Every Curie Counts: Closing the Loop in Radiopharmaceuticals

How AI can turn a scarce, decaying atom into treated patients — a white paper from a collaboration with an emerging biopharma company

Oct 1, 2026 · @Shah Nawaz

A note before you read. As I was transitioning from To64, I had the opportunity to collaborate with several emerging biopharma companies on their technology and AI strategy. This paper is one piece of that work. It was built for a specific use case, a company in targeted radiopharmaceuticals, but the thinking and the approach are universal. Company-specific details have been removed or generalized. I'm sharing it as an invitation: to compare notes, to challenge the ideas, and to explore where this way of thinking applies next.

The story in brief

Most AI strategies in pharma start with the molecule. This one starts with an atom that is disappearing while you watch it.

Targeted alpha therapy is one of the most promising modalities in oncology. A small targeting molecule carries a radioactive isotope directly to the tumor, and the isotope destroys cancer cells from within. The science is extraordinary. The operations are brutal. The isotope is among the scarcest materials on earth. It decays from the moment it is made, and every dose must be manufactured, released, shipped and infused inside a window of a few days, for a specific patient, at a specific site, on a specific day.

The work started with a simple question: in a business like this, what should technology and AI actually be *for*? Its central answer applies well beyond one company: **in radiopharmaceuticals, the competitive advantage is not the molecule, and not even access to supply. It is how efficiently you convert scarce supply into patient outcomes.** That conversion is a control problem, and control problems are what AI is good at.

Today the industry runs this business on human coordination: supply planners, radiopharmacy coordinators and site schedulers working across email and spreadsheets against a clock they cannot stop. The alternative is a closed-loop control system. It sees the patient, computes the dose actually delivered, allocates every curie, automates

manufacture and release, and learns from every outcome, continuously and at machine speed. Done well, it can add roughly half again as many treated patients from the same isotope supply.

The constraint: what a curie actually buys

To understand the strategy, start with the physics. Actinium-225 is the workhorse isotope of targeted alpha therapy, and the whole world's annual supply weighs about half a grain of salt.

Measure	Value	What it means
One curie	37 GBq, or 37 billion decays per second	A curie is a rate, not a quantity. You never *have* a curie; you have one right now, and less of one tomorrow
Mass of one curie of Ac-225	About 17 micrograms	The isotope is almost weightless
Global annual supply	About 34 micrograms	Roughly two curies, gating every alpha-therapy program across the largest companies in the field combined
One patient course	About 40.8 MBq	Sets the theoretical ceiling at around 900 courses per curie
Global courses, theoretical ceiling	About 1,800 a year	Assumes perfect conversion, which nobody achieves
Eligible patients across major indications (US)	Tens of thousands	Demand exceeds the ceiling many times over

At perfect conversion, one curie treats about 900 patients. Nobody converts perfectly. **The gap between what a company buys and what reaches a patient is the entire strategy.** Supply is not the opportunity. Yield is.

This reframes what technology is for. In most of pharma, AI is pitched as a way to discover better molecules or to run trials faster. Here, the most valuable thing software can do is make every atom go further. A software advantage converts directly into a supply

advantage: the company that treats the most patients per curie can afford to pay the most per curie, and so it secures more supply than its competitors.

Uneven frontiers: three shifts that redefine advantage

Three changes are reshaping radiopharmaceuticals at different speeds. Each one moves the source of competitive advantage, and each one points to a capability the winners will need.

The shift	What is happening	The capability that wins
The molecule stops being the moat	AI-driven ligand design is commoditizing binders. Within a few years, credible binders for today's leading tumor targets will exist for anyone with a chemistry group	**Best at converting supply into outcomes.** The durable advantage is the system that turns a decaying atom into a lasting response, and it carries over to every future asset
Supply is the constraint, until it isn't	Global Ac-225 supply covers roughly 2,000 patient courses a year against tens of thousands of eligible patients. New production capacity is arriving from several suppliers	**Best at allocation.** Once access becomes table stakes, efficiency becomes the differentiator: the highest patients-treated-per-curie in the industry
Dosing moves from fixed to personalized	Europe already requires therapeutic radiation to be individually planned and verified. The field is moving from "how much did we inject?" to "how much did the tumor absorb?"	**Best at knowing what was delivered.** Being the only company that can tell a regulator, a payer and a physician exactly how much dose reached the tumor, in every patient and every cycle

The payload is becoming a commodity. The platform is not. And because personalized dosing will eventually be written into labels and regulations, whoever arrives first sets the standard everyone else must meet.

