AI in Drug Discovery and Development

Prof. Rajnish Kumar
Dept. of Pharmaceutical Engineering and Technology
IIT-(BHU), Varanasi

Week-08
Lecture-39

Welcome to the course AI in Drug Discovery and Development. In today's session, we will talk about predicting the outcomes of clinical trials with the help of AI. So, by the end of this lecture, you will be able to understand the rationale behind predicting clinical trial outcomes, recognize the key factors that make outcome prediction complex, and explore how AI can enhance clinical trial prediction. Differentiate between validation strategies used to evaluate predictive AI models and examine real-world case studies demonstrating the impact of AI on trial outcome predictions. And we have seen that AI can be used in all the stages of not only drug discovery and development but also in clinical trials. So let us see how we can use AI to predict the trial outcome.

So if we talk about the possible clinical trial outcomes. So, these clinical trial outcomes refer to the results or endpoints used to determine the effectiveness, safety, and feasibility of a medical intervention like a drug, device, or therapy. So, as you know, the purpose of a clinical trial is to see whether the therapy, whether the drug or a medical device, is being developed for the treatment or prevention of the disease. So, is it effective for the patients? Is it safe for the patients to use or not? And that is what you know; the endpoints of a clinical trial actually mean the outcome of the clinical trial.

So these outcomes can be broadly categorized as primary outcomes, which usually tell about the efficacy, safety, and clinical endpoints; they can also include secondary outcomes, which measure the quality of life. For the functional improvement, as well as the long-term follow-up. And then this can be exploratory; there can be exploratory outcomes as well, like biomarker discovery, genomics, or proteomic data correlation. There can be operational outcomes, such as the success or failure of a clinical trial phase transition, as well as regulatory outcomes, like approval or rejection by the regulatory bodies. So, um, why do we need to predict, or why shall we use, you know, prediction in the clinical trial outcome? So, what it can do is, the first thing it can do is it can, you know, mitigate the risk.

And ensure the patient's safety by forecasting the potential outcomes, researchers can identify early warning signals about a drug's efficacy or adverse effects; thus, this practical approach helps minimize patient exposure to ineffective or harmful treatment. So, if we can see, even before running the clinical trial or during the clinical trial, if we can predict

what will happen, what the endpoint or the result of this trial will be. So, that can help us to reduce you know the risk and ensure the patient safety that is one thing. Another thing is it can help us, you know, utilize the resources. Efficiently reducing the cost.

So it allows researchers to fine-tune trial design by adjusting parameters like sample size, dosing, and endpoints while also efficiently allocating resources to avoid over-investing in trials with low success rates and focusing on promising treatments. Then it can be used in informed ethical decisions, like accurate prediction, to support ethical decision-making by ensuring that patient recruitment minimizes unnecessary risk. and by guiding decisions on whether to continue, modify or halt a trial based on assessment of the benefit versus risk. So, you can use it for, you know, evaluating the risk versus benefit aspect of a clinical trial. And then it can help us with strategic planning and regulatory planning as well.

So, the regulatory bodies and sponsors they rely on robust predictions to guide approval processes such as fast tracking or conditional approvals while also strategizing future research by pinpointing areas for further study or adjustments. And it can enhance the scientific rigor of the study, as well. So those predictive models leverage historical data and advanced analytics, including ML. They enhance hypothesis generation by identifying patients, subgroups that may benefit the most and enable real time adjustments. So overall, AI can reduce the risk.

It can, you know, inform the portfolio decisions and guide the trial design. So now coming to the prediction, what makes it difficult is that, you know, it's not all those AI models; they are not foolproof, actually, and they are not always 100% correct. So, what makes those challenges difficult to predict the outcome of a clinical trial? So, the first thing is the biological complexity because the human body or the animal body is highly complex, and we are still unable to understand all these nitty-gritty details of the pathways and mechanisms of how drugs work. We still have many drugs for which we do not know the exact mechanism of how they actually work. So, a drug may look promising in vitro or in animal models, but it can fail in human trials due to unanticipated mechanisms because it might be working fine in animals.

