DOI: https://doi.org/10.20535/kpisn.2025.4.344350
UDC 519.6; 615.015.8

D.S. Horodetskyi[1*], M.P. Smetiukh[1,2], S.O. Soloviov[1,2]
[1]National Technical University of Ukraine
"Igor Sikorsky Kyiv Polytechnic Institute", Kyiv, Ukraine,
[2]Shupyk National Healthcare University of Ukraine, Kyiv, Ukraine
[*]Corresponding author: dgolodetsky37@gmail.com

DECISION MAKING IN ANTI-CORONAVIRUS DRUG DISCOVERY: MATHEMATICAL MODELLING AND VALUE OF INFORMATION ANALYSIS

Background. The process of preclinical evaluation of antiviral medications typically involves multiple stages, each containing substantial uncertainties. Traditional methods for screening the compounds often lack structured means for optimising the decision-making and calculating the feasibility and risks of transitions between all of the stages. Thus, there appears to be a problem with the inefficient selection of promising antiviral molecules, which subsequently increases the probability of choosing suboptimal research trajectories.
Objective. The paper aims to develop a computational framework for optimising of the transition between stages in preclinical antiviral testing. The system focuses on the integration of decision trees and Markov models in order to include effectiveness, risks and the value of additional information into assessment, supporting an in-depth planning of preclinical research pipelines.
Methods. Experimental data from molecular docking, cytotoxicity CD_{50}, and antiviral activity IC_{50} were used in a multi-stage evaluation system with CTI ≥ 4 being the criterion for progression into further stages. Decision trees provided the explicit rules for advancement of the compounds, while Markov models added context for building sequential strategies under uncertainty and quantified the feasibility of movement to the next stage. Value of information analysis added the assessment of the expected benefit of additional data.
Results. The developed framework consistently produced reliable technical results. The decision used in CTI ≥ 4.0 prediction stage demonstrated a conservative classification pattern, correctly identifying compounds with high therapeutic potential while missing some effective candidates. The Markov model showed steadily increasing state values in docking, cytotoxicity, and antiviral testing phases that confirmed the growth of expected utility. Based on the findings acquired, the most effective solutions were identified for the ongoing investigation into antiviral assays, while the application of value of information analysis indicated that the largest gain occurred after antiviral activity testing, whereas the initial phases serve as filters.
Conclusions. The study showed that both decision trees and Markov models capture different but complementary aspects of the preclinical evaluation process. Decision trees provide an interpretable set of rules that formalise how molecular docking and cytotoxicity measurement influence the progression of compounds, while their limited sensitivity at the CTI threshold highlighted the complexity of predicting the final success of the evaluated compounds. The Markov model simulations showed that the full three-stage pipeline is justified and that progression decisions are influenced by both uncertainty and experimental cost. The value of information analysis clarifies the importance of each stage, helping to emphasise the role of antiviral activity data. These findings support the integration of analytic methods for improving the structure, transparency and efficiency of antiviral preclinical research.
Keywords: coronavirus; drug; preclinical evaluation; decision tree; Markov decision process; value of information.

Introduction

The optimisation of sequential decision-making in preclinical studies of antiviral compounds remains a highly relevant challenge due to the combination of uncertainty, high experimental costs, and limited predictability of candidate efficacy. At each stage of the preclinical pipeline – from in silico screening to cytotoxicity assessment and antiviral activity tests – researchers must make a series of decisions,

Пропозиція для цитування цієї статті: Д.С. Городецький, М.П. Сметюх, С.О. Соловйов, «Прийняття рішень у процесі доклінічної розробки противірусних препаратів проти коронавірусу: математичне моделювання та аналіз цінності інформації», *Наукові вісті КПІ*, № 4, с. 20–30, 2025. doi: https://doi.org/10.20535/kpisn.2025.4.344350

Offer a citation for this article: D.S. Horodetskyi, M.P. Smetiukh, S.O. Soloviov, "Decision making in anti-coronavirus drug discovery: mathematical modelling and value of information analysis", *KPI Science News*, No. 4, pp. 20–30, 2025. doi: https://doi.org/10.20535/kpisn.2025.4.344350

where an inaccurate early-stage choice leads to the loss of time, resources, and potentially promising compounds. This creates the need for systematic approaches capable of increasing the rationality and economic efficiency of the preclinical process.

Despite significant progress in artificial intelligence, current research mainly improves individual steps of drug discovery rather than the full decision-making pipeline. Modern machine learning techniques demonstrate substantial advances in virtual screening, toxicity prediction, and target selection [1]. AI-based integration with organ-on-a-chip platforms and digital twins enhances the accuracy of pharmacokinetic and toxicological modelling [2]. Data-driven design of antiviral peptides using GANs, deep learning and explainable AI demonstrates strong potential for optimising candidate properties [3]. Studies of DHODH inhibitors highlight the complexity of translating promising in vitro results into clinical effects and emphasise the need for step-wise risk assessment [4]. Multi-omics deep learning pipelines accelerate early discovery and facilitate drug repositioning [5]. AI-based prediction of viral mutations supports personalised antiviral strategies and shows the sequential, dynamic nature of decision-making in virology [6]. AI-driven dereplication and classification of natural products further illustrate the need for structured transitions between preclinical stages [7].

However, these advances primarily address predictive accuracy rather than the principled optimisation of decisions across multiple stages. Current research lacks integrated mathematical frameworks that would: formalise transitions between preclinical stages, quantify risks and probabilities of success, incorporate the cost and value of information, and determine when experimental continuation is economically justified. Decision trees and Markov processes are rarely applied specifically to antiviral preclinical pipelines, leaving a methodological gap in modelling sequential choices under uncertainty.