The leaky chain: where the curies go

Seven stages sit between a curie procured and a response delivered. Every one of them leaks, and nobody in the industry measures the whole chain.

Of every curie bought, about a quarter produces a response

Share of procured activity surviving each stage; the small label names the layer that recovers the loss

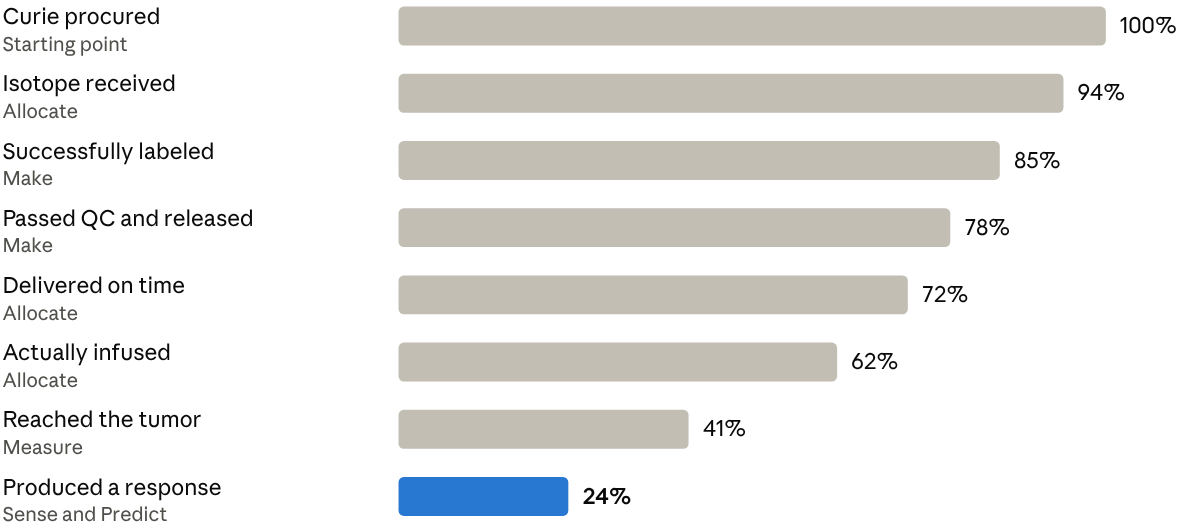

Illustrative stage values. The first phase of the program instruments the real chain.

illustrative stage values from the source strategy · seven stages, curie to response

Each stage loss has an owner in the system described below. Manufacturing intelligence attacks the labeling and release losses. The allocation engine attacks receipt, delivery and infusion. Dosimetry reveals how much actually reached the tumor, and the imaging and prediction layers explain why some doses do not produce a response.

The economics are striking. Raising end-to-end yield from 50% to 75% means **half again as many patients treated from the same supply, with zero additional isotope.** For a company with even a few curies a year, that is on the order of a thousand additional patient courses annually. No discovery program delivers that kind of return on that timeline.

The architecture: one closed loop

The answer to a leaky chain is not six separate tools. It is a single control loop. Each stage of the chain gets a purpose-built AI system, and all of them read from and write to one permanent record of every curie.