But as soon as we go to the clinical trial, it may be ineffective actually because of its unanticipated mechanism. And that is due to the biological complexity of the system. There might be patient heterogeneity as well, where a one-size-fits-all model may overlook subpopulation-specific responses because everyone has, you know, a different genomic makeup. So based on the genomic makeup, one drug may work in one person while it has no effect on another person. And then there is variability in trial design.

So, two trials for the same drug may yield different outcomes due to variations in the trial

structure as well. And then there is the multidimensional and multimodal nature of the data. So, integrating and interpreting heterogeneous, noisy, and often incomplete data sets in a biologically meaningful way demands advanced integration strategies. And that's one of the difficulties with the prediction of a clinical trial. Okay, so how do you know AI can play a role in clinical trial outcome prediction? So, we can learn from the historical data, and that is what predictive analytics actually is.

We can predict, by using AI, whether a phase 2 oncology trial is likely to meet its primary endpoint based on previous similar trials. The data from previous similar trials can be used as training data. And then we can use that to make a decision on the outcome of this trial. And then we can use, you know, the multi-omic data integration, where it can integrate diverse data types like those, you know, omics data such as proteomics, genomics, metabolomics, and epigenomics. And molecular structures, trial protocols, and clinical variables into a unified model, a model may find that trials targeting specific pathways or using adaptive designs have higher success probabilities.

And then we can analyze the probability of success prediction. So, the AI model generates a probability of success score for each trial, guiding portfolio decisions and de-risking development, which helps companies avoid low-likelihood programs early and enables prioritization and optimization of the trial resources. So let us see the process of predicting the clinical trial outcomes using AI. So, the very first step is data collection and integration, where we gather, collect, and curate the data. So, you collect the historical clinical trial data, including patient demographics, treatment regimens, lab results, imaging data, and genetic profiles.

Adverse event records. So, this data may come from past clinical trials; it can be, you know, electronic health records (EHRs); it can be from clinical trial registries; and it can be from the published research as well. And the second thing you know to do is to integrate the diverse sources, consolidating data from various sources and ensuring compatibility. and consistency, which may involve linking disparate databases and ensuring data standardization, because now when we are collecting this data from multiple sources, it might be that the data is not clean. Or it might not be compatible with each other, so we have to make sure that this data is being standardized and treated uniformly.

so that we can integrate them with each other and use them. So, the second step is data preprocessing, where the first thing we do is cleaning. So, we remove the duplicates, fill in or remove the missing values, and resolve the inconsistencies. And the next thing is to do the normalization, standardize the data scales, for example, the lab values or scores, to facilitate model training. And then the next thing to do is the privacy measures because most of the time this data is coming from the patients, and their personal data is also

associated with them.

So, we need to, you know, anonymize or de-identify the data to comply with regulations like HIPAA or GDPR. Once we have prepared the data, the next step is feature engineering, where we try to identify the key variables by determining which variables are likely to influence outcomes. For example, these are now features such as age, biomarkers, disease, and severity. So, these are the features that will lead to, you know, influencing the outcomes, like whether the trial will be successful or not. So that outcome will be relating to these features, actually.

So we need to identify those key features, and then we create, you know, derived features as well. Combine or transform raw data to generate new predictors, such as risk scores or composite scores. So those can be, you know, derived. And then, if we have a load of features where we are unable to decide which features to use, so then we can use techniques for dimensionality reduction like principal component analysis to reduce complexity and eliminate redundant features. So, in this case, for example, we now have 100 features, and we want to identify the top 5 features that correlate well with the outcome.

So, we can use these dimensionality reduction techniques, like principal component analysis. And once the feature engineering is done, the next step is model selection and development. So, we chose an algorithm. So, we evaluate a range of ML algorithms, such as logistic regression, decision trees, random forests, gradient boosting, or deep learning networks based on the data type and prediction goal. And we can try the ensemble methods as well, where instead of using a single, you know, ML method, we try to combine multiple methods and develop the models.

We consider combining multiple models to enhance the predictive accuracy and robustness of the predictive model. and followed by the training and validation where we split the dataset, dividing the data into training, validation, and test sets, or we use cross-validation to ensure the model generalizes well. And we do the hyperparameter tuning where we adjust the model parameters to optimize performance while avoiding overfitting. And we use the performance metrics like we discussed at the beginning of this course, as well as relevant ones. Performance evaluation metrics like accuracy, AUC-ROC curve for classification, and RMSE and MAE for the regression models are used to assess model performance.

