction with the OECD definition, at the moment, we do not elaborate on it and its implications, and maybe that is something we need to add. We need to align to the same understanding throughout the report.

  797
  Kaitlyn McAteer,Merck Animal Health

  Concern or text in question: The earliest examples of real-world use of (simple) AI solutions in PV are from the late 1990s. At this point, disproportionality analysis, first conceptualized in the 1970s, began to be implemented as part of triage algorithms to help direct the attention of PV specialists in their analysis of large national and international collections of individual case reports.

Comment or suggestion for solution: Similar to my previous comment, I don't really think that a disproportionality analysis (DPA) is considered to be AI. These analyses are conducted using mathematical formulas which can be done by a person using a calculator or a program like Microsoft Excel. I do not think that considering simple mathematical calculations as AI aligns with the definition of "artificial intelligence" in the appendix and this is a much broader definition of AI than would be commonly used/understood. Additionally, I do not think that consideration of statistical calculation results by humans as part of a triage algorithm would be considered AI. For example, if a human reviewer sees a calculated PRR >2 for a particular sign and then focuses their time on analysing that sign, I wouldn't consider that to involve AI as the AI only performed the calculation but did not infer anything from it. I would only think of this as AI if the AI program was performing the triage process itself (including decision-making, incorporating other data, etc.).

Further discussion points on this subject:

- The suggestion is to look back to the Introduction and be more precise in terms of how we define AI. It was agreed that in the Scope section, we could elaborate on what is covered and what is not covered in a certain context; be more descriptive.
- What are the implications in practice? RPA is a good example. In terms of the Principles, some of the AI systems have been in place for a long time. For example, there may be an expensive human cost from a governance perspective, e.g. we document what we would not normally view as AI. Even if it is RPA and rule-based, there are checks and balances. The degree may be different to what would be expected for AI.
- There was discussion about when intensive vs routine monitoring can end when we are looking at rule-based components vs AI components.
    - It was put forward that, looking at the lifecycle, when we have a rule-based method, where the rule was defined by humans, we do not have further routine monitoring of the system after validation. We would do post-rollout checks of the AI component of the system, but for the rule-based component, we would not.
    - There was disagreement about this. We still have to monitor because there is a possibility of data drift, even with a simple rule-based system. With the possibility of model drift, which we would have with learning; this introduces a different concern.
    - Active monitoring becomes needed as we move closer to the true AI function/solution. When we have some of these elements: adaptability, ambiguity, and maybe, reproducibility.
    - We do not address the qualification of older/traditional systems and computer system validation in the past. It was decided that the report would not discuss the governance of legacy systems.
- It was decided to use phrasing such as "for the purpose of this report, we have used an inclusive definition, and organisations may set their own boundaries in terms of how to define AI". This is to address the practical challenge of RPA and other types of technologies considered to be AI, which would bring us under the EU AI activities. Otherwise, within some companies, there would be a need to catalogue all AI systems within a registry, which would add an additional level of bureaucracy under the EU AI Act. Using this wording with the definition would avoid a burden not on the PV side, but the other side of governance in terms of ensuring that we have an inventory of AI systems, and have assessed risk according to the EU AI act.

- **Target audience - Ethics committees.** A commentator suggested adding ethics committees to the report's target audience at the scoping section, i.e. to the list: "PV professionals, regulatory authorities and software vendors". There is a similar comment in the Governance &

Accountability chapter, suggesting to add ethics committee to the list of stakeholder bullet points. Similar comments are found throughout the document.

- It was considered that an ethics committee would pertain primarily to the Fairness & Equity and Data privacy chapters, and in these chapters, we delineate those topics that are considered to be routine public health practices - a public health practice as opposed to research applications involving individuals with individual data.
- If our scope is specific to PV systems and development programs, and futuristic development activity, then an ethics committee does play an active role. With any AI activity, even at hospitals today, there are ethics committees overseeing the AI work, including program application validation.
- This may be out of a concern for re-identification, oversight of validation and medical assessment. Ethics committees are involved in research. They do not govern systems – that is what sponsors do – but they have oversight of systems because they are the ones making decisions based on the data that is provided by the systems.
- Ethics committees are involved with reviewing data re-use, protecting patient privacy, etc. But, they can define outcomes where they are not looking at the system per se, they are looking at the process of how we are managing and risk mitigating patient re-identification and related issues. The Ethics committees' interest is more in the scope of data collection in clinical trials. It's not in the purview of AI in PV.
- We may highlight that for certain decisions or steps, e.g. at an escalation point, an ethics committee can get involved.
- This can be addressed in the Data privacy section too.