The study aims to develop and evaluate formalised approaches for optimising sequential decision-making in preclinical antiviral research using decision tree models and Markov decision processes. These models are applied to real-world experimental datasets to quantify transition probabilities, estimate costs, and compare the effectiveness of alternative strategies.

The scientific novelty of the work lies in the integration of an interpretable set of rules provided by the decision trees with globally optimal Markov strategies and value of information analysis. Unlike prior studies, in the proposed framework, predictive patterns, uncertainty quantification, experimental costs and utility maximisation are combined into a unified scheme that supports planning throughout the entire preclinical process.

Problem statement

The object of the study is the process of preclinical evaluation of antiviral drugs, while the subject is mathematical methods for optimising sequential decision-making in this process, in particular decision trees and Markov models. The purpose of the work is to develop and test formalised approaches to assessing the effectiveness, risks and feasibility of transitions between stages of preclinical studies based on real experimental data. The end result is the construction and comparative analysis of two algorithmic models – the decision tree and the Markov process – that demonstrate their ability to support rational, data-driven planning for preclinical testing of antiviral candidates.

Materials and methods

The study uses three interrelated methods: decision trees to formalise the process of selecting compounds, Markov models to describe the sequence of experimental steps over time and value of information analysis metrics to quantify the feasibility of doing additional measurements. This combination allows for moving from the description of individual experiments to a systematic approach where each step is considered to be an element of an optimised decision-making process.

Decision trees act as interpreted classification models that reflect the relationship between a set of input parameters (docking parameters, concentration characteristics, toxicity and antiviral activity indicators) and binary output events (e.g., reaching a chemotherapy index threshold). The decision tree is a hierarchical structure, where each inner node corresponds to a condition of the type "sign ≤ threshold", branches to alternative consequences of this condition, and leaf nodes to result classes. The construction of the tree is carried out by sequentially dividing the feature space to minimise the degree of heterogeneity (for example, the Gini index) in the daughter nodes at each stage. As a result, a set of simple logical rules is formed that allows for explicit interpretation of which combinations of docking, CD_{50}, ID_{50}, and exposure.

Markov chains and Markov decision-making processes are used to describe the evolution of a system in discrete states, taking into account the pro-

babilistic nature of transitions between them. In the simplest case, the Markov chain is given by a set of states and a matrix of transient probabilities, where the probability of moving to the next state depends only on the current state, and not on the complete history. In the context of planning the sequence of experiments, this allows us to consider individual stages (docking, assessment of cytotoxicity, testing of antiviral activity, achievement or failure of therapeutic success) as states of the Markov process, and possible actions of the investigator ("to continue" or "stop" the study at a certain stage) as controlling influences that change the distribution of probabilities of further states. In this formulation, the Markov model of decision-making is used, where each state-action pair corresponds not only to the probability of transition, but also to a certain instantaneous reward or cost, and the optimal strategy is determined by solving the Bellman equations for the value function.

Value-of-information analysis metrics are tools for evaluating the extent to which the anticipated utility of decisions can be enhanced through additional data. Conceptually, the value of information is defined as the difference between the expected utility of an optimal policy given the availability of additional information and the expected utility at the baseline level of uncertainty. The Total Value of Perfect Information (EVPI) reflects a hypothetical increment if the results of the experiments were known in advance without errors; partial value of perfect information (EVPPI) characterizes a similar increase for certain groups of parameters (for example, only for cytotoxicity indicators or only for antiviral activity); The expected value of the sample information (EVSI) assesses how much conducting a realistic incremental experiment with a certain value is able to improve decision-making. In combination with the Markov model of the experimental process, these metrics enable a quantitative comparison of various research design variants, determine the stages at which new measurements give the greatest increase in information about CTI, and justify the optimal balance between the costs of the experiment and the probability of obtaining therapeutically significant candidates.

The study used a multi-level methodology combining experimental data on docking, cytotoxicity, and antiviral activity with mathematical dose-response modelling, decision tree construction, and Markov experiment sequence modelling. The main target characteristic is the CTI chemotherapy index, calculated based on CD_{50} and ID_{50} concentrations for each test sample and cell processing regimen.

Baseline data included energy parameters of molecular docking of a series of candidate compounds to the domains of the spike protein of the virus and the main protease, results of cytotoxicity tests on the cell line at two time points (24 and 48 h) and results of tests of antiviral activity in the therapeutic (L) and therapeutic-prophylactic (LP) modes. For the docking, numerical estimates of binding energy (in conventional units of energy) with individual target sites were considered, which are represented as $SP_1 - SP_5$ for the spike protein and $MP_1 - MP_3$ for the main protease. For each compound, a docking parameter vector was obtained, which was further used as an input trait space in decision tree models.