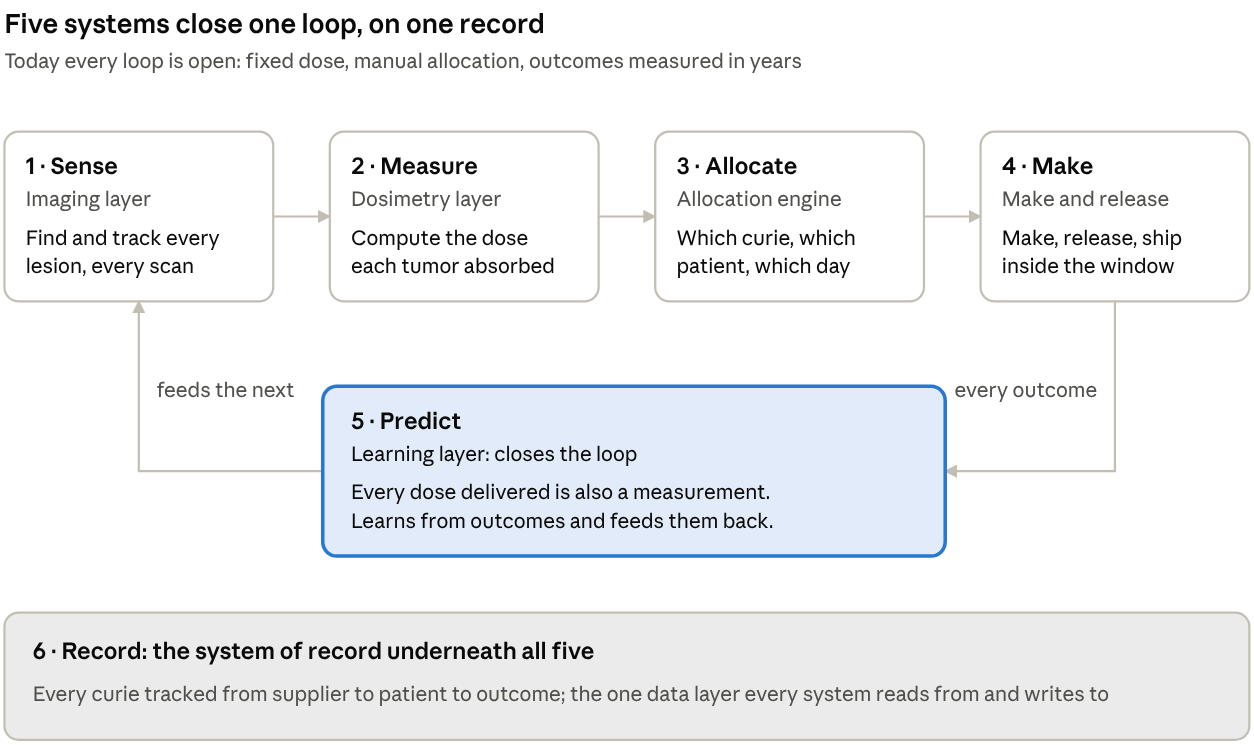

the closed loop · five systems on one system of record

The loop has four forward steps. **Sense** finds and measures every lesion. **Measure** computes the radiation actually absorbed. **Allocate** decides which curie goes to which patient on which day. **Make** manufactures, releases and ships the dose inside its window. The fifth system, **Predict**, closes the loop, because in this modality every dose delivered is also a measurement. It learns how delivered dose relates to outcome and feeds that back into the next scan, the next allocation and the next dose. **Record**, the system of record, sits underneath all five.

Today every one of these loops is open across the industry: doses are fixed, allocation is manual, and outcomes take years to measure. Closing them is the strategy.

Inside the six capabilities

Each capability is a distinct AI system with its own inputs, models and outputs. The numbers below are illustrative, modeled on published dosimetry and on how such a network typically runs. Replacing them with measured actuals is the first deliverable of the program.

1. Sense: every scan becomes a structured map of disease

The imaging layer reads every diagnostic PET, every CT or MR and every post-therapy SPECT scan. It returns a lesion-level record in which each tumor keeps its identity across months, modalities and scanners.

- **Inputs:** diagnostic receptor PET, anatomical CT/MR, and post-therapy SPECT.
- **Where AI is applied:** learned reconstruction that recovers signal from photon-starved alpha-therapy SPECT; automated delineation of every receptor-avid lesion; and matching that gives each lesion a persistent identity across modalities and timepoints.
- **Outputs:** a lesion graph (persistent IDs, volume, uptake), whole-body burden metrics, and tumor-to-kidney and tumor-to-marrow ratios.

Why this matters: standard response criteria track a handful of target lesions and return one verdict per patient. Track every lesion and a different picture appears. In one illustrative patient followed for four cycles at the same dose, eight lesions behaved three ways:

Lesion	Site	Change in uptake	Call
L01	Soft tissue	−76%	Responding
L02	Soft tissue	−67%	Responding
L03	Lymph node	−44%	Responding
L04	Soft tissue	−26%	Stable
L05	Lymph node	−14%	Stable
L06	Bone	+35%	Progressing
L07	Bone	+51%	Progressing
L08	Lung	+29%	Progressing

So there are three responding, two stable and three progressing lesions, all in one patient. Lesion-level tracking is what makes dose–response learnable at all.

