

AI in Drug Discovery and Development
Prof. Rajnish Kumar
Dept. of Pharmaceutical Engineering and Technology
IIT-(BHU), Varanasi
Week-11
Lecture-54

Welcome to the course "AI in Drug Discovery and Development." In today's session, we will talk about the regulatory considerations for AI implementation. As we have seen recently, a lot of government agencies, world leaders, and people who have played an important role in the development of AI have shown concerns regarding the misuse of AI. So, in this lecture, our objective will be to understand what those regulatory considerations are and what regulatory bodies are making rules or regulations regarding the implementation of AI in drug discovery and development. So, by the end of this lecture, you will be able to understand the regulatory significance of AI implementation across the drug development life cycle. Compare global regulatory frameworks, including those of the FDA, EMA, PMDA, and CDSCO.

Examine key regulatory considerations for AI-based software as a medical device. Explore best practices such as GMLP and regulatory sandboxes for building compliant and scalable AI systems. As well as reflecting on ethical and legal implications surrounding AI adoption in clinical and pharmaceutical settings. So, if we talk about it, we have seen a lot of, you know, uses of AI in drug discovery and development.

So, now why do those regulations matter? Because AI is influencing key decisions across the R&D pipeline, from discovery to clinical trials, as well as post-market surveillance. So, the regulatory oversight is critical to ensure, first of all, safety, where preventing harmful or inappropriate AI-driven decisions is a major objective. Another one is the validity; it requires transparency, reproducibility, and scientific rigor, and fairness is another important aspect, where it addresses bias impacting underrepresented populations as well as accountability, which defines responsibility and enables traceability. So, as AI scales, clear regulations are essential to safeguard patients and ensure ethical effective implementation. So, this is a kind of recap of how AI has been used in the drug discovery and development cycle.

So, starting from the drug discovery process where AI is being used to predict novel drug targets using deep learning and to screen millions of compounds via QSAR and molecular docking, It could be structure-based methods or ligand-based methods, where we have been using them for identifying hit compounds and designing novel molecules with the help of generative AI. and it enables rapid accurate protein structure prediction as well which is

absolutely used in the structure based drug design. And then, in preclinical development stages, AI is helping to predict ADMET toxicity properties using ML models and AI models. It simulates efficacy and toxicity in virtual animal models as well, and it also assists in formulation selection and delivery optimization. While in clinical development, it is being used to optimize trial design, which could include virtual trial adaptive models or dropout predictions.

It enhances patient recruitment through EHR mining. And then it also enables risk-based monitoring and anomaly detection, as well as creating synthetic control arms using real-world data. And this we have already seen, like how AI has been instrumental in designing and executing clinical trials. And then in pharmacovigilance and post-marketing surveillance, AI is helping to detect adverse events via NLP and ML from diverse sources, including social media content as well. And then it also helps to prioritize safety signals using Bayesian and deep learning models, as well as supports labeling updates and REM's decisions through AI-driven analysis.

So, let us see what those regulatory bodies and regulatory agencies are saying about the implementation of AI. So, we talk about the global regulatory framework coming from the FDA. The FDA adopts a risk-based lifestyle approach to regulating AI in drug development. It emphasizes scientific rigor, transparency, and continuous oversight of AI systems as AI systems evolve. It also supports adaptive regulation to match the dynamic nature of AI and ML technology because, you know, these AI and ML technologies are evolving very fast; like, every month, new technologies are coming up.

So, catching up on those developments in technology and implementing those regulations. So, that is again a big challenge. That is why the FDA is using adaptive regulation. So, where those can be, you know, changes in the technologies as well. So, if you look at the oversight of AI use across the drug development life cycle that the FDA is implementing.

So, in the discovery where the AI applies in the you know target identification molecule generation. So, the FDA's regulatory focus is, you know, at least the discovery phase, which is not regulated yet unless it is used in the regulatory submission. And then for preclinical studies like toxicity prediction or PK/PD modeling, the FDA regulation says that the use of AI must align with GLP if it is part of the safety package. And then for clinical trials where it is being used in the design, recruitment, or synthetic control arms. So, it must meet good clinical practices, and you need to disclose its use in the investigational new drug applications.

