Artificial Intelligence in Oncology

From organ segmentation to outcome prediction

Nikos Paragios, Distinguished Professor

CentraleSupélec

université PARIS-SACLAY

THERAPANACEA

Artificial intelligence in medicine: between hype and reality

Nine years ago, one of the founders of modern AI predicted the end of an entire medical speciality. What happened?

“Medical schools should stop training radiologists now.”

— Geoffrey Hinton, Nobel laureate, 2017

What AI is good at

Spotting patterns in images, text and numbers — sometimes more reliably than human experts.

Doing repetitive tasks tirelessly, in seconds rather than hours.

Discovering signals in data that are simply too complex for the human eye.

What AI still cannot do

Understand a patient as a whole person, with their history and context.

Take responsibility for a decision, or explain it to a worried family.

Work safely outside the conditions it was trained on — a different scanner, a different population.

Radiologists, pathologists and radiation oncologists are still here. The question is no longer “replace or not” — it is “how do we use AI well?”

What is artificial intelligence?

Behind the buzzword, three concrete things a machine can be taught to do.

A computer programme that succeeds at tasks normally requiring human intelligence.

1

Reason

Mimic how a human thinks through a problem — weighing options, applying rules.

Example

Suggesting a likely diagnosis from a patient's symptoms.

2

Act

Reproduce a specific task by learning from thousands of examples.

Example

Drawing organ contours on a CT scan, like a radiologist would.

3

Discover

Find patterns in data that a human would never notice.

Example

Predicting which patients will respond to a new treatment.

Why is AI changing medicine now, and not 20 years ago?

The mathematical ideas behind modern AI are decades old. Three things changed.

Compute

Computers are millions of times faster than in 1990.

Graphics chips originally built for video games turned out to be ideal for training AI models. Calculations that once took years now take days.

Why it matters

AI can now learn from millions of medical images instead of a few hundred.

Data

Every hospital is now a digital archive.

Imaging, lab results and patient records are stored electronically and can be linked together — something that was paper-based a generation ago.

Why it matters

AI learns by example: the more good examples, the more reliable the prediction.

Algorithms

Better learning recipes — not magic.

A handful of mathematical tricks (notably “back-propagation”) let computers automatically improve themselves from their mistakes.

Why it matters

Models can now learn complex tasks no engineer could programme by hand.

Historical Statements

1950

In (2000) 70% chances for an average human of not being able to distinguish human answers from computer generated answers in a conventional dialogue

Alan Turing

1960

In (1970) computers exploring AI will be able to become world champions of chess

1997

2015

2025?

Cancéropôle Île-de-France

RÉPUBLIQUE FRANÇAISE

inca Institut National du Cancer

AI healthcare landscape

Health system & regulations authorities

Patients

Oncologists

Physicians

Flagship cancer centers

Public hospitals

Private clinics

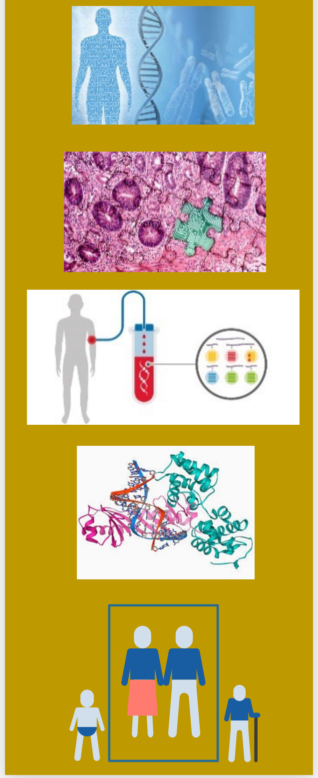

Radiotherapy

Interventional Radiology

(Neuro) Surgery

Drug Development

Precision Medicine

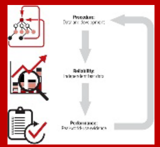

Treatment Companion

PART 2

Pre-operative classification of lung lesions

LUCA-PI — sparing patients with benign nodules from unnecessary surgery.

