What are we actually getting from AI?

Seven years, 100+ centers, one honest answer.

Dr. Sergey Morozov

Helping hospitals deploy AI in diagnostics

MD, PhD, MPH (Harvard) · Founder Medlogic / 3R Swiss Imaging Network · EuSoMII Past President & Honorary Member

HLTH.rad Summit, Amsterdam

The paradox of medical AI implementation

Eric Topol, Ground Truths, May 2026.

🔍 Proven, yet under-deployed

Imaging AI is validated at scale (mammography, colonoscopy RCTs), yet rarely runs in routine care.

= missed opportunities

💬 Adopted, yet unproven

Generative AI is used daily by millions and most physicians, with little clinical validation.

= missing evidence

Both gaps are about everything around the model: infrastructure, orchestration, monitoring.

Topol's prescription: invest in infrastructure and monitoring now, before scaling AI.

More than 1400 FDA-approved AI tools, tons of hype, one fair question:

“What are we actually getting from *hlthcr* AI?”

A CT scanner you can see. A workstation you can see. AI is mostly invisible.

3R Swiss Imaging Network

NETWORK TODAY

22 CENTERS

60 RADIOLOGISTS

>150 MR, CT, XR, US

2 LANGUAGES

7 SWISS CANTONS

AI adoption in 2020-2026

3R GROUP SWITZERLAND

Safe (Sn)
Precise (Sp)
On Time
Convenient

10 AI solutions
~100K exams / year

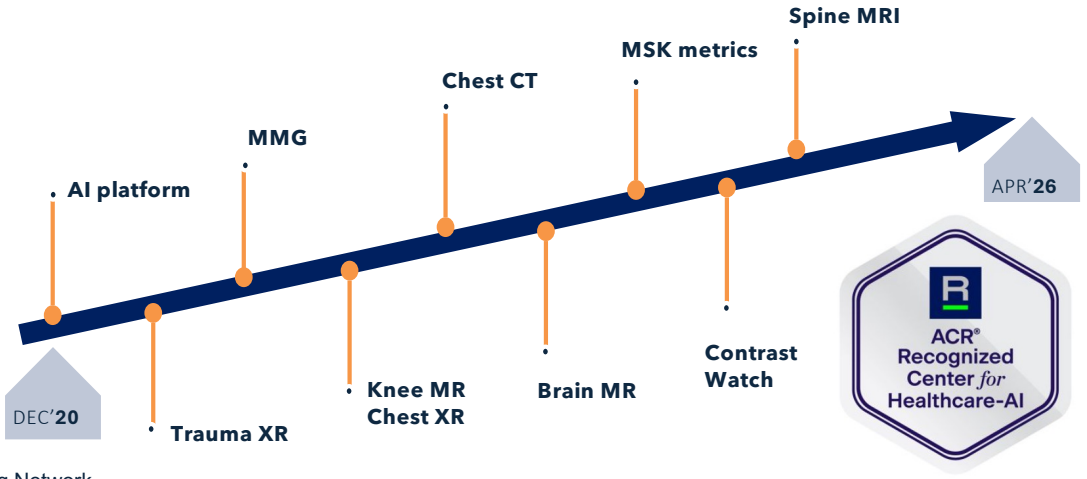

HOW I INVESTIGATED

Three lenses, triangulated

Radiologist survey

Surveys and interviews: what the doctors actually say

AI dashboard

Exams actually processed: what the AI is really doing

Operational data

Report turnaround: what happens in the live pipeline

Triangulated truth

what doctors say, what the AI does, and what actually happens, cross-checked

THE MONITORING APPROACH

One dashboard, seven signals

Aspect	Quality Assurance	Performance Tracking
Purpose	Ensuring quality during development	Monitoring performance in real-world scenarios
Timing	Pre-deployment	Post-deployment
Focus	Standards compliance	Long-term reliability
Activities	Testing, validation, verification	Monitoring, feedback, retraining
Example	Validating a tumor detection model	Tracking tumor detection rates


International Journal of Computer Assisted Radiology and Surgery (2022) 17:1969–1977
https://doi.org/10.1007/s11548-022-02669-1

SHORT COMMUNICATION

Changes in software as a medical device based on artificial intelligence technologies

Victoria Zinchenko[1] · Sergey Chetverikov[3] · Ekaterina Akhmad[1] · Kirill Arzamasov[3] · Anton Vladzymyrskyy[4] · Anna Andreychenko[2,3] · Sergey Morozov[5]

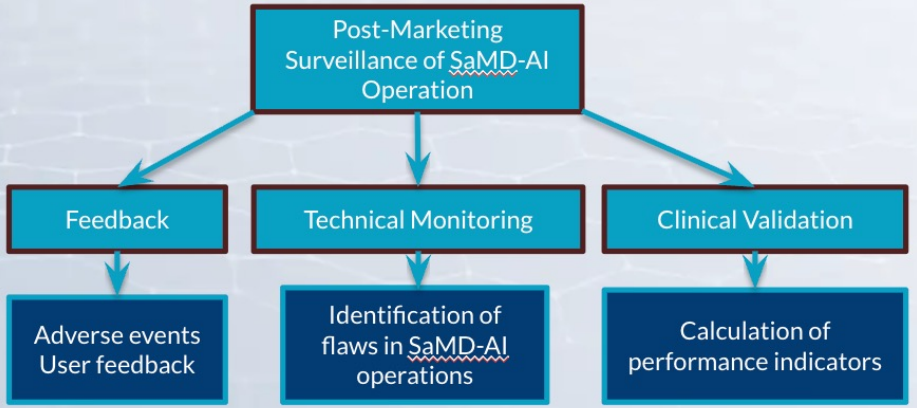

https://doi.org/10.17691/stm2022.14.5.02; https://doi.org/10.1007/s11548-022-02669-1