2. Measure: from activity injected to dose absorbed

Activity is an input the company controls. Absorbed dose, the radiation that actually reaches the tumor and the organs at risk, is an output the industry does not routinely measure. The dosimetry layer closes that gap in under a minute, for every patient in every

cycle.

- **Where AI is applied:** a Monte Carlo surrogate model, where seconds of inference replace hours of simulation; decay-chain modeling that tracks daughter isotopes as they redistribute after injection; and single-timepoint dose reconstruction, which needs one scan plus a learned prior instead of four imaging visits.
- **Outputs:** a 3D voxel dose map, absorbed dose in Gy for every tracked lesion, and cumulative kidney and marrow dose across cycles.

The finding that justifies the whole program: **identical dose in, wildly different dose delivered.** In an illustrative cohort of 32 patients, each given the same fixed activity every cycle, absorbed tumor dose varied about **seven-fold**. Nine of the 32 fell below the assumed therapeutic threshold, and four sat at or above the kidney constraint. Under fixed dosing, none of this is visible today.

3. Predict: learning what a delivered dose actually buys

With lesion identity from the imaging layer and delivered dose from the dosimetry layer, it becomes possible to do something no oncology company has done routinely: **fit a dose–response curve from routine care.**

- **Where AI is applied:** a hierarchical lesion-control model (lesions nested within patients); an organ-toxicity model for cumulative kidney and marrow dose; and causal inference that treats operational noise (shipping delays, decay differences between batches) as a natural experiment in dose.
- **Outputs:** probability of control for each lesion at the delivered dose, probability of toxicity for each organ, and a continue, stop or re-treat recommendation. Every prediction is locked before each readout and scored afterwards.

The data advantage comes from scale at the lesion level. A trial of a few hundred patients, with roughly 20 lesions each, yields thousands of dose–outcome observations. That is how a small trial produces a large learning signal.

This curve decides everything. In the illustrative version, half of lesions are controlled at about 27 Gy absorbed. Below 15 Gy, control is under 20%, which means those lesions are effectively untreated. Above 40 Gy, control exceeds 80%, yet 9 of 32 patients never reach that level under fixed dosing. The lever is to move dose, not activity: the same curie, better placed.

The strategy is honest about this risk. If the curve turns out to be flat on retrospective data, personalized dosing is dropped and all investment goes into allocation. Reaching that decision costs **a small fraction of a late-stage program**, because it uses data the company already has.

4. Allocate: air traffic control for a decaying inventory

Think of air traffic control: scarce landing slots, aircraft that cannot circle forever, and weather that keeps changing. The allocation layer continuously decides which curie goes to which patient, at which site, on which day, and it prices the opportunity cost of every decision.

- **Inputs:** a supply forecast by isotope lot, manufacturing and QC slots, courier windows, site capacity, and organ-dose limits from the dosimetry layer.
- **Where AI is applied:** probabilistic demand forecasting that turns referral signals into expected patient arrivals; a rolling-horizon stochastic optimizer that re-solves the whole network every few hours; and a learned value function that estimates what a curie is worth if it is held for later.
- **Outputs:** a daily schedule (who is dosed, where and when), the prescribed activity for each patient within specification, and a batch and routing plan.

The binding constraint is time. Ac-225 has a half-life of 9.92 days, but a labeled batch is usable only while it stays within its activity specification, a window of roughly 72 hours. Every upstream step eats into that window. When a patient cancels, the system reassigns the dose to another patient within the same window, so no activity is discarded, and it learns from every disruption.

5. Make: every hour removed is product recovered

Decay cannot be slowed. Latency everywhere else can be removed: in the recipe, in the reaction, and above all in the wait for release testing.

- **Where AI is applied:** Bayesian optimization that learns the formulation recipe (concentration, radioprotectant, reaction time) from a small number of experiments; a release soft-sensor that predicts radiochemical purity from in-process signals before the assay finishes; and supplier-quality models, because each production route for the isotope behaves differently.
- **The human stays in charge:** the model prepares the release decision, and a qualified person approves it.

