Mathematics

HYBRID FUZZY–EVOLUTIONARY FRAMEWORK FOR MULTI-CRITERIA DRUG SELECTION IN PERSONALIZED MEDICINE

Javid Yashar Abbasli
Azerbaijan Technical University, Baku, Azerbaijan: cavid.abbasli@aztu.edu.az
https://orcid.org/0009-0005-8982-0164

Abstract. Personalized drug selection requires balancing therapeutic efficacy, adverse events, drug interactions, and cost–objectives that are inherently conflicting. This study proposes a hybrid architecture integrating Mamdani fuzzy inference, ensemble machine learning, NSGA-II optimization, and Fuzzy DIBR II ranking for multi-objective pharmacotherapy. The Mamdani system encodes eight clinical safety rules to filter unsuitable patient-drug pairs based on renal and hepatic function. Random Forest and XGBoost models trained on 44,692 pairs achieve AUC-ROC of 0.871 and 0.907 for efficacy and adverse event prediction, respectively. NSGA-II generates nine unique Pareto-optimal prescriptions over 30 generations. Fuzzy DIBR II elicits criterion weights from physician preferences modeled as triangular fuzzy numbers ($\sigma = 0.7$), yielding $w_1 = 0.387$ (efficacy), $w_2 = 0.302$ (safety), $w_3 = 0.205$ (interactions), $w_4 = 0.105$ (cost). The top-ranked prescription achieves 84.4% efficacy, 39.4% adverse events, 11.6% interaction severity, and 608.60 AZN cost. Results demonstrate clinically interpretable multi-criteria recommendations with explicit trade-off quantification.
Keywords: *multi-objective optimization, fuzzy inference, NSGA-II, fuzzy DIBR II; personalized medicine*.

Introduction. Clinical pharmacotherapy is a complex multi-criteria decision problem that requires clinicians to simultaneously balance therapeutic efficacy, adverse event risk, drug–drug interactions, and treatment cost while respecting patient-specific physiological constraints [1]. Existing drug selection practices rely mainly on clinical guidelines and physician experience, typically optimizing one dominant criterion and handling the others implicitly or via ad hoc rules [2].

Multi-objective optimization offers a principled way to generate sets of Pareto-optimal prescriptions that represent distinct trade-off strategies instead of enforcing fixed preferences a priori. Evolutionary algorithms such as NSGA-II provide efficient search and diversity maintenance, but by themselves do not address two key clinical requirements: enforcing patient-specific safety constraints and translating Pareto sets into actionable recommendations aligned with uncertain clinician preferences. Fuzzy logic and fuzzy MCDM methods are well-suited to encode linguistic clinical knowledge and model uncertainty, while ensemble machine learning can provide robust data-driven predictions of efficacy and adverse events.

This study proposes a hybrid five-stage decision-support architecture that integrates Mamdani fuzzy safety filtering, ensemble machine learning (Random Forest and XGBoost) for patient-specific outcome prediction, NSGA-II-based multi-objective optimization, and Fuzzy DIBR II for uncertainty-aware preference modeling [3]. Personalized drug selection is formulated as a constrained four-objective problem that maximizes efficacy and minimizes adverse events, drug–drug interaction severity, and cost under fuzzy safety constraints [4]. Experimental evaluation on a synthesized data yet clinically realistic dataset with 1,000 patients and 50 drugs shows that the framework can generate diverse Pareto-optimal multi-drug prescriptions and rank them into clinically interpretable strategies, with the top solution achieving high predicted efficacy while maintaining acceptable safety and cost.

Related works. Multi-objective evolutionary algorithms have been used in healthcare for problems such as radiotherapy planning, chemotherapy scheduling, and dose optimization, where trade-offs between tumor control and toxicity are critical. These studies demonstrate the value of Pareto-based optimization but typically treat continuous dosing variables, rely on deterministic objective functions, and do not incorporate patient-specific pharmacological safety constraints or discrete multi-drug selection. Fuzzy logic has a long history in medical diagnosis and risk assessment, enabling robust handling of imprecise laboratory values and linguistic clinical rules [5]. Mamdani-type fuzzy systems and neuro-fuzzy models have achieved high diagnostic accuracy [6, pp. 177–179]. Several studies have shown that Random Forest, gradient boosting and graph-based models can accurately predict drug response and adverse events from clinical and omics data.

