



Operationalizing Clinical AI in Diagnostic Imaging


Turning validated models
into bedside impact

How Innova Solutions partners with academic medical centers to bring research-validated imaging AI into governed, integrated, everyday clinical use—and is building a marketplace to scale it.

WHITE PAPER · CLINICAL AI · DIAGNOSTIC IMAGING



At a Glance

- **The Gap.** Validated imaging AI rarely reaches routine care — the barrier is operational, not scientific.
 - **The Partnership.** Validated model science from medical research and academic institutions paired with Innova's enterprise deployment, integration, and governance.
 - **The Model Classes.** Three investigational, "For Research Use Only" CT model classes—cardiovascular aging, liver radiomics, and liver-fat quantification—on one shared backbone.
 - **The Approach.** Establishment of value on your own data first, then scaling across models, sites, and modalities.
 - **The Marketplace.** A growing, curated catalog of validated clinical AI models, delivered on one governed backbone.
- 

1

EXECUTIVE SUMMARY

THE SCIENCE EXISTS. THE CHALLENGE IS TRANSLATING IT INTO ROUTINE CLINICAL PRACTICE.

Medical imaging now produces more clinically meaningful signals than the workforce can extract by eye. Volume and complexity are climbing far faster than the capacity available to interpret each study, and much of the deeper quantitative signal inside every scan—including patterns associated with cardiovascular, hepatic, and cardiometabolic risk, years before symptoms—goes unmined.

Artificial intelligence can surface that latent signal automatically, from scans already acquired for other reasons, without requiring additional imaging or radiation exposure. A growing body of peer-reviewed research demonstrates the potential of these approaches, and related technologies from academic research programs have already reached regulatory clearance. Yet relatively few imaging AI models progress from research validation to routine clinical

use. The barrier is not the algorithm; the challenge lies in translating a research-validated model to a governed, integrated, continuously monitored clinical service.

This paper describes how **Innova Solutions** partners with **medical research and academic institutions** to address that gap. Partner institutions bring clinical and scientific rigor to model development and validation, with research programs whose track records extend to FDA-cleared products. Innova brings in enterprise engineering, PACS and EHR integration, local validation, and production governance, supported by reusable accelerators that onboard each new model faster and at lower marginal cost. Neither capability succeeds alone; together they form a repeatable pathway from validated science to the bedside. Beyond any single deployment, Innova is building a marketplace for clinical AI

models, a curated catalog through which validated models reach providers on one governed backbone.

The approach is illustrated through three classes of investigational imaging-AI models derived from routine CT: a cardiovascular-aging model (which returns two readouts, a vascular age and a 10-year MACE risk score), a liver-radiomics model, and a liver-fat quantification model. All three are currently **investigational and “For Research Use Only,” and are on a governed path toward clinical deployment**. Delivered on a common integration-and-governance backbone and subject to local validation on the client’s own data before clinical use, these models can unlock new clinical and economic value from imaging data already being generated as part of routine care.

2

THE IMAGING BURDEN

IMAGING IS OUTGROWING THE PEOPLE WHO READ IT

Demand for medical imaging is rising on every front—more studies, more slices per study, and more subtle findings expected from each read—while the workforce available to interpret them is barely growing. The result is a widening structural gap between the signal captured in a scan and the human capacity to extract it.

3–4%

Annual Imaging Growth

Volume rises 3–4% a year; the radiologist workforce grows only ~1%¹.

1,000+

Images Per CT Study

Modern scans can exceed 1,000 slices per study—a new data points generated every few seconds.

+50%

Rise in Attrition

Radiologist attrition is up 50% since 2020; ~1 in 3 is now aged 55+².

The demand side is structural, not cyclical

The forces driving imaging volume are durable. An aging population with increasing chronic disease burden generates more cross-sectional imaging per person. Screening guidelines continue to broaden—lung, cardiac, and colorectal—and each new indication adds recurring volume. Cross-sectional modalities also keep getting richer—thinner slices, multi-phase protocols, and dual-energy acquisition all multiply the number of images produced per encounter. These factors are likely

to sustain imaging demand even as temporary backlogs are reduced

The supply side cannot respond quickly

Radiology capacity is slow to expand. Training a radiologist takes well over a decade, so the pipeline cannot quickly adjust to meet the demand spikes. At the same time the existing workforce is aging toward retirement, while attrition and burnout add further pressure. The consequence is not simply a longer queue. When every study must be interpreted under given time constraints, the focus naturally remains on the

clinical question that prompted the scan while deeper quantitative information may go unmeasured. While marginal minutes are spent answering the referral question—the reason the scan was ordered—the deeper quantitative signals are necessarily deprioritized.

