Council for International Organizations of Medical Sciences

12th meeting of the CIOMS Working Group WG XIV on Artificial Intelligence in Pharmacovigilance

8th of September 2025, virtual meeting

Meeting Minutes

Overview:

Summary

The CIOMS Working Group (WG) XIV on Artificial Intelligence (AI) in Pharmacovigilance (PV) held its 12th and final meeting virtually on the 8th of September 2025. The purpose of the meeting was for the Chapter and Section Leads to consult the full WG about any final unresolved comments from the Public Consultation and for the WG members to question how the comments had been addressed up until this point. The meeting was also used for discussing initial report awareness raising activities following publication.

Minutes of discussion

1. Welcome and opening remarks

- Lembit Rägo, CIOMS Secretary General, opened the meeting as Chairman.
- Lembit proposed putting forward an abstract to promote the WG XIV report for a session at the DIA Global Annual Meeting in Philadelphia, US, over 14th-18th June 2026. The application submission deadline is the 18th of September. Lembit welcomes volunteers.
- Two CIOMS webinars have been scheduled to promote the CIOMS WG XII report on *Benefit-risk balance for medicinal products* for the 10th of September and the 23rd of October 2025. The two events are targeting different time zones. It may be helpful to arrange similar webinars for the CIOMS WG XIV once the report is ready.
- Sanna Hill was rapporteur.

2. Draft report – Overview of chapters and sections

- Every Chapter and Section Lead and co-Lead made a point to thank their fellow WG members for their fruitful team work.
- Every Chapter and Section had editorial changes made.

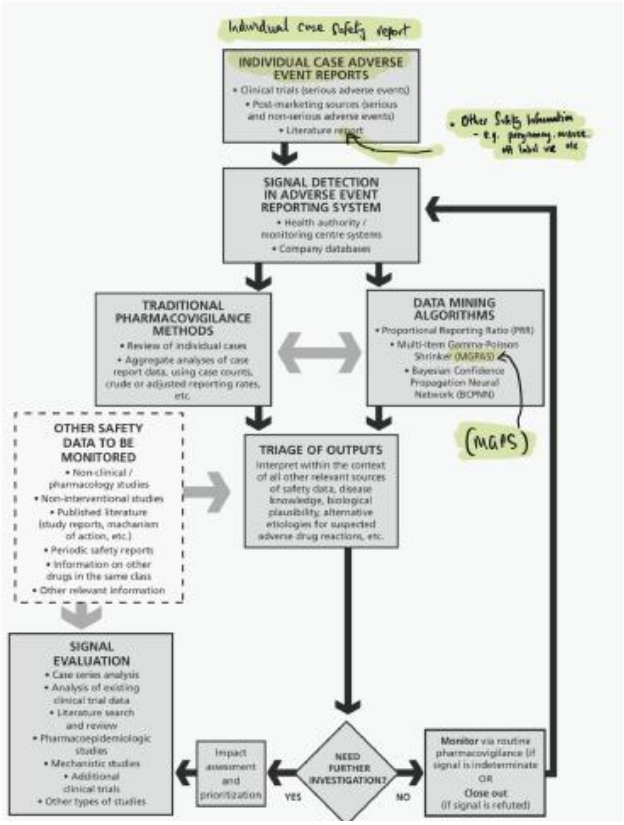

Chapter 1: Introduction - presented by Vijay Kara

- The Figure 1 has been re-named as *Traditional Representative signal management process* and will be revised at the graphic design layout stage, as shown here on the right.
- The AI systems definition has been broadened in line with the definition in the Glossary.

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

- The section on *Early application of generative AI in pharmacovigilance* has been re-named to *Early application of generative LLMs in pharmacovigilance*. Similarly, the same section has been revised as a commentator read the heading as *generative applications* and took it to mean that one could generate text, images and videos; and therefore, the examples given felt out of scope. If we use generative LLMs and constrain their output either through prompting or through post-processing, such that they perform classification, this is not GenAI. We did not want to talk about generative applications per se but about generative LLMs as a new technology.
- Some content has been moved from the Chapter 2 to the chapters on Transparency and Validity & Robustness.
- The Table 1 on *Examples of deployed artificial intelligence solutions in pharmacovigilance described in the public domain* was criticised for not having industry content. Vijay proposed adding a recent scoping review by Maurizio Sessa, which includes a landscape analysis of AI-enabled signal management / signal detection tools across board. This includes vendor solutions but has been peer reviewed and is in print. Niklas will review the article.
- In the regulatory considerations section, one of the more significant questions was about how we chose the guiding principles and established the Table 2 on *Comparison of CIOMS Working Group XIV guiding principles for artificial intelligence across regional and country government institutions, and international organizations*. The US Bill of Rights that had informed the US column is no longer valid due to the change of administration. The new US administration has an action plan but it has not yet updated its guiding principles. The US Bill of Rights has been archived and is provided as the WG XIV report reference.
- It is noted that in drafting our guiding principles, we did not refer to current AI policies from several other countries, e.g. from China, and regulatory/non-regulatory bodies – our list is not exhaustive – but instead, we have included an AI policy tracker.