Dose-effect modelling for cytotoxicity and antiviral activity was carried out using a four-parameter sigmoidal model. For each compound, exposure time, and treatment regimen, a set of concentrations was given $\boldsymbol{x}_i$ and the corresponding measured values of relative cell viability (for cytotoxicity) or relative viral activity (for antiviral action), normalised to control in the interval [0; 1]. As a model function, the expression

$$f(x)=\frac{a}{1+e^{(-bx+c)}}+d \qquad (1)$$

where x is the concentration of the compound, $f(x)$ is the expected relative value of the indicator (viability or activity), a, b, c, d are unknown parameters of the curve describing the amplitude of the effect, the steepness of the transition, the shift along the concentration axis and the baseline, respectively. The estimation of the parameters was carried out by the method of least squares by minimising the root mean square error

$$MSE(a,b,c,d)=\frac{1}{N}\sum_{i=1}^{N}(y_i - \\ - f(x_i,a,b,c,d))^2 \qquad (2)$$

where y_i – experimental values of the relative viability of cells or the relative activity of the virus, N is the number of points of the curve. Optimisation was carried out by the numerical method of nonlinear regression with constraints on parameters to avoid unrealistic decisions; in cases where numerical optimisation did not match, stable heuristic initial approximations were used, providing a smooth monotonic curve within the studied concentration range.

Based on the fit of the sigmoidal model, the characteristic concentrations of CD_{50} and ID_{50} were determined. The concentration of CD_{50} was determined as the solution of the equation

$$f_{cyto}(x) = 0.5 \quad (3)$$

that is, the concentration at which the relative viability of cells is 50 % of the control. Similarly, the concentration of ID_{50} was defined as the solution

$$f_{virus}(x) = 0.5 \quad (4)$$

corresponding to a 50 percent level of residual virus activity. For a given four parameters (*a*, *b*, *c*, *d*), the analytical expression for such a concentration was obtained from the equation x^*

$$y_{\text{target}} = \frac{a}{1+e^{(-bx^*+c)}} + d \quad (5)$$

by algebraic transformation:

$$y_{\text{target}} - d = \frac{a}{1+e^{(-bx^*+c)}}, \quad \frac{a}{y_{\text{target}} - d} - 1 = e^{(-bx^*+c)}; \quad (6)$$

$$-bx^* + c = \ln\left(\frac{a}{y_{\text{target}} - d} - 1\right), \quad x^* = \frac{c - \ln\left(\frac{a}{y_{\text{target}} - d} - 1\right)}{b}. \quad (7)$$

In cases where the expression under the logarithm was incorrect (negative or zero) or the parameter *b* was close to zero, the value of CD_{50} or ID_{50} was considered uncertain (no intersection with the level of 50 % in the studied range).

The chemotherapeutic index for each combination "compound – time – treatment regimen" was calculated according to the standard ratio

$$\text{CTI} = \frac{CD_{50}}{ID_{50}} \quad (8)$$

which is interpreted as a safety margin: the larger the CTI, the wider the therapeutic interval between cytotoxic and antiviral concentrations. For further classification analysis, CTI was converted to a binary trait by threshold: the value of $\theta = 4{,}0$ class 1 denoted combinations with CTI ≥ 4, and class 0 – CTI < 4, which made it possible to interpret the problem as a two-class problem of "promising / unpromising" candidates.

To investigate the relationship between the docking profile of compounds, cytotoxicity parameters and the probability of obtaining a high CTI, the decision trees method of the CART (Classification and Regression Trees) type was used. In the first model, the decision tree described the probability of obtaining a determined CD_{50} based on docking indicators. The trait vector included energy parameters of interaction with different regions of the spike protein and the main protease (SP_1 – SP_5, MP_1 – MP_3), as well as a coded timestamp of cell exposure (time_class, where 0 corresponded to 24 hours, 1 to 48 hours). The target variable class_CD_{50} took a value of 1 if CD_{50} was defined for the corresponding compound-time combination, and 0 in the opposite case. Thus, the first model evaluated which docking profiles are associated with the presence of a correct dose-appropriate cytotoxicity curve.

In the second model, the decision tree modeled the dependence of the "connection-time-mode" combination belonging to the CTI class ≥ 4 on the combination of docking characteristics, CD_{50} and ID_{50} parameters, and the treatment regimen. In addition to SP_1 – SP_5, MP_1 – MP_3 and time_class, the numerical values of CD_{50}, ID_{50}, CTI itself, as well as the encoded trait treatment_type_class (0 for mode L and 1 for LP ≥ class_CTI) were added to the trait vector. Both models were built as binary trees with a Gini index division criterion that minimises class heterogeneity in nodes.

Decision trees were trained according to the scheme of dividing the sample into training and test subsamples in the ratio of 70 % / 30 % with a fixed random number generator to ensure reproducibility. In the presence of both classes, a stratified division was carried out to preserve the proportions of the classes in the training and test parts. The depth of the trees was limited to a predetermined maximum to avoid overtraining, and the number of leaf nodes and the structure of the resulting rules were analysed to control the complexity of the model. The text representation of the tree in the form of nested "if" rules was obtained by traversing the structure of the tree, where each inner node specifies a condition of the form "sign ≤ threshold", and the leaf node – belonging to class 0 or 1.

Assessment of the quality of classification models was carried out on test subsamples using a set of standard metrics. Accuracy was defined as the proportion of correctly classified examples:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (9)$$

where TP (true positives) is the number of true positive classifications, TN (true negatives) is true negative, FP (false positives) is false positive, FN (false negatives) is false negative. Sensitivity (or recall for a positive class) was defined as

$$Sensitivity = \frac{TP}{TP+FN} \quad (10)$$

which reflects the model's ability to detect positive cases. Specificity was calculated as

$$Specificity = \frac{TN}{TN+FP} \quad (11)$$

characterising the ability of the model to correctly cut off negative cases. To assess the balance between sensitivity and accuracy of positive classifications, the F1 measure was used:

$$\text{Precision} = \frac{TP}{TP+FP}, \quad F1 = \frac{2 \cdot \text{Precision} \cdot Sensitivity}{\text{Precision} + Sensitivity}. \quad (12)$$

In the case of probabilistic model outputs (predict_proba), the area under the ROC curve (ROC-AUC) was additionally calculated, which characterises the trade-off between sensitivity and specificity when varying the classification threshold. In cases where all observations belonged to the same class and the ROC curve was incorrectly determined, the ROC-AUC was not interpreted.