Multi-criteria decision-making methods such as AHP, TOPSIS, PROMETHEE, and their fuzzy extensions have been applied to drug and health technology evaluation, showing how structured criteria weighting can support formulary and reimbursement decisions [7]. However, these methods usually rely on precise or simple fuzzy weights and do not explicitly represent uncertainty in preference ratios. The DIBR family introduces ratio-based weight elicitation with a stronger theoretical basis and recent fuzzy extensions, but has not yet been applied to pharmacotherapy to jointly capture uncertainty in fuzzy preference ratios.

Mathematical modelling of multi-criteria drug selection. Let $D = \{d_1, d_2, \dots, d_n\}$ denote the set of available drugs and p a fixed patient. A prescription is represented by a binary decision vector $\mathbf{x} = [x_1, x_2, \dots, x_n]^T \in \{0,1\}^n$, where $x_i = 1$ indicates that the drug d_i is included in the regimen and $x_i = 0$ otherwise. The number of selected drugs is constrained by the cardinality condition

$$1 \le \|\mathbf{x}\|_1 = \sum_{i=1}^{n} x_i \le 5 \tag{1}$$

reflecting clinical practice in which only a limited number of concurrent medications is acceptable. For each patient-drug pair (p, d_i) a fuzzy safety system produces a scalar safety score $S_{\text{fuzzy}}(p, d_i) \in [0,1]$. This score combines age, renal and hepatic function, and drug-specific toxicity into a continuous measure of suitability. The following hard safety constraint is imposed, $S_{\text{fuzzy}}(p, d_i) \ge \tau_{\text{safety}}, \forall i$ such that $x_i = 1$, where $\tau_{\text{safety}} = 0.5$ is a clinically chosen threshold.

Drugs violating this inequality are excluded from any admissible prescription. The quality of a prescription $\mathbf{x}$ is assessed through four criteria: expected therapeutic efficacy, risk of adverse events, severity of drug-drug interactions, and total treatment cost. These are modeled as objective functions $f_j(\mathbf{x}), j = 1, \dots, 4$, to be minimized simultaneously.

Machine learning models provide patient-specific probabilities $\hat{P}_{\text{eff}}(p, d_i)$ of achieving sufficient clinical response and $\hat{P}_{\text{se}}(p, d_i)$ of experiencing a clinically relevant adverse event. Let $\mathcal{J}(\mathbf{x}) = \{i : x_i = 1\}$ denote the index set of selected drugs. The first two objectives are defined as

$$f_1(\mathbf{x}) = -\frac{1}{|\mathcal{J}(\mathbf{x})|} \sum_{i \in \mathcal{J}(\mathbf{x})} \hat{P}_{\text{eff}}(p, d_i), \tag{2}$$

$$f_2(\mathbf{x}) = \frac{1}{|\mathcal{J}(\mathbf{x})|} \sum_{i \in \mathcal{J}(\mathbf{x})} \hat{P}_{\text{se}}(p, d_i), \tag{3}$$

so that minimizing f_1 corresponds to maximizing mean efficacy, whereas minimizing f_2 reduces the average probability of adverse events across all prescribed drugs. Potential pharmacological conflicts between drugs are modeled through a symmetric interaction matrix – drug-drug interaction - DDI $\in [0,1]^{n \times n}$, where DDI_{ij} quantifies the normalized severity of the interaction between drugs d_i and d_j. For any prescription containing at least two drugs, the interaction objective is

$$f_3(\mathbf{x}) = \frac{2}{|\mathcal{J}(\mathbf{x})|(|\mathcal{J}(\mathbf{x})| - 1)} \sum_{\substack{i,j \in \mathcal{J}(\mathbf{x}) \\ i<j}} \text{DDI}_{ij}, \tag{4}$$

representing the mean pairwise interaction severity; for single-drug prescriptions, $f_3(\mathbf{x})$ is set to zero.