Part 2: should this lung nodule become surgery?

Lung cancer is a dominant cause of death in modern society. Life expectancy at 5 years after diagnosis is 50%.
When a small spot is found on a chest CT, the only way to know if it is cancer is to remove it surgically. But most nodules turn out to be benign — meaning many patients undergo a major operation they did not need.

The clinical question

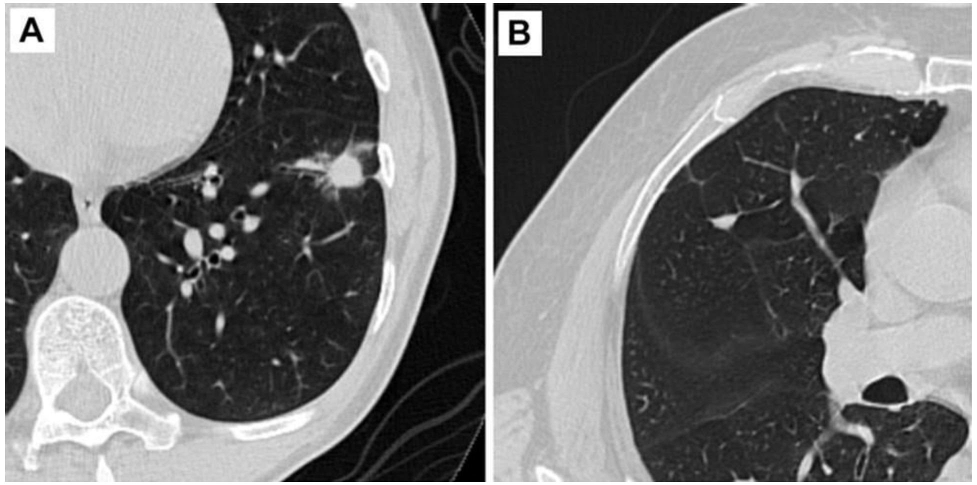

Our approach: LUCA-PI, APHM, IGR, INSERM, …

Use information already available before surgery — patient history, scan features — to predict whether a nodule is cancer.

Cohort: EPITHOR registry

3,032 patients operated for a lung nodule across French thoracic surgery departments (89% malignancy, median age 66).

Inputs

Smoking history, FEV1 lung function, comorbidity score, clinical TNM stage, nodule morphology (ground-glass / solid), spiculation, emphysema.

Model

Gradient-boosted decision trees trained to mimic the surgical-pathology ground truth, with built-in interpretability (SHAP).

Goal: spare patients with clearly benign nodules from a thoracic operation they do not need.

What this means for patients

On the held-out validation set, here is what LUCA-PI delivers — in plain numbers and in patients.

100 patients with a benign nodule:

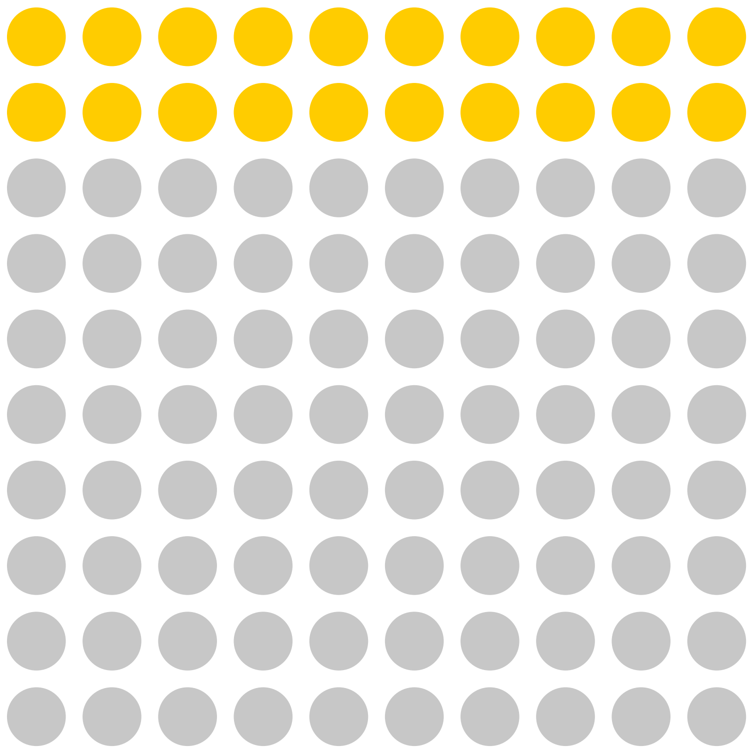

20 spared from surgery

80 still operated (safety margin)

Model performance

0.99
Catches 99% of cancers

20%
of benign patients flagged safe

0.85
Overall discrimination (AUC)

What “sensitivity 99%” means

Almost no cancer is missed. The trade-off is that we are conservative on the benign side: we only spare patients whom the model is very confident about.

Why it matters

Lung surgery has real costs: chest pain for weeks, reduced breathing capacity, possible complications. Sparing one in five benign patients means thousands of avoided operations every year in France alone.

LUCA-PI — Lung Cancer Screening

Recherche Hospitalo-Universitaire programme: a national, multi-disciplinary alliance across clinical, academic and industrial partners.

3,032

patients from EPITHOR

1 in 5

benign patients spared surgery

Partners

Clinical leadership

Société Française de Chirurgie Thoracique et Cardio-Vasculaire (EPITHOR registry); national network of French thoracic surgery departments, APHM, IGR, Inserm, Inria

Academic

Université Paris-Saclay (CentraleSupélec, INRIA); CHU thoracic surgery and radiology departments; INSERM clinical research units.

Industry

TheraPanacea (project coordination, machine learning, data engineering, regulatory).

Funding

France 2030 — Programme d'Investissements d'Avenir, Agence Nationale de la Recherche (ANR), RHU funding instrument.

Cancer, radiotherapy and why this is urgent

One in two of us will face a cancer diagnosis. The numbers explain why AI is no longer optional.

30M

new cancer cases per year worldwide by 2030

15M

annual cancer deaths projected in 2030

60%

of cancer patients will need radiotherapy at some point

Today, treating one patient with radiotherapy means…

Hours of expert time

A senior radiation oncologist spends 3 to 12 hours per patient just preparing the plan — outlining each organ on every CT slice, by hand.

Different care, different outcomes

The same tumour treated in two different centres can lead to two different outcomes — not because of biology, but because of practice.

One-size-fits-all dose

Today's treatment cannot adapt to how a tumour shrinks or how a patient's anatomy changes during a 6-week course of radiotherapy.

AI can save expert time, harmonise care across centres, and personalise dose for each patient.

Radiotherapy : Challenges and Limitations

1 Planning

Imaging + Physics + Optimization + Simulation

2 Treatment

Imaging + Positioning + Radiation

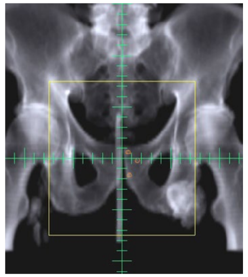

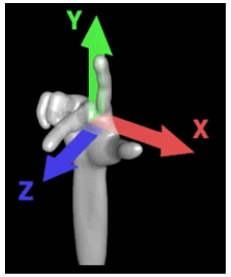

Time

Tedious, time-consuming manual steps

Expert bias

Quality of treatment heavily depends on expertise

Anatomy's evolution

Lack of handling anatomy's evolution and physiological changes

One target equal to One dose

Inability to account for local tumor proliferation

Exploiting outcomes?

Inability to connect treatment choices with outcomes

Facts

Human expertise

3-12 h human expertise required per patient in preparation leading to sub-optimal workflows

Unequal access to treatment

Same patient treated by 2 different centers will see different clinical outcomes

Increased toxicity and Side effects

Adjusting treatment through delivery could allow decreased margins, personalization and eliminate side effects

PART 1

Auto-contouring of organs at risk

Replacing 3–12 hours of manual contouring per patient with reproducible AI segmentation.