Those deeper signals lie in the tissue volumes, densities, and textures embedded within the scan. This information requires quantitative analysis and is difficult to extract during the radiologist's read. As a result, clinically meaningful information contained within the images may remain unmeasured.

Sources

1. Radiologist Shortage: A National Healthcare Challenge, Anderson Hospital, 2025
2. Projected US Radiologist Supply, 2025 to 2055, Harvey L. Neiman Health Policy Institute, 2025

3

THE OPPORTUNITY

THE SIGNAL IS ALREADY IN YOUR SCANS

Because the signal has already been captured, the opportunity is additive rather than disruptive. Clinical AI can analyze each scan automatically, in the background, from images taken for entirely different reasons. It then surfaces additional risk information without a new scan or added radiation, and without changing the clinician's workflow at the point of care.

1.

Use of Already Available Scan

Uses existing scans. No additional imaging acquisition, radiation exposure, or scan-related cost.

2.

AI-Read Hidden Signal

Quantifies patterns the human eye cannot measure.

3.

Risk Surfaced Early

Flags elevated risk years before symptoms appear.

Opportunistic imaging is an established paradigm

Extracting additional clinical value from a scan acquired for another indication is not a novel idea; it is a growing, evidence-backed practice. Bone-density estimates and coronary-calcium scores are already being read opportunistically from routine chest CT. Rather than depending on a radiologist to detect and measure a secondary finding, AI analyzes every eligible scan automatically and consistently. The suite described here extends this paradigm to cardiovascular and cardiometabolic risk.

Why CT is a uniquely rich substrate

CT is volumetric, quantitative, and

reproducible. A single study contains tissue-level information—including the volume, density, and radiomic texture of fat depots, organs, and vessels—that is fully present in imaging data but is not routinely quantified during the standard visual read. These are precisely the features that quantitative models are designed to analyze. The economic logic follows naturally: the imaging has already been acquired, the patient bears no additional burden, and the analysis runs in the background. The incremental cost is primarily compute and integration, rather than additional imaging, allowing the same infrastructure to support multiple models and extend the value of each scan beyond its original purpose.



THE SCIENCE EXISTS. THE DEPLOYMENT DOESN'T.

Validated models rarely reach the bedside. Between a research-validated algorithm and its daily clinical use stand four distinct barriers—none of them scientific, all of them operational. Understanding each one clarifies precisely what a credible deployment partner must solve.

Barrier	What Stands in the Way
Research-to-clinical translation	Models validated for research use need a rigorous, documented path to clinical deployment.
No local validation	Performance is unproven on the client's own patients, scanners, and imaging protocols.
No workflow integration	Results live outside PACS and EHR, so clinicians never see them in the context of the read.
No governance	Without monitoring, drift detection, and audit, models cannot satisfy clinical and compliance review.

Translation. A model that performs well on a curated research cohort may perform differently on real-world data from a population or environment it was not developed for—a problem of distribution shift. Closing this gap means documenting intended use, characterizing the performance envelope, and defining known failure modes, so clinicians understand not only what the model does, but where it should not be trusted.

Local validation. Scanner make and model, reconstruction kernels, slice thickness, contrast protocols, and patient mix all vary from site to site. A

model that performs at one institution may not perform identically at another without site-specific validation against local reference standards—which is why proving value on the client's own retrospective data is non-negotiable before anything goes live.

Workflow integration. A result that lands in a separate portal may be missed at the moment of decision. Real integration means routing studies to inference automatically, orchestrating the worklist, and writing structured results back into the PACS viewer and the EHR through DICOM, HL7, and FHIR—so the output appears in

the report, in clinical context, without requiring the clinician to check another screen.

Governance. A production model is not a finished artifact; it is a system that requires ongoing oversight. Performance can change as scanners, protocols, and populations change; subgroup performance must be monitored for bias; every inference needs an audit trail, and every model change needs version control. Without these controls, no clinical-governance or compliance committee can responsibly approve routine use.

THE THROUGH-LINE

Each barrier is an engineering and governance challenge rather than a scientific one. That is why effective deployment requires both validated model science and enterprise-grade engineering—a winning combination at the heart of this partnership.

ILLUSTRATIVE CLASSES OF RESEARCH-VALIDATED CLINICAL AI MODELS

Each model reads a routine CT to surface risk that standard tools miss, and all of them run on the same integration and governance backbone. This makes onboarding the second and third model faster and more efficient than building each deployment independently. The illustrative suite comprises three models focused on cardiovascular and cardiometabolic risk, using imaging that patients are already receiving.