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

- One of the most visually significant changes was to isolate the examples of risk mitigation approaches. There was a discussion about moving these to the Validity & Robustness chapter but it was decided to keep them in the Risk-based approach chapter.
- Some comments proposed a more near-prescriptive approach with greater detail and they were declined.
- There was a suggestion to include a diagram from the GAMP guidance but it was felt that it was not clearly enough connected to the chapter content and so this was declined.
- Multi-agent systems are considered out of scope is some chapters, e.g. in the Governance & Accountability chapter, but not in others, e.g. in the Transparency chapter, where this features in one of the tables, although not in an in-depth manner. Multi-agent systems were discussed - we are aiming for perhaps a sentence or two at most but not entire sections on the topic.

Chapter 4: Human oversight - presented by Julie Durand

- One comment regarding patient involvement was declined as in our experience patients are not widely involved with most PV systems.
- There were a few commentators who questioned whether the section on *Transformation of traditional roles* was appropriate for the chapter but it was decided that this would be kept.

Chapter 5: Validity & Robustness - presented by Niklas Norén

- A section was added on *Generative output* to distinguish between a generative LLM and a generative application, which generates free form such as text in most cases. We write about performance evaluation and most of the applications we focus on are of this nature - typically classification tasks. It was important to note that this is only one type of application, and that going forward, there will be other types of tasks, where the output will not be simply yes/no in nature, but may be something different. We acknowledge this is a fast-moving field; we do not go in too deep/authoritatively into the subject. Niklas talked through the examples and references used. Vijay offered to propose some additional references.
- The above section, *Generative output*, is distinguished from another section, *Unsupervised learning*, which addresses other applications different from the classification modality.
- The title *Performance evaluation* was re-named to *Performance evaluation for classification tasks.* There was a discussion about whether readers will expect to read about performance evaluation for e.g. other tasks. The surrounding sections address other tasks. Reverting back to the previous title *Performance evaluation* would require re-working the section content. Niklas and Andrew to discuss off-line.
- The title *Reproducibility* was re-named to *Non-deterministic systems* in response to Public Consultation comments, which focused besides the point (reproducibility, replicability, repeatability), rather than on the non-deterministic nature of some systems and what that means.

Chapter 6: Transparency - presented by Niklas Norén

- One of the Key messages was amended as follows to make it clear that the PV system is the focus of the message (with the underlined text added): "To give a comprehensive picture of an AI system's effectiveness and limitations in a pharmacovigilance application, performance results for the specific task should be presented and describe the scope and nature of the test set(s), including definitions of their reference standards and sampling strategies". Andrew proposed changing "comprehensive" to "clear".

- At the end of the section on *Inherent vs post hoc explainability*, a paragraph has been added on generative LLMs. Niklas and Andrew discussed some edits during the meeting and this will be finalised off-line.

Chapter 7: Data privacy - presented by Walter Straus

- The historical background on bioethics has been shortened by two thirds in line with suggestions from a number of commentators.
- Hyperlinks can be provided for key references if this is helpful but this is not done throughout the report largely because many sources can become outdated fairly quickly. The print version would need to be adapted.
- The sections on HIPAA and GDPR have been reduced, i.e. points related to specificity.
- A reference has been added on an article that was written this year relating to issues on ethical considerations in the use of AI in PV.
- Walter requested for help with verifying the content in Tables 6 and 7 on *Data privacy regulations for using secondary data* in Japan and China.
- There was a comment that Figure 6 was out of date but Walter proposed to keep it. Taxiarchis offered to create an up-to-date version.
- Walter was appealing to Taxiarchis for help regarding the interpretability of the green area in the figure.
- In Figure 6, Niklas questioned whether we can connect the models directly with a certain level of data privacy. Taxiarchis originally proposed the figure because it discusses how specific models tackle security issues. Walter concluded that the figure shows that as the size of the models increases, there may be increasing challenges related, and how rapidly this accelerates. Niklas made the point that the bigger risk is leakage through the prompts, not the network size. Vijay suggested that we keep the diagram. The purpose of the diagram is explained in the text underneath. The authors correlated the amount of information that was available in the literature in relation to privacy, and they noted that as the size of the language models have increased, there has been less discussion of privacy and privacy concerns associated with those language models. It was an inference of the developers that they are not discussing privacy as a focus point of the models that are now being released.