Step 1: teaching AI to outline organs on CT scans

Before radiotherapy, every organ near the tumour must be drawn by hand, slice by slice. AI can now do this for the radiation oncologist.

What contouring looks like

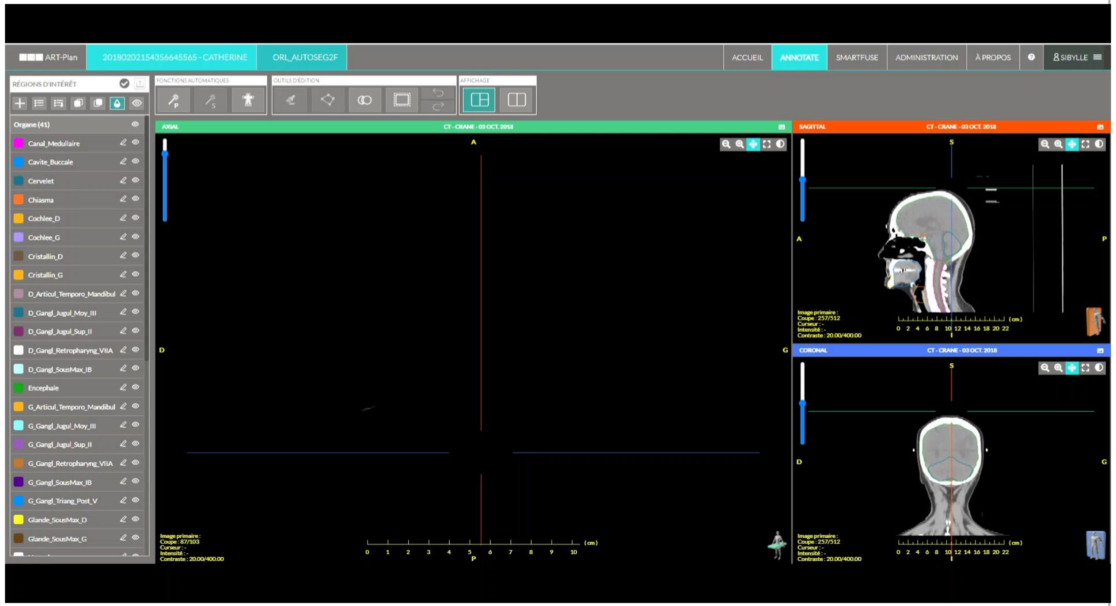

Results

187 min

by hand, per patient

6–41 min

with AI — up to 97% faster

How good are the AI contours?

Blind evaluation, 8 experts, 30 cases:

AI — **79% acceptable**

Human expert — **69% acceptable**

Same quality as a senior radiologist, in a fraction of the time — freeing experts for the cases that really need them.

PART 3

Predicting response to immunotherapy

TRIAGE-IO — identifying patients at risk of early death under first-line single-agent immunotherapy in advanced NSCLC.

Part 3: which patients will benefit from immunotherapy?

Immunotherapy has transformed lung cancer treatment — but it does not work for everyone. About 4 in 10 patients die within 6 months despite having all the markers that usually predict benefit.

The clinical problem

Today there is no reliable test to tell, before treatment, which patients will not respond. Patients spend their last months on a treatment that is not working — instead of receiving the more aggressive combination they need.

TRIAGE-IO

A model that combines 6 routine variables — ECOG performance status, glucose, haemoglobin, bone metastases, STK11 and TP53 mutations — to flag high-risk patients before the first injection.

