Council for International Organizations of Medical Sciences

10th meeting of the CIOMS Working Group WG XIV on Artificial Intelligence in Pharmacovigilance

23rd April 2025, virtual meeting

Meeting Minutes

Overview:

Summary

The CIOMS Working Group (WG) XIV on Artificial Intelligence (AI) in Pharmacovigilance (PV) held its 10th meeting virtually, on the 23rd of April 2025. The purpose of the meeting was to present the draft report in its entirety following a WG-internal review and comment period, and to offer an opportunity for everyone to discuss any outstanding issues as we prepare for a Public Consultation. The Public Consultation is scheduled for 29 April – 6 June. The next in-person meeting will be held in Geneva, Switzerland, over 25-26 June 2025, when the WG will meet to discuss the Public Consultation comments.

Minutes of discussion

1. Welcome and opening remarks

- Lembit Rägo, CIOMS Secretary General, opened the meeting as chairman.
- For a list of participants see Annex 1.
- Sanna was rapporteur.

2. Draft report – Overview of updates resulting from WG-internal review

Preface - presented by Walter Straus

- The Preface underscores the formative role that CIOMS has played in PV.
- It highlights examples of CIOMS's contributions through its WGs.

- It makes the points that AI represents a major technology change with significant implications for PV, and that the report authors have taken a focused approach on this broad, rapidly evolving field.
- The Preface identifies the audience for the report.

Executive Summary - presented by Manfred Hauben

- The Executive Summary emphasises the following points:
    - AI in PV is uniquely cross-disciplinary;
    - It has indications, a method of administration, adverse effects, precautions, and warnings for use;
    - The field is dynamic – this is not a detailed, technical exposition of AI in PV but elaborates methodological, ethical, and regulatory principles, which will guide the development and will further engage the field.
- It lists the key points from each chapter.
- The Executive Summary will be useful as a tool to present the report to member organisations, external forums and stakeholders.

Chapter 1: Introduction - presented by Andrew Bate

- The Introduction describes what is changing in PV to set out the scope.
- A change was made recently to a figure and text that focus on how fundamentally generative AI is changing how we think about AI as this part received a lot of questions and comments.

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

- Since the 9th meeting in Darmstadt, input was received from Bob, Vijay, Taxiarchis, and Andrew, and others added use cases, which has made the coverage of early work more complete.
- References have been added.
- Some content has been moved from the Future Vision chapter on robotic process automation, referring to the TransCelerate survey.
- With regard to Table 1 on *Published examples of deployed artificial intelligence solutions in pharmacovigilance*, a note will be added in the caption to say that many pharmaceutical companies also have deployed AI solutions but typically this has been done through vendors and the descriptions of these may not be in the public domain. We will cite the TransCelerate survey as many use this.
- The regulatory consideration section has incorporated the comments that were received.
- Updates to Table 2 on *Comparison of the CIOMS Working Group XIV guiding principles*:
    - References have been added;
    - The original documents were revisited and more check marks were added to the principles, because even though some of the documents did not have these headings, the principles were nevertheless mentioned in the documents.
    - The authors may need to edit in the table header some of the names: 'US' (FDA?), Canada (Health Canada?), 'UK' (MHRA?), 'Australia' (TGA?). We could move the international organisations e.g. WHO, PAHO, and OECD to the right, rather than in between the regulatory authorities.
    - There was a suggestion to add a cross-reference to the use cases appendix.

- Some of the documents were updated e.g. the EMA Reflection Paper on the Use of Artificial Intelligence in the Medicinal Product Lifecycle.
- More references have been added e.g. the FDA draft guidance has been added: FDA Draft Guidance on Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products.

Chapter 3: Risk-based approach - presented by Julie Durand

- The comments and suggestions have been implemented including to the Principle description.
- Amongst the Key messages, there were comments to stress that the risk-based approach applies across principles including Data privacy and Fairness & Equity; and this has been implemented.
- The body of the chapter has not changed much.

Chapter 4: Human oversight - presented by Julie Durand

- Most of the suggested edits have been implemented, which were not substantial.
- There were a few requests to expand on how the different oversight mechanisms are described in the European Commission's ethics guideline for trustworthy AI and how this would apply in the PV space; and so some sentences have been added to illustrate this.
- At the 9th meeting in Darmstadt, it was agreed that we should discuss in one place in the report the impact of the rollout of AI systems in PV on the human work force; and a section has been added to the Human oversight chapter on the 'Transformation of traditional roles'.

