



+ Precision Medicine

GENETIC DIAGNOSIS OF PROSTATE CANCER

Benchmark vs. Logistic Regression, Random Forest, Boosted Tree & Neural Network

UC#08 – 2025/06 (v4.0)

xtractis.ai

PROBLEM DEFINITION

GOAL Design an AI-based decision system that accurately and instantly makes a rational medical diagnosis of prostate cancer from genetic sequencing of prostate tissue.

PROS & BENEFITS

- ▶ Identify the genes involved in cancer and enhance medical knowledge by helping urologists and oncologists understand the causal relationships between specific genes, their combination, and the presence of cancer.
- ▶ Help the medical profession to make earlier and more personalized decisions through rapid, systematic, and explainable diagnoses.
- ▶ Contribute to improving patient care (pain, survival, duration of treatment) and extend access to high-level diagnoses even in medical deserts.

REFERENCE DATA

Source:
D. Singh & al., Department of Adult Oncology, Brigham and Women's Hospital, Harvard Medical School.

Dataset:
www-genome.wi.mit.edu /mpr/prostate (2014)

Variable to Predict: The model diagnoses the sampled prostate tissue as **NORMAL | TUMOR**

Potential Predictors: 12,600 variables are the level of expression of genes characterizing each patient, normalized to the median.

Observations: 136 genetic sequencing of prostate tissue from patients with or without cancer. 102 cases compose a Learning Dataset for model induction using Training and Validation Datasets. 34 samples from a different experiment compose an External Test Dataset to check the top-model's performance on real unknown data and for benchmarking.

Learning Dataset: 102 patients 75.00% 80% for Training, 20% for Validation		External Test Dataset: 34 patients 25.00%	
NORMAL	TUMOR	NORMAL	TUMOR
50 49%	52 51%	9 26.47%	25 73.53%

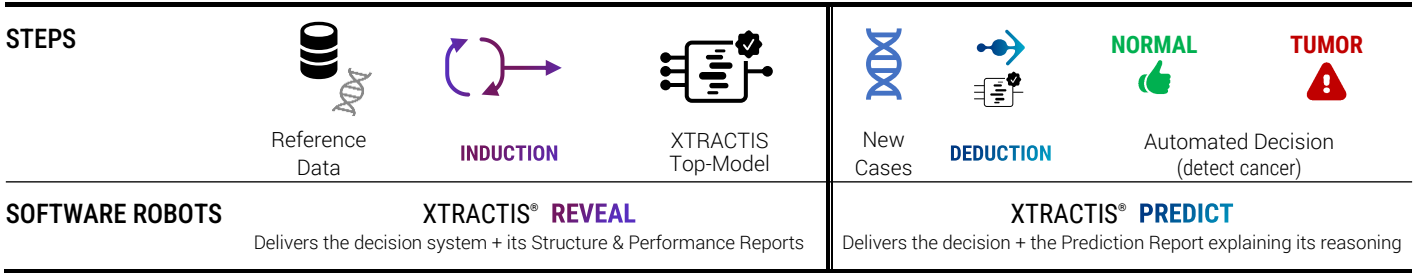
MODEL TYPE

Regression Multinomial Classification **Binomial Classification** Scoring

XTRACTIS-INDUCED DECISION SYSTEM

- Intelligible Model, Explainable Decisions**
 - ▶ The top-model is a decision system composed of 4 gradual rules without chaining aggregated into 2 disjunctive rules.
 - ▶ Each rule uses from 2 to 4 predictors among the 7 variables that XTRACTIS automatically identified as significant (out of the 12,600 level of genes expression describing each patient).
 - ▶ Only a few rules are triggered at a time to compute the decision.
- High Predictive Capacity** It has an excellent Real Performance (on unknown data).
- Ready to Deploy** It computes real-time predictions up to 70,000 decisions/second, offline or online (API).

XTRACTIS PROCESS



TOP-MODEL INDUCTION

INDUCTION PARAMETERS

Powered by:



1. We launch 100 inductive reasoning strategies; each strategy is applied to 40 different 5-fold-partitions of the Learning Dataset to get a reliable assessment of the descriptive and predictive performances, respectively from Training and Validation Datasets.
2. Each strategy thus generates 200 unitary models called **Individual Virtual Expert (IVE)**, whose decisions are aggregated with 3 possible operators into a **College of Virtual Experts (CVE)**.
3. Among the 300 induced CVEs, the top-CVE with the best predictive performance remains complex: 658 rules share 471 predictors.