And in the CMC, in the manufacturing optimization or real-time release, it is being used. According to the FDA, it requires validation under the CGMP. And then for regulatory

submission, document summarization, or automated analysis, the output must be reviewable and auditable. So, it means that the AI tool we have been using should be transparent and explainable as well. So, talking about the EMA, the European Medical Agency, EMA adopts a cautious yet supportive stance on AI in drug development.

It prioritizes transparency, ethical AI use, and patient centrality. It strongly aligns with the EU principles of safety, data protection, the GDPR, and public trust. So, MS's focus is, for example, on the discovery stage where it is being used for target prediction or identification. So, MS's focus is on low regulatory interest unless it is used in submission. In the preclinical stage, it must be robust and GLP-aligned.

In the clinical trials, it shall again be GCP-compliant, explainable, and disclosed in the protocol. And then for CMC, the transparency in signal logic and GDPR compliance is mandatorily required; only then can they approve its use. Talking about the PMDA in Japan, the Pharmaceutical and Medical Device Agency of Japan. So, it emphasizes safety efficacy and quality while actively supporting AI driven healthcare innovation. So, the subcommittee on SAMD software as a medical device reviews AI/ML-specific regulatory issues, including bias, data reuse, and database development.

So, the key regulatory focus areas are plasticity, managing AI's adaptability over time, predictability, which ensures consistent and explainable outputs, and data quality using clinically relevant and unbiased datasets. Autonomy defines acceptable levels of self-learning versus human oversight, as well as life cycle management where oversight beyond initial approval is needed. So, coming to the requirement of the Indian regulatory agency, which is CDSCO. So, there are no dedicated AI guidelines yet for drug development, and AI tools are assessed under existing frameworks for drugs and medical devices. And it aligns with the international standards, including those from the IMDRF (International Medical Device Regulators Forum).

And the upcoming digital drug regulatory system DDRS, it aims to modernize regulatory workflows and may integrate AI in the future. So, if you look at the differences in regulatory maturity across the regions, coming again to the FDA, it has reached an advanced level where they have published an AI/ML action plan in 2021. They also have this pre-certified program for software as a medical device, and they have a clear pathway for adaptive AI models. And then they have drafted guidelines on real-world evidence and algorithm transparency as well. While EMA is kind of between the developing and advanced stages, they have published a reflection paper on the AI in medicine life cycle in 2023.

And they emphasize explainability, data quality, ethical AI, and AI oversight integrated

across clinical and pharmacovigilance stages. And then the PMDA Japan is well-structured and progressive. So, they have this science board, AI-specific subcommittees, and then horizon scanning and post-market adaptability reviews as well. For the Indian regulatory agency, it is still evolving and at the national stage. So, the regulatory focus on AI plasticity and transparency, along with increasing digitalization, lacks AI-specific regulations yet, but those are coming very soon.

Okay, talking about the software as a medical service. So, SaMD, as defined by the International Medical Device Regulators Forum, is software intended for a medical purpose that functions independently of a physical medical device. So, it can operate on general-purpose platforms like smartphones or the cloud and is used for diagnosis, monitoring, decision support, and therapy recommendations. So, the AIML-based software as a medical service can analyze imaging patient data or clinical inputs to aid diagnosis, recommend treatments, and assess trial eligibility. So, the regulatory concerns include the opacity of black-box models and the need for continuous model updates to prevent performance degradation.

So, these are some examples of FDA-approved AI/ML-based tools as software as medical devices. So, IDXDR is one of the tools that is coming from the digital diagnostics company. So, its use case is that it is kind of an autonomous AI for diabetic retinopathy detection, and it is the first FDA-cleared autonomous diagnostic AI, which was cleared back in 2018. And then you have Viz.ai, which is an AI-based stroke detection tool for CT scans.

So, it is FDA cleared for time-sensitive trials. And then you have the Caption Guidance from Caption Health. So, it is AI-guided cardiac ultrasound acquisition. It is being approved as a real-time imaging support tool. And then you have the GI Genius, which is coming from Medtronic.

