HMA/EMA Multi–stakeholder workshop on Artificial Intelligence (AI)

2025 Workshop Report

Day 1

Welcome and vision: Setting the Stage for AI in Medicines lifecycle

The workshop was opened by **Emer Cooke**, Executive Director of the European Medicines Agency (EMA), **Karl Broich**, President of the Bundesinstitut für Arzneimittel und Medizinprodukte (BfArM), and **Marco Marsella**, Director at the European Commission DG SANTE.

Emer Cooke highlighted that while previous AI workshops explored the *promise* of AI, this workshop focuses on the *practice* of AI. With these new opportunities come responsibilities. We must shape this transformation of our work in an ethical way and establish guardrails for industry. This workshop is about the responsible development of frameworks that foster innovation. All stakeholders have a role to play. To summarise the essence of the NDSG's work this year: **"we're all in this together."** The Network has accomplished much through intentional collaboration.

Karl Broich noted that the workshop marks a milestone—it is the first under the steering of the Network Data Steering Group (NDSG), which was established earlier this year following the transition from the BDSG. The NDSG is uniquely positioned to enable better medicines regulation through better data. The co-chairing of Peter Arlett from EMA and Karl Broich from HMA demonstrates the profound partnership between EMA and the national medicines agencies. Every day, we see both the opportunities and challenges presented by AI. Several NCAs have already developed AI tools. The spirit of this workshop is about **working together, learning from each other, and building foundations** for a future where data and AI serve public health.

Marco Marsella reiterated that AI in healthcare is here already. The benefits are evident—not only for patients but also for regulators and policymakers. We are introducing AI-driven technology into a sector with a high degree of regulation. The EC has been working to support efforts to use data to accelerate both the development and adoption of innovative tools. The legislative ecosystem includes three key aspects:

1. **Data access and sharing** (regulation on the European Health Data Space),
2. **Ensuring security and trust of access** (which involves both technological elements and equitable access to benefits), and
3. **Safety and robustness** (including the AI Act, clinical trial regulations HTA).

Several strategic actions are underway: creating infrastructure conducive to AI development, holistic skill-building, and support across the value chain.

Keynote speeches

Alain Cornet, from **LUPUS Europe**, introduced LUPUS GPT, an AI tool developed to improve patient access to reliable high-quality information, aiming to improve treatment responses through better education. Built with strong patient involvement, the tool provides accurate and empathetic answers in 100 languages without storing personal data. Since its launch, it has answered over 60,000 questions and continues to evolve with features like an "Easy Lupus" mode for simplified responses and plans for a spoken version for visually impaired or lower literacy users.

The initiative began in 2018 with LUPUS100, a multilingual resource created through collaboration between patients and doctors. In 2024, discussions on AI led to rapid development: a beta model was built in one month and validated by GPs. By September, LUPUS GPT was live in 100 languages, trained on just 119 documents selected by patients and clinicians.

Alain stressed that while AI can enhance healthcare, it cannot replace patient expertise or peer review. He warned against misuse, noting that technology must complement - not substitute - the human touch. "We cannot let AI erase the unique experience patients bring."

Regulator AI use cases: improving efficiency and enhancing decision making

Moderators: Zaide Frias (EMA) and Georg Neuwirther (Austrian Medicines and Medical Devices Agency - AGES)

AI@MPA toolbox - AI for EU medicines regulation from the Swedish Medical Products Agency,

Gabriel Westman, Seamus Doyle, Samuel Fransson (Swedish Medical Products Agency - MPA)

The **AI@MPA toolbox** is a suite of services designed to support AI regulation. It includes both encoder-based and generative AI tools, hosted on a secure cloud-based platform. Each month, the system processes approximately 15,000 regulatory requests. The most widely used application is **REGULUS**, which provides regulatory support to multiple use cases. REGULUS is a self-deployed MPA service featuring three assistant modules based on open-weight LLMs. It can process text documents, images, and audio.

RADAR is an upcoming tool that aligns with the current pilot on analysing patient-level clinical study data in medicine evaluation. RADAR enables quick and easy analysis of raw data, transforming an analysis prompt into actionable outputs. Users can work with existing datasets or upload their own. The app is scheduled for release before the end of the year.

SCHEMA (*Semantic Chemical Embeddings Map*) assists in predicting expected properties of new substances, such as adverse reactions, interaction pathways, and chemical similarity. SCHEMA has a database with approximately 600,000 substances vectorized using a pretrained chemical language model. Users upload a chemical substance file, and the tool returns the most similar substances along with detailed metadata.