FINDING 1: IT WAS NEVER RESISTANCE

Adoption was not the problem

>90%
of radiologists use AI

2/3
use it regularly

0
evidence of luddites

How often radiologists use AI

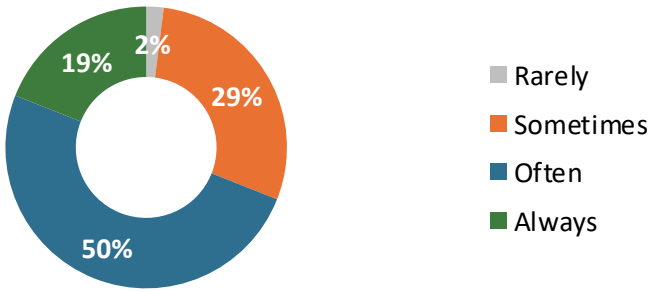

Adoption was never the barrier

Two-thirds of radiologists use AI often or always. The friction is downstream, in the workflow, not in the people.

FINDING 2: WHAT DOCTORS VALUE

A safety belt first, then time. Not speed for its own sake.

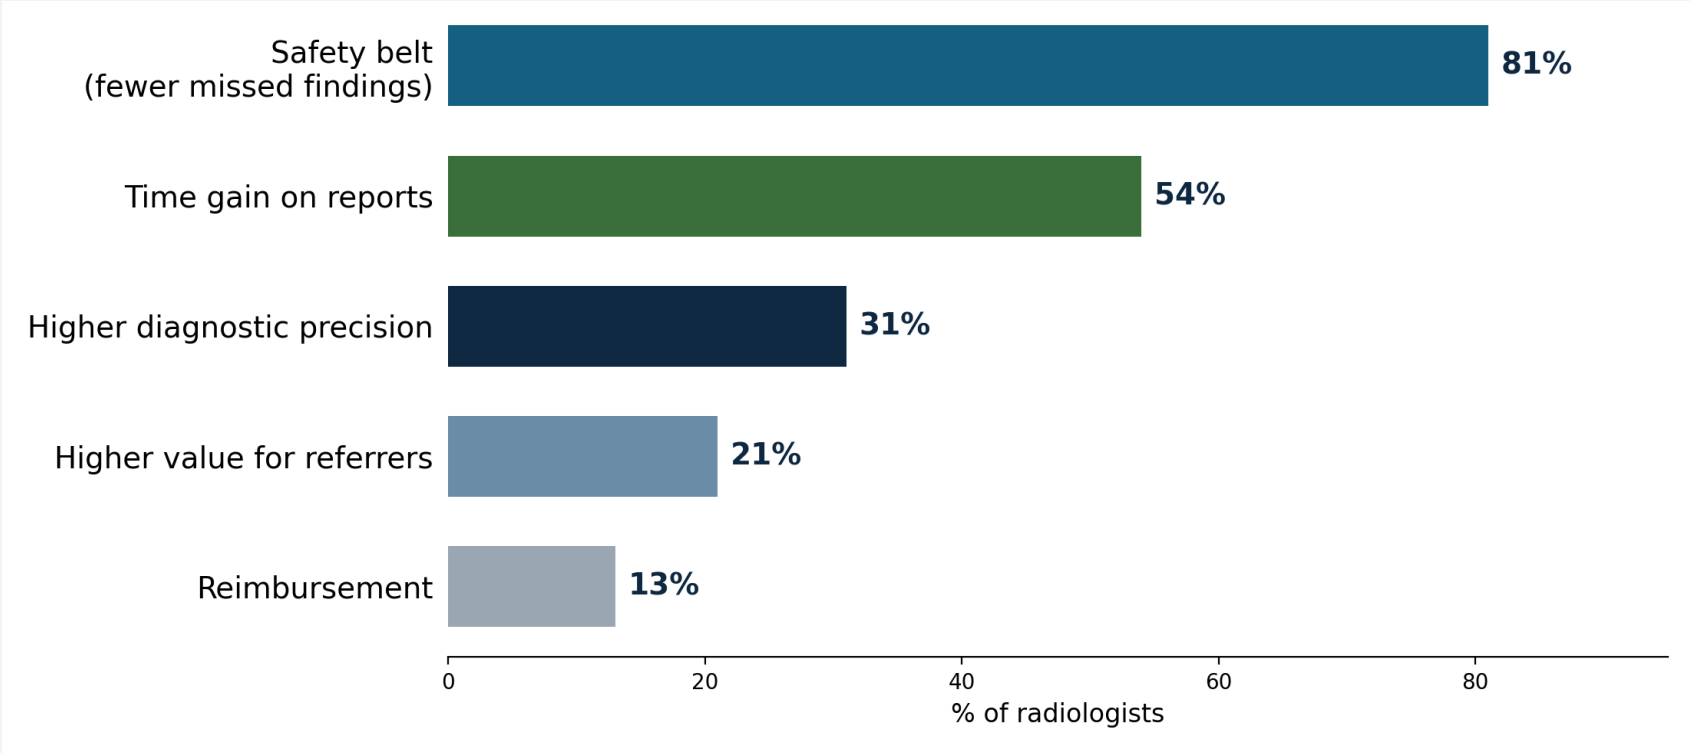

N=48; POLL 1 reveal: the safety belt (81%) beat time and precision

NOT ALL AI IS EQUAL

Radiologists love some tools and reject others

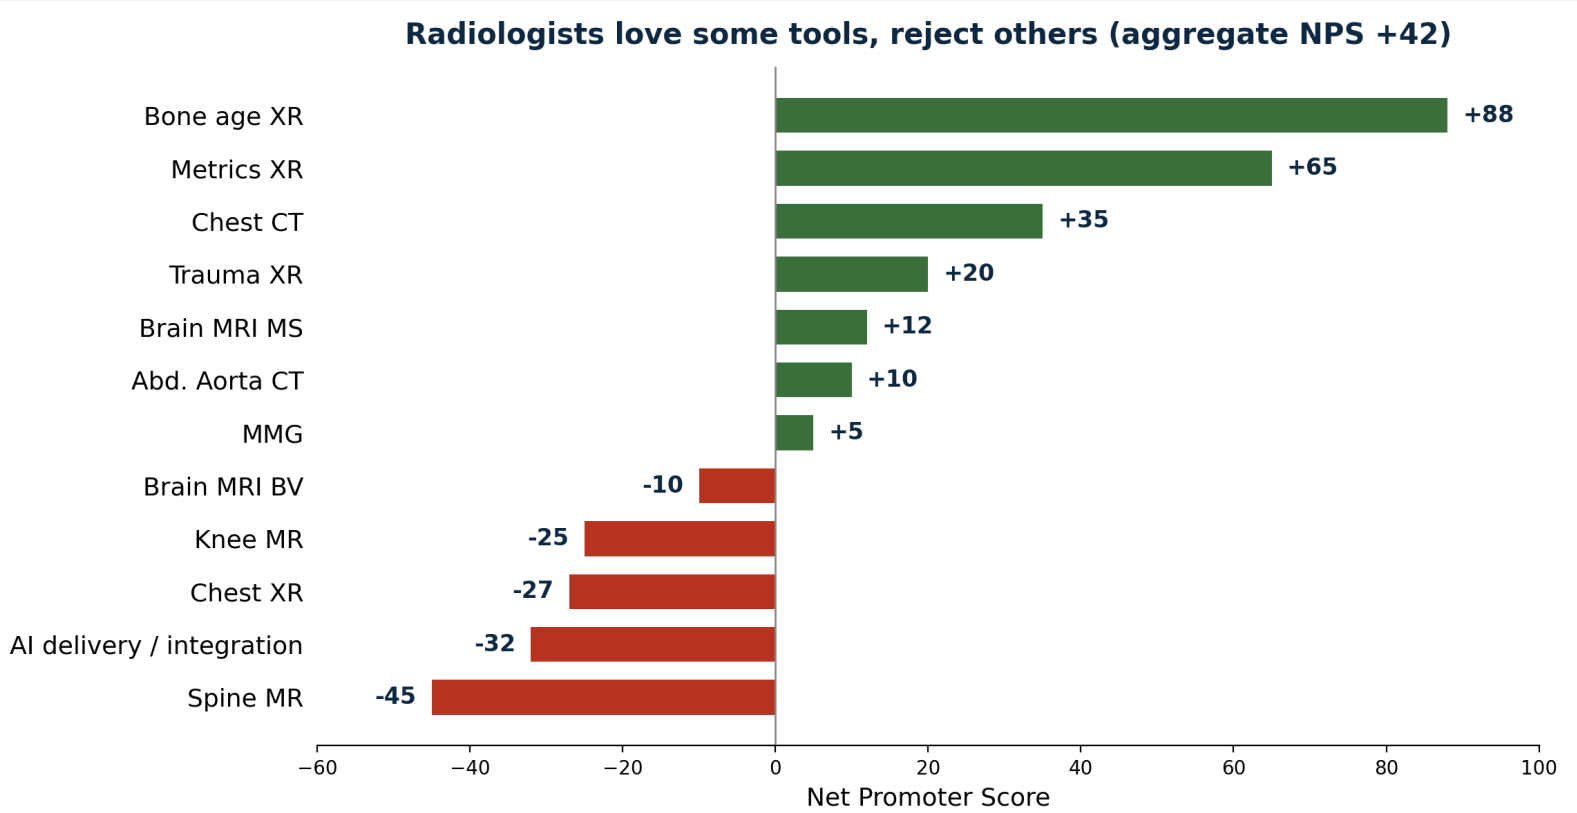

Net Promoter Score per application; aggregate +42

FINDING 3: IT IS PLUMBING

Most of the delay was not inference. It was moving data around.