Model	What It reads → Output	What the Validation Work Suggests
 Cardiovascular Aging	 Peri-aortic fat on routine CT → cardiovascular age (CV-Age) and 10-year MACE risk score. One model, two readouts.	 Captures early vascular aging and inflammation before symptoms; a high PARC has been associated with a multi-fold higher rate of MACE, including in patients rated low or intermediate by PREVENT.
 Liver Radiomics	 The liver on a routine cardiac / CAC-CT → a heart-failure risk score.	 Flags heart-failure risk even when the calcium score is zero and in patients rated low-risk by PREVENT.
 Liver-Fat Quantification	 Quantitative liver fat on routine CT → cardiometabolic risk and MASLD.	 Detects fatty liver even when CT attenuation looks normal and links hepatic fat burden to downstream cardiovascular events.

STATUS · INVESTIGATIONAL · IN VALIDATION

All three models are “For Research Use Only” and on a governed path to clinical deployment. They are not currently cleared for diagnostic use. The associations above come from validation cohorts while the clinical claims are made only after local validation and completion of the appropriate regulatory pathway. Full methods, cohorts, and results are detailed in the originating programs’ published materials.

Cardiovascular and cardiometabolic diseases are high-prevalence, high-cost, and frequently silent—making them strong cases for earlier quantitative risk stratification. The relevant tissue is already captured in routine chest and cardiac CT: peri-aortic fat sits in the field of view of countless studies, while the liver appears on cardiac scans (acquired

for entirely unrelated reasons). The recurring theme across all three models is their ability to find risk that may not be captured in a conventional read: elevated risk even when the calcium score is zero, or hidden fatty-liver disease when attenuation looks normal. These are hypotheses the respective validation studies are designed to test, rather than settled clinical claims.

The strategic value lies not in any single biomarker but in the shared platform beneath them. Because every model consumes the same imaging inputs and returns structured results through the same connectors, the core integration validation, and governance work is built once and reused—turning what could otherwise be a series of one-off deployments into a platform whose cost per additional model falls over time.

Every model in the suite follows the same basic pattern: start from a scan already acquired for another reason, quantify a tissue signature the standard read never measures, and return a single interpretable result into the report. What differs is the tissue each one analyzes and the risk it surfaces. The three walk-throughs below show that pattern in each model.




Age gap (Δ Age) = Cardiovascular Age – Chronological Age.

A positive gap indicates accelerated aging and elevated risk; a negative gap signals resilience. Expressing risk as an age is a deliberate choice. A 52-year-old with a vascular age of 65 grasps the message without interpretation. Because the cardiovascular age can also be recomputed on future scans, it has the potential to provide a longitudinal measure of change. Alongside the cardiovascular age, a high peri-aortic risk score has been associated with a multi-fold higher rate of MACE, including among patients rated low or intermediate by PREVENT.



Liver Radiomics— Heart-Failure Risk, From a Cardiac CT

The model analyzes the liver texture on a routine cardiac calcium (CAC) scan to identify future heart-failure risk, years before symptoms, turning a scan acquired for coronary assessment into an early-warning signal for the heart.



1




Routine CAC-CT

A cardiac calcium scan already obtained in CV care.



2



Liver radiomic signature

AI reads liver texture reflecting metabolic & hemodynamic stress.



3



Liver Radiomics Score

It outputs a single heart-failure risk score.


The signal.

A high score is associated with higher future heart failure risk. In validation work, the association persisted even when conventional markers were reassuring: materially elevated risk when the calcium score is zero, and in patients PREVENT rates as low-risk. The findings suggest that liver texture may provide information related to cardiovascular risk that is not captured by the coronary calcium score alone.



Liver-Fat Quantification— Liver Fat & Cardiometabolic Risk From a Routine CT

The model quantifies liver fat from a CT already obtained for cardiovascular indications, catching hidden MASLD and elevated cardiometabolic risk that a qualitative read would pass over.



1



Routine CT

A CT is already obtained for CV indications (e.g., CAC-CT).



2



Quantitative Liver-Fat Analysis

AI estimates actual liver-fat %, catching subtle, uneven fat.



3



Liver-Fat Score

The resultant liver-fat burden linked to cardiometabolic risk.

The signal.

The model detects true liver fat even when the CT looks normal: attenuation can sit in the normal Hounsfield range while clinically significant fat is still present, and that hidden burden signals MASLD and elevated cardiovascular (MACE) risk standard tools miss. As with the other two models, the value comes from quantifying a tissue the eye cannot measure precisely: the same existing-scan, quantified-signature, interpretable-output pattern, applied to a different organ and a different risk.

