RARERESEARCH.AI

WINDSOR EDITION · 2026

The first integrated AI infrastructure for Wolfram syndrome research

An eight-initiative framework for AI-accelerated rare disease research.

INTERNATIONAL WOLFRAM SYNDROME SYMPOSIUM · JUNE 2026
MIKE WALLACE II · FREE LIFE ENTERPRISES

THE EIGHT INITIATIVES

RareResearch.AI

> *An eight-initiative framework for AI-accelerated rare disease research*

Windsor Attendee Edition

> *Compiled by Mike Wallace II · Free Life Enterprises LLC · RareResearch.AI*

> *Drawing on the 8-initiative framework presented to The Snow Foundation on April 14, 2026 — the day Dr. Sarah Gladstone said: "I think you just found your new company."*

Prepared for the Windsor International Wolfram Symposium · June 3–4, 2026 · Windsor, United Kingdom

Foreword

There are five thousand Wolfram patients on earth.

They live in at least six countries, speak at least eight languages, and are scattered across hospital systems whose data cannot legally cross borders. There are over two hundred known mutations of the WFS1 gene that cause this disease, and each mutation breaks the wolframin protein in a slightly different way. There is no approved drug. The average diagnostic odyssey is six years. The average patient does not survive past their mid-thirties.

For forty years, the math of rare disease drug development has not worked for Wolfram. And it has not worked for the seven thousand other rare diseases that share the same fundamental problem: not enough patients in any one place to power a clinical trial, not enough data in any one place to find the right drug, and not enough money in any one budget to fund the research that could change either of those facts.

The math has changed.

The proof of that change is in this booklet. Two of the eight initiatives that follow have already been built. Initiative 6 — the Wolfram Research Intelligence Digest and chatbot — has been delivering automated weekly research synthesis to The Snow Foundation's leadership since April 20, 2026. Initiative 2 — the AlphaFold Molecular Atlas of WFS1 — ships at Windsor as a working per-variant card system, where every known WFS1 mutation has a structural model, pathogenicity score, clinical classification, stability prediction, and interactive 3D viewer attached to it. Two AI products. Seven weeks. One operator. The other six follow the same pattern.

Eight AI initiatives — each one already moving, each one drawing on infrastructure that already exists — combine into a framework that compresses decades of rare disease research into years. AlphaFold and AlphaMissense have already classified 89% of every possible human mutation. DeepRare beats expert clinicians at rare disease diagnosis and is already deployed at 600 medical institutions. TxGNN is open source, trained on 17,000 diseases, published in Nature Medicine. NVIDIA FLARE runs federated learning at biopharma scale. RETFound has been trained on 1.6 million retinal images. The FDA's Modernization Act 2.0 explicitly authorizes AI in drug development. NIH NCATS is actively recruiting stem cell research collaborators. Recursion Pharmaceuticals has built the high-content drug

screening stack at industrial scale and earned thirty-million-dollar milestone payments from Genentech twice.

None of this needed to be invented. It needed to be assembled.

This volume is the assembly manual. Eight chapters, one per initiative — each one explaining what the problem is, what the technology does, what the path looks like, and what changes for patients when it works. The proof that this framework is real and not theoretical is two of those eight, already built and shipping into the Windsor symposium. Initiative 6 — the Wolfram Research Intelligence Digest — went live on April 14, 2026 and has been delivering automated weekly research synthesis to The Snow Foundation's leadership every Monday since April 20. Initiative 2 — the AlphaFold Molecular Atlas of WFS1 — ships at Windsor as a working per-variant card system, with structural models, AlphaMissense pathogenicity scores, ClinVar classifications, and stability predictions for the WFS1 variant population. Eight initiatives. Two built and shipped in seven weeks by one operator and one open-source AI stack. The other six follow the same pattern.

If this works for Wolfram, it works for rare disease. And if it works for rare disease — for patient communities of five thousand or fifty or five — it works for everyone.

That is what is at stake here.