Additionally, two complementary procedures were used to analyse the contribution of individual traits. First, standard estimates of the importance of traits in the tree were used based on a decrease in the Gini criterion when splitting according to the corresponding trait. Secondly, the sensitivity analysis of the "drop-one-feature" sensitivity was performed: for each feature, a new decision tree was built without this feature in the feature vector, and then the model metrics were compared with the base variant. A significant degradation of sensitivity or specificity in the exclusion of a certain trait was interpreted as an indicator of its critical importance for decision-making.

To formalise the sequence of decision-making on the continuation or termination of laboratory tests at different stages (docking, assessment of cytotoxicity, testing of antiviral activity), the Markov Decision Process (MDP) was used. The state space described the main stages of the study: post-docking baseline (S_0), post-doc status (S_1), post-antiviral activity (S_2), post-doctrinal status, and two absorption states – success ($S_{success}$, candidate acquisition with CTI ≥ 4) and completion without success (S_{fail}). In each of the non-absorption states, two actions were considered: "stop" – stop further experiments, and "continue" – move to the next stage of the study.

Transitions between states were described by probabilities that were estimated on the basis of empirical frequencies in the population of the compounds studied. For example, the probability of transition from $SP(s'|s,a)_0$ to S_1 under the action "continue" was estimated as the proportion of compounds for which CD_{50} could be determined; the probability of transition from S_1 to S_2 as the proportion of combinations for which CTI was determined; the probability of transition from S_2 to $S_{success}$ as the proportion of cases with CTI ≥ 4 among those who have passed to the stage of antiviral activity tests. The reward system $R(s, a)$ included negative contributions in the form of docking costs, cytotoxicity and antiviral activity tests, as well as a positive reward for achieving a state of $S_{success}$ corresponding to obtaining a promising candidate. To assess the long-term usefulness of action policies, a discount factor $\gamma \in (0;1)$ was used, which takes into account the decrease in the "value" of time-distant results.

The optimal policy that maximizes the expected discounted total reward was determined by the value iteration method. At each iteration, the value of the utility function $\pi^*(s)V(s)$ was updated according to the rule

$$V_{k+1}(s) = \max\left[R(s,a) + \gamma \sum_{s,a} P(s' \mid s,a) V_k(s')\right] \quad (13)$$

until the changes in $V(s)$ for all states become less than the predetermined error. The choice of action in state s was carried out as an argument to the maximum on the right side of the Bellman equation.

Interpretation of the supplementary experiments as sources of information about the probability of success and the associated costs allowed us to integrate the concept of information value analysis into the Markov formulation. The total expected value of perfect information (EVPI) is the difference between the expected utility from having complete, error-free information about the outcome (e.g., CTI for each compound before the experiments were performed) and the expected utility at the current level of uncertainty. The expected value of partial perfect information about a subset of parameters (EVPPI) marks a similar difference, but only for information about a particular block of parameters (for example, only about CD_{50} or only about ID_{50}). The expected value of sample information (EVSI) determines the increase in expected utility obtained by making additional, but not error-free measurements (for example, additional experiments on cytotoxicity

or antiviral activity). In mathematical formulation, these indicators are calculated as the difference between the maximum expected utility according to the refined probability distribution (after taking into account the new data) and the maximum according to the original distribution; within the framework of MDP, it comes down to comparing the values of $V(s)$ under different assumptions about the information state of the system.

Formally, let θ denote the vector of uncertain parameters (e.g., transition probabilities and success rates), d a decision or policy, and $U(d, \theta)$ the total discounted utility under this decision. The baseline expected value at the initial state s_0 is

$$V_{base}(s_0) = \max_{d \in D} E_\theta[U(d,\theta) \mid I_0] \quad (14)$$

where I_0 denotes the current information set. The expected value of perfect information (EVPI) is defined as

$$\text{EVPI} = E_\theta[\max_{d \in D} E_\theta[U(d,\theta) \mid I_0] - V_{base}(s_0). \quad (15)$$

It quantifies the maximum gain in expected utility that could be achieved if θ were known without uncertainty before any decision is made.

For a subset of parameters $\varphi \subset \theta$ (for example, only cytotoxicity or only antiviral activity parameters), the expected value of partial perfect information (EVPPI) is given by

$$EVPPI(\varphi) = E_\varphi \times$$

$$\times[\max_{d \in D} E_\theta[U(d,\theta) \mid I_0]] - V_{base}(s_0). \quad (16)$$

Here, only the subset φ is assumed to be known perfectly, whereas the remaining parameters θ\φ remain uncertain.

The expected value of sample information (EVSI) associated with a realistic additional experiment y is defined as

$$\text{EVSI}(y) = E_y[\max_{d \in D} E_\theta[U(d,\theta) \mid I_0]] -$$

$$- V_{base}(s_0) - C(y) \quad (17)$$

where y denotes the possible outcomes of the new experiment, and $C(y)$ is the cost of collecting this information. Within the MDP formulation used in this work, all three quantities can be evaluated as differences between optimal state values $V(s_0)$ computed under different information scenarios (baseline information, perfect information on θ or φ, and posterior distributions updated by sample data y).