Because on the basis of these evaluation matrices, we decide whether the model we have developed is good or not. Okay, once we have trained the model and validated it, the next step is interpretation, model interpretation, and explainability because now we are talking about AI a lot. And we have talked about AI, which is kind of a black box, especially deep

learning. Can we use some interpretability tools? So, for example, we use the techniques like SHAP or LIME to understand feature contribution and ensure that the model's decision-making process is transparent and it's not working like a black box actually. And then we also need to define the clinical relevance of those features.

So we work with clinical experts to validate that the predictive factors align with biological plausibility and clinical experience. And this we have seen in the QSAR session as well, where we need to make sure that the features we are identifying have a causality relationship, not just a correlation relationship. Okay, once we have got a nice model and have explained the applicability domain and all those things, the next step is to integrate it into the clinical trial design. So, we can use, you know, the adaptive trial design where we use the model predictions to optimize trial parameters such as patient selection criteria, dosing strategies, and endpoint definitions. We can use it for resource allocation, adjusting resource allocation by focusing on patient subgroups that are more likely to benefit, thereby reducing unnecessary risks and costs.

And we can use it for decision-making, where we can set the criteria for early stopping or modification of trials based on real-time predictive insights. Once we have integrated into a trial design and have used all those, you know, well-validated predictive models. The next step is the deployment and continuous learning, where we can use real-time monitoring. We can integrate that AI model into the trial process to monitor outcomes and adjust predictions as new data become available. And then we can, you know, update the model as well.

We can regularly retrain the model with new trial data to ensure ongoing accuracy and relevance because each time we are using it for prediction. We are generating new data, and with the clinical trial as well, we are generating new data so that it can be further used to update the model, improving the accuracy and relevance of those models. And we can use a feedback loop in which we establish a system for continuous feedback from clinical operations to refine both the model and the trial design. So, this can be an exemplary pipeline. So, where you use multimodal inputs, such as those from the drug molecule, you can use the structure of the drug molecule, as well as the disease pathology and biomarkers.

And all that stuff from the patient, and then you can use the information from the existing literature from the clinical trial, dot go dot in, so like the trial protocol. For the registries, you embed those as features, then make the models, and use the trial outcome to predict using that model. And then you can use this, you know, to integrate this into the ongoing clinical trial, generate the data, and then update the model and use that model for further analysis. Predicting the trial outcomes. So, this is how we can use it for trial outcome prediction.

Okay, this is, you know, kind of a similar workflow where you can see that we can use the data from the AACT database, which is the aggregate analysis of clinicaltrials.gov. And then there is another database, you know, which is known as CHIA. It's a novel, large, annotated corpus of patient eligibility criteria. So, you get the patient eligibility criteria from the CHIA data set, and you can get the clinical trial data from the, you know, AACT database.

So, doing the feature engineering and, of course, using NLP because NLP is very good at handling all those text-based inputs. By using that, we obtain an enhanced data set, which is then used for data analytics and for drafting the protocol for a new interventional study. Okay, and then further, this data is used for training a model, and this model is being evaluated, and then you get a final model. And with the help of this final model you evaluate, you know, or you polish your protocol draft from the new interventional study, and then you study the termination probability. This is kind of, you know, the outcome, actually.

So you predict the outcome of whether this model is giving a score less than the threshold. If it is less than the threshold, then if it is not, you finalize the protocol. And if it is yes, then you interpret the prediction with this app and try to identify the features that are giving this negative outcome. And then, based on those features, you try to edit the protocol, and then you finalize the protocol. So, there are multiple model validation strategies, such as using retrospective validation, which employs historical clinical trial data to assess model performance.

and then the model is trained and tested on past trials with known outcomes. It is really good for initial proof of concept, study, and model tuning. And then you can use the quasi-prospective validation, where you simulate a real-world scenario by training the model on older data and testing it on more recent trials not seen during training. And so, it mimics a kind of time-forward validation thing. So, it evaluates the generalizability to the future and, you know, predicts the unseen conditions.