Figure 6: Relationship of timing of large language model introduction, parameter size and attentiveness to data privacy
Source: [265]

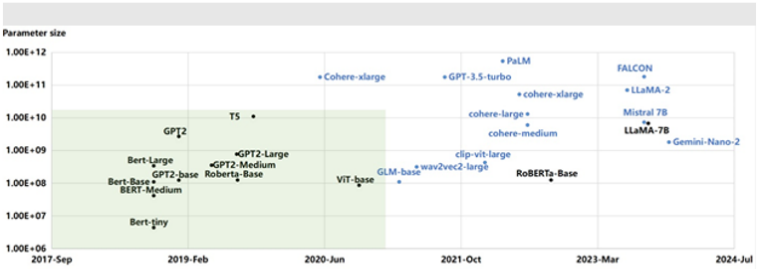

The horizontal axis represents the time of LLMs release, while the vertical axis represents the size of model parameters. Blue dots signify LLM instances not addressed in the literature pertaining to privacy protection, whereas black dots indicate those that have been examined in such literature. The green backdrop delineates the central cluster zone of LLMs with the potential to facilitate privacy protection.

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

- Elizabeth thanked the technical experts on the Editorial Team who have been helpful in getting some of the chapter paragraphs framed correctly with the appropriate references.

Chapter 9: Governance & Accountability - presented by Hua Carroll

- The last column of the Governance grid has been removed and its contents have been transferred to the section underneath describing each lifecycle phase.
- A lifecycle diagram has been added (integrated) to the top of the table.

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

- Tom has added some references to the chapter.
- Content has been added on the topic of the ongoing evolution of AI infrastructure.
- There was a comment suggesting to add principles on Human factors and ergonomics (HFE), and although Tom did not feel this was necessary, as he has provided a reference and added a descriptive sentence about the HFE scientific discipline.
- The Black Swan event content has been moved from the Fairness & Equity section in Chapter 10 to the Risk-based approach section in Chapter 10, as this seems to be a more appropriate place. As Black Swan events are addressed in another chapter earlier in the report, Tom proposed removing the text in this chapter and making a cross-reference to the earlier place. Andrew expressed a preference for mentioning the Black Swan event in the previous chapter, Chapter 3.
- There was a comment that human oversight may not be sufficient for certain types of highly complex applications. Tom has added text about not only human-in-the-loop, but about how we may need to develop additional tools and training to maintain proper oversight.
- One open comment relates to cross-functional, multidisciplinary PV teams. Tom has added text about meeting demands of validating AI-enabled systems.
- There has been a discussion about what PV is and what PV is not, and Tom decided to include nanotechnology and smart organs as examples.
- On the subject of data privacy, some examples were requested of federated learning and blockchain, and additional information with references have been added.
- Additional content has been added on heterogeneity of regulatory authorities on data privacy. This concerns restrictions and data sharing across regions.
- A commentator requested more information on the metrics and measuring of effectiveness of what we consider as success to be added to the conclusion section, but it was felt that this was not the appropriate place.
- Niklas suggested adding a sentence at the end of the section on Validity & Robustness, in the light of the discussions on the section on *Generative applications*, to mention how (along the lines of) “this area is evolving rapidly, it is in its early days, these are the current common practices, and it will deserve more thought”.
- Similarly, when we move away from simple, compartmentalised AI solutions, and the producers test their performance, when they are going to be much more integrated, when we are interacting with an AI solution to perform a specific task, and the human is heavily engaged in the exchange, how do we ensure validity and robustness in that context? Should PV individuals develop new skillsets which fit better with human oversight?

Appendix 1: Glossary: presented by Douglas Domalik

- Older terms that were no longer applicable have been removed.
- Some of the remaining terms have been revised based on the feedback received at the last WG meeting in Geneva, and based on comments from the Editorial Team members. The last comments were made based on Validity & Robustness and Transparency.
- The final Glossary changes will be circulated among the Chapter / Section Leads to ensure consistency throughout the report. No major changes are expected.

Appendix 3: Use Cases: presented by Vijay Kara

- The Team has added some pre-text to describe how to utilise the Use Cases, i.e. that they provide practical illustrations of the considerations that have been discussed earlier in the document.

- In response to feedback from the Public Consultation and discussions at the previous WG meeting, a new section has been added to each of the Use Cases, the *Challenges and Lessons Learned.* This extends the previous scope of the Use Cases which was to give the practical considerations of the Use Cases in addition to the Governance Framework.
- In the majority of the Use Cases, where applicable and if relevant, additional details have been added around the results and performance characteristics. Some of the feedback received was that that section was a little bit limited in terms of the information provided.

Appendix 4: Content related to Explainability and to Fairness & Equity - presented by Elizabeth MacEntee Pileggi and Niklas Norén

- There is one reference to add on the Fairness & Equity content from Beth and Monica.
- The Explainability content is yet to be reviewed by Niklas and Vijay.