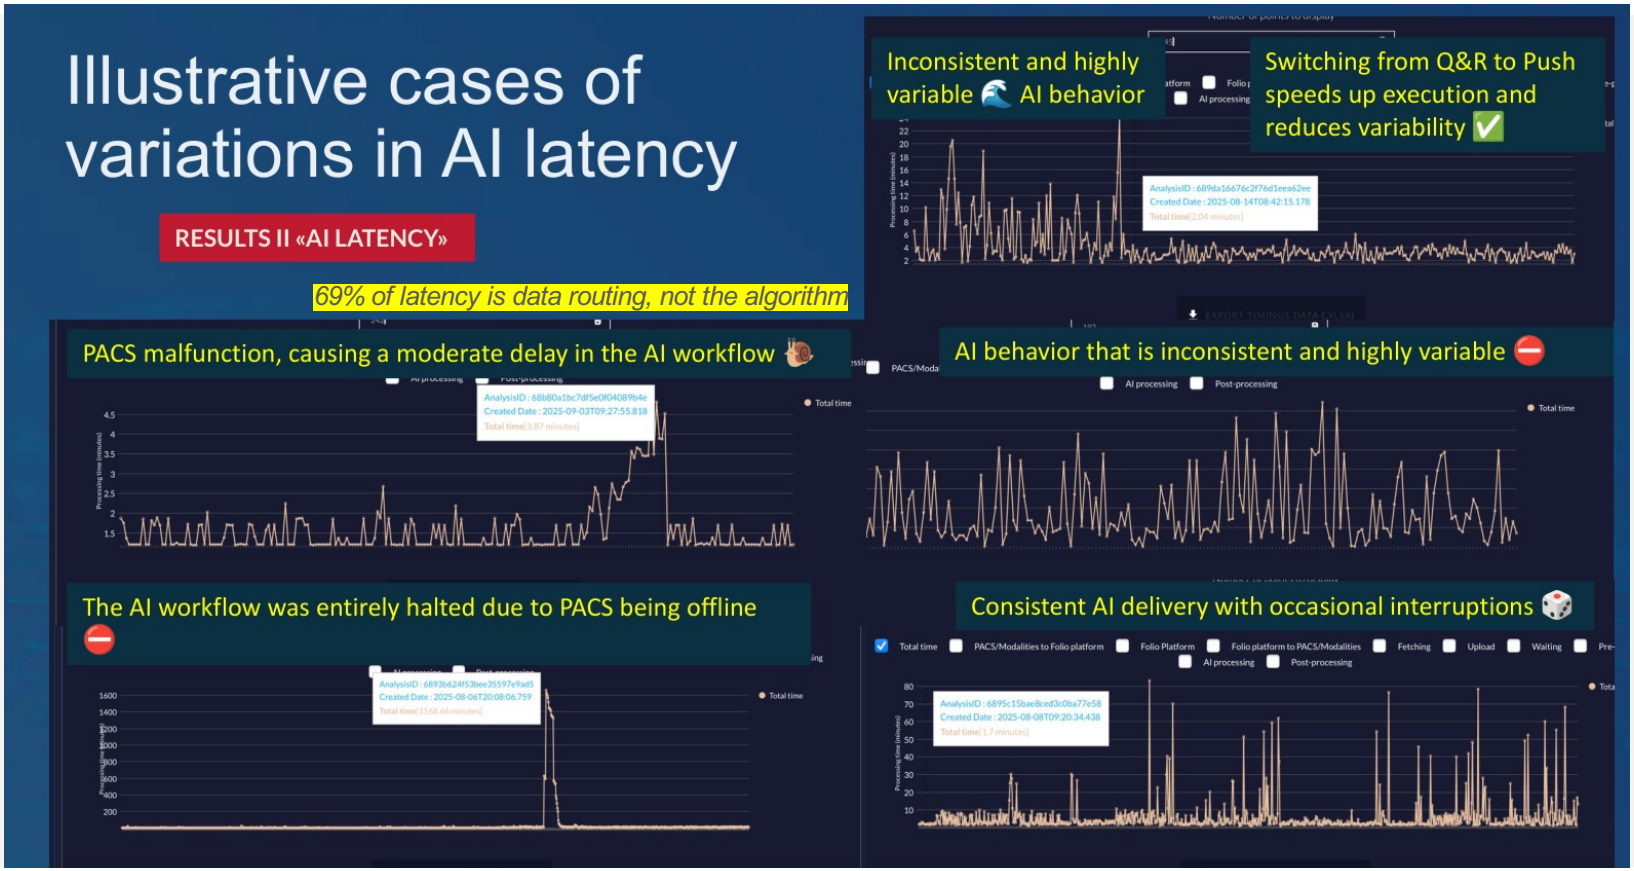

When the result lands after sign-off, you paid for intelligence nobody saw.

An expensive screensaver.

Accuracy without timing is a metric. Accuracy plus timing is a workflow.