Let $C(d_i) > 0$ denote the monetary cost of the drug d_i and C_{ref} a reference value used for normalization (in the experiments, $C_{\text{ref}} = 200$ AZN). The cost objective is given by

$$f_4(\mathbf{x}) = \frac{1}{C_{\text{ref}}} \sum_{i \in \mathcal{J}(\mathbf{x})} C(d_i), \tag{5}$$

which scales the total cost to a dimensionless quantity of order one, facilitating joint optimization with the other criteria. The multi-criteria drug selection problem for a patient p can therefore be written compactly as the constrained multi-objective program,

$$\min_{\mathbf{x}\in\{0,1\}^n} \mathbf{F}(\boldsymbol{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), f_3(\mathbf{x}), f_4(\mathbf{x})]^T \tag{6}$$

$$1 \le \|\mathbf{x}\|_1 \le 5, S_{\text{fuzzy}}(p, d_i) \ge \tau_{\text{safety}}, \forall i \in \mathcal{I}(\mathbf{x}). \tag{7}$$

A prescription $\mathbf{x}^{(1)}$ is said to Pareto-dominate another prescription $\mathbf{x}^{(2)}$ if

$$f_j\left(\mathbf{x}^{(1)}\right) \le f_j\left(\mathbf{x}^{(2)}\right), \forall j \in \{1,2,3,4\}, \tag{8}$$

and the inequality is strict for at least one objective. The set of all non-dominated feasible prescriptions constitutes the patient-specific Pareto set. Subsequent sections describe how this formal model is instantiated within a hybrid fuzzy-evolutionary decision-making framework that generates, evaluates, and ranks Pareto-optimal multi-drug prescriptions.

Hybrid fuzzy–evolutionary decision-making framework. The fuzzy safety module employs triangular membership functions defined as:

$$tri(x; a, b, c) = \begin{cases} 0, & \text{if } x < a; \\ \frac{x-a}{b-a}, & \text{if } a \le x \le b; \\ \frac{c-x}{c-b}, & \text{if } b \le x \le c; \\ 0, & \text{if } x > c; \end{cases} \tag{9}$$

where $a \le b \le c$ represent the lower bound, modal value, and upper bound, respectively. Three input linguistic variables are defined as $x_{\text{age}} \in [18,90]$ years for a lifetime, $x_{\text{GFR}} \in [15,120]$ mL/min/1.73m² for **Glomerular Filtration Rate (GFR)** and, $x_{\text{AST}} \in [10,150]$ U/L for **Aspartate Aminotransferase (AST).** Each variable uses three fuzzy sets with triangular membership functions in Table 1.

Table 1

Partitioning of variables into three fuzzy sets using triangular membership functions

Lifetime	GFR	AST
$\mu_{\text{young}}(x) = \text{tri}(x; 18,18,40)$	$\mu_{\text{poor}}(x) = \text{tri}(x; 15,15,45)$	$\mu_{\text{normal}}(x) = \text{tri}(x; 10,10,40)$
$\mu_{\text{middle}}(x) = \text{tri}(x; 30,55,70)$	$\mu_{\text{moderate}}(x) = \text{tri}(x; 35,60,85)$	$\mu_{\text{elevated}}(x) = \text{tri}(x; 30,60,90)$
$\mu_{\text{elderly}}(x) = \text{tri}(x; 60,90,90)$	$\mu_{\text{normal}}(x) = \text{tri}(x; 75,120,120)$	$\mu_{\text{high}}(x) = \text{tri}(x; 80,150,150)$

The output variable, safety score $s \in [0,1]$, is partitioned into three fuzzy sets, $\mu_{\text{unsafe}}(s) = \text{tri}(s; 0,0,0.4)$, $\mu_{\text{marginal}}(s) = \text{tri}(s; 0.3,0.5,0.7)$ and, $\mu_{\text{safe}}(s) = \text{tri}(s; 0.6,1.0,1.0)$. The knowledge base consists of eight IF-THEN rules encoding clinical safety constraints:

1. **R_1**: IF age is young AND GFR is normal AND AST is normal THEN safety is safe
2. **R_2**: IF age is elderly AND GFR is poor THEN safety is unsafe
3. **R_3**: IF GFR is poor AND AST is high THEN safety is unsafe
4. **R_4**: IF age is middle AND GFR is moderate AND AST is elevated THEN safety is marginal
5. **R_5**: IF age is young AND GFR is poor THEN safety is marginal
6. **R_6**: IF age is elderly AND GFR is normal AND AST is normal THEN safety is safe
7. **R_7**: IF age is middle AND GFR is normal THEN safety is safe
8. **R_8**: IF age is elderly AND AST is high THEN safety is unsafe

For the input tuple $(x_{\text{age}}, x_{\text{GFR}}, x_{\text{AST}})$, the firing strength α_k of rule R_k is computed using the minimum t-norm:

$$\alpha_k = \min\{\mu_{A_{k1}}(x_{\text{age}}), \mu_{A_{k2}}(x_{\text{GFR}}), \mu_{A_{k3}}(x_{\text{AST}})\} \tag{10}$$

where A_{ki} denotes the antecedent fuzzy set for the variable i in rule k. The aggregated output membership function is obtained via maximum aggregation:

$$\mu_{\text{agg}}(s) = \max_{k=1\ldots8}\{\min(\alpha_k, \mu_{C_k}(s))\} \tag{11}$$

where C_k is the consequent fuzzy set of the rule k. Centroid defuzzification produces the crisp base safety score:

$$s_{\text{base}} = \frac{\int_0^1 s\cdot\mu_{\text{agg}}(s)\,ds}{\int_0^1 \mu_{\text{agg}}(s)\,ds} \approx \frac{\sum_{j=1}^{N} s_j\cdot\mu_{\text{agg}}(s_j)}{\sum_{j=1}^{N} \mu_{\text{agg}}(s_j)} \tag{12}$$

where $N = 1000$ discretization points are used for numerical integration. The base safety score is adjusted for drug-specific toxicity profiles:

$$S_{\text{fuzzy}}(p,d) = \max\{0, s_{\text{base}} - \Delta(p,d)\} \tag{13}$$

where the penalty term is:

$$\Delta(p,d) = \beta_1 T_{\text{nephro}}(d)\cdot\mathbb{I}(x_{\text{GFR}} < 60) + \beta_2 T_{\text{hepato}}(d)\cdot\mathbb{I}(x_{\text{AST}} > 60) + \beta_3 T_{\text{terato}}(d)\cdot\mathbb{I}_{\text{pregnant}} \tag{14}$$

with penalty coefficients, $\beta_1 = 0.5$, $\beta_2 = 0.6$, $\beta_3 = 1.0$ and $T_{\cdot}(d) \in \{0,1\}$ indicating drug toxicity characteristics. For each safety-eligible patient–drug pair (p,d), a 10-dimensional feature vector is constructed:

$$\mathbf{z} = [\text{age,BMI,GFR,AST,HbA1c,severity}, E_{\text{base}}(d), R_{\text{base}}(d), C(d), \text{class}(d)]^T \tag{15}$$

where patient features include age, body mass index (BMI), glomerular filtration rate (GFR), aspartate aminotransferase (AST), glycated hemoglobin (HbA1c), and disease severity grade. Additionally, drug features comprises on baseline efficacy $E_{\text{base}}(d)$, baseline adverse event risk $R_{\text{base}}(d)$, and cost $C(d)$. A Random Forest classifier with $T = 100$ decision trees and maximum depth $h_{\max} = 20$ predicts the probability of high efficacy:

$$\hat{P}_{\text{eff}}(\mathbf{z}) = \frac{1}{T}\sum_{t=1}^{T} h_t(\mathbf{z}) \tag{16}$$

where $h_t: \mathbb{R}^{10} \to \{0,1\}$ is the t-th decision tree trained on a bootstrap sample via Gini impurity minimization:

$$\text{Gini}(S) = 1 - \sum_{c=0,1}\left(\frac{|S_c|}{|S|}\right)^2 \tag{17}$$

for a node sample set S with class subsets S_0 and S_1. An XGBoost classifier with $K = 100$ boosting iterations, maximum depth 6, and learning rate $\eta = 0.1$ models adverse event probability:

$$\hat{P}_{\text{se}}(\mathbf{z}) = \sigma\left(\sum_{k=1}^{K} f_k(\mathbf{z})\right) \tag{18}$$

where $\sigma(\cdot)$ is the sigmoid function, and f_k minimizes the regularized objective:

$$\mathcal{L}^{(k)} = \sum_{i=1}^{n_{\text{train}}} \ell(y_i, \hat{y}_i^{(k-1)} + f_k(\mathbf{z}_i)) + \gamma T_{\text{leaves}} + \frac{\lambda}{2}\sum_{j=1}^{T_{\text{leaves}}} w_j^2 \tag{19}$$

with logistic loss:

$$\ell(y,\hat{y}) = -y\log(\sigma(\hat{y})) - (1-y)\log(1-\sigma(\hat{y})) \tag{20}$$

complexity penalty γ, and L^2 regularization parameter λ. For two candidate solutions $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$, solution $\mathbf{x}^{(1)}$ Pareto-dominates $\mathbf{x}^{(2)}$ (denoted $\mathbf{x}^{(1)} \prec \mathbf{x}^{(2)}$) if and only if:

$$f_j(\mathbf{x}^{(1)}) \le f_j(\mathbf{x}^{(2)}) \text{ for all } j \in \{1,2,3,4\} \tag{21}$$

$$\exists k \in \{1,2,3,4\}: f_k\left(\mathbf{x}^{(1)}\right) < f_k\left(\mathbf{x}^{(2)}\right) \tag{22}$$

Let $\mathcal{F}_r$ denote the set (front) of solutions that share the same non-domination rank r, where rank is computed with respect to all objective values $\{f_1(\mathbf{x}), f_2(\mathbf{x}), f_3(\mathbf{x}), f_4(\mathbf{x})\}$ in the current population. The first front $\mathcal{F}_1$ thus contains all non-dominated solutions, i.e., all $\mathbf{x}$ for which there exists no other solution $\mathbf{y}$ such that $\mathbf{y} \prec \mathbf{x}$. Non dominated sorting iteratively identifies $\mathcal{F}_1, \mathcal{F}_2 \cdots$ by counting, for each solution, how many other solutions dominate it and grouping solutions with the same rank into the corresponding front $\mathcal{F}_r$. For solution i in front $\mathcal{F}$, the crowding distance CD_i quantifies how isolated this solution is from its neighbors in objective space:

$$CD_i = \sum_{j=1}^{4} \frac{f_j(i+1) - f_j(i-1)}{f_j^{\max} - f_j^{\min}} \tag{23}$$

where, within front $\mathcal{F}$, solutions are sorted by objective j, $f_j(i+1)$ and $f_j(i-1)$ denote the objective values of the immediate neighbors of solution i along dimension j, and boundary solutions in $\mathcal{F}$ receive $CD_i = \infty$ to ensure their preservation. For this, NSGA-II operates with a population size $N_{\text{pop}} = 50$ and evolves over $G = 30$ generations [8]. Each individual is represented as a binary vector of length equal to the number of safety-eligible drugs for the patient. Each individual is initialized by randomly selecting 1–5 drugs from the safety-eligible set. Binary tournament selection based on non-domination rank and crowding distance [9, pp. 45–52]. Bit-flip mutation with probability $p_m = 0.3$. If the mutation violates cardinality constraints ($\|\mathbf{x}\|_1 \notin [1,5]$), the offspring is rejected. Combined parent and offspring populations are sorted by non-domination rank and crowding distance; the top N_{pop} solutions survive to the next generation.