- **GenAI test sets and transparency.** One reviewer suggested that GenAI systems cannot be used for PV-related use cases. This comment below in blue was made in the Transparency chapter, and the same comment seems to have been given in multiple places. In the opinion of the CIOMS WG XIV, GenAI is in scope.

  788
  Rory Raftery, Tomasz Dyszynski, Bayer

  Concern or text in question: To give a comprehensive picture of an AI solution's effectiveness and limitations, the presentation of performance evaluation results should describe the scope and nature of the test set(s) used including definitions of their reference standards and sampling strategies.

  Comment or suggestion for solution: This section's requirements again suggest that GenAI/LLMs cannot be used in PV use cases because they cannot fulfil these criteria. Can CIOMS advise what information technology companies would need to provide regarding transparency to fulfil these requirements?

  Further discussion points on this subject:

  - There was some confusion about this comment. The nature of the confusion was this: when the commentator was speaking of not knowing the nature of the test set, they were thinking of the tests that OpenAI may have done when they developed ChatGPT or Google may have done when they developed BERT, but we are concerned with tests done on our specific application. This point can be clarified.
  - It was agreed that this point would be addressed:
    - In the performance section of the Validity & Robustness chapter;
    - In the Transparency chapter where the comment occurs, from the angle of being transparent about test sets;
    - In the Governance & Accountability chapter.

- It would be difficult to address this point in the Introduction before the Principles have been met and it would be too late to write about this in Chapter 10 because the issue is already occurring today.

4. Public Consultation – Specific Comments

- The WG members worked most of the day in their Chapter/Section Teams addressing the Public Consultation comments.

5. Closing remarks Day 1

- The CIOMS WG XIV consensus report is keenly awaited by PV organisations.
- There is a fair amount of expectation in the field for the CIOMS report.
- The CIOMS report will offer an opportunity to encourage good practice in the PV space, which we do not see collaboratively often with other AI and ML functions. This report provides guidance to take into consideration.
- This report will help to eliminate some of the confusion, given the vast amount of choice available from vendors.
- It will help to clarify some of the terms in the field.
- Hervé would like to keep the CIOMS WG XIV open beyond the publication of the report to enable the WG to review issues in the future.
- Hervé congratulated the WG on its achievements.

Day 2

6. Welcome back and opening remarks

- Lembit Rägo, CIOMS Secretary General, opened the meeting as Chairman for Day 2.
- Lembit informed the WG that Robert Ball, the WG member from the US FDA, has had to leave the CIOMS WG XIV due to FDA's current administration policy and we will have to remove his name from the report. This was the case with the CIOMS WG XII report too. In the case of the WG XII report Acknowledgements section, a note was added as below, and the same can be envisaged for the WG XIV report too:

  "The Council for International Organizations of Medical Sciences (CIOMS) is thankful to the experts from drug regulatory authorities, including ICH founding members, […]
  Finally, we would like to warmly thank those who kindly gave their time and contributed their expertise in key roles but who unfortunately cannot be named."

- We highly appreciate the valuable Public Consultation comments provided by both US FDA CDER and CBER.
- The below timeline was drafted to enable the WG XIV report to be tentatively published in early/mid-December 2025:

  - w/c 18 Aug: Editorial Team meeting to discuss Public Consultation comments and to align across Chapters/Sections

[Post-meeting comment: this may be a good deadline for proposing final glossary terms, definitions and references to allow the Glossary Team to review and the Chapter/Section Leads to ensure the terms are applied consistently throughout the report.]
    - *w/c 25 Aug: Chapter/Section Leads send content to Sanna for compiling into updated draft report [Post-meeting comment: this line was added since 11th WG meeting.]*
    - w/c 1 Sept: Share updated draft report with full WG
    - w/c 8 Sept: Virtual WG meeting; Chapter/Section Leads make final edits
    - w/c 15 Sept: Sanna to finalise references, check hyperlinks, check against editorial style guideline, start process to pay for external intellectual license fees if any and to obtain high resolution artwork
    - w/c 22 Sept: WG's 2-week no-objection check of document (there is no action necessary), in parallel proofread doc
    - w/c 6 Oct: Finalise resulting edits
    - w/c 13 Oct: Start graphic design of document
    - w/c 8 Dec: Electronic publication, start print