Given the small number of reference cases in the reference dataset, the XTRACTIS **CVE→IVE** Reverse-Engineering process is necessary to get a more intelligible model:

4. We build a synthetic dataset composed of 20,400 new cases simulated by deduction from the top-CVE, around the 102 original learning cases but distinct from them.
5. We apply 2,000 induction strategies to the same single 70% Training | 15% Validation | 15% Test partition of this new dataset: XTRACTIS induces 2,000 IVEs.
6. The top-IVE selected is the one that is the most intelligible while being as efficient as the top-CVE.


Total number of induced unitary models	Criterion for the induction optimization	Validation criterion for the top-model selection	Duration of the process (Induction Power FP64)
22,000 IVEs	F₁-Score	F₁-Score	17 days (1.13 Tflops)

TOP-MODEL STRUCTURE

The top-IVE model has an excellent intelligibility as it combines 7 predictors into 4 rules with 3 predictors per rule on average. Its Structure Report reveals all the internal logic of the decision system and ensures that the model is understandable by the human expert. It is a transparent model that can be audited and certified before deployment to end-users.

PREDICTORS

- 7 genes identified out of 12,600
- Ranked by individual contribution (2 strong, 3 medium & 2 weak signals):
#1 [gene 36883_at](#) / #2 [gene 37639_at](#) / ...
- Labeled by fuzzy classes
Example: **fuzzy interval** "inferior to about 5"



RULES

- 4 connective fuzzy rules without chaining (aggregated into 2 disjunctive fuzzy rules)
- 2 to 4 predictors per rule (on average, 3 predictors per rule)
- Example: fuzzy rule **R4** uses 4 predictors and concludes **TUMOR**. 3 other rules complete this model.


```

IF gene 39939_at IS inferior to ~5
AND gene 35178_at IS inferior to ~-2
AND gene 36883_at IS inferior to ~87
AND gene 40282_s_at IS inferior to ~77
THEN Diagnosis IS TUMOR
    
```

Literally, the sampled prostate tissue gets a tumor diagnosis if the level of expression of gene #39939 is under around 5, and that of gene #35178 is under around minus 2, and that of gene #36883 is under around 87, and that of gene #40282_s is under around 77.

TOP-MODEL PERFORMANCE

The top-IVE performances, measured in Training/Validation/Test on synthetic data, then in External Test on reference data, guarantee the model's predictive and real performances.



EXPLAINED PREDICTIONS FOR 3 UNKNOWN CASES

CASE

(from the External Dataset, i.e., not included in the Learning Dataset)

PATIENT #1

actual value = TUMOR

gene 39939_at	5
gene 33792_at*	1.7
gene 35178_at	2
gene 36883_at	39
gene 37639_at	162
gene 37367_at	114
gene 40282_s_at	26




DEDUCTIVE INFERENCE OF RULES

For this patient, 3 rules are triggered:

R4 is fired at 0.940 to conclude TUMOR, R1 at 0.117, and R2 at 0.022 to conclude NORMAL.

R3 is not activated.



AUTOMATED DECISION

NUMBER OF TRIGGERED RULES

3 / 4

FUZZY PREDICTION

{ TUMOR | 0.940,
NORMAL | 0.117 }

FINAL PREDICTION

{ TUMOR }

The system delivers a correct diagnosis of cancer compared to that given by the genetic oncologist:

TUMOR 

PATIENT #30

actual value = NORMAL

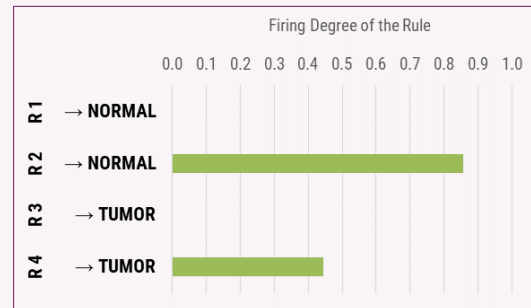
gene 39939_at	24
gene 33792_at	296.9
gene 35178_at	2
gene 36883_at	21
gene 37639_at	33
gene 37367_at	92
gene 40282_s_at	60



For this patient, 2 rules are triggered:

R2 is fired at 0.857 to conclude NORMAL, and R4 at 0.445 to conclude TUMOR.