This approach provides a holistic methodological framework: sigmoidal dose-response modelling allows for stable estimates of CD_{50}, ID_{50}, and CTI; decision trees formalise the logic of the selection of compounds according to the profiles of docking, cytotoxicity and antiviral activity; The Markov model with associated information value metrics allows you to assess the feasibility of continuing or stopping experiments at different stages, taking into account costs and the likelihood of achieving chemotherapy success.

The study used a holistic experimental dataset which covered seven test samples (T1 – T7), each having obtained both in silico and in vitro characteristics. The generalised structure of this set is provided in Table 1.

Results of molecular docking to the spike protein and the main protease of the virus are available for each sample and presented in the form of binding energy ranges. In particular, for the spike protein, the minimum and maximum values of the docking energy for a set of domains (SP_1 – SP_5) were taken into account, and for the main protease – the minimum and maxi-

***Table* 1.** Generalised characteristics of the experimental dataset

Test Sample	Min. Docking energy (spike protein), kcal/mol	Max. Docking energy (spike protein), kcal/mol	Min. Docking energy (basic protease), kcal/mol	Max. Docking energy (basic protease), kcal/mol	Summarised CTI in 24 hours	Summarised CTI in 48 hours	Cytotoxicity data (24/48 hours)	Antiviral activity data (L, LP; 24/48 hours)
T1	–6.26	–2.26	–5.68	–4.04	24	16	yes	yes
T2	–6.32	–3.3	–7.34	–4.07	–	4	yes	yes
T3	–5.24	–2.4	–7.12	–4.04	3	2	yes	yes
T4	–5.74	–2.26	–7.34	–4.04	–	–	yes	yes
T5	–4.96	–2.37	–7.34	–4.04	4	3	yes	yes
T6	–6.26	–2.4	–7.34	–4.04	6	6	yes	yes
T7	–5.64	–2.26	–7.34	–4.21	–	–	yes	yes

mum values of the interaction energy with individual functional regions (MP_1 – MP_3). Table 1 shows two aggregate intervals for each test sample: minimum and maximum docking energy for the spike protein and major protease, reflecting the spectrum of possible binding configurations within the docking protocols used.

Results

The results obtained demonstrate a consistent chain of transition from in silico characteristics to integral assessment of the chemotherapeutic index and optimisation of the sequence of experiments. First of all, the analysis of the initial CD_{50}, ID_{50} and CTI values for seven test samples showed significant inter-sample variability: for some of the sample-time-mode combinations, CD_{50} or ID_{50} could not be correctly estimated at all within the studied concentration range (the curve did not cross the 50 % level), while others showed well-defined half-inhibition points and high CTI values. It is important to note that this variability turned out to be structured – it is related to the profile of the docking to the main protease and to the exposure mode, and is not random noise; This is confirmed by the construction of the first decision tree and the Markov model.

In the "docking → cytotoxicity" model, the decision tree gives a compact but meaningful structure (Fig. 1).

The root node separates all observations by the MP_2 parameter, which characterises the binding energy of compounds to one of the functional areas of the main protease. If MP_2 is found to be above the threshold value of approximately – 7.0 kcal/mol (i.e., binding is weaker), the model, without further branching, assigns the corresponding sample-time combination to a class for which CD_{50} is not determined: all cases where the dose-response of cytotoxicity remains flat or monotonically low are concentrated in this leaf node. Instead, for compounds with a more favourable MP_2 value (≤–7.0 kcal/mol), the tree moves on to the second critical parameter – MP_1. For MP_1 values ≤–5.97 kcal/mol, all examples fall into the sheet with the class "has CD_{50}", i.e. a sufficiently strong interaction with two regions of the main protease is a reliable predictor of the presence of correctly defined CD_{50}. In the intermediate zone, where MP_2 still indicates a fairly strong binding, but MP_1 is already closer to the threshold, the model additionally takes into account the incubation time: for 24 h, some of the combinations remain in the class without a defined CD_{50}, while at 48 h the tree leans towards a class with a defined CD_{50}. Thus, exposure time acts as a secondary, modulating factor that can compensate for the insufficient "force" of docking, but only in a narrow sub-range of MP_1 and MP_2 values. On the test sample, this model demonstrates classification performance with all metrics equal to 1.0, indicating that the hierarchical combination of MP_2, MP_1, and incubation time provides deterministic rules for predicting CD_{50} availability. Trait weights and analysis of the exclusion of individual predictors confirm the dominance of MP_2 and MP_1 parameters: they provide the main contribution to the reduction of the Gini index, while the time factor affects the quality of classification much weaker.

The results for the "cytotoxicity → CTI ≥ 4.0" model turned out to be fundamentally different and highlighted the structural features of the existing dataset. Applying the decision tree to an extended trait vector that included docking scores, CD_{50} numerical values, ID_{50}, CTI itself, encoded treatment time and regimen, resulted in a moderately complex tree structure (Fig. 2).