Additional upcoming projects were also discussed.

PHAIR - PV by AI in Real time - preliminary results from Danish consortium and potential for the future,

Claus Møldrup (Danish Medicines Agency - DKMA)

PHAIR is a project sponsored by the Danish Innovation Fund, universities, and DKMA. Its goal is to proactively identify new drug reactions and validate adverse drug reactions (ADRs). The current approach relies heavily on spontaneous reports, which are often incomplete due to underreporting. PHAIR uses AI combined with real-world data to match medicine use with symptom reports - even when healthcare professionals have not initially made the connection. The process spans data preparation, analysis, communication, and implementation with relevant stakeholders. The team is developing a minimal viable product (MVP) to validate outputs from unstructured and registry data, aiming to accelerate the time from case detection to verification. They leverage "cradle-to-grave" data to capture confounding variables and map patterns of causes, effects, events, and outcomes.

An example shared was the study of long COVID using MRI brain scans. The scans detected microbleeds, and by comparing images before and after COVID infection, researchers identified an increased risk of microbleeds post-infection. This illustrates how AI-driven analysis can uncover previously unseen associations.

How AI creates consistency across assessment reports

Paul Zabbal (EMA)

Ensuring appropriate consistency in assessment reports presents several challenges: varying writing styles and levels of detail among assessors, repeated information and overlapping sections, time spent on parts of the assessment that add little assessment value but are still required, and quality control that often occurs late in the process and is fully manual. Consistency is critical to regulatory integrity, building reliability, enhancing efficiency, and ensuring compliance. AI can support this goal as a partner—not as a replacement for assessors. It can be used through structured prompts and checklists to guide assessors through key sections, surface similar past assessments or regulatory wording, and help draft summaries or checkpoints based on content.

The plan is to create a **prompt library** that enables assessors to follow a structured approach, perform cross-checks on content, and validate their work. A suggested quality assurance use case involves assessors accessing different product or project documents and applying prompts tailored to this scenario. This exploratory work is currently being developed at EMA.

If successfully implemented, the prompt library could serve multiple use cases: drafting early summaries of assessments to support committee discussions, reviewing scientific advice letters to ensure structure and tone follow common patterns, performing consistency checks across assessment reports, and assisting in generating factual texts based on essential checklists. Ultimately, the goal is to deliver solutions that support a wide range of regulatory tasks.

Industry AI use cases: from innovation to impact

Moderators: Hilmar Hamann (EMA) and Pelle Persson (MPA)

GxP Validation of AI Sub-Systems used in clinical trials (CTs) and pharmacovigilance (PV): Ensuring Compliance in pharmaceutical R&D Regulated Environments

Alexander Kunz (European Federation of Pharmaceutical Industries and Associations - EFPIA)

Industry operates under various international regulations and guidance documents. Starting with a computerized information system, the next step may involve an AI sub-system - for example, for data extraction - followed by an AI model. These three levels form the **"onion model"** developed for this purpose. While the context of use can vary, it must be clearly defined during development. Creating a validation plan, validation report, and verification activities is essential in the planning and execution phase. During the maintenance phase, ongoing monitoring ensures that all AI tool operations remain within acceptance criteria.

A **Proof of Concept** serves as an evaluation activity to determine feasibility and apply a risk-based approach. Sub-system verification must ensure fit-for-purpose data, predictive power, effective human-AI interaction, stability and robustness, and mitigation of bias.

The ongoing monitoring phase follows several guiding principles. Human oversight should not be symbolic; it is required during operation and monitoring to identify when intervention is needed. Operators must receive effective training. The request to EMA is for **risk-based recommendations** on clinical development and pharmacovigilance. In areas without current limitations, industry seeks additional guidance.

Pre-MA development strategies and regulatory registration

Stéphane Bellec (International Plasma and Fractionation Association - IPFA)

IPFA identified a key business need: regulatory strategy departments must review the drug regulatory and competitive environment, including guidance from authorities, publicly available data, and product information. Another need is an automated form of regulatory intelligence to assess new guidance and compare it with previous versions.

To address this need, IPFA considered two options: an in-house solution versus an out-of-the-box solution. A decision criteria list was developed, including factors such as implementation timeline and expected investment from end users for system deployment. In IPFA's case, they opted for an out-of-the-box solution.

The process involved defining business needs, establishing test cases, creating three internal tester groups, and collaborating with a vendor during a three-month pilot. The pilot revealed significant time-saving opportunities and demonstrated that the system could retrieve highly specific information - cases where manual efforts would likely fail due to the complexity or "needle in a haystack" nature of the data.