R1 and R3 are not activated.



NUMBER OF TRIGGERED RULES

2 / 4

FUZZY PREDICTION

{ NORMAL | 0.857,
TUMOR | 0.445 }

FINAL PREDICTION

{ NORMAL }

The system delivers a correct diagnosis of cancer compared to that given by the genetic oncologist:

NORMAL 

PATIENT #5

actual value = TUMOR

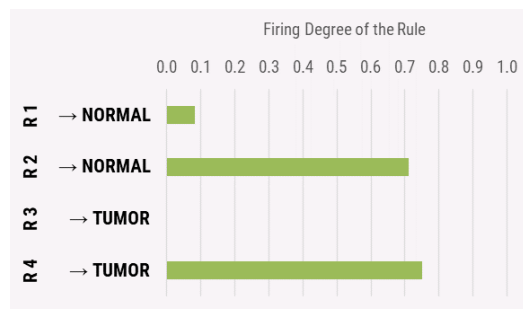
gene 39939_at	14
gene 33792_at	20.6
gene 35178_at	4
gene 36883_at	20
gene 37639_at	55
gene 37367_at	75
gene 40282_s_at	46



For this patient, 3 rules are triggered:

R4 is fired at 0.751 to conclude TUMOR, R2 at 0.711, and R1 at 0.082 to conclude NORMAL.

R3 is not activated.



NUMBER OF TRIGGERED RULES

3 / 4

FUZZY PREDICTION

{ TUMOR | 0.751,
NORMAL | 0.711 }

FINAL PREDICTION


{ TUMOR }

The system delivers a correct diagnosis of cancer compared to that given by the genetic oncologist, despite uncertainty/hesitation:

TUMOR 

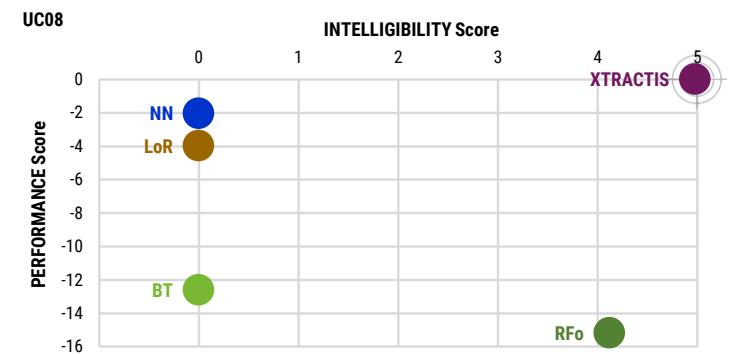
*Predictor value outside the variation range of the model but inside the allowed extrapolation range. XTRACTIS will refuse to give a result for an extrapolation far from the allowed extrapolation range. It is one situation of the "Refusal" prediction.

TOP-MODELS BENCHMARK: DECISION STRUCTURE & INTELLIGIBILITY × PERFORMANCE SCORES

	XTRACTIS 	LOGISTIC REGRESSION	RANDOM FOREST	BOOSTED TREE	NEURAL NETWORK	
MODELING PARAMETERS	MODELS RELEASE	2021/06	2022/10	2021/08	2021/04	2022/03
	ALGORITHM VERSION	XTRACTIS REVEAL 11.2.38531	Python 3.9.12 Scikit-Learn 1.0.2	Python 3.6 LightGBM 2.2.2	Python 3.6 LightGBM 2.2.2	Python 3.6 TensorFlow 2.6.2 Keras 2.6.0
	CROSS-VALIDATION TECHNIQUE	40x5 folds for each CVE model. Then 1-Split Validation for each IVE model: 70% Training 15% Validation 15% Test	40x5 folds for each CVE model	40x5 folds for each CVE model	40x5 folds for each CVE model	40x5 folds for each CVE model
	NUMBER OF EXPLORED STRATEGIES⁽¹⁾	100 induction strategies for the CVE on Training / Validation data. 2,000 induction strategies for the IVE on synthetic data	300 data analysis strategies on Training / Validation data	300 ML strategies on Training / Validation data	300 ML strategies on Training / Validation data	300 ML strategies on Training / Validation data
	TOP-MODEL SELECTION⁽²⁾	Top-CVE among 300 CVEs. Then Top-IVE among 2,000 IVEs	Top-CVE selected among 300 CVEs, then single model obtained by applying best CVE strategy on 100% of the Learning Dataset			