So, it's an AI-assisted polyp detection system during colonoscopy. It is the first FDA-cleared GI support AI. And then you have the DermaSensor coming from DermaSensor Inc., and it is a skin cancer detection aid still pending wider release, and it has received breakthrough device designation. So, you can see these are some of note; this list is not exhaustive, but these are some examples that have been approved as software as a medical device (SaMD).

So, talking about the AI-specific regulatory challenges, the problem is the lack of transparency; the black box nature of these models, many AI models, especially those deep learning systems, operates opaquely. So, we see the input and the output, but we do not know how the model got there. So, we do not know how it is coming to that output. So, that is a little bit, you know, difficult, and the implications of that reduce clinician trust,

hinder validation of AI-driven decisions, and lead to conflicts with regulations like GDPR, which upholds the right to explanation for automated decisions. So, the regulatory considerations are that you need to have explainable AI to enhance interpretability, and you also mandate transparent documentation of the AI decision-making process.

Another thing is model drift and lifecycle management. So, model drift is when an AI model may lose accuracy over time as input data shifts or clinical patterns evolve, leading to degraded performance. So, the implications of model drift are reduced diagnostic accuracy, flawed treatment recommendations, and increased patient safety risk if models are not retrained or recalibrated. So the regulatory consideration states that we need to implement continuous monitoring and periodic revalidation of those AI models, as well as define standardized protocols for model updates and regulatory re-approvals. Then there is a problem with data bias and representation in AI.

So, the bias training data the AI system trained on non-representative datasets may reinforce existing disparities in the healthcare leading to marginalized populations may receive inaccurate or substandard care as well as public trust in AI driven healthcare solution may erode as well. So, the regulatory consideration states that we need to enforce the use of diverse, inclusive data sets during AI development, as well as require regular bias audits to ensure fairness and equity in the outcomes. Coming to another consideration, which is validation and reproducibility in AI, the inconsistent performance across settings is one of the challenges. So, the AI system must prove they can deliver reliable reproducible results across diverse populations and clinical environments. So, the implications of it are that variability in performance may lead to inconsistent or unsafe patient outcomes, and the lack of reproducibility slows regulatory approvals and adoption.

So, the regulatory consideration says that require standardized validation protocol across multiple healthcare settings as well as promote the transparency in publishing validation results and methodology. Another thing is auditability and traceability in AI; it has, you know, limited transparency in the model life cycle. So, tracking the development deployment and update history of AI model is critical for safety and accountability. So, the implication of this can be it is difficult to investigate AI driven errors without proper traceability. So, if you are not able to trace what changes are made in that model, we cannot trace the error, actually.

And then outdated models may be used without clear version control if we do not have proper version control. So, those outdated models can be used, which might not be correctly predicting the outcomes. So, the regulatory consideration states that we need to have a comprehensive mandate for documentation of development workflows and version histories, as well as establish audit trails to track model performance updates and usage

over time. And then coming to human oversight and responsibility in AI-driven care. So, now if the AI is you know taking care of the diagnosis and recommendations.

So, who will take responsibility? So, there is a kind of ambiguous accountability; clarifying who is responsible when AI is involved in clinical decision-making remains a significant concern. And the implications of it are that unclear accountability may lead to legal and ethical challenges, and over-reliance on AI could weaken clinicians' judgment and autonomy. And then the regulatory consideration states that we need to define explicit roles and responsibilities for developers, clinicians, and the institutions. As well, we shall enforce human-in-the-loop frameworks to ensure clinicians retain final decision authority, and this we have seen in the earlier sessions as well. Okay, another consideration is the security and privacy in healthcare AI.

So, the sensitive data requirement AI models depend on a large amount of patient data, raising serious security and privacy concerns. And the implications of this are that the risk of data breaches compromising confidential health information, as well as non-compliance with regulations like HIPAA and GDPR, may also lead to legal penalties. So, the regulatory consideration states that we need to enforce robust data production practices like encryption or access controls, as well as ensure compliance with privacy regulations and require regular security audits. So, coming to some of the regulatory initiatives and guidance documents. So, the FDA's AI/ML-based SaMD action plan is one of them, actually.

