

+ Precision Medicine

SEROLOGICAL DIAGNOSIS OF CHRONIC KIDNEY DISEASE (CKD)

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

UC#14 – 2025/06 (v3.0)

xtractis.ai

PROBLEM DEFINITION

GOAL Design an AI-based decision system that accurately and instantly makes a rational medical diagnosis of chronic kidney disease, from the patient record and its blood measures.

- PROS & BENEFITS**
- ▶ Identify the parameters involved in the kidney disease and enhance medical knowledge by helping nephrologists understand the causal relationships between these parameters, their combination, and the disease.
 - ▶ Help the medical profession to make earlier and more personalized decisions through rapid, systematic, and explainable diagnoses.
 - ▶ Use a model with few predictors to limit medical data that can be expensive to collect.

REFERENCE DATA

Source:
Dr. P.Eswaran, Dept of Computer Science and Engineering, Alagappa University, Karaikudi, Tamilnadu, India

Dataset:
[<http://archive.ics.uci.edu/ml>] (2015)

- Variable to Predict:** The model diagnoses the patient record as **No CKD | CKD**
- Potential Predictors:** 24 variables characterize each patient [Age, appetite, blood measures (Pressure, Specific Gravity, Albumin Sugar, Red Blood Cells, Potassium, ...), Hypertension, Diabetes, Coronary Artery Disease, Pedal Edema, Anemia].
- Observations:** 400 records from patients with or without Chronic Kidney Disease. 340 cases compose a Learning Dataset for model induction using Training and Validation Datasets. 60 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: 340 patients 85.00% 80% for Training, 20% for Validation	
No CKD	CKD
128 37.7%	212 62.3%

External Test Dataset: 60 patients 15.00%	
No CKD	CKD
22 36.7%	38 63.3%

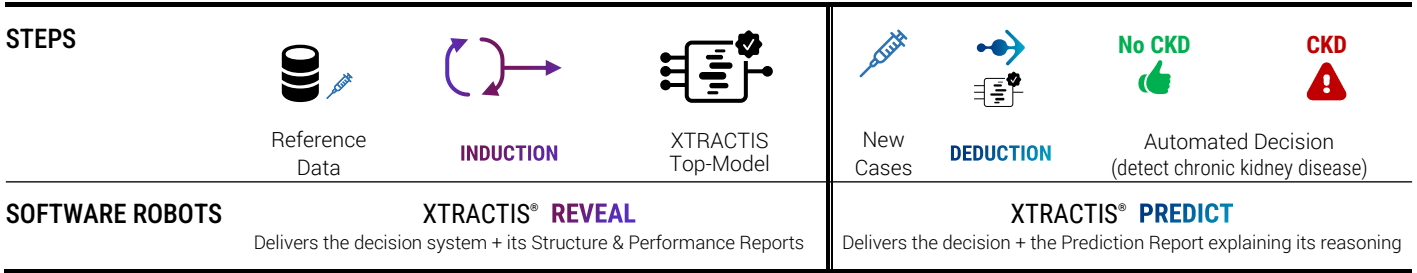
MODEL TYPE

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 6 predictors among the 8 variables that XTRACTIS automatically identified as significant (out of the 24 features describing each patient).
 - ▶ Only a few rules are triggered at a time to compute the decision.
- High Predictive Capacity** It has a perfect 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:



- We launch 1,000 inductive reasoning strategies; each strategy is applied to 20 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.
- Each strategy thus generates 100 unitary models called **Individual Virtual Expert (IVE)**, whose decisions are aggregated with 3 possible operators into a **College of Virtual Experts (CVE)**.
- Among the 3,000 induced CVEs, the top-CVE with the best predictive performance remains complex: 200 rules share 12 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:

- We build a synthetic dataset composed of 34,000 new cases simulated by deduction from the top-CVE, around the 340 original learning cases but distinct from them.
- We apply 2,000 induction strategies to the same single 34% Training | 33% Validation | 33% Test partition of this new dataset: XTRACTIS induces 2,000 IVEs.
- 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)
102,000 IVEs	F₁-Score	F₁-Score	2 hours (1.13 Tflops)

TOP-MODEL STRUCTURE

The top-IVE model has an excellent intelligibility as it combines the 8 predictors into 4 rules with 3.0 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

- 8 features out of 24.
4 are continuous predictors and 4 are nominal.
- Ranked by individual contribution (1 strong, 1 medium & 6 weak signals):
#1 **Specific gravity** / #2 **Hemoglobin** / ...
- Labeled by modalities or fuzzy classes
Example: **fuzzy interval** "inferior to about 1.0210"



RULES


- 4 connective fuzzy rules without chaining (aggregated into 2 disjunctive fuzzy rules)
- 2 to 6 predictors per rule (on average, 3.0 predictors per rule)
- Example: fuzzy rule **R2** uses 2 predictors and concludes **CKD**. 3 other rules complete this model.

```
IF Specific gravity IS inferior to ~1.0210
AND Hypertension IS Yes
THEN Diagnosis IS CKD
```

Literally, the patient gets a Chronic Kidney Disease diagnosis if the specific gravity (urine density) is under around 1.0210 and if the patient is prone to hypertension.