TOP-MODEL STRUCTURE	NUMBER OF PREDICTORS (out of 12,600 Potential Predictors)	7	120	19	24	12,600
	AVERAGE NUMBER OF PREDICTORS PER RULE OR EQUATION	3.0 per rule	120 per equation	1.8 per rule	1.9 per rule	11,700.9 per equation
	STRUCTURE OF THE DECISION SYSTEM	4 fuzzy rules without chaining (aggregated into 2 disjunctive fuzzy rules) Only some rules are triggered at a time to compute a prediction	1 linear equation	15 trees without chaining 50 binary rules	1 chain of 14 trees 48 binary rules Tree #N corrects the error of the N-1 previous trees	1 hidden layer 13 hidden nodes 14 equations 13 unintelligible synthetic variables

	Random ⁽³⁾	XTRACTIS	LoR	RFo	BT	NN
INTELLIGIBILITY Score⁽⁴⁾		4.98	0.00	4.12	0.00	0.00
CVE Real Performance (F ₁ -Score) in External Test		100.00	97.96	87.50	88.00	97.96
Gap to CVE Leader in External Test		0.00	-2.04	-12.50	-12.00	-2.04
IVE Real Performance (F ₁ -Score) in External Test	92.00	100.00	94.11	82.14	86.79	97.96
Gap to IVE Leader in External Test		0.00	-5.89	-17.86	-13.21	-2.04
Average Real Performance in External Test	92.00	100.00	96.04	84.82	87.40	97.96
PERFORMANCE Score⁽⁴⁾		0.00	-3.97	-15.18	-12.61	-2.04



(1) For all algos: on the same Learning Dataset. All Models are optimized according to their Validation F₁-Score.

(2) All top-models are selected according to their Validation F₁-Score while checking that it remains close to their Training F₁-Score.

(3) Baseline performances that models must exceed to perform better than chance (P-value = 0.001; 100,000 models generated by random permutation of the output values). The value of each performance criterion is generally achieved by a different random model.

(4) See Appendices for explanations and detailed results. Performance Scores are calculated on all available unknown data. XTRACTIS's perfect results on External Test could be explained by a low number of reference points compared to the very large number of potential predictors.

More Use Cases:
xtractis.ai/use-cases/

APPENDIX 1 – Calculation of the Intelligibility × Performance Scores

AI Technique #i	T _i	i ∈ [1 ; n] n = number of AI Techniques benchmarked in terms of data-driven modeling = 5
Benchmark #k	B _k	k ∈ [1 ; p] p = number of Benchmarks for the Use Case ∈ {1, 2, 3}

Remarks:

- In case of a small number of reference data, a CVE model (College of Virtual Experts) is generated by each explored strategy of T_i, generally via an N×K-fold cross validation. In this case, a Benchmark is led with the top-CVE on the External Test Dataset (ETD, composed of unknown reference cases). Then, a top-IVE model (Individual Virtual Expert) is generated from the top-CVE, through the XTRACTIS® reverse-engineering process, or for the other T_i, by applying the top-strategy, which has generated the top-CVE, on the Training and Validation Datasets. And a second Benchmark is led with this top-IVE on the same ETD.
- In case of a huge number of reference data, an IVE is generated by each explored strategy of T_i, via a 1-split validation. In this case, Benchmarks are led with the top-IVE on the Test Dataset (TD, composed of unknown reference cases) and on the available ETDs.
- Each Benchmark uses the latest versions of the following algorithms available at the date of the benchmark. XTRACTIS®: REVEAL; Logistic Regression: Python, Scikit-Learn; Random Forest & Boosted Tree: Python, LightGBM; Neural Network: Python, TensorFlow, Keras.
- Each B_k uses exactly the same TD and ETD for each T_i model.
- No Regression models can be obtained by Logistic Regression. So, this Data Analysis technique is benchmarked only for Classification or Scoring problems.
- The Holy Grail for critical AI-based decision systems is to obtain a model with the highest Performance and the highest Intelligibility scores (top-right corner of the graph).

PERFORMANCE Score

For each B_k, we calculate the values of the Performance Criterion (PC) on the same ETD for all the T_i top-CVEs; and on the same TD and ETDs for all the T_i top-IVEs. The PC is: RMSE in percentage for a Regression; F₁-Score for a Binomial Classification; Average F₁-Score or Average F₂-Score for a Multinomial Classification; Gini index for a Scoring. Then, we compare the value of the PC of each T_i top-CVE (resp. top-IVE) to the best value of this PC reached by the best T_i top-CVE (resp. top-IVE) on ETD (resp. on TD and ETDs).