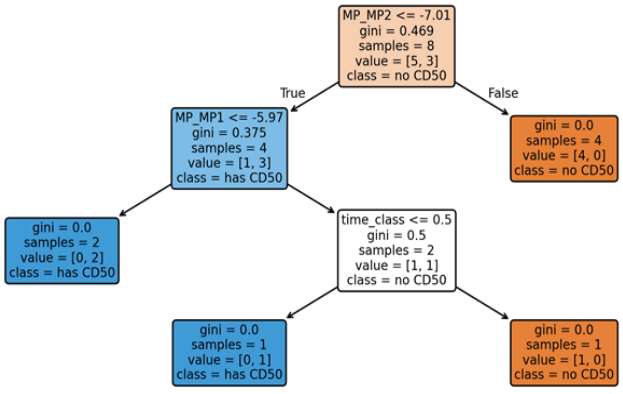

Fig. 1. Decision tree for the transition "docking → cytotoxicity" using the parameters of interaction with the main protease (MP_1, MP_2) and exposure time (24/48 hours) as predictors of the presence of correctly determined CD_{50}

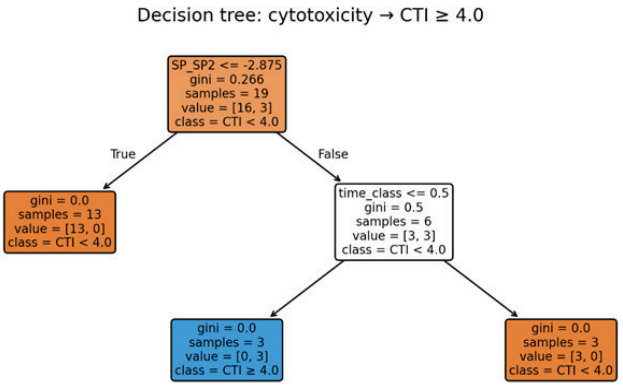

Fig. 2. Decision tree for the transition "cytotoxicity → CTI ≥ 4.0"

Unlike the first model, this tree exhibits branching that attempts to separate compounds achieving CTI ≥ 4.0 from those with lower therapeutic indices. However, the classification performance reflects the challenging nature of this prediction task. The model achieved an overall accuracy of 0.89, indicating that 89 % of test cases were correctly classified. Notably, sensitivity was substantially lower at 0.50, meaning that only half of the compounds that actually achieved CTI ≥ 4.0 were correctly identified by the model. In contrast, perfect specificity at 1.00 indicates that all compounds predicted to have CTI < 4.0 were indeed below this threshold – the model made no false positive predictions. This imbalance between precision and recall is reflected in the F1-score of 0.67, while the ROC-AUC of 0.75 suggests moderate discriminative ability.

This performance pattern reveals an important characteristic of the current model: it adopts a highly conservative strategy and predicts high CTI rarely unless multiple favourable conditions are met. Although this approach does eliminate false optimism, it often results in missed opportunities, as approximately half of promising compounds are not identified. The Markov model of the sequence of experiments made it possible to quantify how expedient it is to continue the study at each stage, taking into account the costs and probability of obtaining at least one candidate with a CTI ≥ 4.0. For states S_0 (after docking), S_1 (after assessment of cytotoxicity), and S_2 (after assessment of antiviral activity), values of the utility function $V(s)$ were calculated, which increase from about 9.7 for S_0 to 22.5 for S_1 and 84.0 for S_2 as shown in Table 2.

***Table* 2.** The value of the utility function $V(s)$ for Markov model states and the optimal policy $\pi^*(s)$

	S_0	S_1	S_2	$S_{success}$	S_{fail}
$V(S_i)$	9.688	22.5	84	0	0
$\pi(S_i)$	continue	continue	continue	stop	stop

Such monotonous growth means that each subsequent block of experiments significantly increases the expected "cost" of the candidate portfolio: at the docking stage, information about the potential of compounds is still very uncertain; obtaining CD_{50} adds an important layer of safety assessment and cuts off clearly toxic variants; completion of antiviral activity tests virtually determines whether a portfolio has a chance of containing at least one drug with an acceptable CTI. The calculated optimal policy of π^*(s) turned out to be unambiguous: for all three non-absorption states, the action "continue experiments" is recommended, while in the absorption states of success or failure – "stop". On the one hand, this is consistent with the high CTI values among the samples that have passed the previous filters: the projected benefit from the complete passage of all stages exceeds the total costs. On the other hand, such a policy indicates that the structure of the experimental program does not currently contain "redundant" stages: each of them significantly changes the expected utility, and therefore makes a non-trivial contribution to reducing uncertainty about CTI.

Interpreting these results in terms of information value shows that the greatest gain in utility is given by the transition from a post-cytotoxicity state to a post-antiviral state. This means that it is the ID_{50} results and the associated CTI values that are key to the final decision on the feasibility of promoting the compound; information about the docking profile and CD_{50} plays mainly the role of a pre-filter. The tree "docking → cytotoxicity" clearly shows that already at the stage of in silico evaluation, a combination of MP_2 and MP_1 parameters can be distinguished, which, with a high probability, leads to the formation of a correct CD_{50} curve. Further, the Markov model demonstrates that, despite the costs, continuing the studies to the stage of antiviral activity is economically justified, since the expected gain from the potential detection of at least one candidate with a high CTI significantly exceeds the alternative of "stopping" in the early stages. At the same time, the moderate performance of the tree "cytotoxicity → CTI ≥ 4.0" (Table 3) signals that in the current dataset, predicting final therapeutic success from intermediate parameters remains challenging. While the model's high specificity ensures that no unpromising compounds are incorrectly advanced, its lower sensitivity indicates that approximately half of actually promising candidates are not recognised by the current decision rules, suggesting that more sophisticated stratification approaches or additional predictive features may be needed to improve early identification of therapeutic candidates.