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 2 UNKNOWN CASES

CASE

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

DEDUCTIVE INFERENCE OF RULES

AUTOMATED DECISION

PATIENT #30

actual value = CKD

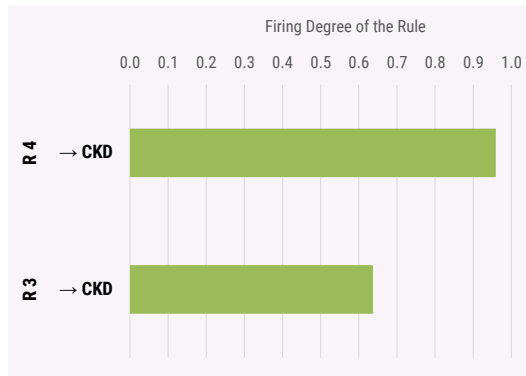
Specific gravity *	1.0050
Blood Glucose Random (mg/dL)	Missing Value
Hemoglobin (g/dL)	12.9
Packed Cell Volume (%)	38.0
Hypertension	No
Diabetes Melitus	No
Appetite	Yes
Pedal Edema	No



For this patient, 2 rules are triggered:

R4 is fired at 0.959 and R3 at 0.637 to conclude CKD.

R1 and R2 are not activated.



NUMBER OF TRIGGERED RULES

2 / 4

FUZZY PREDICTION

{ CKD | 0.959 }

FINAL PREDICTION

{ CKD }

The system delivers a correct diagnosis compared to that given by the nephrologists:

Chronic Kidney Disease



PATIENT #305

actual value = No CKD


Specific gravity *	1.0250
Blood Glucose Random (mg/dL)	128
Hemoglobin (g/dL)	13.1
Packed Cell Volume (%)	45.0
Hypertension	No
Diabetes Melitus	No
Appetite	Yes
Pedal Edema	No



For this patient, 2 rules are triggered:

R1 is fired at 1.000 to conclude No CKD, and R3 at 0.332 to conclude CKD.

R2 and R4 are not activated.



NUMBER OF TRIGGERED RULES

2 / 4

FUZZY PREDICTION

{ No CKD | 1.000, CKD | 0.332 }

FINAL PREDICTION

{ No CKD }


The system delivers a correct diagnosis compared to that given by the nephrologists:

No Chronic Kidney Disease



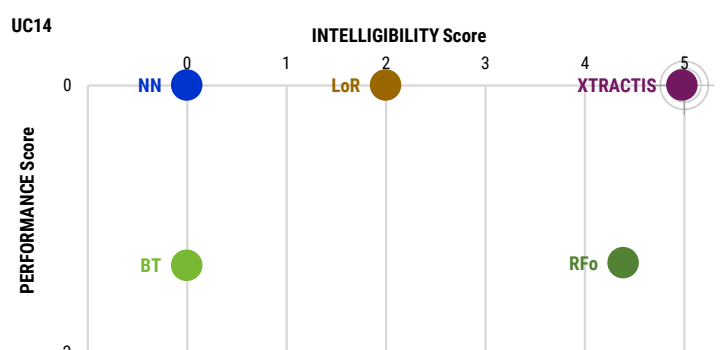
*Predictor value outside the variation range of the model but inside the allowed extrapolation range. XTRACTIS will refuse to give a result for an extrapolation far from the allowed extrapolation range. It is one situation of the "Refusal" prediction.

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

	XTRACTIS 	LOGISTIC REGRESSION	RANDOM FOREST	BOOSTED TREE	NEURAL NETWORK	
MODELING PARAMETERS	MODELS RELEASE	2023/02	2023/02	2023/02	2023/02	
	ALGORITHM VERSION	XTRACTIS REVEAL 13.0.44983	Python 3.9 Scikit-Learn 1.1.2	Python 3.9 LightGBM 3.3.2	Python 3.9 LightGBM 3.3.2	Python 3.9 TensorFlow 2.10.0 Keras 2.10.0
	CROSS-VALIDATION TECHNIQUE	40x5 folds for each CVE model. Then 1-Split Validation for each IVE model: 34% Training 33% Validation 33% Test	20x5 folds for each CVE model	20x5 folds for each CVE model	20x5 folds for each CVE model	20x5 folds for each CVE model
	NUMBER OF EXPLORED STRATEGIES⁽¹⁾	1,000 induction strategies for the CVE on Training / Validation data. 2,000 induction strategies for the IVE on synthetic data	1,000 data analysis strategies on Training / Validation data	1,000 ML strategies on Training / Validation data	1,000 ML strategies on Training / Validation data	1,000 ML strategies on Training / Validation data
	TOP-MODEL SELECTION⁽²⁾	Top-CVE among 3,000 CVEs. Then Top-IVE among 2,000 IVEs	Top-CVE selected among 1,000 CVEs, then single model obtained by applying best CVE strategy on 100% of the Learning Dataset			