Opinion piece on what was nearly the Appendix entitled "Beyond Vigilance: Navigating AI's Potential Dominion of Medicine"

- This text will not be a part of the CIOMS WG XIV report.
- Lembit is exploring avenues within CIOMS for publishing this opinion piece as a stand-alone CIOMS publication.
- Ravi is sending an updated version.

3. Communication strategy after publication

- CIOMS webinars
    - Lembit proposed holding two webinars targeting two different time zones.
    - We recommend recording presentations in advance and hosting a live Q&A session with a panel.
    - Please see an example of a webinar to promote the report of the CIOMS WG XIII on RWD/RWE.
- DIA Global Annual Meeting in Philadelphia, US, over 14th-18th June 2026.
- Scientific article to be submitted in an open access publication.
    - CIOMS would pay publication fees.
    - The WG discussed the possibility of approaching *Drug Safety*.
    - The article would be a summary of the report, which is slightly unusual to submit for publication.
    - We could consider packaging the article such that the core of the article is about the report with a surrounding story about something more.
    - It tends to be easier to get this type of article submitted if the WG has contacts within the editorial committee of a publication and we have a prior consultation.
    - There are several WG members on the editorial board of *Drug Safety* and the Editor-in-Chief would make the decision.
    - Andrew questioned to what extent the recommendations and guidance we provide generalises beyond PV because the broader medical journals may not be interested. If we go too far out, we may be out of our comfort zone / area of expertise.
    - Andrew suggested an article angle: "How to think about the use of AI in high-risk application-type areas?".
    - From an impact perspective, the higher profile medical journal we go to, the more visibility there will be, but there is also a correlation with workload.
    - Andrew is on an editorial board of a few journals and offered to enquire.

- Lembit felt the WG's report is applicable: 1) to the broad medical context, especially in the US, including in everyday medical practice where doctors use AI-assisting tools; and 2) to regulatory affairs, as PV is partly regulatory business.
- Andrew agreed that a piece around PV and to what extent it can or should generalise wider could certainly be done. The question is, we need to be very careful to be clear where we do not think things may be not generalised, but they are lower priority, for example, certain parts of patient engagement, the nature of data, and what can be shared or not shared will vary by application and things.
- Further ideas are welcome.

4. Closing remarks

- Lembit thanked everyone for their time and warmly commended all for their contributions.

5. Annex 1: List of participants

Andrew Bate (GlaxoSmithKline), Adrian Berridge (Takeda), Taxiarchis Botsis (Johns Hopkins University School of Medicine, US), Hua Carroll (Biogen), Mariane Diniz (Bayer), Douglas Domalik (AstraZeneca), Julie Durand (EMA), Kirsten Egebjerg Juul (Danish Medicines Agency), Thomas Henn (Unither), Sanna Hill (CIOMS), Vijay Kara* (GlaxoSmithKline), Dieter Kempf (Genentech, Roche), Benny Ling (Health Canada), Denny Lorenz (independent expert), Elizabeth MacEntee Pileggi (Johnson & Johnson), Yusuke Matsunuga (PMDA, Japan), Eva-Lisa Meldau* (UMC/ WHO), Dirk Mentzer (PEI), Flávia Moreira Cruz (ANVISA, Brazil), Jensen Morten (Danish Medicines Agency)*, Niklas Norén (UMC/ WHO), Lembit Rägo (CIOMS), Hans-Jörg Römming (Merck Group), Panos Tsintis (CIOMS Senior Adviser), and Walter Straus (Moderna).

Apologies
Russ Altman (Stanford University, US), Justyna Amelio (AbbVie), Luisa Barrios (Merck Sharp & Dohme), Arvind Bellur (CSL Behring), Brian Buch (MHRA, UK), Yauhenia Cherkas* (Johnson & Johnson), Selin Cooper* (AbbVie), Jean-Michel Dogné (University of Namur, Belgium), Piero Francesco Franco (Pfizer), Julie Girod* (Sanofi), Neal Grabowski (Sanofi), Kendal Harrison (MHRA, UK), Manfred Hauben (Merck Group)*, Stephen Heaton* (CIOMS), Satoko Hirokawa (Health and Youth Care Inspectorate, Netherlands), Yuki Kikuchi* (PMDA), Hervé Le Louët (CIOMS), Monica da Luz Carvalho Soares (ANVISA, Brazil), Ravi Patel (Unither), Stephen Rosenfeld (North Star Review Board), Irene Scholz (Swissmedic), Thomas Stammschulte (Swissmedic), Phil Tregunno (MHRA, UK), James Whitehead* (AstraZeneca), and Brian Yau (WHO).

* Alternate WG member