For Regression, we calculate for each T_i top-model (CVE and IVE): PS(T_i, B_k) = Best_PC(B_k) - PC(T_i, B_k).

For Classification and Scoring, we calculate for each T_i top-model: PS(T_i, B_k) = PC(T_i, B_k) - Best_PC(B_k).

$$\text{Performance Score of } T_i$$

$$\text{PS}(T_i) = \text{Mean} (\text{PS}(T_i, B_k))_{k \in [1 ; p]}$$

Remark:

- Each PS varies theoretically from -100 (Lowest Score) to 0 (Highest Score), but practically between -50 and 0.

INTELLIGIBILITY Score

We consider the T_i top-IVE. Its Intelligibility Score IS(T_i) is valued from 0.00 to 5.00 regarding the structure of the model: number of predictors, classes, rules, equations, trees, synthetic variables, modalities to predict for classifications (or numeric variables to predict for regressions or scoring). The more compact the model, the higher its IS.

The IS of each T_i is obtained by accumulating the following five penalty values to the ideal IS value of 5.00 (each penalty has a null or a negative value):

- Penalty 1 (logarithmic penalty regarding the number of predictors):

$$\text{Pen1}(T_i) = \min(0, 1 - \log_{10} \text{number of predictors})$$

Examples: Pen1 = 0.00 for up to 10 predictors
Pen1 = -3.00 for 10.000 predictors

- Penalty 2 (linear penalty regarding the average number of rules or equations per variable or modality to predict):

$$\text{Pen2}(T_i) = \min\left(0, \frac{1 - \text{average number of rules or equations per variable or modality to predict}}{40}\right)$$

Examples: Pen2 = 0.00 for 1 rule or equation per variable or modality to predict on average
Pen2 = -3.00 for 121 rules or equations per variable or modality to predict on average

- Penalty 3 (linear penalty regarding the average number of predictors per rule or equation):

$$\text{Pen3}(T_i) = \min\left(0, \frac{9 - 3 \times \text{average number of predictors per rule or equation}}{7}\right)$$

Examples: Pen3 = 0.00 for up to 3.0 predictors per rule or equation on average
Pen3 = -3.00 for 10.0 predictors per rule or equation on average

- Penalty 4 (linear penalty regarding the number of chained trees, here for BT only):

$$\text{Pen4}(T_i) = \min(0, 1 - \text{number of chained trees})$$

Examples: Pen4 = 0.00 for 1 tree
Pen4 = -3.00 for 4 chained trees

- Penalty 5 (maximum penalty due to unintelligibility of synthetic variables, here for NN only):

$$\text{Pen5}(T_i) = -5$$

Intelligibility Score of T_i

$$\text{IS}(T_i) = \max(0.00, 5.00 + (\text{Pen1} + \text{Pen2} + \text{Pen3} + \text{Pen4} + \text{Pen5}))$$

Remarks:

- For the difference between the Intelligibility and the Explainability of a model, please see the XTRACTIS® Brochure, page 7.
- The real complexity of the process/phenomenon under study is intrinsic, i.e., it could not be reduced or simplified, but only discovered; thus, the top-model will be complex if the process/phenomenon turns out to be complex [Zalila 2017]. Consequently, for some complex process/phenomenon, IS can be equal to 3.00 or less, even if T_i natively produces intelligible models (XTRACTIS, Random Forest).
- For similar structures, the Boosted Tree model is always less intelligible than the Random Forest one, as it is composed of chains of trees, instead of a college of trees (see Penalty 4).
- Neural Network model has always the lowest IS of 0.00, because it uses synthetic unintelligible variables (hidden nodes) in addition to all the potential predictors (see Penalty 5).