TOP-MODEL STRUCTURE	NUMBER OF PREDICTORS (out of 24 Potential Predictors)	8	10 2 nominal variables are decomposed into 12 binary variables	9	9	34 2 nominal variables are decomposed into 12 binary variables
	AVERAGE NUMBER OF PREDICTORS PER RULE OR EQUATION	3.0 per rule	10.0 per equation	2.6 per rule	2.3 per rule	26.3 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	9 trees without chaining 51 binary rules	1 chain of 11 trees 45 binary rules Tree #N corrects the error of the N-1 previous trees	1 hidden layer 3 hidden nodes 4 equations 3 unintelligible synthetic variables

TOP-MODEL SCORES		Random ⁽³⁾	XTRACTIS	LoR	RFo	BT	NN
	INTELLIGIBILITY Score⁽⁴⁾		4.98	2.00	4.39	0.00	0.00
	CVE Real Performance (F ₁ -Score) in External Test		100.00	100.00	98.67	97.30	100.00
	Gap to CVE Leader in External Test		0.00	0.00	-1.33	-2.70	0.00
	IVE Real Performance (F ₁ -Score) in External Test	87.58	100.00	100.00	98.67	100.00	100.00
Gap to IVE Leader in External Test		0.00	0.00	-1.33	0.00	0.00	
Average Real Performance in External Test	87.58	100.00	100.00	98.67	98.65	100.00	
PERFORMANCE Score⁽⁴⁾		0.00	0.00	-1.33	-1.35	0.00	



(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.

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								
<i>Number of Random Permutations (P-value) = 100,000 (0.001%)</i>								
<i>Performance against chance (External Test)</i>								
	23.33%	68.18%					81.58%	
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)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	0 (0.00%)
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)	0.11%	99.84%	99.84%	100.00%	100.00%	99.61%	99.92%	0 (0.00%)
IVE - Predictive Performance (Validation)	0.09%	99.89%	99.89%	99.97%	99.99%	99.72%	99.94%	0 (0.00%)
IVE - Real Performance (Test)	0.14%	99.81%	99.81%	99.97%	99.99%	99.53%	99.90%	0 (0.00%)
IVE - Real Performance (340 original points)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	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)	0.88%	97.66%	100.00%	97.66%	98.60%	100.00%	99.30%	
CVE - Predictive Performance (Validation)	0.59%	98.44%	100.00%	98.44%	99.07%	100.00%	99.53%	
CVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
IVE - Descriptive Performance (Training)	0.88%	97.66%	100.00%	97.66%	98.60%	100.00%	99.30%	
IVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
RANDOM FOREST TOP-MODEL								
CVE - Descriptive Performance (Training)	0.29%	99.53%	99.53%	100.00%	100.00%	99.22%	99.76%	
CVE - Predictive Performance (Validation)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
CVE - Real Performance (External Test)	1.67%	97.37%	97.37%	100.00%	100.00%	95.65%	98.67%	
IVE - Descriptive Performance (Training)	0.29%	99.22%	100.00%	99.22%	99.53%	100.00%	99.76%	
IVE - Real Performance (External Test)	1.67%	97.37%	97.37%	100.00%	100.00%	95.65%	98.67%	
BOOSTED TREE TOP-MODEL								
CVE - Descriptive Performance (Training)	0.29%	99.53%	99.53%	100.00%	100.00%	99.22%	99.76%	
CVE - Predictive Performance (Validation)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
CVE - Real Performance (External Test)	3.33%	94.74%	94.74%	100.00%	100.00%	91.67%	97.30%	
IVE - Descriptive Performance (Training)	0.29%	99.53%	99.53%	100.00%	100.00%	99.22%	99.76%	
IVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
NEURAL NETWORK TOP-MODEL								
CVE - Descriptive Performance (Training)	1.18%	98.44%	99.06%	98.44%	99.06%	98.44%	99.06%	
CVE - Predictive Performance (Validation)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
CVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
IVE - Descriptive Performance (Training)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	
IVE - Real Performance (External Test)	0.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	

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 (2022-2025). XTRACTIS® the Reasoning AI for Trusted Decisions. Use Case #14 | Precision Medicine: Serological Diagnosis of Chronic Kidney Disease – Benchmark vs. Logistic Regression, Random Forest, Boosted Tree & Neural Network. IDAGRAI LABS, June 2025, v3.0, Compiègne, France, 6p.