THE WORKFLOW EPIPHANY

World-class algorithm + broken timing = zero ROI

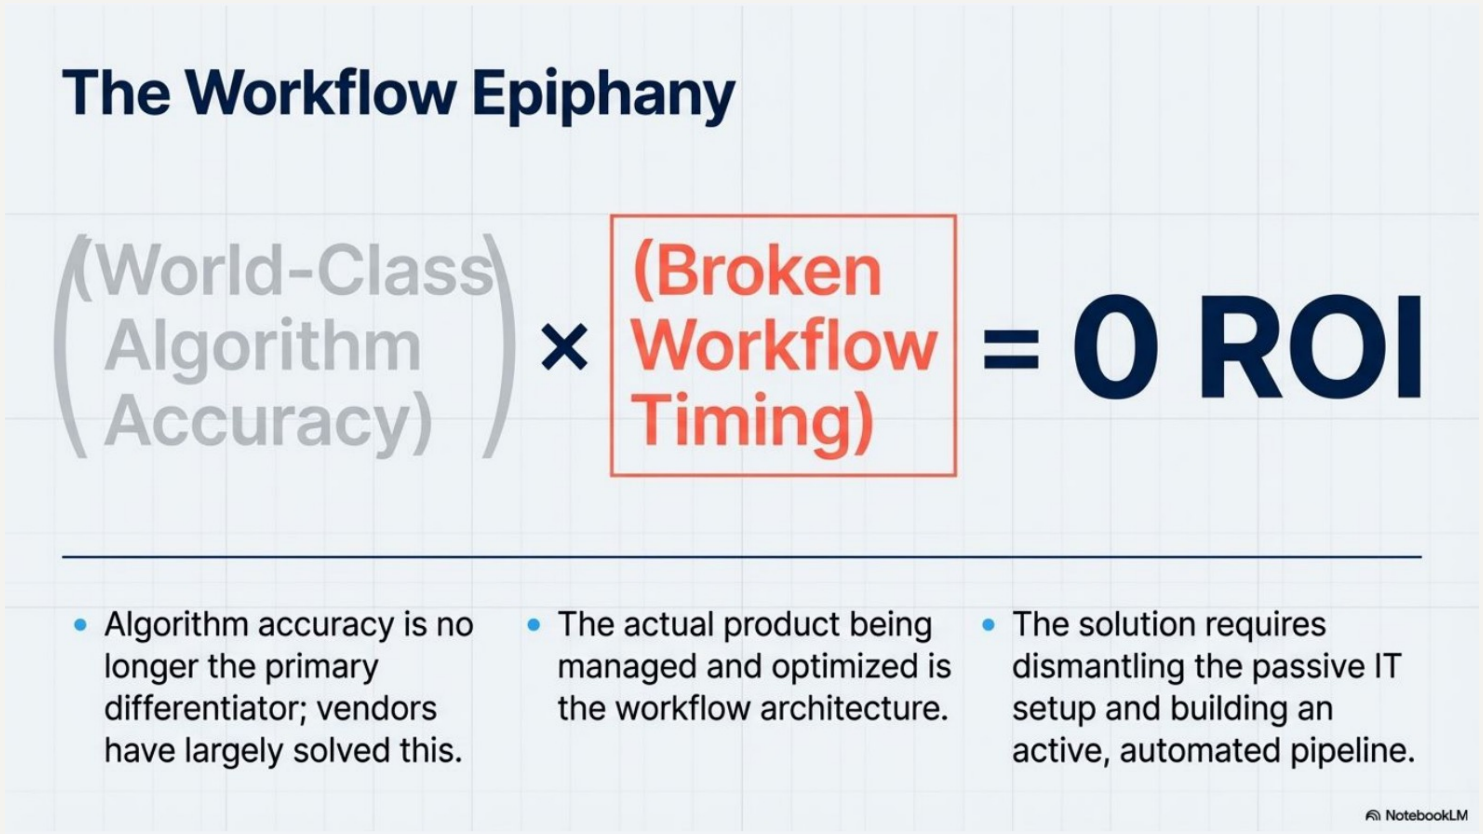

THE PAYOFF

Reporting time fell from about 25 minutes to about 14

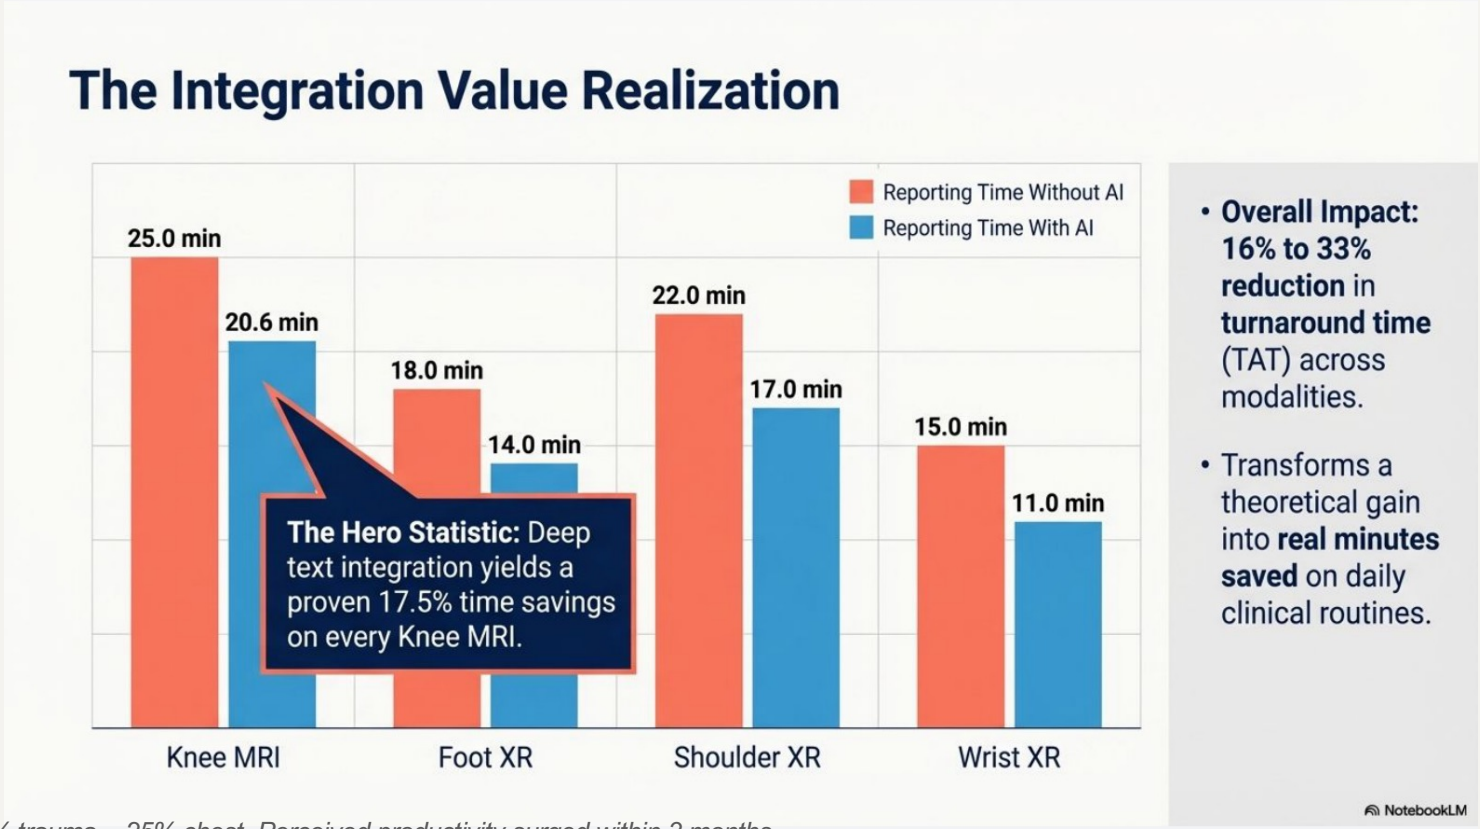

TAT down ~33% trauma, ~25% chest. Perceived productivity surged within 3 months.