Chapter 5: Validity & Robustness - presented by Niklas Norén

- Minor feedback has been consolidated since the 9th meeting in Darmstadt.
- We have added one Key message, which was already presented in draft form in Darmstadt, and this has not changed since then.
- For all of the chapters, helpful editorial suggestions were received from Mariane for the Principles and most of these were accepted.

Chapter 6: Transparency - presented by Niklas Norén

- The explainability content was incorporated into the Transparency chapter as a section, and therefore explainability is no longer a Principle in its own right. This was suggested in the FDA comments prior to the meeting in Darmstadt. The decision was not made at the meeting in Darmstadt but the Editorial Team made the decision in the autumn 2024.
- A Key message has been added for explainability to the Transparency chapter.
- The remaining explainability content has been placed in Appendix 4.

Chapter 7: Data privacy - presented by Walter Straus

Walter presented slides with the following content:

Ethical and Regulatory Foundations for Data Privacy in AI-Driven Pharmacovigilance

Ethical Principles (Belmont Report) – Developed for clinical research; applied to PV/public health

- Respect for Persons: Control over personal data, informed consent.
- Beneficence: Maximize benefit, minimize harm in PV activities.
- Justice: Fair data use and equitable access to PV benefits.

- Applies to both Data Privacy & Fairness and Equity

Key Regulations/Legislation:

- HIPAA (US): PHI protection, minimum data use, de-identification; exemptions for PV with safeguards.
- GDPR (EU): Broader scope, includes "right to be forgotten", strict consent rules, limits on data transfer.

Global Considerations:

- Varied country-specific rules (e.g. Germany, Japan).
- Regulatory alignment essential for global PV and AI deployment.
- Academic PV requires IRB approval and robust data safeguards.

AI Risks and Practical Safeguards for Data Privacy in PV

AI-Driven Privacy Risks (primary focus – GenAI):

- Generative AI may link disparate data, breaching privacy (intentional or not)
- Re-identification risk with minimal data points.
- Open LLMs riskier due to weaker controls on data intake/oversight

Risk Mitigation Measures:

- Maintain and monitor de-identification standards.
- Restrict AI model access and ensure internal controls (closed LLMs)
- Use contractual and technical safeguards with partners.
- Audit data use, reinforce minimum data principles.
- Train PV staff; ensure oversight of AI in PV.

Takeaway: Ethical AI use in PV demands vigilant governance and evolving privacy safeguards.

*For team – discuss preferred figure.

Feedback from WG members:

- It may be possible to cross-reference to the ethics content in the CIOMS WG XIII report on RWD and RWE as there is some overlap. In answer - this content was already reduced and is considered to be necessary for both the Data privacy and the Fairness & Equity chapters.
- There was concern over the wording of open LLMs being riskier than closed LLMs; this is almost an artificial distinction. This language could be softened. We could add some text specifically about the risk of compromising patient confidentiality via prompts.

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

- This chapter is closely linked with the chapter on Data privacy and with the chapter on Human oversight regarding changing roles.
- Fairness & Equity focuses on identifying and addressing discrimination with the biggest area of focus being underserved populations.
- The chapter is composed of sources of potential threat to Fairness and Equity with the use of AI in PV and key mitigation strategies. Further editing was carried out to clarify the message: mitigate unfair bias in PV and avoid discrimination.
- Most changes to the chapter were editorial, i.e. there were suggestions for better phrasing, but the chapter structure has not changed. The Principle description was refined with regard to its relevance to PV.
- We also clarified that important topics outside of scope of PV will not be covered in depth.
- Beth appealed for a reference for an example that a commentator had provided of a prediction of patients with chronic illness who would benefit from a proactive intervention.

Chapter 9: Governance & Accountability - presented by Hua Carroll

- The majority of the proposed edits were accepted and this includes updates to the Governance framework grid.
- The Governance & Accountability chapter was moved from its position as the first chapter to the last chapter. It was felt that this chapter was a good anchor to end the series of chapters.
- A paragraph on explainability was moved to the Transparency chapter.
- A Section on Traceability and version control was moved to the Governance & Accountability chapter from the Transparency chapter.
- The conclusion section was re-purposed as an Introduction section with some light editing in order to be consistent with the chapters structure throughout the report.