Stage	Today	With manufacturing intelligence
Elution and purification	6 h	6 h
Labeling reaction	4 h	3 h
QC and release testing	14 h	4 h
Transit to site	10 h	8 h
Site preparation to infusion	6 h	5 h
Total, and activity remaining	**40 h · 89%**	**26 h · 93%**

Removing 16 hours of process latency recovers about 5% of every batch and widens the window the allocation layer has to work with. There is one important caveat. Predicting a release assay is a multi-year analytical validation project, not a three-month modeling exercise, and that is exactly why it has to start now.

6. Record: one curie, end to end, permanently

The system of record tracks every curie from supplier to patient to outcome. It is the single governed data layer that all the other systems read from and write back to. Operators (radiochemists, medical physicists, clinical operations, supply planners and site coordinators) and their AI agents work through it. It connects to supplier systems, manufacturing execution, LIMS, EDC, imaging archives, courier APIs, ERP and collaboration tools.

Here is what one illustrative lot looks like in that ledger:

Time	Event	Activity remaining	Written by
T+00:00	Isotope lot received and assayed	100.0%	Record
T+06:12	Purification complete	98.2%	Make
T+10:40	Labeled (radiochemical purity 99.1%, yield 94%)	96.9%	Make
T+14:05	Release predicted, then confirmed by assay	95.9%	Make
T+24:30	Shipped, calibrated for T+62h	92.9%	Allocate

Time	Event	Activity remaining	Written by
T+61:50	Infused at the prescribed activity	83.7%	Allocate
T+85:00	Post-therapy SPECT acquired	—	Sense
T+96:00	Absorbed dose: 34.1 Gy tumor, 11.2 Gy kidney	—	Measure
Cycle 4	Partial response confirmed at the lesion level	—	Predict

For the first time, the company can state what share of a lot reached a patient: 83.7% in this example. A 34.1 Gy dose is now linked to a response, and repeating that thousands of times gives the prediction layer its curve. Each row is written automatically by the system that produced it, with no manual reconciliation and no lost provenance.

This record exists nowhere else in the industry, and it cannot be reconstructed later. A competitor that starts two years from now cannot recover the data from the years it missed.

Why the advantage compounds

Traditional radiopharma treats each dose as an independent event. But a radioligand has an unusual property: it images itself. Every administration is also a quantitative record of where the drug went, which is the closest thing oncology has to a direct measurement of exposure. **Every treatment is also an experiment, and every experiment makes the system better.**

The flywheel works like this:

1. Better allocation means
2. more patients treated per curie, which means
3. more economic value per curie, which means
4. the company can outbid others for scarce supply, which means
5. more curies and more patients treated, which generates
6. more dose-and-outcome data, which produces better allocation.

Each layer improves as the loop turns. The imaging layer gets better at finding and tracking disease with every patient scanned. The dosimetry layer learns what is actually delivered, not what was intended. The allocation layer learns site reliability, courier performance and no-show risk from every shipment.

Operating data no one else has

The most defensible asset in this strategy is not a model. It is operating data that does not exist outside the company running the network: yield curves, courier reliability, site no-show rates, supplier quality by production route, and lesion-level dose–response. None of it can be bought.

A competitor can license a binder and build a plant. It cannot buy ten years of data on how to turn a decaying atom into a response.

The operating model: from coordination to control

Today, dose operations in this industry look like a coordination organization. A supply planning lead, radiopharmacy coordinators, a courier desk, batch schedulers, a site coordinator for each region, plus clinical operations, supply analysts, QA release, logistics and scheduling: roughly fifteen or more people working reactively, with about two weeks of visibility, coordinating by email and spreadsheet.

	Coordination model (today)	Control model (proposed)
How decisions are made	Manually, by people negotiating across functions	By one optimization engine running continuously across the whole network, with human override
Visibility	About two weeks	Continuous, network-wide
When a patient cancels or a shipment slips	Escalation by email; activity often wasted	Automatic reallocation within the same window
Learning	Lives in individual heads; lost when people leave	Captured from every disruption; the system improves itself
Staffing	15+ people in reactive coordination roles	Smaller teams supervising, handling exceptions and improving the system

This is not about removing people from a regulated, safety-critical process. Radiochemists, physicists, QA and clinical teams remain accountable for every release and every dose. What changes is where their time goes: away from reconciling

spreadsheets against a decaying clock, and toward the decisions only humans should make.