Physician preferences over criteria $\{C_1, C_2, C_3, C_4\} = \{\text{Efficacy, Safety, Interactions, Cost}\}$ are modeled via pairwise preference ratios represented as triangular fuzzy numbers:

$$\tilde{\theta}_{ij} = (\theta_{ij}^{l}, \theta_{ij}^{m}, \theta_{ij}^{u}) \tag{24}$$

with confidence-dependent bounds:

$$\theta_{ij}^{l} = \theta_{ij}^{m}(1 - 0.2\sigma), \theta_{ij}^{u} = \theta_{ij}^{m}(1 + 0.2\sigma) \tag{25}$$

where $\sigma \in [0,1]$ quantifies preference certainty. The preference ordering Efficacy $\succ$ Safety $\succ$ Interactions $\succ$ Cost is encoded with modal values $\theta_{12}^{m} = 1.3$, $\theta_{23}^{m} = 1.5$, $\theta_{34}^{m} = 2.0$, and confidence parameter $\sigma = 0.7$. For DIBR II weight computation, cumulative preference ratios are computed via fuzzy multiplication [10, pp. 72–96]:

$$\tilde{\theta}_{13} = \tilde{\theta}_{12} \otimes \tilde{\theta}_{23}, \tilde{\theta}_{14} = \tilde{\theta}_{13} \otimes \tilde{\theta}_{34} \tag{26}$$

Fuzzy weights are derived from the preference structure:

$$\tilde{w}_1 = \left[\tilde{1} \oplus \tilde{\theta}_{12}^{-1} \oplus \tilde{\theta}_{13}^{-1} \oplus \tilde{\theta}_{14}^{-1}\right]^{-1} \tag{27}$$

$$\tilde{w}_j = \tilde{w}_{j-1} \oslash \tilde{\theta}_{(j-1)j}, j = 2,3,4 \tag{28}$$

where ⊘ denotes fuzzy division. Graded mean defuzzification converts fuzzy weights to crisp values:

$$\mathrm{GM}(\tilde{A}) = \frac{a_l + 4a_m + a_u}{6} \tag{29}$$

Normalized crisp weights are obtained via [11]:

$$w_j = \frac{\mathrm{GM}(\tilde{w}_j)}{\sum_{k=1}^{4} \mathrm{GM}(\tilde{w}_k)} \tag{30}$$

For each Pareto solution i, objectives are normalized via min–max scaling [12]:

$$\bar{d}_{ij} = \frac{d_{ij} - d_j^{\min}}{d_j^{\max} - d_j^{\min}} \text{ if } C_j \text{ is benefit criterion} \tag{31}$$

$$\bar{d}_{ij} = \frac{d_j^{\max} - d_{ij}}{d_j^{\max} - d_j^{\min}} \text{ if } C_j \text{ is cost criterion} \tag{32}$$

where d_{ij} is the value of the criterion j for solution i. The final aggregated score is:

$$\mathrm{Score}_i = \sum_{j=1}^{4} w_j \bar{d}_{ij} \tag{33}$$

As a result of DIBR II weight ranking, solutions are ranked in descending order of Score_i, with the top-ranked prescription representing the optimal balance of all criteria under the specified physician preferences [13].

Experimental results. The proposed hybrid framework was evaluated on a synthetic clinical dataset constructed to emulate realistic pharmacotherapy scenarios under full experimental control. The dataset comprises 1,000 patients with clinically plausible distributions for age, BMI, GFR, AST, and HbA1c, and 50 drugs defined by baseline efficacy, adverse event risk, cost, binary toxicity indicators, and a symmetric drug–drug interaction matrix representing low to moderate interaction severity. Applying the Mamdani fuzzy safety filter with threshold $\tau_{safety} = 0.5$ to all 50,000 patient–drug pairs yielded 44,692 safety-eligible combinations, which were split 80/20 into 35,754 training and 8,938 test instances for ensemble modelling, with predictive performance summarized in Table 2.