APPENDIX 2 – Use Case Results (all Performance criteria of all Top-Models)

Performance Criterion	Classification Error	Min. Sensitivity Specificity	Sensitivity	Specificity	PPV	NPV	F ₁ -Score	Refusal
-----------------------	----------------------	------------------------------	-------------	-------------	-----	-----	-----------------------	---------

RANDOM MODEL

Nb of Random Permutations (P-value) = 100,000 (0.001%)

Performance against chance (External Test)	11.76%	0.698					92.00%	
--	--------	-------	--	--	--	--	---------------	--

XTRACTIS TOP-MODEL

CVE - Descriptive Performance (Training)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	0 (0.00%)
CVE - Predictive Performance (Validation)	1.98%	97.96%	98.08%	97.96%	98.08%	97.96%	98.08%	1 (2.94%)
CVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	0 (0.00%)
IVE - Descriptive Performance (Training)	0.65%	99.26%	99.43%	99.26%	99.28%	99.42%	99.36%	0 (0.00%)
IVE - Predictive Performance (Validation)	0.49%	99.40%	99.61%	99.40%	99.42%	99.60%	99.52%	0 (0.00%)
IVE - Real Performance (Test)	0.36%	99.27%	100.00%	99.27%	99.30%	100.00%	99.86%	0 (0.00%)
IVE - Real Performance (102 original points)	1.96%	98.00%	98.08%	98.00%	98.08%	98.00%	98.08%	0 (0.00%)
IVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	0 (0.00%)

LOGISTIC REGRESSION TOP-MODEL

CVE - Descriptive Performance (Training)	1.96%	98.00%	98.08%	98.00%	98.08%	98.00%	98.08%	
CVE - Predictive Performance (Validation)	2.94%	96.15%	96.15%	98.00%	98.04%	96.08%	97.09%	
CVE - Real Performance (External Test)	2.94%	96.00%	96.00%	100.00%	100.00%	90.00%	97.96%	
IVE - Descriptive Performance (Training)	0.98%	98.00%	100.00%	98.00%	98.11%	100.00%	99.05%	
IVE - Real Performance (External Test)	8.82%	77.78%	96.00%	77.78%	92.31%	87.50%	94.11%	

RANDOM FOREST TOP-MODEL

CVE - Descriptive Performance (Training)	3.92%	94.23%	94.23%	94.23%	98.04%	96.08%	96.08%	
CVE - Predictive Performance (Validation)	1.96%	98.00%	98.08%	98.00%	98.08%	98.00%	98.08%	
CVE - Real Performance (External Test)	17.65%	77.78%	84.00%	77.78%	91.30%	63.64%	87.50%	
IVE - Descriptive Performance (Training)	0.98%	98.00%	100.00%	98.00%	98.11%	100.00%	99.05%	
IVE - Real Performance (External Test)	29.41%	11.11%	92.00%	11.11%	74.19%	33.33%	82.14%	

BOOSTED TREE TOP-MODEL

CVE - Descriptive Performance (Training)	2.94%	96.15%	96.15%	96.15%	98.04%	96.08%	97.08%	
CVE - Predictive Performance (Validation)	1.96%	98.00%	98.08%	98.00%	98.08%	98.00%	98.08%	
CVE - Real Performance (External Test)	17.65%	66.67%	88.00%	66.67%	88.00%	66.67%	88.00%	
IVE - Descriptive Performance (Training)	1.96%	96.00%	100.00%	96.00%	96.30%	100.00%	98.11%	
IVE - Real Performance (External Test)	20.58%	44.44%	92.00%	44.44%	82.14%	66.67%	86.79%	

NEURAL NETWORK TOP-MODEL

CVE - Descriptive Performance (Training)	0.98%	98.08%	98.08%	100.00%	100.00%	98.04%	99.03%	
CVE - Predictive Performance (Validation)	1.96%	98.00%	98.08%	98.00%	98.08%	98.00%	98.08%	
CVE - Real Performance (External Test)	2.94%	96.00%	96.00%	100.00%	100.00%	90.00%	97.96%	
IVE - Descriptive Performance (Training)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
IVE - Real Performance (External Test)	2.94%	96.00%	96.00%	100.00%	100.00%	90.00%	97.96%	

The entirety of this document is protected by copyright. All rights are reserved, particularly the rights of reproduction and distribution. Quotations from any part of the document must necessarily include the following reference:
Zalila, Z., Idagrai Labs & Xtractis (2018-2025). XTRACTIS® the Reasoning AI for Trusted Decisions. Use Case #08 | Precision Medicine: Genetic Diagnosis of Prostate Cancer – Benchmark vs. Logistic Regression, Random Forest, Boosted Tree & Neural Network. IDAGRAI LABS, June 2025, v4.0, Compiègne, France, 6p.