ENTERPRISE DEPLOYMENT, BUILT TO REUSE

Every model is first validated on your existing scans, then embedded into the workflows clinicians already use. Reusable accelerators connect the models to the PACS viewer and EHR, so results appear in the context of the read rather than in a separate application. This creates a reusable deployment foundation: once the core integration and validation processes are established, additional models can be onboarded without rebuilding the platform from scratch.

1. CT / PACS

Existing scans & storage.

2. Model Validation

Validated on your existing images first.

3. AI Inference

Validated models in the cloud.

4. FHIR / PACS Connectors

Reusable integration layer.

5. Radiologist + Physician

Results in the PACS viewer & EHR.

Clinical models & enterprise engineering

The models are born in an Academic Medical Center; the delivery is enterprise-grade. Bringing the two together is what lets validated science reach patients—combining clinical credibility with the engineering discipline to run models safely at scale.

What Academic Medical Center Brings	What Innova Brings
Model development, validation, & regulatory alignment	Enterprise engineering & PACS / EHR integration
Conventional rigor and clinical credibility	Local validation, governance, & monitoring at scale
Research pipelines with track records extending to FDA-cleared products	Reusable accelerators to onboard each new model fast

Why the split matters

These capabilities rarely exist within the same organization. Academic Centers bring clinical research expertise and credibility but typically do not operate as enterprise software deployment organization across multiple health systems. Technology vendors bring engineering and integration capabilities, but not the underlying model science. The partnership assigns each responsibility to the party best equipped to own it: academic-medical-center owns model science, validation methodology, and the regulatory pathway; Innova leads deployment, integration, local validation at scale, and production monitoring—with feedback loops connecting the two.

The reference architecture

A study originates at the imaging modality or PACS. Before any model is trusted in production, it is validated on the client's own retrospective images to establish its performance across local scanners, protocols, and patient mix rather than relying solely on published results. Once validated, studies are routed securely and automatically to a cloud inference environment provisioned for the client, where each model runs as an isolated container against the incoming images. An orchestration layer determines which models apply to each study and manages processing and scale. Results are then returned through the integration layer to the PACS viewer for the radiologist and the EHR for the referring physician. For clinicians, no separate application or workflow is required; results appear within the systems they already use.

The accelerator and connector layer

The connectors form a reusable core of the platform. Standards-based integration components for DICOM,

HL7, and FHIR handle the movement of images in and results out and can be reused across models. This reduces the integration effort required for subsequent models. Once the core connectors and validation framework are established, the additional models can use the same infrastructure while undergoing their own validation.

Cloud, scale, and security

Packaging each model as a container makes it portable and independently deployable. This keeps all the models isolated from one another, allowing the environment to scale horizontally as the imaging volume grows. Because the workload runs on provisioned cloud infrastructure, capacity can expand with demand rather than being constrained by on-premise hardware. Protected health information is encrypted in transit and at rest, with access controls and audit logging appropriate to the clinical data. Security and compliance are incorporated into the architecture from the outset rather than being added later.

Why architecture matters

- 1



Validate first, then deploy.

Every model proves itself on your existing scans before it ever runs live.
- 2



Onboard once, reuse everywhere

New models plug into the same connectors: integration is built once, not rebuilt per model.
- 3



Fits the read, doesn't add to it

No new screen or workstation for clinicians; results appear where they already work.
- 4



Lower marginal cost per model:

Platform economics, not project economics: the cost of each additional model falls over time.



7

THE MARKETPLACE

ONE BACKBONE, MANY MODELS: A MARKETPLACE FOR CLINICAL AI

The backbone described above points toward something larger than a single suite. Innova is building a marketplace for clinical AI models: a curated catalog through which validated models reach providers on one governed deployment platform. Model developers gain an enterprise route to the bedside—integration, local validation, monitoring, and commercial infrastructure they would otherwise have to build themselves. Providers gain one integration effort and one governance framework across every model they adopt, rather than a new vendor relationship for each algorithm.

Every model listed in the marketplace clears the same bar: documented intended use, characterized performance, a defined regulatory status, and readiness for local validation. Admission is defined by that evidence standard rather than by institution type; qualifying models may come from academic medical centers, independent research institutes, or health-system innovation programs. The three model classes profiled here—spanning cardiovascular, hepatic, and cardiometabolic risk—illustrate the kinds of entries the catalog is designed to host, with additional tissues, modalities, and risk domains to follow.