- Before the Editorial Team meeting in the w/c 18 Aug, it would be helpful to collate all the chapters and sections so that the Editorial Team can review how all edits affect the various sections.
- In the w/c 22 September, we would like to invite an external proofreader to read the draft report, and there was a suggestion to invite this person to review earlier in the process regarding overlaps or opportunities to consolidate the document.
- It was felt that the time between the virtual WG meeting in the w/c 8 Sept and Sanna finalising the references in the w/c 15 Sept is too tight for the Chapter/Section Leads to do their work, and we discussed various solutions, e.g. reducing the WG's 2-week no-objection check of the document and the graphic design of the document, although these would seem difficult to shorten too. [Post-meeting note: the timeline was brought back to the original proposed version as above but all efforts will be made to maximise the time available to the Chapter/Section Leads. For example, if there are several options for a virtual WG meeting in the w/c 8th of Sept that are equally work-able, then we will favour the option that offers the most time to the Chapter/Section Leads after the meeting.]

7. Public Consultation – Specific Comments summary

- Any comments that have already been made in the minutes have not been repeated below.
- Almost all Chapter/Section Leads reported addressing editorial Public Consultation comments. This type of comments have not been repeated for each chapter/section summary.

Executive Summary

- We had a brief hybrid session to make a start on the comments on the Executive Summary.
- To give an example of one of the comments, there was a suggestion to add "mathematics" to the list of fields in the sentence: "The CIOMS report on artificial intelligence (AI) in pharmacovigilance (PV) addresses a rapidly emerging cross-disciplinary field that is at the intersection of PV, computer science, mathematics, regulation, law, medicine, human rights, psychology and social science".

Chapter 2: Landscape - presented by Niklas Norén and Benny Ling

- There was a specific section on Common Data Models (CDMs) that referred to Observational Medical Outcomes Partnership (OMOP) and Fast Healthcare Interoperability Resources (FHIR) in the context of ICSR processing, and it came across as a little off-topic and incomplete. To talk

about a CDM for ICSR processing, it would need to be E2B, but there are other comments suggesting that we should not really emphasise a CDM, and so our proposal is to omit that text.

- In the Comparison Table, the US data is out of date as new AI measures and guiding principles are being developed. It was suggested to edit the table to say that it was correct on the given date and consider it as a historical document. We will provide a date for the hyperlinked document.

Chapters 3 and 4: Risk-based approach and Human oversight - presented by Julie Durand

- The comments have been largely requests to provide more detail and examples, e.g. types of risk. Many of these have been declined as the report is not intended to be exhaustive.
- There was a request for a checklist for developing a risk strategy and this was declined.
- So far, there have not been comments that would necessitate substantial changes to either chapter.

Chapter 5: Validity & Robustness - presented by Niklas Norén

- There were many comments, several quite detailed, and some recurring.
- More than one organisation requested pedagogical presentation of performance evaluation metrics. We had content on this in the report at one stage and we removed it because we felt that it was off topic. Our solution now is to cite a paper by the Uppsala Monitoring Centre that covers this topic and not go into detail in the report.
- A commentator asked about the validity and robustness of GenAI solutions, and we feel most of the principles apply. There are some aspects of performance evaluation that we could cover. We are considering drafting one or two paragraphs towards the end of the chapter over the summer, citing any available papers. The request was specifically about how to optimise and train GenAI, and we think that is out of scope because we do not focus on anything that technical or so much on training.
- We will add a few sentences when we think there are specific considerations related to agents in general or multi agents in particular.

Chapter 7: Data privacy - presented by Walter Straus

- The FDA provided a number of helpful comments related to HIPAA.
- Walter requested assistance with verifying the content in Tables 5-7 on Data privacy regulations for using secondary data in China, Germany and Japan.
- Walter requested for examples of data privacy breaches that have occurred in PV in the public domain. So far in the chapter, there are examples of non-PV-specific data privacy breaches.
- We discussed inherent limitations in safeguarding data breaches involving LLMs and how to mitigate the risk of data breaches, focusing on the roles of training and restricted access. It is easier to assure proper use by individuals with good intentions than improper use by people with bad intentions. How do we to mitigate the risks associated with data breaches involving LLMs?

Chapter 8: Fairness & Equity - presented by Elizabeth MacEntee Pileggi

- Some changes were proposed to the references and these need to be evaluated offline.

Chapter 9: Governance & Accountability - presented by Hua Carroll

- Three quarters of the comments require discussion by the chapter team and the team anticipates continuing working virtually over the summer after the in-person WG meeting.
- Light editing of the chapter will be needed.