THE POPULATION WIN

AI cut mean reporting time about 20%, and ~30% in the 14-month update

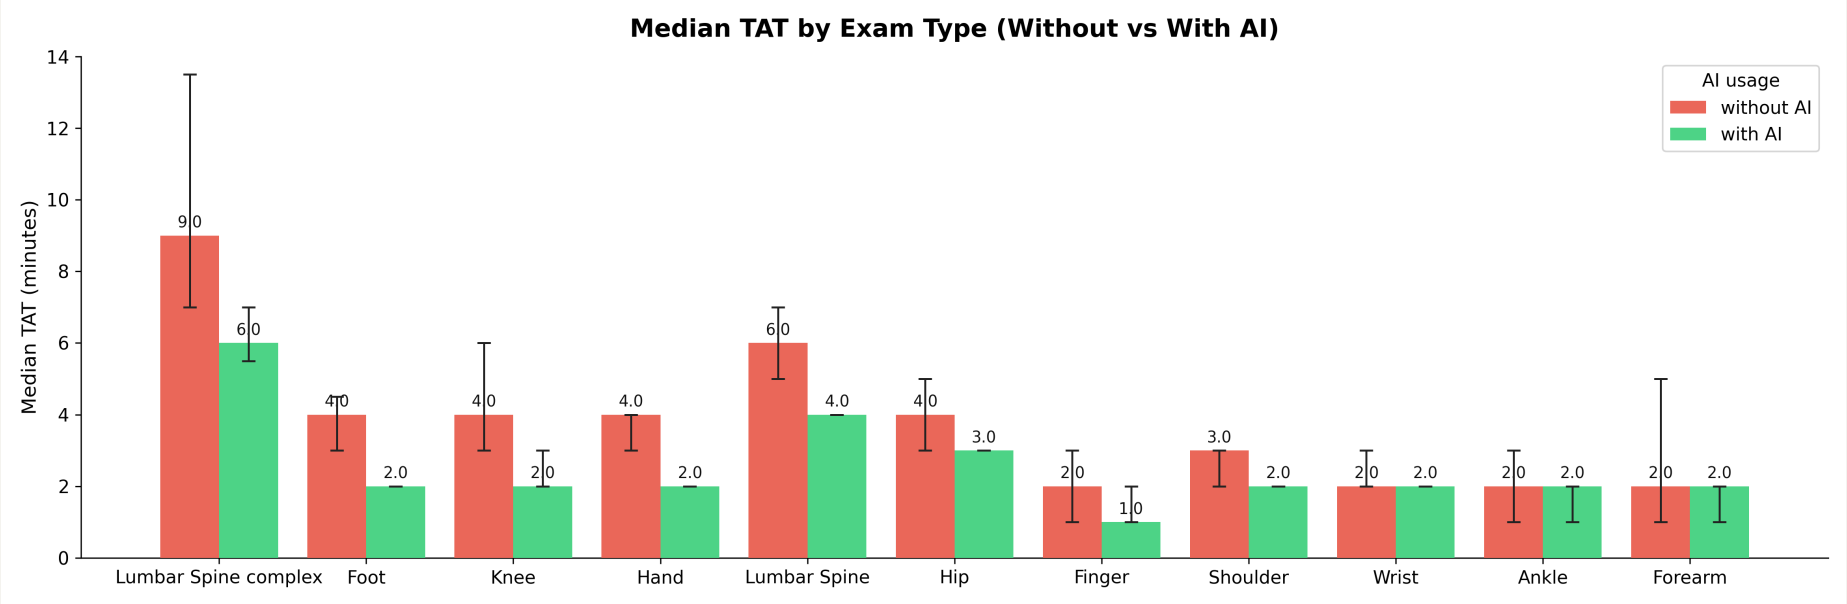

37% of radiologists were individually faster. The average hid who benefits.

3R, nearly 20,000 trauma X-rays

THE 3R INTEGRATION FRAMEWORK

Three moves from screensaver to safety belt

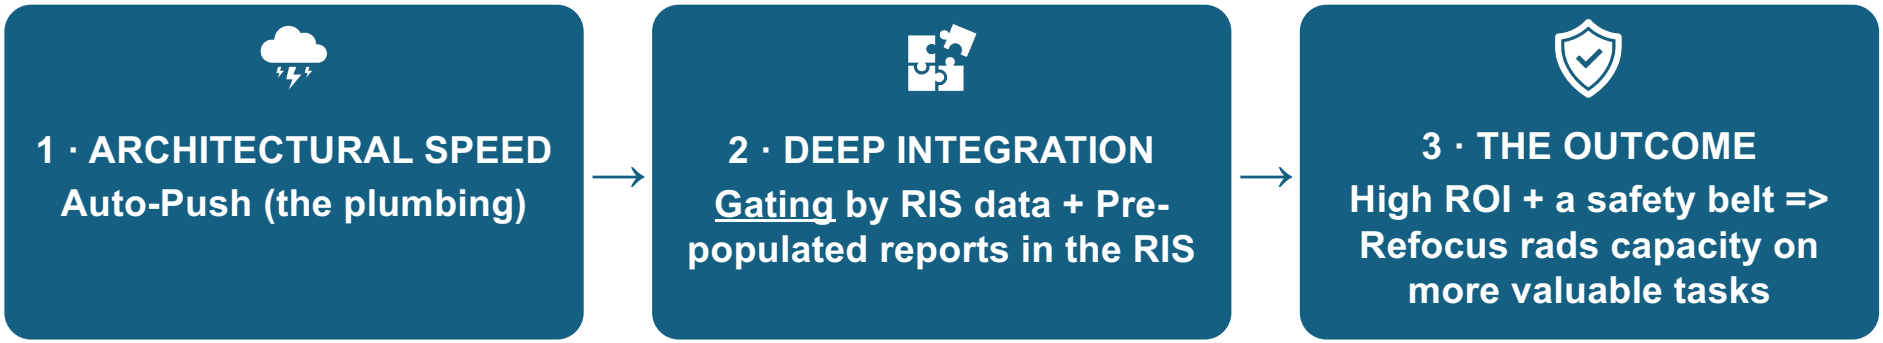

Eliminating latency is the non-negotiable foundation; integration turns speed into adopted value; the outcome is durable ROI.