A representative test patient with age 62.5 years, GFR 63.1 mL/min/1.73m², and AST 10.0 U/L was selected for detailed analysis. Fuzzy safety assessment indicated that all 50 drugs satisfied the safety threshold for this patient, enabling exploration of the full drug space. NSGA-II was executed with population size $N_{pop} = 50$ for $G = 30$ generations. The evolutionary process exhibited consistent improvement in Pareto front quality, with the hypervolume indicator increasing from 0.412 (generation 1) to 0.584 (generation 30), indicating effective convergence and diversity maintenance. The final population contained 20 non-dominated solutions in the raw Pareto front. Application of duplicate removal based on binary prescription vector comparison reduced this to 9 unique Pareto-optimal prescriptions, representing a 55% duplication rate typical of evolutionary algorithms on discrete combinatorial problems.

Table 2

Predictive Performance of Ensemble Machine Learning Models

Model	Task	AUC-ROC	Accuracy	Precision	Recall
Random Forest	Efficacy	0.871	0.794	0.782	0.801
XGBoost	Adverse Events	0.907	0.829	0.815	0.838

In the figure given below, illustrates key trade-offs on the Pareto front between efficacy, safety, and cost. High-efficacy prescriptions (efficacy > 0.80) show increased adverse event risk (≈0.35–0.42), whereas ultra-safe options (adverse events < 0.15) reach only moderate efficacy (≈0.45–0.50). High-efficacy regimens are also more expensive (≈550–610 AZN), while lower-cost alternatives (<250 AZN) do not systematically worsen safety, and all prescriptions keep drug–drug interaction severity below 0.15, reflecting the moderate interaction matrix.

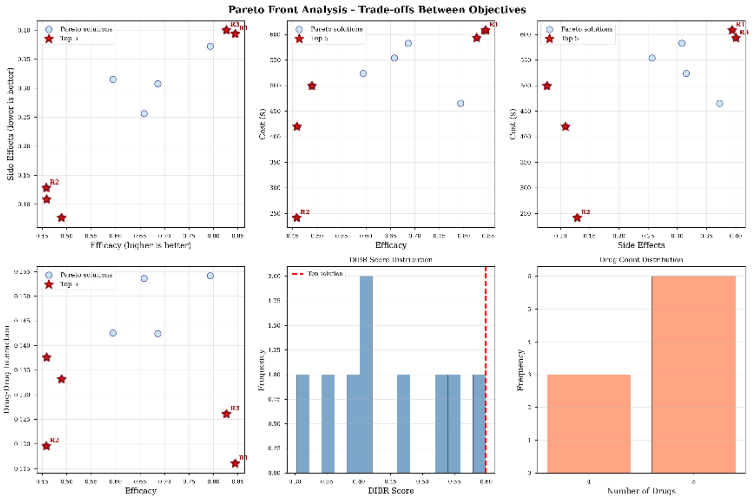

Pareto Front Analysis – Trade-offs Between Objectives

Drug count analysis shows that 4- and 5-drug regimens dominate the Pareto set (7 of 9 solutions), indicating that multi-drug combinations provide better coverage of the objective space. Physician preferences were modeled via fuzzy triangular numbers with the ordering Efficacy ≻ , Safety ≻ , Interactions ≻ Cost, and confidence parameter $\sigma = 0.7$, yielding crisp weights $w_1 = 0.387$, $w_2 = 0.302$, $w_3 = 0.205$, $w_4 = 0.105$, which emphasize efficacy and safety while assigning lower importance to cost. The top five Pareto-optimal prescriptions, together with their objective values and qualitative clinical interpretations, are summarized in Table 3.