- One area to clarify will be how the governance structure will tie into the current existing Quality Management System. We are calling for a separate stream of work.
- A second area to clarify is the deconditioning of the AI system and how it ties into the current process.
- Eventually, we would need the Editorial Team, and then the full WG, to review changes made to the governance framework grid. We received comments on the grid, and the Chapter Team has provided some suggestions for changes too.
    - We discussed the idea of potentially shorting the grid to fit on one page for easy viewing and use. If this is not possible, then we would like to shorten the grid to make it fit on two pages side-by-side.
    - The Chapter Team is considering removing the last column entitled "General considerations". The little content in this column would be incorporated into the preceding columns or into the text before or after the grid. The "General considerations" column is not a development phase, and if this grid were to be used as a template of activities, this column would not be filled out.

  For the grid to be valuable, it should not be reduced too much.
- The Chapter Team is also considering adding a graph to help readers visualize the whole change process from phase-to-phase.

Chapter 10: Considerations on future development and deployment of AI in pharmacovigilance – presented by Thomas Henn

- Some comments were asking for more granularity.
- There was a comment concerning black swan events, which has been moved to Chapter 3.
- Regarding smart organs and nanotechnology, we received a comment saying that these were too advanced to include in the report; however, in the framework of future considerations, and as one of the directions of AI deployment, it was felt that these topics could be left in.
- It came to light that within the WG, we have different understandings about what we mean by the term "Autonomy", and this may be useful to include in the Glossary.
- There was a comment concerning federated learning, and it may be that this may be best moved to the Data privacy chapter.

Glossary: presented by Douglas Domalik

- Douglas has reviewed all the comments but the Glossary Team has not met during the in-person meeting because the Glossary Team members were involved in other Chapter/Section Team meetings.
- Some comments were suggestions for using alternative definitions. Where CIOMS definitions were proposed to be replaced, these comments were declined, but other comments will be discussed over the coming two or three weeks.
- Some comments have the potential to add clarity but none of the comments will bring significant changes.
- The WG members have suggested we add the terms "Autonomy" and "Agent".
- Douglas requested that the Editorial Team members contact him if they have terms, definitions and references to add over the summer for adding to the glossary and for the Glossary Team to work on.

Use Cases: presented by Vijay Kara

- As the Governance & Accountability chapter grid was being developed at the same time as the Use Cases were being discussed, there was a disconnect between the governance grid principles and their implications within the Use Cases. The comments show that readers are trying to get a better understanding of the thought processes that went behind defining each of the positions within the governance grid. Therefore, the Use Case Team will go back to think about how to integrate the grid thought processes better into the Use Cases.
- There was feedback questioning the utility of some of the Use Cases:
    - Use Case H, which is entitled "Artificial intelligence to support diagnosis and prediction of (hydroxy)chloroquine retinopathy". This Use Case is to do with PV in the clinic. We felt that this Use Case has a purpose in terms of demonstrating that we have PV processes and that they can also extend into the clinic. We will think about what changes can be made to provide the thought processes to show the PV perspective.
    - Use Case G, which is entitled "Generative artificial intelligence for enhanced and intelligently structured outputs from large pharmacovigilance document libraries". In its current format, this Use Case is to do with how one might develop a GenAI solution overall. There were references on its applicability to a PV system, but these also tied to some comments where the readers were trying to get an idea of how aggregate reporting might be supported by the GenAI solutions, so it gives us an opportunity to reposition this particular Use Case and think about the principles and what considerations would need to be given from a requirements gathering phase and what would need to be considered to support particular Use Cases. We will try to incorporate components around what is available currently in the public domain as a precursor to thinking about how that might apply to reporting. It may provide us with a slightly more useful, rounded Use Case.
- The Governance & Accountability chapter team requested to have visibility of the Use Case edits once completed in case there are any counter corrections required in the Governance & Accountability chapter. This could be in terms of developing or qualifying a system, setting expectations for implementation and usage.
- There were many comments with requests for more examples throughout the report, and in the Use Cases, we can address this kind of questions: how do you test the performance of outputs of LLMs; and what do we define as the threshold for acceptability? These are things that we cannot address within the body of the document but we can attempt to address in the Use Cases by saying that this performance threshold was acceptable because of the kind of risk assessment frameworks and the quality management system that supported the process, and we can go into a little bit more detail.

Appendix 4: Content related to explainability - presented by Niklas Norén

- Some content in the appendix contradicts content in Chapter 6 on Transparency and this will need to be adjusted. This content was drafted by Satoko, who is not present at the moment, and this content needs to be reviewed. Our thinking about explainability has evolved over the years of writing the CIOMS WG XIV report and some content may have become outdated. Niklas invited others to join him in reviewing.