MAKING THE NUMBER REAL

A business case speaks money, volume, and capacity

- Translate AI into the board's language: hours, FTE, CT/MRI redeployment
- Demand vendor data, build on published evidence, align CEO/COO/CMIO
- Revenue is a bonus; the safety net and recovered time are the case

From measured time-savings to board-level value

−19.6%
mean reporting time (TAT)

→

355 hrs
recovered per year

→

≈0.69 FTE
redeployable capacity

→

CT / MRI
redeploy freed time to higher-acuity reads

Capacity gained: ≈ 0.69 FTE
Platform cost: ≈ 0.23 FTE of one radiologist

≈ 3 : 1
return on high-volume radiography

SIGNAL: CONCORDANCE

AI vs radiologist concordance at scale. Below 90%, trigger a review.

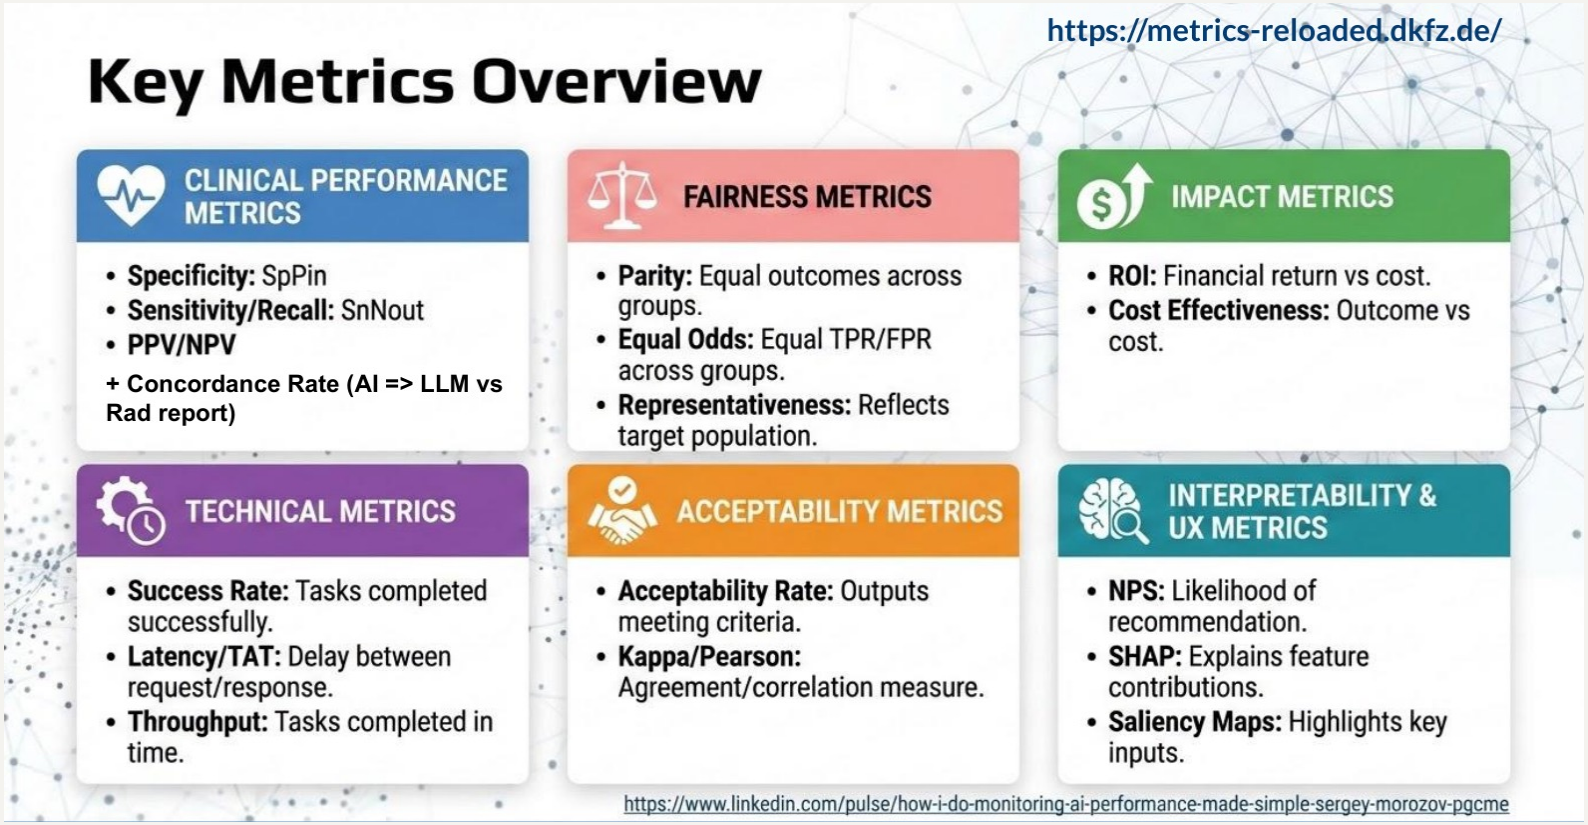

WHERE 3R SITS INTERNATIONALLY

Different systems, same lesson: cost and integration gate adoption

Survey Parameter	US Healthcare Survey	ESR Survey	French CHU Survey	3R Imaging Survey (in publishing)
Organization	Scottsdale Institute	EuroAIM/EuSoMII ESR	French University Hospitals	Groupe 3R Network
Survey Period	2024	2024	2025	2025
Completed Responses	43 C-suite executives	572 rads, residents, radiographers, physicists	80% of CHU radiologists	53 radiologists
Response Rate	64.2%	2.0%	~80%	91.4%
Geographic Scope	United States	Europe	France	Switzerland
AI Usage Rate	100% Ambient Notes, 90% imaging	47.9% currently using AI	80% adoption rate	91.4% using AI tools
Top Success Area	Clinical documentation (53% success)	Breast imaging subspecialty	Improved pathology detection quality; MSK, Thorax, and Neuro leading.	MSK X-ray (NPS +86.7)
Main Barrier	Immature tools (77%), financial (47%)	Immature tools, costs, legal issues	Cost and complexity, limited workload reduction	Latency issues (48%), limited integration
Adoption Status	Variable, generative AI leading	Moderate adoption with concerns	High adoption with mixed satisfaction	Early Adopter (Index 73/100)
Key Limitation	Large health systems only	Low response rate, self-selection bias	Academic setting bias, implementation challenges	Single network, social desirability bias
Main Recommendation	Tool maturity, evaluation frameworks	Education, regulation compliance	Address cost-benefit ratio, improve integration	Speed improvements, workflow integration

THE EUROPEAN REALITY

Europe has algorithms. The bottleneck is infrastructure.

Current EU grant programmes fund algorithms. The integration layer that connects them to care is unfunded.

Data normalisation

Protocol mapping, terminology harmonisation, DICOM tag standardisation

Pseudonymisation

GDPR-compliant token management across sites and vendors

Orchestration

Multi-vendor AI routing, priority queuing, load balancing

Context-aware delivery

Urgency-based AI result routing – urgent / pre-dictation / async

Open infrastructure lowers the barrier: Orthanc · FHIR APIs (AWS/GCP) · Mistral · SNOMED on FHIR – but EU grant programmes must fund the deployment gap.

Fund the glue – not only the models