Chapter 10: Considerations on future development and deployment of AI in pharmacovigilance (previously known as 'Future Vision') – presented by Thomas Henn

- This chapter had not yet been seen by the WG members.
- This chapter was scripted by Ravi, edited by Tom in line with feedback from the Editorial Team, it was further fine-tuned by Panos and Walter, and then checked for transferring 'current' content to the Landscape chapter by Niklas.
- Linkage has been strengthened with the preceding chapters.
- Some of the content relating to the most advanced technologies, e.g. esoteric examples of AI in general medicine, have been moved to an accompanying appendix, leaving the chapter to focus on AI in PV aspects alone.
- The title has been changed from 'Future Vision' to 'Considerations on future development and deployment of AI in PV'. The 'vision' is now rather captured in the accompanying appendix.
- The chapter has three sections:
  - The first section now speculates on where AI and PV may take us in the future;
  - The second section considers and gives examples of expert systems such as fuzzy logic;
  - The main section focuses on the evolution of AI and the guiding principles.
- The references still need to be inserted in the right places.
- The WG members are invited to review the chapter and provide feedback by 2 pm on the 25th of April.

Accompanying appendix: Beyond Vigilance: Navigating AI's Potential Dominion of Medicine – presented by Thomas Henn

- This appendix had not yet been seen by the WG members.
- The appendix's primary author was Ravi, co-author Tom, and it was fine-tuned by Panos and Walter.
- This appendix covers advanced technological potential of AI usage and includes some speculative medical subjects such as robotics, nanotechnology, gene editing, nanomedicines, and smart organs.
- The idea of lifting this content out of chapter 10 was supported overall.
- The current title may be edited: 'Beyond vigilance: navigating AI's potential dominion of medicine'.

- If we move outside of PV, the WG members may need to consult with other colleagues within their organisations for input to verify alignment with e.g. medical device groups.
- We need to be careful that some content that we consider as the future is not already current day and vice versa.
- In the title, the element 'Beyond vigilance' may give the impression that this is outside of PV, whereas we mean that in the long-term future, the lines will start to blur. We could remove the 'Beyond vigilance' part. The question is not about the definition but about involvement. PV in clinical practice accommodates these concepts.
- 'Beyond vigilance' is broadening the topic (it is an extension of PV) and yet we are providing a link.
- If this appendix is saying that we are blurring the lines of what PV is, then this is fine, but we do not want to focus on AI in drug development.
- We are addressing product development in a wider sense. This is from early conception of ideas until a product is withdrawn from the market.
- Is this about how we perform PV in the future? It is also asking the question: what is PV? The piece is not exhaustive. We cannot know where AI and advanced technologies may take us.
- The WHO definition for PV includes the word prevention, which includes primary, secondary and tertiary prevention; and so PV naturally has a component in the clinic, as we see in some of the use cases that have been added.
- The clinic is in scope of PV and will be increasingly so as we move to a learning healthcare (PV?) system. However, we need to be careful about AI's impact and what is broader or not. Clinical support systems sometimes push back on adding safety flags for fear of clinician overload. But AI might make this more 'user friendly', e.g. less alerts, and therefore more implementable. It is not that AI enables PV in the clinic per se, as it can and should be done now.
- Other potential titles put forward included:
    - 'Navigating AI's potential during the product lifecycle';
    - 'Navigating AI's potential dominion of medicine and future PV'.
- The WG members are invited to review the appendix and send in feedback by 2 pm on the 25th of April.

Glossary – presented by Douglas Domalik and Niklas Norén

- The WG XIV Glossary Team's own drafted definition received the most votes in the WG-internal survey but the Editorial Team had some concerns, mainly that we did not want to re-create new guidance around AI.
- The Editorial Team feels the best compromise would be to choose the updated definition from the OECD, as below, which would be a more global approach and a little clearer, and which has also been adopted by the EU AI Act recently. The EU AI Act definition is the published definition that received the most votes in the WG-internal survey.

> "An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment."