So, the US FDA has introduced a strategic framework to regulate AI/ML-based software as a medical device, ensuring safety without stifling innovation. So, it ensures patient safety and regulatory clarity while enabling the responsible evolution of AI- and ML-driven healthcare tools. So some of the key principles are transparency, which promotes clear disclosure of the algorithm's purpose, training data, and limitations to build trust and support regulatory oversight. Another principle is real-world performance monitoring, which emphasizes post-market surveillance to assess ongoing safety and effectiveness in real-world use. and support for the continuous learning systems that encourage adaptive AI with safe update mechanisms through a predetermined change control plan where updates and conditions for change are pre-specified.

Okay, looking at the EMA's reflection paper on AI in the medicinal product life cycle. So, it is issued by EMA. This paper addresses the challenges and best practices for using AI throughout the entire drug life cycle from development to the post marketing. So, it promotes the safe and ethical integration of AI in medicinal development while safeguarding patient safety and data quality.

So, the key principle is GXP compliance. Where the AI systems must comply with good practice standards like GMP, GCP, and GLP, ensuring data integrity. And traceability, explainability, and transparency are another, you know, key principles which avoid black box models in critical decisions, and outputs must be interpretable to support trust and accountability. And then we have, you know, another key principle: the algorithm validation, where models must undergo rigorous training, validation, and testing with representative datasets, and the entire AI lifecycle should be documented. Okay, coming to the ICH guidelines and discussions. So, the International Council for Harmonization sets global standards for pharmaceutical development and regulatory approval.

So, it fosters structured data collection and analysis, ensuring consistency across regions for AI application in clinical and quality domains. So, the relevant guidelines are ICH E6 R3, which is the Good Clinical Practice. So, it promotes modernized clinical trials integrating digital technologies and decentralized models, potentially supported by AI, and then you have the ICH Q12, which is a product lifecycle management guideline. It enables regulatory flexibility for post-approval changes in AI-driven systems using manufacturing or quality control. Documents are like public consultations and white papers.

So, WHO guidance on the ethics and governance of AI for health is a 2021 document. So, which highlights human rights, inclusivity and accountability and provides a framework for ethical AI use addressing risk like bias and inequality. And then you have the MHRA AI Assurance Framework coming from the UK. So, it's a detailed guiding principle for AI product development, validation and regulation, focusing on transparency, safety and lifecycle assurance of AI models. So, these documents drive global dialogues and foster consensus on AI ethics and regulatory governance.

And then they are, you know, the open collaborations with the regulators. So, that is one of the, you know, important aspects of implementing AI and using it to its fullest extent, actually to help with drug discovery and development. So, it bridges the gap between AI innovation and regulatory oversight through transparency, open data, and early engagement with developers. So, some examples of these collaborative initiatives include the FDA, NIH, and academia. So, collaborative efforts to develop regulatory grade datasets for AI model validation, fostering more robust and fair algorithms.

OpenFDA is a public access platform offering datasets and adverse event reports to enhance transparency and innovation. And then the CDER AI pilot program, which provides early-stage feedback for sponsors submitting AI-based tools during development. Talking about some of the best practices for compliance is important. So, regarding the documentation and model transparency, So, clear documentation is essential for making AI models understandable, auditable, and trustworthy, particularly for regulatory

submission. And then the transparency reduces black boxes, addresses concerns, and supports regulatory reviews.

So some of the key tools are model cards, which summarize model purpose, training data, intended use, limitations, and performances across subgroups. helping regulators and clinicians understand when to trust the AI model data sheets for data sets provide metadata on data set provenance composition, labeling and potential biases. And then you have the transparent reporting where you clearly state assumptions, known risks, and validation boundaries to ensure reproducibility, regulatory auditability, and user trust. Then we have these good machine learning practices as well. So, what are the 10 foundational principles for the development and deployment of AI in healthcare as per the FDA? So, these 10 principles you know are, for example, the first one is that multidisciplinary expertise is leveraged through the total product life cycle.

Another one is that good software engineering and security practices are implemented. Clinical study participants and datasets are representative of the intended patient population. Training datasets are independent of the test datasets. Selected reference datasets are based on the best available methods. Model design is tailored to the available data and reflects the intended use of the device.