Standards exist (FHIR, IHE AIW-I/AIR/AIRAI). EU AI Act Art 9, 13, 72.

FOUR POLICY ACTIONS

For the system builders in the room

1 **Mandatory IHE profiles for ALL diagnostic specialties**

AIW-I / AIR / AIRAI in CE marking. Extends beyond radiology → pathology, cardiology, nuclear medicine. Structurally eliminates PACS lock-in.

2 **Clinical context as MDR performance criterion**

Require indication-aware metrics in post-market surveillance (Art. 84). Incentivises vendors to build AI that reads the referral.

3 **Minimum AI Result API – mandatory**

DICOM SR TID 1500 or FHIR DiagnosticReport + confidence scores + model version. Eliminates bespoke per-vendor integration.

4 **Standard for workflow-adjusted AI delivery**

IHE (+ EuSoMII) to specify urgency-based result delivery – critical / pre-dictation / async. Closes the timing integration gap.

What's arriving: agents, foundation models, and local AI

SIIM26's clearest signal: the bottleneck moved from algorithms to integration

AI agents in radiology

- Agents chain steps: pull priors, run a model, format the report
- Not single-task classifiers anymore
- The harness wins, not the LLM: tools, context, guardrails
- New vendor question: do you expose an MCP?

Foundation models for reporting

- Harrison.rad.1.5 drafts reports from images, priors, context
- Only model to pass the FRCR 2B Short Case
- Microsoft MedImageInsight + CXRReportGen, open-source
- Shift from 'detect a nodule' to 'generate the report'

Local-first AI, by design

- Run models locally, keep PHI off public clouds
- Ollama + Open WebUI; LM Studio on Apple Silicon
- Open-weight clinical models ready: Harrison.rad, MS Foundry
- No vendor lock-in

The model matters less than the harness, the delivery, and keeping data in the building.

3R reinvents diagnostics with AI

3R GROUP SWITZERLAND

3R Radiologists / AI users:

Roger AEBI
Marcelo AGUILAR
Alessandro ALIMENTI
Mourad AMOR
Marcela ANCHANTE
Anne-Catherine BAFORT
Anastasia BARRAS
Eugénie BARRAS
Hachem BEN BOUZID
Pierre BENEDICT
Olivier BERREBI
Iheb BOUGAMRA
Johann CARRARD
Ana CARUSO
Theofilos CHRISTOFORIDIS
Monica Oana DEAC
Sofiane DERROUIS
Amira DHOUIB
Luca DUC
Sylvain DUC
Victor FERNANDES
Frank-Günther FÜCHSEL
Maria KEBETS
Dominique FOURNIER
Julien GALLEY
Matteo GANDALINI
Marc GIRAUD
Cécile GRANDIN
Arnaud GREGOIRE
Enrico GUIDETTI
Catrina HANSEN-PHAM
Peter KELEMEN
Romain KOHLER
Amine KORCHI
Georges KROMPECHER
Jean-François LANDEL
Vincent LENOIR
Gibran MANASSEH
Marc MAZILU
Benoît MOREL
Patricia NIN
Mehmet O. OKSUZ
Kahina OUAMER
Alain PELLATON
Jacques PERRIN
Bahar POPAL
Lawrence PUPULIM
Miriam PYKA
Diana RIBEIRO
Guillaume RIITANO
Benoît RIZK
Anna Maria ROSANÒ
Diego SAN MILLAN
Patrique SANTOS OLIVEIRA
Abdulhakim SARRAJ
Anne Laure SAVEROT
Sandro SCHERRER
Karen SCHUSTER
Caroline SCHÜTZ SCHWEIZER
Georgios SGOURDOS
Jean-Marc STEITY
Aphrodite SYROGIANNOPOULOU
Catherine WAEBER VON DER WEID
Oguzhan ZAGRALIOGLU
Lorena ZAMORA
Tihana ZVONAREVIC MISURA

Thank you to the R&D team of 3R