- By adopting an existing definition, we enhance the likelihood of the report being read and accepted.

- At the provided link above for the OECD definition for AI, there is information about why this definition was written and the different considerations that were taken into account.
- The WG XIV's own definition puts AI into the context of PV and we propose providing this underneath as a corollary note:
  "Artificial intelligence (AI) is a science and technology that enables computers and machines to emulate human learning, comprehension, problem solving, decision making, creativity and autonomy. In the context of pharmacovigilance, the use of artificial intelligence with systems and activities is aimed at enhancing drug safety monitoring, patient safety and regulatory compliance."
  Proposed definition from the Glossary Working group

Use cases – presented by Taxiarchis Botsis

- The set of use cases in the report is an opportunity to be little technical without too much risk to outdate the document.
- Most use cases are supported by a journal or another publication.
- The use cases focus on the following areas: ICSR Processing, ICSR Reporting, Safety Analysis, Causality Assessment, Signal Detection, PV document retrieval, and PV in The Clinic.
- Each use case follows the same structure:
  - Business rationale and challenges;
  - Solution;
  - Results;
  - Compliance with the governance framework.
- Each use case includes a table to show how the use case aligns with the governance framework, including the principles described in the report and how the principles are met in the different phases of the development cycle (collection of specifications/requirements, development and change management, pre-deployment & post-change sign-off, post-deployment & post-change hyper-care, routine use).
- There is an 'activities column' which describes how the principles are met or supported. In some use cases, the principles are applicable/not applicable in the lifecycle phases.

Public consultation process – presented by Sanna Hill

- The Public Consultation is scheduled to start on the 29th of April.
- The draft report will be posted on the CIOMS website on the page dedicated to the WG XIV, together with a comment form – a prepared form was shared – for a duration of six weeks.
- The comments received will be inserted by the CIOMS Secretariat into the draft report for easy reviewing by the WG members.

3. Annex 1: List of participants

Andrew Bate (GlaxoSmithKline), Robert Ball (US FDA), Luisa Barrios (Merck Sharp & Dohme), Adrian Berridge (Takeda), Taxiarchis Botsis (Johns Hopkins University School of Medicine, US), Hua Carroll (Biogen), Mariane Diniz (Bayer), Douglas Domalik (AstraZeneca), Julie Durand (EMA), Piero Francesco Franco (Pfizer), Neal Grabowski (Sanofi), Manfred Hauben (Merck Group), Thomas Henn (Unither), Sanna Hill (CIOMS), Vijay Kara* (GlaxoSmithKline), Benny Ling (Health Canada), Denny Lorenz (independent expert), Monica da Luz Carvalho Soares (ANVISA, Brazil), Elizabeth MacEntee Pileggi (Johnson & Johnson), Yusuke Matsunuga (PMDA, Japan), Jensen Morten (Danish Medicines Agency), Niklas Norén (WHO/UMC), Lembit Rägo (CIOMS), Hans-Jörg Römming (Merck Group), Irene Scholz (Swissmedic), Walter Straus (Moderna), Phil Tregunno (MHRA, UK).

Apologies

Justyna Amelio (AbbVie), Russ Altman (Stanford University, US), Arvind Bellur (CSL Behring), Brian Buch (MHRA, UK), Yauhenia Cherkas* (Johnson & Johnson), Shaun Comfort* (Genentech, Roche), Selin Cooper* (AbbVie), Jean-Michel Dogné (University of Namur, Belgium), Kirsten Egebjerg Juul (Danish Medicines Agency), Julie Girod* (Sanofi), Kendal Harrison (MHRA, UK), Stephen Heaton* (CIOMS), Satoko Hirokawa (Health and Youth Care Inspectorate, Netherlands), Dieter Kempf (Genentech, Roche), Yuki Kikuchi* (PMDA), Hervé Le Louët (CIOMS), Eva-Lisa Meldau* (WHO/UMC), Dirk Mentzer (PEI), Flávia Moreira Cruz (ANVISA, Brazil), Ravi Patel (Unither), Stephen Rosenfeld (North Star Review Board), Thomas Stammschulte (Swissmedic), Panos Tsintis (CIOMS Senior Adviser), James Whitehead* (AstraZeneca), and Brian Yau (WHO).

*Alternate