Focus is placed on the performance of the human-AI team. Testing demonstrates device performance during clinically relevant conditions, users are provided clear essential information, and deployed models are monitored for performance while retraining risks are managed. Some of the algorithm validation approaches are as follows: robust validation is essential to ensure the generalizability, safety, and efficacy of algorithms, particularly in clinical or trial-related applications. So, they are, for example, internal validation, external validation, real-world evidence, and pre-specified endpoints. So, the purpose of internal validation is to hold out the test sets from the same data source to test for overfitting. And then external validation is used to test on independent data sets or to assess the generalizability, and the real-world evidence is that its purpose is evaluation using post-deployment data like EHRs or registries.

And the pre-specified endpoints, the purpose is to avoid data dredging and ensure alignment with the clinical goals. So, regulatory sandboxes are another tool, you know, a controlled environment where innovative technology such as AI-based systems can be tested in real-world conditions under the supervision of regulators without the full regulatory burden typically required for commercial deployment. So, it allows developers to experiment with new technologies while ensuring they meet safety, efficacy and ethical standards before full scale implementation. So, the AI use case could be, for example, a target identification algorithm, and the sandbox purpose would be to assess the

reproducibility and traceability of the model logic. And for the compound generation, like those you know, generative models, the purpose of the sandbox could be to explore toxicity safeguards and filtering criteria.

For the virtual screening platforms, the purpose of the Sandbox could be to validate the predictive performance on novel scaffolds. And for the AI-based SAR models, the purpose of the Sandbox could be testing documentation standards and explainability tools, and then integrating with lab automation like robotics. So Sandbox could evaluate the GXP applicability and the audit trail. So, these are some of the, you know, real-world examples of AI tools under regulatory review. So, you have the unlearn AI, which is kind of a digital twin, and then the function is to create patient-specific models to simulate outcomes in the trial, and it is using deep learning and time series modeling.

And then the FDA is collaborating, you know, so they have pre-submission meetings as well with the FDA. And then the status is active in the clinical trial design for neurodegenerative diseases. Then you have PATH-AI, which is used for histopathology slide analysis to assist in drug development and diagnosis. So, it is based on CNN. At PATH-AI, we are working with FDA and pharma partners to validate the endpoints.

It is used in biomarker validation under multiple ongoing reviews. And then you have OWKIN, which is a predictive modeling tool for patient stratification and survival in oncology trials. So, it is using federated learning and predictive ML, and it is in the pilot phase for regulatory-grade deployment. And then you have the MediData synthetic control arm, which uses RWD to simulate control arms in oncology trials. It uses matching algorithms and Bayesian inference, and it is currently used in multiple sponsored trials with FD engagement. And then you have Aiforia, which is an AI-assisted tissue analysis for drug discovery and companion diagnosis.

So, it is using deep learning on the histological data. Currently, it is being used for, you know, pathology endpoints in the trial. And then you have IBM Watson for clinical trial matching, which matches patients to eligible trials using the EHR data. So, it is using NLP rule-based methods, ML-based methods, or hybrid methods. And then it is currently deployed in large hospital systems that have discontinued it commercially, but you know the framework is still cited. And then you have Tempus, which is the AI for genomic profiling and clinical trial matching.

So, again it is using ML on genomics and EHR data; it is in commercial use again at present. And then you have the in-situ, which is an AI-driven target discovery using cellular phenotyping and genomics. Using deep learning and representational learning in preclinical to clinical transition monitoring for regulatory inputs. Okay, coming to the

summary, AI is transforming drug discovery but presents regulatory, ethical, and operational challenges. The regulatory bodies like the FDA, EMA, PMDA, and CDSCO are developing risk-based frameworks for the safe and effective use of AI.

The key concerns include data bias, model drift, lack of explainability and auditability, and the oversight measures include lifecycle management and human-in-the-loop systems. And software as a medical device adds regulatory focus on transparency, real-world validation, and continuous monitoring. And the best practices, like GMLP model documentation and regulatory sandboxes, promote responsible innovation and patient safety. So, in the end, I have a small activity for you. So, just try to make a list of five AI-based tools that are being used for drug discovery and development and that are currently under consideration by worldwide regulatory agencies.

So, these agencies could be any of the following: FDA, EMA, PMDA, or CDSCO. And in the end, I have some suggestions for further reading where you can go through these articles to learn more about this topic. And with that, thank you.