Survival in the validation cohorts

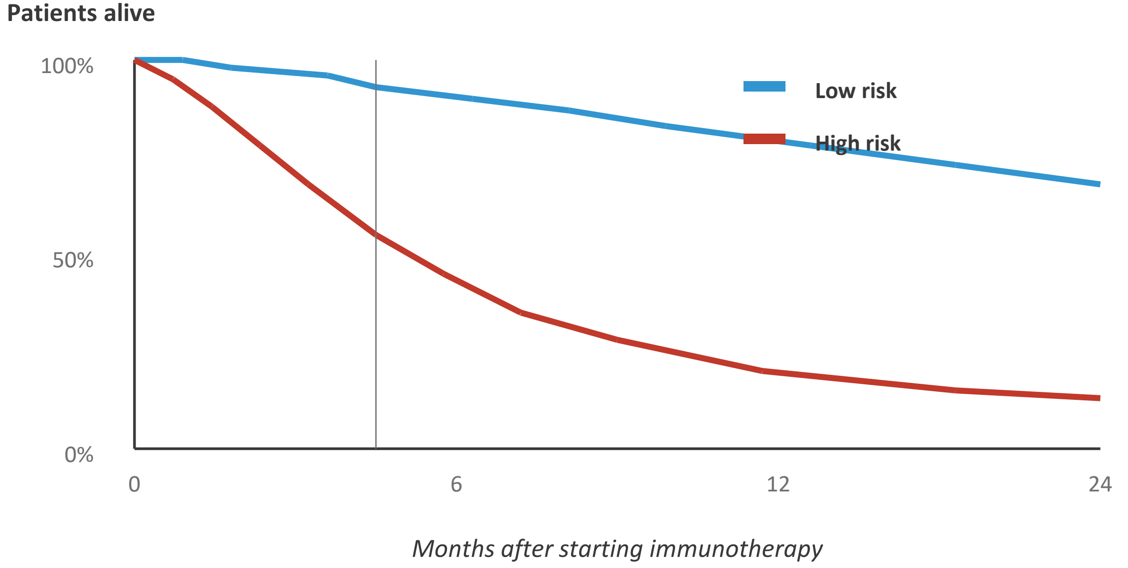

External validation — 3 cohorts, 105 patients

93%
Negative predictive value

88%
Sensitivity

Twice the risk of early death

Patients flagged “high-risk” died from any cause at twice the rate of low-risk patients (HR 1.99, p=0.001).

TRIAGE-IO: external validation in three cohorts

Three independent prospective cohorts — 105 patients, 25 early-death events.

Barcelona
n = 25

NPV

1.00

No early deaths among low-risk patients.

Marseille
n = 30

NPV / Sensitivity

0.89 / 0.71

HR 2.48 (95% CI 1.20–5.15; p=0.015).

Paris
n = 50

NPV / Sensitivity

0.93 / 0.92

Decision-curve analysis: net clinical benefit.

Partner sites and institutions

Spain

Hospital Universitario de Navarra (HUN) — Navarrabiomed–IdiSNA, Pamplona • Hospital Clinic de Barcelona • Complejo Hospitalario de Navarra — Royal Navarre Hospital

France

AP-HP Hôpital Bichat – Claude-Bernard, Paris • Centre Jean Perrin, Clermont-Ferrand • AP-HM, Marseille • Université Paris-Saclay • TheraPanacea

Challenges of AI in healthcare

Standardisation

Clinical practices are different, data as well

→ *Handling Data / Practices heterogeneity*

Correlation ? Causality

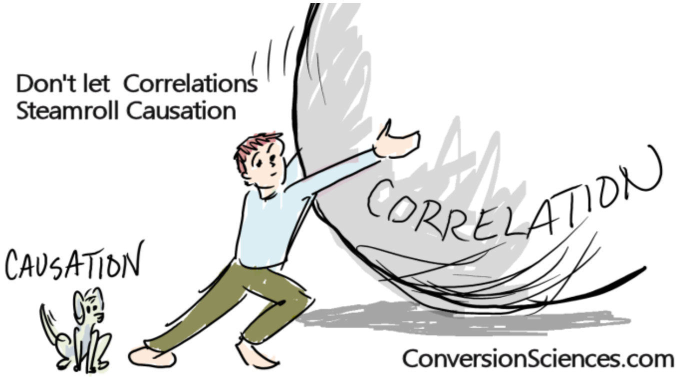

Too many variables to interpret

→ *Random Correlations*

Bias? Generalization

Unrepresentative training and testing sets

→ *Methods that do not generalize*

Responsible AI in oncology is a co-design problem — clinical, technical, legal and ethical at the same time.