Appendix: Beyond Vigilance: Navigating AI's Potential Dominion of Medicine (not in the Public Consultation)

- This content was not part of the Public Consultation.
- It was proposed that this content be published as a separate opinion piece in a journal, i.e. not as part of the CIOMS WG XIV report and not as a consensus paper.
- Lembit is investigating into the possibility of publishing the content under CIOMS.

- This content started being shaped before the CIOMS WG XIV was formed. There were conversations about how the future development of medicine and the way the delivery of medicine is going to impact how PV is going to be adopted into clinical practice. This was not limited to 10, 20, or 30 years, and beyond. Ravi and Tom started drafting the text, and it was then realised that this was too far extended out in time, and this is one of the reasons why the WG could not reach consensus on the text. A lot of the principles and guidance of the report are thinking about how to manage things today and in the near future, and not necessarily 30 years down the road.
- There was a suggestion to focus the paper around transportation technologies or broaden the scope to include drug delivery, devices, Internet of Things, and all of the other kinds of fora of technologies that could impact PV. The scope has not been decided yet. We would like to receive feedback from Lembit. More detail will surely be added.
- Other volunteers will be able to join to help finalise the paper.
- Would we consider suggesting joining forces with public-private partnerships, e.g. Innovative Health Initiative (IHI), who are already working in this space to avoid duplication of efforts? We would first need to determine the scope of the paper we are writing. Also, we are not necessarily writing a consensus paper.

8. Closing remarks Day 2

- Closer to the publication of the report (perhaps at the next WG meeting), we can discuss awareness raising activities to promote the CIOMS WG XIV report such as:
    - WG members speaking at key conferences;
    - Scientific article in open access journal;
    - Webinar with WG members as speakers.
- The IHI topic that is in the pipeline for future calls for proposals: Topic 3: AI-powered signal detection in pharmacovigilance (https://www.ihi.europa.eu/apply-funding/future-opportunities) will be included towards the next WG meeting agenda.
- Lembit thanked everyone for their time and commended all for their contributions.

9. Annex 1: List of participants

Justyna Amelio (AbbVie), Adrian Berridge (Takeda), Hua Carroll (Biogen), Mariane Diniz (Bayer), Douglas Domalik (AstraZeneca), Julie Durand (EMA), Kirsten Egebjerg Juul (Danish Medicines Agency), Piero Francesco Franco (Pfizer), Neal Grabowski (Sanofi), Thomas Henn (Unither), Sanna Hill (CIOMS), Vijay Kara* (GlaxoSmithKline), Dieter Kempf (Genentech, Roche), Benny Ling (Health Canada), Denny Lorenz (independent expert), Monica da Luz Carvalho Soares (ANVISA, Brazil), Elizabeth MacEntee Pileggi (Johnson & Johnson), Yusuke Matsunuga (PMDA, Japan), Niklas Norén (UMC/ WHO), Ravi Patel (Unither), Irene Scholz (Swissmedic), and Walter Straus (Moderna).

Attended one day
Hervé Le Louët (CIOMS), and Lembit Rägo (CIOMS)

Attended virtually during a hybrid session
Manfred Hauben (Merck Group)*, Jensen Morten (Danish Medicines Agency)*, Hans-Jörg Römming (Merck Group), and Phil Tregunno (MHRA, UK).

Apologies
Russ Altman (Stanford University, US), Luisa Barrios (Merck Sharp & Dohme), Andrew Bate (GlaxoSmithKline), Arvind Bellur (CSL Behring), Taxiarchis Botsis (Johns Hopkins University School of Medicine, US), Brian Buch (MHRA, UK), Yauhenia Cherkas* (Johnson & Johnson), Selin Cooper*

(AbbVie), Jean-Michel Dogné (University of Namur, Belgium), Julie Girod* (Sanofi), Kendal Harrison (MHRA, UK), Stephen Heaton* (CIOMS), Satoko Hirokawa (Health and Youth Care Inspectorate, Netherlands), Yuki Kikuchi* (PMDA), Eva-Lisa Meldau* (UMC/ WHO), Dirk Mentzer (PEI), Flávia Moreira Cruz (ANVISA, Brazil), Stephen Rosenfeld (North Star Review Board), Thomas Stammschulte (Swissmedic), Panos Tsintis (CIOMS Senior Adviser), James Whitehead* (AstraZeneca), and Brian Yau (WHO).

*Alternate