So, another way is to use prospective validation, which tests the model's predictions in real time on ongoing or newly launched clinical trials. However, it requires time to observe the actual outcomes because we will not be able to know the actual outcomes until the ongoing study is completed. So here you can see the use cases like those in retrospective validation. So, it's used in model development or benchmarking. And the example can be like predicting outcomes of a phase two trial completed between 2010 and 2020.

The quasi-prospective validation is being used for robustness testing. And then you can

see that it can be used for a model trained on trials before 2017 and tested on those from 2018 to 2020. So, you are using the historical data, and then you are using it for the already available clinical trial data. You predict the outcomes. And then there is prospective validation where you use real-world validation and deployment, where the model makes predictions for the active phase of the trial, and accuracy is assessed once outcomes are published, because until we get the results from the study, we do not know the accuracy, and we cannot compare experimental outcomes versus the Predicted outcome.

Okay. So, what are those, you know, AI/ML models that can be used for this task? So, you can use traditional statistical models like logistic regression and Cox proportional hazard models, which offer interpretability for binary outcomes and survival analysis, respectively. Then you can use ensemble methods like random forest and gradient boosting machines, XGBoost or LightGBM, which provide robustness and insight into feature importance, often excelling with structured clinical data. And then you can use deep learning models like neural nets, including deep feedforward networks, CNNs for imaging, and RNNs and LSTMs for sequential data. You can capture the complex, nonlinear relationships. And then you can use transformer-based models like BERT.

So, it is tailored for biomedical texts, which are useful for processing unstructured clinical notes and literature. And then you can use the specialized survival models where you use the random survival forests and deep learning, which are tailored for time-to-event data. These are crucial for understanding the long-term outcomes. So let us have a look at the real-world examples where people have been using these AI and ML-based tools for predicting clinical trial outcomes.

So inCliniko is one of the tools. So, which can predict the transition from phase two to phase three. So, it's a transformer-based AI platform integrating omics, as I said, like multi-omics data, metabolomics, proteomics, epigenomics, trial design, and drug properties and features. and it could predict the outcomes for melanoma and alzheimer's trial enabling early resource allocation because if we can you know predict Whether the trial is going to be successful, so it can help us in the early allocation of resources for that trial, and what were the impacts of using this inclinico tool on the accuracy. Accuracy-wise, it was 79% accurate in prospective prediction for phase 2 trial outcomes.

The AUC, you know, was around 0.88 in quasi-prospective validation for the phase 2 to phase 3 transition, and the ROI generated, you know, at least a 35% return in a 9-month virtual investment portfolio. So, this is kind of a return on investment, actually. So, you can see that it can be really useful for cutting the cost of these studies. So, this is InClinico's AI framework. So, the platform integrates two core predictive models focused on trial design and target selection.

So each model independently estimates the probability of success for a submitted trial and rank the most influential feature driving the predictions. You can see here that you have all the data, including trial design data such as the number of patients, clinical sites, number of arms, comparator choice, and blinding. And the target choice data, like clinical data on drugs, disease oncology, drug-target interaction, omics data, protein-protein interaction, and signaling pathways. And then you use an ensemble model, and you can get a competency report that can tell you about the probability of success. and also identify the features that impact the probability of success of this trial.

Then there is another tool called Lifted. So, it's a multimodal outcome prediction. So, it's a mixture of experts model that unifies electronic health records, omics, and trial protocols into natural language for predictions. So, it improved predictions for rare disease trials by integrating social determinants of health. So, the impact was, you know, it outperformed the baseline in predicting outcomes across phases 1, 2, and 3 trials. and dynamic weighing, where it automatically prioritized the critical modalities, such as biomarkers for the oncology trials.

And then we have the MIT Sloan Project Alpha, which is a risk-stratified trial forecasting where they have used an ML model analyzing over 140 features. And those features include, for example, sponsor track record, trial duration, accrual rate, etc. And it guided investment decisions in cardiovascular drug trials, reducing capital waste. and the impact was, you know, so failure prediction, so identify high risk trials with greater than 80% accuracy means it was very good at identifying which trials might go wrong or might fail. So, that's the advantage of using this tool, and that is how it could save a lot of course.