8

WHO BENEFITS

VALUE FOR EVERY STAKEHOLDER

The same background analysis—designed to minimize workflow disruption—creates a distinct win for everyone the scan touches, from the patient on the table to the health system that owns the imaging assets. Because the analysis is automatic and universal rather than selective, these benefits accrue at population scale rather than case by case.

The wins are not competing claims to be traded off against one another. Rather, they are the same event seen from four vantage points. A single background inference simultaneously advances a clinical goal (earlier detection), an operational one (no added read burden), and a financial one (better performance against value-based-care and quality metrics). That alignment is what makes the model sustainable: the intervention that helps the patient is the same one that strengthens the health system's economics.

Stakeholder Outcomes

Stakeholder

The Win

Patients



- Hidden cardiovascular risk caught earlier.
- No extra scan, radiation, or cost.
- Time to act before symptoms start.

Physicians



- Sharper, objective risk stratification.
- Borderline patients reclassified for earlier action.
- A repeatable measure to track over time.

Radiologists



- Automated quantitative results on scans already read.
- Signal captured beyond the referral question.
- No added read time or new workstation.

The Hospital



- Opportunistic screening at population scale.
- Alignment with value-based care & quality goals.
- New value from imaging assets already owned.

BUILT TO EARN CLINICAL AND REGULATORY CONFIDENCE

Because these are clinical models operating on patient data, trust is engineered in from the start rather than added afterward. Four commitments define how models move from research use toward clinical use, with each supported by a defined process.

Validation before go-live

Published performance alone is not sufficient to support clinical deployment. Each model is evaluated on the client's own retrospective images against appropriate reference standards, with performance assessed across clinically relevant subgroups. Only once performance has been evaluated on local data and reviewed with the institution's clinical leads does a model move toward live use. This provides institution-specific evidence to inform the decision to proceed toward clinical use.

Monitoring and equity in production

Once in production, model performance and incoming data are monitored continuously. The platform tracks the distribution of incoming data and model

Commitment	What It Means in Practice
 Local Validation First	Every model is proven on the client's own retrospective imaging data before moving toward live use.
 Continuous Monitoring	Performance, bias, and drift are tracked in production with full audit trails.
 Regulatory Alignment	A clear, governed path from research use toward clinical clearance: models remain "For Research Use Only" until that path is complete.
 Feedback Loops	Clinical usage feeds back improving both the models and the integration over time.

outputs, monitors for drift as scanners, protocols, and patient populations change, and raises alerts when performance or data characteristics move outside predefined thresholds.

Subgroup performance is monitored separately so that strong overall performance does not obscure underperformance in particular patient populations. Model inferences and versions are logged and traceable, providing an auditable record for clinical governance and ongoing oversight.

The regulatory pathway, stated plainly

The three models currently operate

under "For Research Use Only" status. The path toward clinical clearance follows established frameworks for software as a medical device, and nothing in this document should be read as a clinical or regulatory claim about the current suite. The partnership provides a governed pathway from research use toward clinical use, supported by local validation, monitoring, documentation, and regulatory alignment. FDA clearances achieved within contributing research programs provide relevant precedent for that pathway, while the current models remain under "For Research Use Only" status.

10

CONCLUSION

FROM PROVEN SCIENCE TO EVERYDAY CARE

The potential clinical value of imaging AI is increasingly supported by research; the remaining challenge is translating validated models into safe, integrated clinical workflows. What separates a promising algorithm from a working clinical service is execution: local validation, deep workflow integration, and production-grade governance, delivered on a backbone that makes each additional model cheaper than the last.

Innova's enterprise engineering and its Academic partners' clinical and scientific rigor turn research-validated models into governed, integrated, monitored care, unlocking new value from imaging assets providers already own. The

approach is designed to be repeatable: validate one model on local data, establish the integration and governance foundation, and extend that foundation across additional models, service lines, and modalities. The marketplace widens that same path — as the network of contributing research partners grows, each new validated model reaches providers over infrastructure that is already in place. The next natural step is a focused engagement using your own retrospective data. This, inevitably, offers a practical way to move from a compelling thesis to evidence from your own patient population.

IMPORTANT NOTICE

The cardiovascular-aging, liver-radiomics, and liver-fat quantification models described in this paper are investigational and For Research Use Only. They are not cleared or approved for diagnostic use, and nothing here constitutes a clinical claim. Any clinical use follows local validation and the appropriate regulatory pathway.





innova[®]
SOLUTIONS

www.innovasolutions.com