Numerical evaluation of the information-value metrics within the MDP formulation yielded an expected value of perfect information (EVPI) equal to zero, as well as a zero partial value of perfect information (EVPPI) for the antiviral stage, while the expected value of sample information (EVSI) for an additional antiviral experiment was negative and equal to −10 in the adopted arbitrary utility units. This pattern indicates that, under the current estimates of transition probabilities and reward structure, even hypothetically perfect knowledge of the

***Table* 3.** Comparative metrics of decision tree models and Markov model

Model	Positive class	Accuracy	Sensitivity (Recall)	Specificity	F1-score	ROC-AUC
Docking → CD_{50} (presence of CD_{50} defined)	has CD_{50}	1	1	1	1	1
Cytotoxicity → CTI ≥ 4.0	CTI ≥ 4.0	0.89	0.50	1.00	0.67	0.75

antiviral success probability would not change either the optimal policy or the expected value at the initial state, so that additional information of this type has no incremental decision value. At the same time, the negative EVSI reflects the fact that, in the simplified scenario considered, a realistic extra experiment on antiviral activity does not lead to a sufficient increase in the expected utility to compensate for its cost, implying that resources would be more efficiently allocated to adjusting earlier-stage selection criteria or expanding the candidate set rather than intensifying measurements at the final antiviral stage.

The comparative values of classification metrics for both decision trees, as well as the utility functions of the states of the Markov model, are summarised in Table 3, allow you to quantitatively compare the accuracy of predictive decisions and the expected effectiveness of various strategies for conducting experiments.

Collectively, this indicates that further experimental studies should be directed either to expand the sample set (to include more examples with intermediate and low CTIs) or to clarify dosing regimens and time regimens, where the gap between efficacy and toxicity will be less obvious and, accordingly, will provide a richer structure for building more complex but informative decision trees.

Conclusions

This research shows that decision trees and Markov decision processes are used for complementary aspects in preclinical antiviral research optimisation, and their performance is fundamentally shaped by the structure of the available data. The decision-tree model linking docking parameters to the presence of a well-defined CD_{50} achieved classification success on the test set, with all metrics equal to 1.0. This indicates that strong binding to the main protease, in particular, favourable values of MP_1 and MP_2 combined with incubation time provides deterministic rules for predicting the formation of a stable dose-response curve. While these results imply robust predictive ability within the current dataset, the limited sample size requires validation on larger, more diverse compound libraries in order to confirm generalizability.

In contrast, the tree built for the transition from cytotoxicity parameters to CTI ≥ 4.0 exhibited substantially different performance characteristics. The model achieved 89 % overall accuracy but demonstrated an asymmetric error profile: sensitivity of only 0.50 (correctly identifying half of compounds with CTI ≥ 4.0) combined with specificity of 1.00 (no false positives). This classification strategy ensures that compounds predicted to achieve high CTI are indeed therapeutically promising. Still, it results in missed opportunities, as approximately half of actually promising candidates are not recognised by the current decision rules. The F1-score of 0.67 and ROC-AUC of 0.75 show that the relationship between intermediate experimental parameters and final therapeutic success is more complex than can be captured by simple threshold-based rules with the current feature set and sample size.

The Markov decision process provided a global, quantitatively interpretable view of the same experimental pipeline, explicitly integrating transition probabilities, experimental costs and the probability of achieving at least one candidate with CTI ≥ 4.0. The estimated state values $V(S_0)$, $V(S_1)$ and $V(S_2)$ increased monotonically during the stage "after docking → after cytotoxicity → after antiviral testing", confirming that each successive block of experiments substantially raises the expected utility of the candidate portfolio. The optimal policy consistently recommended the continuation of the preclinical program from docking through cytotoxicity to antiviral assays, and that the stop is required only in the absorbing states of success or failure. Despite the second decision tree performing the imperfect classification, the Markov model still indicated that, under the assumed costs and probabilities of success, the full three-stage pipeline is an economically justified solution and does not contain redundant experimental steps.

The value of information analysis shows how much each stage of the pipeline contributes to reducing uncertainty and improving decision quality. The results suggest that the largest incremental gain in

expected utility appears when the compound moves from the state after cytotoxicity assessment to the state after antiviral activity testing. This highlights the decisive role of ID_{50} and derived CTI values in confirming or rejecting candidates. Information on docking and CD_{50} acts primarily as a preliminary filter that shapes the distribution of outcomes that are observed at later stages. Decision trees are most effective as local, interpretable tools for formalising the cut-off rules based on docking and toxicity profiles at the early stages, whereas the Markov model outperforms them in terms of providing a globally optimal, cost-aware strategy for navigating the entire preclinical pipeline. Together, these results show that further development should focus on expanding and rebalancing the experimental dataset (to enable more informative tree-based models at the CTI level) and on refining Markov and value of information formulations to incorporate richer biological and economic parameters in the optimisation of antiviral preclinical programs.