Risks of AI in healthcare

Performance is necessary, but not sufficient. Deployment must address structural risks.

Equal access

AI tools risk widening the gap between flagship centres and under-resourced hospitals — unless deployment is intentionally inclusive.

Fairness and bias

Models trained on majority populations can underperform for women, minorities and rare phenotypes. Fairness needs to be a design constraint.

Explainability

Black-box models limit clinician trust and informed consent. Interpretability tools (SHAP) help but are not a substitute for understanding.

Healthcare professionals

Automation may erode hard-won expertise, displace skills, and shift accountability. Training and role design must evolve with the tools.

Privacy and data governance

Patient data is sensitive, valuable and increasingly aggregated. GDPR compliance is the floor, not the ceiling.

Liability and regulation

When an AI recommendation contributes to harm, who is responsible? CE marking and FDA clearance are necessary but only frame the question.

Responsible AI in oncology is a co-design problem — clinical, technical, legal and ethical at the same time.

Maximizing impact of AI in healthcare

- ✓ able to solve full-scale problems
- ✓ where domain expertise is encoded into the solution
- ✓ and continuous access to unique data is guaranteed
- ✓ and algorithmic supremacy is part of the solution

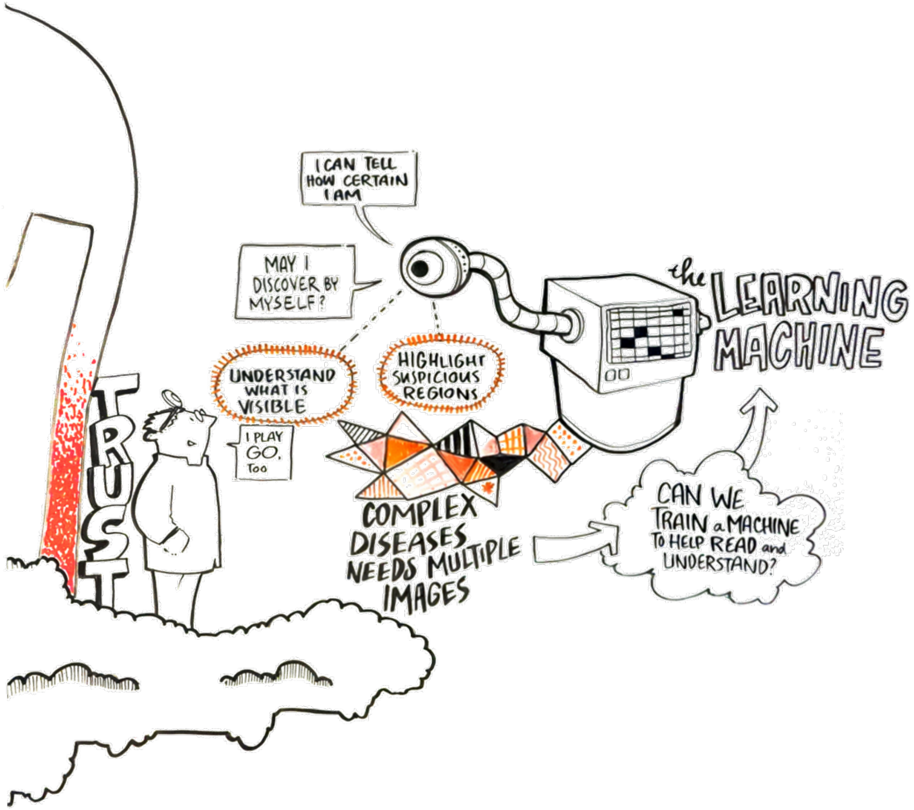

Responsible AI in oncology is a co-design problem — clinical, technical, legal and ethical at the same time.

5

continents

250+

clinical sites

250,000+

targeted oncology patients in 2025

Nikos Paragios, Professor : n.paragios@therapanacea.com