A key takeaway: initiate early internal discussions to ensure the chosen solution complies with the AI Act and organizational policies; prepare the pilot phase with well-defined test cases; recognize that testing requires time; and ensure the solution provides constant access to a hotline that is both technically proficient and understands business needs.

Using AI as quality control to accelerate patient access to ATMPs

Toni Manzano (Parenteral Drug Association - PDA)

AI can significantly enhance quality control in cell therapy. **ATMPs** (Advanced Therapy Medicinal Products) present a strong use case for AI because they are inherently variable, and robust quality control requires integrating knowledge from development through to commercial manufacturing.

Humans are not naturally equipped to compare numerous variables simultaneously - but AI is. The audience was walked through the operational procedure for ATMP manufacturing, emphasizing that when developing an AI tool for ATMPs, the problem statement must be clear and comprehensive, requiring interdisciplinary input. Toni also provided an overview of ATMP unit operations and their components.

Several AI use case opportunities were highlighted, including an example demonstrating how AI can advance personalized medicine. By using patient biomarkers as data inputs and evaluating clinical efficacy, the tool can help determine which therapies are best suited for patients with different biomarker profiles.

Spotlight talk: Opportunities for international convergence

Moderator: Karl Broich (BfArM)

Joerg Zinserling (BfArM) presented on behalf of the European Medicines Regulatory Network (EMRN). The EMRN has three key AI drivers: regulate AI in the medicines lifecycle, increase efficiency, gain deeper insights into health data.

Joerg presented the **AI workstream** from the NDSG's AI workplan. Collaborations are a core part of the workplan and help expedite learnings and promoting certainty and predictability in a rapidly evolving environment.

In practice, collaborations with multiple stakeholders include engagement with the EU Agencies Network (i.e. between different EU-level agencies), knowledge sharing on regulatory use of AI and on guidance development at the International Coalition of Medicines Regulatory Authorities and bilateral collaborations such as the one between EMA and the U.S. FDA.

Collaboration with the FDA has included discussions on common AI guiding principles and mapping terminology. These efforts are valuable because international collaboration promotes regulatory certainty and supports accelerated innovation. It is also easier to explore convergence opportunities before guidance drafting begins, paving the way for broader discussions on global regulatory standards for AI.

Looking ahead, planned actions include publishing **EMA-FDA joint Guiding Principles for Good AI Practice**, developing an AI glossary, and exploring further convergence opportunities.

Anindita Saha (FDA) added that FDA is seeing AI leveraged across discovery, nonclinical research, clinical research, manufacturing, and post-market safety monitoring. FDA has already received over 800 submissions involving AI.

CDER's core principles include ensuring adaptive, risk-based regulation rather than a one-size-fits-all approach and fostering collaboration across a broad set of stakeholders. Next steps include strengthening collaboration within FDA and with international regulators, harmonizing glossaries, guiding principles, and standards, and advancing regulatory approaches that support innovation.

Day 2

Spotlight talk: The Algorithm will see you now: Rethinking Clinical Trial matching using AI

Moderator: Karl Broich (BfArM)

Julián Isla (European Dravet Syndrome Federation and Foundation 29) presented the, in development, digital tool – **TrialGPT**. Julián began his presentation by recalling how the discovery of DNA in 1952 transformed science. Working with language has always been a computational challenge with limitations. Similar to the DNA breakthrough, the development of transformers changed everything - yet language remained a bottleneck until the advent of Large Language Models (LLMs), which marked the next major leap.

Foundation 29 is developing TrialGPT, a platform designed to retrieve, match, and sort clinical trial information. A key challenge today is finding patients for trials, even when they are registered in databases like CTIS. TrialGPT allows users to upload a medical record and select demographic identifiers. The platform then identifies relevant clinical trials in the user's region. Its LLM simplifies trial descriptions and provides inclusion and exclusion criteria, along with a recommendation on patient eligibility. The tool can even draft an email to the study contact, including relevant patient details.

Development of the tool was fast. However, several challenges were considered, including bias and fairness, transparency and explainability of results, and evaluation of outputs to ensure accuracy. An evaluation framework is critical during planning. Developers must also guard against overfitting; insufficient training in certain areas can lead to incorrect results due to low coverage. Compliance with GDPR, AI Act, MDR, and other regulations is essential, particularly when software supports clinical decisions under medical device rules. Adaptive LLMs currently fit heterogeneously into existing regulations, creating overlapping requirements that are especially challenging for small companies—but a concern for all. Julián emphasized the need to balance regulation with innovation by aligning requirements while remaining agile and responsive.