References

[1] U.E. Ogbonna *et al.*, "Advances in machine learning for optimizing pharmaceutical drug discovery", *Curr. Proteomics*, Vol. 22, no. 2, p. 100015, Apr. 2025. Retrieved from doi: https://doi.org/10.1016/j.curpro.2025.100015

[2] A. Gangwal and A. Lavecchia, "Artificial intelligence in preclinical research: enhancing digital twins and organ-on-chip to reduce animal testing", *Drug Discov. Today*, Vol. 30, no. 5, p. 104360, May 2025. Retrieved from doi: https://doi.org/10.1016/j.drudis.2025.104360

[3] M. Mashhadi Abolghasem Shirazi *et al.*, "Next-generation antiviral peptides: AI-driven design, translational delivery platforms, and future therapeutic directions", *Virus Res.*, Vol. 361, p. 199642, Nov. 2025. Retrieved from doi: https://doi.org/ 10.1016/j.virusres.2025.199642

[4] A. Luganini *et al.*, "DHODH inhibitors: What will it take to get them into the clinic as antivirals?", *Antiviral Res.*, Vol. 236, p. 106099, Apr. 2025. Retrieved from doi: https://doi.org/10.1016/j.antiviral.2025.106099

[5] N. Vora *et al.*, "Artificial intelligence and multi-omics in drug discovery: A deep learning-powered revolution", *Cure Care*, p. 100011, Nov. 2025. Retrieved from doi: https://doi.org/10.1016/j.ccwv.2025.100011

[6] K.O. Oyediran *et al.*, "Artificial intelligence in human immunodeficiency virus mutation prediction and drug design: Advancing personalized treatment and prevention", *Pharm. Sci. Adv.*, Vol. 3, p. 100080, Dec. 2025. Retrieved from doi: https://doi.org/10.1016/j.pscia.2025.100080

[7] A. Gangwal and A. Lavecchia, "Artificial Intelligence in Natural Product Drug Discovery: Current Applications and Future Perspectives", *J. Med. Chem.*, Vol. 68, no. 4, pp. 3948–3969, Feb. 2025. Retrieved from doi: https://doi.org/10.1021/acs.jmedchem.4c01257

Д.С. Городецький, М.П. Сметюх, С.О. Соловйов

ПРИЙНЯТТЯ РІШЕНЬ У ПРОЦЕСІ ДОКЛІНІЧНОЇ РОЗРОБКИ ПРОТИВІРУСНИХ ПРЕПАРАТІВ ПРОТИ КОРОНАВІРУСУ: МАТЕМАТИЧНЕ МОДЕЛЮВАННЯ ТА АНАЛІЗ ЦІННОСТІ ІНФОРМАЦІЇ

Проблематика. Доклінічне оцінювання кандидатів як противірусних препаратів – це багатоетапний процес, який супроводжується значною невизначеністю і потребує формальних інструментів підтримки прийняття рішень. Традиційні підходи до скринінгу сполук зазвичай не містять структурованих методів оптимізації для послідовного вибору, а також оцінювання доцільності й ризиків переходів між етапами. Це призводить до неефективності під час відбору перспективних молекул і підвищує ймовірність вибору субоптимальних дослідницьких траєкторій.

Мета дослідження. Розробити й обґрунтувати формалізований підхід, щоб оптимізувати переходи між етапами доклінічного тестування противірусних препаратів. Цей підхід інтегрує дерева рішень і марковську модель для оцінювання ефективності, ризиків і цінності додаткової інформації, що забезпечить раціональне планування послідовності доклінічних досліджень.

Методика реалізації. Експериментальні дані з молекулярного докінгу, цитотоксичності CD_{50} та антивірусної активності IC_{50} були інтегровані в каскадну систему оцінювання із критерієм переходу XTI ≥ 4. Дерева рішень забезпечили інтерпретовані правила просування сполук, а за допомогою марковської моделі було змодельовано послідовні стратегії в умовах невизначеності та оцінено доцільність переходів між етапами. За допомогою аналізу цінності інформації було оцінено очікувану користь додаткових експериментальних даних.

Результати дослідження. Описаний підхід дав узгоджені технічні результати. Дерево рішень для прогнозування CTI ≥ 4,0 показало консервативний шаблон класифікації, правильно визначаючи сполуки з високим терапевтичним потенціалом, але пропускаючи частину ефективних кандидатів. Марковська модель допомогла оцінити стан системи на етапах докінгу, цитотоксичності й антивірусного тестування, що показало зростання очікуваної корисності. Ґрунтуючись на отриманих результатах, було визначено оптимальні рішення щодо продовження досліджень до антивірусних тестів, тоді як за допомогою аналізу цінності інформації було встановлено, що найбільший приріст очікуваної корисності досягають після тестування антивірусної активності, коли ранні етапи виконують роль фільтрів.

Висновки. Дослідження показує, що дерева рішень і марковські моделі відображають різні, але взаємодоповнювальні аспекти доклінічного оцінювання. Дерева рішень допомагають структурувати правила на ранніх етапах дослідження, показуючи, як етапи докінгу та цитотоксичності впливають на просування сполук. Водночас їх обмежена чутливість підкреслює складність передбачення кінцевого противірусного успіху на основі проміжних показників. Марковський процес дає ширший погляд на послідовність експериментів і демонструє виправданість вибору повного трирівневого дослідження та впливу невизначеності й витрат на рішення щодо прогресії сполук. Результати аналізу цінності інформації уточнюють важливість кожного етапу, підкреслюючи ключову роль даних про антивірусну активність. Разом ці результати показують важливість впровадження методів прийняття рішень для підвищення структури, прозорості та ефективності доклінічних досліджень противірусних препаратів.

Ключові слова: коронавірус; препарат; доклінічне оцінювання; дерево рішень; марковський процес прийняття рішень; цінність інформації.

Рекомендована Радою
факультету прикладної математики
КПІ ім. Ігоря Сікорського

Надійшла до редакції
18 жовтня 2025 року

Прийнята до публікації
15 грудня 2025 року