The same governed data layer also makes everyone a builder. Operators get self-service tools for analysis and automation. Agents mine radiology reports for referral signals and draft documentation and batch records. Non-engineers can build their own workflows on governed connections, under one consistent set of controls.

The roadmap: staged, falsifiable, metric-driven

The plan is built so that the biggest scientific risk can be tested in the first 18 months, before any large commitment, and so that the largest near-term prize does not depend on that risk at all.

Phase	Months	What happens	Gate to pass
0 · Instrument	0–6	Build the system of record and measure the curie waterfall end to end for the first time	Deliver the number no one in the industry has: true end-to-end yield
1 · Validate	6–18	Run the imaging, dosimetry and prediction layers on retrospective trial data. Run allocation in shadow mode alongside the human schedule	Is there a real dose–response signal? If not, stop personalized dosing and keep allocation
2 · Operate	18–36	Allocation takes live scheduling authority. Dosimetry is deployed as clinical decision support	+25 points of end-to-end yield; above 95% ship-to-infuse
3 · Differentiate	36–60	Run a personalized-dose trial against the fixed-dose regimen	A differentiated label, at which point the moat becomes regulatory

The investment: a focused cross-functional team of a few dozen people (data scientists, medical physicists, radiochemistry and operations experts, and engineers), with GxP validation staffed from day one. Retrofitting compliance onto research code is where pharma AI programs die.

Two design choices make this plan credible to a board.

- **It is falsifiable.** Each phase has a measurable gate and an explicit kill condition. The most expensive scientific bet, personalized dosing, gets an early and cheap test.
- **The floor is high.** Even if the dose–response curve turns out to be flat, the allocation and manufacturing gains alone recover most of the value. Raising end-to-end yield from 50% to 75% means half again as many patients treated, with no additional isotope.

Lessons that travel beyond radiopharma

This story is specific to a decaying isotope, but the principles carry over to any company where a scarce input meets a complex delivery chain: cell and gene therapy, personalized vaccines, rare-disease supply, even hospital operations.

1. **Find the conversion problem.** Most AI strategies optimize the asset. The larger prize is often the yield between what you have and what reaches the patient.
2. **Instrument before you optimize.** Phase 0 here delivers no model, only a measurement. Without it, every later claim is a guess.
3. **Measure the output, not just the input.** Moving from "activity injected" to "dose absorbed" changed the whole problem. Every field has a similar gap between what is controlled and what is achieved.
4. **Learn at the finest grain the data allows.** Lesion-level data turned a few hundred patients into thousands of observations.
5. **Design the kill switch first.** A strategy that names the condition under which it stops is one leadership can fund with confidence.
6. **Make the system of record the moat.** Data that is generated only by running the operation, and cannot be reconstructed later, is the one advantage a competitor cannot buy.
7. **Validate from day one.** In a regulated setting, compliance is architecture, not paperwork.

Closing

Fixed dosing for every patient is the last artifact of a pre-AI era in this modality. The isotope will always be scarce, and it will always be decaying. What can change is how much of it becomes a treated patient, and how precisely each patient's dose is known.

Whoever closes that loop first will not just treat more patients with the same supply. They will set the standard that the rest of the field has to meet.

An invitation

This paper started as one company's strategy. I'm sharing it because the questions behind it travel well. Where does value leak between what we have and what reaches the patient? What do we control, and what do we actually achieve? What would it take for the system to learn from every outcome?

If you are building in a field where a scarce or perishable input meets a complex path to the patient (radiopharmaceuticals, cell and gene therapy, personalized medicines), or you are working out how to make AI compound rather than pile up as pilots, I'd welcome the conversation. The specifics will differ. The approach carries over.

About this paper: it draws on strategy work done in collaboration with an emerging biopharma company in targeted radiopharmaceuticals. Company, partner, supplier, trial and program details are omitted or generalized. Figures marked illustrative are modeled on published dosimetry and typical network operations; they show structure and scale, not reported results.