Regulating AI: Updates and Perspectives

Moderators: Peter Arlett (EMA) and Georgia Gavriilidou (EMA)

AI Act updates

Tatjana Evas (EU AI Office)

Tatjana emphasized that regulation should be seen as a benefit rather than an obstacle. Two words form the foundation of EU AI policy: **excellence** and **trust**. She presented the **AI Continent Vision**, which includes five components: computing infrastructure, data, skills, algorithm development and adoption, and regulatory simplification. The EU is funding initiatives across these areas. There are parallel and converging strategies in the AI space broadly, as well as within healthcare and medicine. One notable example is the Virtual Human Twins Initiative, which offers significant benefits for drug development.

The **AI Act** was developed with core principles aligned to a lifecycle approach, a risk-based framework, and existing product regulations. It does not only regulate LLMs but also systems

trained through machine learning, making its scope much broader. The Act promotes safety and trust, fosters innovation, complements other EU laws, and provides an adaptable, future-proof approach. Importantly, it focuses on applications rather than specific technologies.

Tatjana outlined the AI Act's risk levels and explained how they align with the **Medical Device Regulation** and other sectoral legislation. For **high-risk AI systems**, mandatory requirements must be met before market entry, including an EU declaration of conformity and CE marking by the provider. Upcoming compliance tools include harmonized standards and Commission guidelines. She also presented the AI Act timeline, noting that some requirements are already in effect, with additional obligations phased in through 2027. Current priorities include establishing governance structures, issuing practical guidance, and engaging stakeholders to support compliance. Tatjana introduced the AI Act Service Desk and encouraged participants to explore it.

AI Guidance update – focus on GMP Annex 22

Ib Alstrup (DKMA)

In traditional software, inspection processes are well defined and cover all GxP areas. However, for machine learning models, development occurs through learning rather than programming. To inspect these models, one must review the tests and the test data used to document model performance.

Members of the **GMP Annex 11 and 22 Drafting Group**, which includes representatives from the EU/EEA and several other countries, identified the need to update Annex 11 and to include a section on AI. A public consultation on a concept paper highlighted 33 points, including topics related to AI. The outcome was that the AI section would become a standalone document—**GMP Annex 22**—to allow future updates without revising Annex 11 entirely. Both Annex 11 and Annex 22 have been through public consultation, and feedback is now being implemented. The annex defines the scope as models used in critical applications with direct impact on patient safety, product quality, or data integrity. Within this scope, it explicitly excludes Generative AI and LLMs in its current draft.

The draft focuses on recommendations for documenting intended use, specifying the data used for testing, and ensuring that test results meet predefined criteria.

AI Guidance development

Tobias Fellinger (AGES)

The AI Act and the **EMA Reflection Paper** on AI are key guiding documents anchoring guidance development. The reflection paper acknowledges the potential for various use cases for AI throughout the medicinal product lifecycle. Each use case may be associated with different risk levels, applicable standards and regulations, and specific documentation requirements. It also outlines technical aspects and guiding principles.

Additional guidance is under development at both European and international levels, including the draft guideline **ICH M15**, **GMP Annex 22**, **Guiding Principles for Good AI Practice**, a Glossary of AI Terminology, and guidance from **PRAC** and **MWP**.

Several platforms enable engagement with regulators, such as **Portfolio and Technology Meetings**, the EMA/HMA **Focus Group on AI with industry stakeholders**, Innovation Taskforce Meetings, and Scientific Advice/Qualification of Novel Methodologies.

Enabling responsible innovation with AI

Moderated by: Emmanuel Cormier (EMA) and Konstantina Boumaki (European Patients' Forum)

AI Ethics

Alessandro Blasimme (ETH Zürich University)

The projected market size for AI use in clinical trials indicates significant growth in the near future, offering potential benefits for productivity in clinical development for sponsors, regulators, and trial participants. While these gains are promising, they come with considerable ethical uncertainties.

Even seemingly simple applications, such as site selection, can raise fairness concerns by prioritizing certain populations for access to treatment. Historical biases may influence these decisions. Similarly, matching trials and patients' aims to balance risks and benefits, but risks may be overlooked or benefits overstated. Subpopulation inclusion remains a challenge, as historic underrepresentation—such as women in trials—could persist.

AI can perform well in these areas, but ethical quality must be tested, and integration should proceed cautiously. Decentralized clinical trials are particularly suitable for AI-driven design, and currently, 25% of global Phase II and III decentralized trials integrate AI. However, regulators report limited disclosure of AI use, raising questions about transparency and creating an ethical imperative to document AI involvement in trial databases.