So it enables sponsors to reallocate resources from trials likely to fail. And then you have the Cleveland Clinic and Hint model, which could, you know, again indicate or predict the early termination. So, SHAP analysis of 420,268 trials flags risky eligibility criteria like stringent BMI thresholds, etc. So, you know if we could, the beauty of this tool is that it is analyzing the eligibility criteria of a large number of trials, actually. And based on that, it could predict which of those eligibility criteria can lead to the failure of a clinical trial, so it could identify, you know, features that actually identify high dropout risk in a pediatric asthma study.

And the impact was that the area under the curve was 0.80 in predicting early terminations; you can see that it is quite good at it. Predicting early termination and how it can be helpful in risk mitigation. So, reduce protocol amendments by 30% in a diabetes trial as well. And then there is another, you know, example: Lindus Health's Citrus Platform.

So it is being used for real-time outcome monitoring. So, it automate it is an automated anomaly detection in decentralized trial using variables and NLP. So, it flagged the adverse events in a Phase 3 hypertension trial, enabling early intervention. So, the impact was that it achieved almost 95% in maternal health trial data completeness and retention improvement. So, it reduced dropouts by 25% through predictive engagement as well. Okay, so those are the examples where we have now seen how we can use the AI tools to predict the clinical trial outcome.

So let us talk about some of the limitations of AI-enabled prediction. So, the first thing is the data quality and availability, and that is, you know, a general limitation of using AI in drug discovery and development because the biological data is really complicated. It is heterogeneous, and the availability is very low as well. So, clinical trial datasets often contain missing values, noise, and may be limited in size or scope, which can compromise model training and accuracy. And then there is a lot of bias and generalizability as well; like the historical data may reflect bias in patient selection or treatment protocols, leading to a model that might not generalize well across diverse populations or novel scenarios.

And then interpretability and transparency are another challenge. Like many advanced models, especially deep learning approaches, they can function as a black box, making it difficult for clinicians and regulators to understand and trust the decision-making process. Integrating AI models into existing clinical trial processes and the integration into clinical workflow can be challenging, requiring alignment with operational practices, clinical training, and continuous model updates to accommodate the evolving data. And regulatory and ethical challenges are also present, where navigating regulatory approval requires comprehensive documentation, validation, and ethical oversight, and the lack of standardized guidelines for AI in clinical research can slow adoption and trust. Okay, so what is in the future? Like, you know the integration with real-world evidence, combining clinical trial data with real-world data sources such as electronic health records and insurance claims. and patient registries, which further enhance model generalizability and relevance to broader patient populations, bridging the gap between controlled trial settings and everyday clinical practice.

And then you can use federated learning for data privacy. So, which enables model training across multiple institutions without transferring sensitive data. So, it preserves patient privacy, complies with the data sharing regulations, and allows learning from distributed data sets, such as data that can come from hospitals or pharma partners, while maintaining confidentiality. And then we need to use explainable AI (XAI), which improves the interpretability of AI models by highlighting the features driving predictions. So, it builds trust with clinicians, regulators, and pharma stakeholders, and it facilitates transparent decision-making in trial design and portfolio management.

Okay, let's move on to the summary. So, you know, predicting clinical trial outcomes is a growing priority in drug development due to the high attrition rate, especially in phase two and phase three studies. And we need to understand that if a drug is failing in phase three studies, then the company that is running that clinical trial is losing billions of US dollars. So, they are losing a huge amount of money that they have invested in that study or in that project, actually. Because it did not start with the clinical trial, it started with target identification, validation, hit identification, and optimization. And finally, the clinical trial shows that a decade of work and a huge amount of monetary investment go to waste if a drug fails in phase three of the clinical trial.

So that is why it is really important to make some models we can use for machine learning. and AI to predict the trial success or the trial outcome so that will be one of the best things we can do So the AI enables data-driven forecasting by learning patterns from vast multimodal datasets, ranging from omics and molecular data to trial design and real-world evidence. However, key challenges remain in prediction, including biological complexity, patient variability, and diverse trial designs. And there are future advancements, such as XAI federated learning and integration with real-world data, that can enhance the predictive power and transparency of AI in clinical research.

So, in the end, I have an open question for you. If an AI model predicts that a well-designed phase 2 clinical trial is likely to fail, Should the trial be stopped before it begins? Just think it over, and then I have some, you know, some publications. Here are some references that you can go through to enhance your understanding of this topic, and with that, thank you.