Table 3

Top Five Prescription Recommendations with Clinical Profiles

Rank	DIBR Score	Efficacy	Adverse Events	DDI	Cost (AZN)	Drugs	Clinical Profile
1	0.599	84.4%	39.4%	11.6%	608.6	5	High-efficacy, aggressive
2	0.545	45.8%	12.9%	12.0%	241.2	4	Budget-constrained
3	0.525	82.6%	40.0%	12.6%	593.3	5	High-efficacy alternative
4	0.478	48.9%	7.7%	13.3%	499.2	5	Ultra-safe, elderly
5	0.417	45.9%	10.9%	13.8%	419.8	5	Balanced, general use

Quantitative analysis of the Pareto front confirms fundamental trade-offs in multi-criteria drug selection. The efficacy–safety relationship shows a strong negative correlation (Spearman $\rho = -0.78, p < 0.01$), where increasing mean efficacy from 45% to 85% leads to a 3.1-fold rise in adverse event probability (from 12.9% to 40.0%). The efficacy–cost correlation is moderate ($\rho = 0.61, p = 0.03$), with high-efficacy solutions (>80%) costing on average 2.4× more than moderate-efficacy regimens (<50%), whereas safety and cost are only weakly correlated ($\rho = 0.23$, $p = 0.18$), indicating that budget limits do not inherently require higher adverse event risk. Drug–drug interaction severity remains relatively uniform (coefficient of variation 18.3%), suggesting that the interaction matrix design and cardinality constraints effectively restrict harmful combinations across Pareto-optimal prescriptions.

Conclusion and future works. This study presented a hybrid decision-support framework for multi-criteria drug selection that jointly optimizes efficacy, safety, drug–drug interactions, and cost at the individual patient level. The architecture combines a Mamdani fuzzy safety filter, ensemble machine learning for outcome prediction, NSGA-II for multi-objective optimization, and Fuzzy DIBR II for preference-based ranking, yielding interpretable Pareto-optimal prescriptions instead of a single opaque recommendation. Experimental evaluation on a synthesized data but clinically realistic dataset demonstrated that the framework can generate diverse treatment options with clearly quantified trade-offs aligned with physician priorities. Future work will include validation on real-world electronic health record and prescription data to better capture rare adverse events and complex interaction patterns, as well as the development of interactive tools for dynamic physician preference elicitation. Further research directions involve integrating uncertainty-aware or robust optimization techniques and extending the framework to high-risk, polypharmacy-intensive domains such as oncology.

REFERENCES

1. World Health Organization. Medication Safety in Polypharmacy, Technical Report WHO/UHC/SDS/2019.11, 2019.
2. Bates D.W., Cullen D.J., Laird N., Petersen L.A., Small S.D., Servi D., Laffel G., Sweitzer B.J., Shea B.F., Hallisey R. Incidence of adverse drug events and potential adverse drug events, JAMA, 1995, vol.274(1), pp.29–34.
3. Topol E.J. High-performance medicine: the convergence of human and artificial intelligence, Nature Medicine, 2019, vol.25(1), pp.44–56.
4. Shortliffe E.H., Sepúlveda M.J. Clinical decision support in the era of artificial intelligence, JAMA, 2018, vol.320(21), pp.2199–2200.
5. Zadeh L.A. Fuzzy sets, Information and Control, 1965, vol.8(3), pp.338–353.
6. Ross T.J. Fuzzy Logic with Engineering Applications, 3rd ed., John Wiley & Sons, 2010, 607 p.
7. Innocent P.R., John R.I. Computer aided fuzzy medical diagnosis, Information Sciences, 2004, vol.162(2), pp.81–104.
8. Deb K., Pratap A., Agarwal S., Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II, IEEE Transactions on Evolutionary Computation, 2002, vol.6(2), pp.182–197.
9. Coello C.A.C., Lamont G.B., Van Veldhuizen D.A. Evolutionary Algorithms for Solving Multi-Objective Problems, 2nd ed., Springer, 2007, 800 p.
10. Hwang C.L., Yoon K. Multiple Attribute Decision Making: Methods and Applications, Springer-Verlag, 1981, 269 p.
11. Breiman L. Random forests, Machine Learning, 2001, vol.45(1), pp.5–32.
12. Chen T., Guestrin C. XGBoost: A scalable tree boosting system, Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp.785–794.
13. Mardani A., Jusoh A., Zavadskas E.K. Fuzzy multiple criteria decision-making techniques and applications – Two decades review from 1994 to 2014, Expert Systems with Applications, 2015, vol.42(8), pp.4126–4148.

Accepted: 18.12.2025