Recent developments in **AI ethics** include reviews of ethical principles and frameworks for human-centered oversight. In Europe, **the right to explanation** was reinforced by a recent CJEU ruling, and the **Code of Practice for General-Purpose AI Models** provides guidance on documentation. Explainability approaches have limitations: human-centered methods often leave decisions as black boxes, while machine-centered approaches risk confirmation bias. ETH Zurich's **"ethics-in-the-loop"** framework addresses acceptable bias levels for performance benefits. Guidance suggests informed consent models are evolving toward AI-ready consent, but opting out after data has trained a model remains nearly impossible for now.

Research priorities update

Marie Orre (EMA)

Preliminary results from an AI research priorities survey were presented. The survey sought broad input from stakeholders following earlier discussions during the 2024 AI workshop and aimed to inform regulatory science research on AI in the medicinal product lifecycle. It covered seven domains, each with four research questions, and respondents were asked to rank these questions by importance. The domain ranked highest overall was **"accuracy and reliability of AI tools,"** followed by **"data governance, confidentiality, and consent."** Rankings of individual research questions varied across stakeholder groups and levels of AI experience.

Next steps for analysis include conducting a secondary review to stratify results by different groups and refine priorities based on stakeholder input. The results will be reviewed and endorsed by NDSG, with a final publication of research priorities expected in **Q1 2026**.

How EMRN is advancing AI implementation: AI use cases roadmap update

Joaquim Berenguer Jornet (EMA)

One of the highlights from last year's AI workshop was that several NCAs were unaware of existing AI solutions available within the network. The **NDSG Data and AI Workplan** aims to identify use cases and foster collaboration on both existing and new solutions.

The AI use case roadmap includes a collection of use cases from EMRN. The process began with assessing the status quo, including compiling AI use cases, followed by evaluating them based on business criticality and technical feasibility. This was complemented by an **AI landscape deep-dive workshop** to discuss solutions with IT offices across the network.

The second step involved synthesizing EMRN-wide needs, including clinical trials, marketing authorisation, and pharmacovigilance. Next steps will require consolidating knowledge and expertise, identifying tools needed for solution delivery, and creating a delivery plan. The proposal will be presented to NDSG for approval in the coming weeks.

From Plan to Practice: Executing the Data and AI Workplan

Moderated by: Peter Arlett

The final session of the workshop was a **panel discussion** focused on the execution of the HMA-EMA NDSG workplan—particularly its AI workstream—and on collecting input from stakeholders. The session was moderated by **Peter Arlett (EMA)**.

Luis Pinheiro (EMA) opened the session by presenting an update on the AI workstream, outlining progress for 2025 and planned next steps for 2026. This was followed by a panel discussion featuring:

- **Thomas Brookland (EFPIA)**
- **Stephen Gilbert (Technische Universität Dresden)**
- **Alice Accorroni (European Academy of Neurology)**
- **Gregorio Sambataro (European Heart Network)**
- **Diana Teixeira (AnimalhealthEurope)**

The panel explored key topics related to AI in the regulatory and healthcare space. Stakeholders highlighted themes to reinforce or further explore in the coming year, including:

- **Explore the use of sandboxes** – consider creating regulatory sandboxes for the use of AI
- **Regulatory guidance** – progress in guidance development and increase transparency in the development planning
- **EMRN AI transparency** – maintain visibility on AI activities under the workplan, including regulatory AI use cases
- **International harmonisation** – continue exploring opportunities for global regulatory convergence
- **AI for patient safety** – support initiatives that provide reliable information and enhance safe use of medicines

- **Health data infrastructure** – enable research and clinical applications
- **AI literacy** – promote understanding across stakeholders
- **Veterinary specificities** – ensure veterinary medicines are systematically considered for inclusion in AI initiatives
- **Expand access to real-world data sources** – include patient-recorded data
- **Ethical AI governance** – address bias mitigation and fairness
- **Responsible experimentation** – enable safe and ethical experimentation
- **Enhanced collaboration** – foster cooperation and a common use of terminology across stakeholders

Audience input was gathered via **Slido**, revealing the top three priority actions:

- **International harmonisation** – pursue global regulatory convergence
- **Regulatory guidance** – advance guidance development and increase planning transparency
- **Explore the use of sandboxes** – consider creating AI regulatory sandboxes

A **fireside chat** followed to reflect on these priorities, and the session concluded with an interactive **Q&A**, allowing participants to engage directly with the speakers.