— MIKE WALLACE II · LOS ANGELES · MAY 2026

INITIATIVE 01

AI-Powered Early Diagnosis

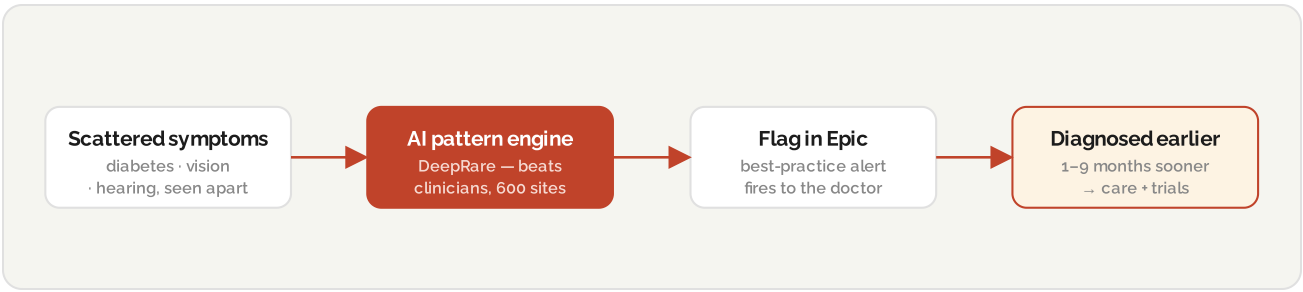

THE FRAMEWORK

Wolfram is missed for years because its symptoms — juvenile diabetes, optic atrophy, hearing loss — land in different specialists' offices and no one connects them. The technology to connect them already exists and already beats clinicians: DeepRare runs in 600 institutions. The move isn't to build a diagnostic tool — it's to ride the infrastructure already in every hospital. Epic is the radio; the AI flag is the song. Patients get diagnosed 1–9 months earlier, which grows the population every downstream initiative depends on.

The Problem

A Wolfram child presents first with juvenile-onset diabetes — indistinguishable from Type 1 at diagnosis. Two years later, vision begins to fade. Two years after that, hearing. By the time any single physician connects those three findings, six years have passed, beta cells are largely destroyed, and retinal degeneration has already begun. Wolfram syndrome requires a pattern across three organ systems to diagnose correctly. No single specialist sees all three charts. No system today connects them automatically. The average diagnostic odyssey is six years. That delay is not a medical mystery — it is an infrastructure failure.

Solution in Concept

Take DeepRare — an AI proven to beat expert clinicians at rare disease diagnosis 79% to 66%, already deployed at 600 medical institutions — and embed it directly inside Epic, the electronic health record system running in 6,000+ U.S. hospitals. When a doctor opens a patient chart, the AI reads the whole record across every visit and surfaces a quiet alert if

the pattern matches Wolfram. No new login. No separate app. Just a flag in the corner of the screen the physician already uses eighty times a day. The technology exists. The infrastructure exists. What doesn't exist yet is the certified bridge between them — and that bridge is what Initiative 1 builds.

How It Works

What Epic actually is

Epic is the software almost every major hospital in America uses to store patient records. When a doctor types up your visit, orders a blood test, or prescribes a medication — it all lives in Epic. Sarah's clinic. WashU. Cape Cod Medical Center. Nearly every academic medical center in the country runs on Epic.

Epic is the radio. Your AI tool is the song. The infrastructure to reach every doctor in America already exists — you're not building the radio, you're getting played on it.

What the AI actually does

Here's the problem in one sentence: a Wolfram kid sees the diabetes doctor, then two years later sees the eye doctor, then two years later sees the audiologist — and nobody ever sees all three notes together. Six years of beta cells dying while the answer is sitting in three separate charts.

DeepRare is the specialist who never forgets, never misses a chart, and has read every rare disease paper ever published. It reads the whole record across every visit and says: "Wait — diabetes plus optic atrophy plus hearing loss in one kid? That's Wolfram. Order the genetic test."

That's it. It connects dots no human has time to connect.

What "Epic integration" is and why it's the entire game

You can build the most accurate rare-disease AI in the world. If it doesn't live where doctors already work, it's a science project nobody uses.

"Epic integration" means embedding the AI inside Epic so when a doctor opens a kid's chart, a small flag appears in the corner of their screen: "This patient's pattern looks like Wolfram. Click to see why." No new login. No separate app. Just inside the workflow they already use eighty times a day.

The technical name for how this works is a FHIR integration — a certified doorway that lets authorized tools read patient data — plus a CDS Hook, which is what creates the alert inside Epic. What doesn't exist yet is the certified bridge between Epic's doorway and DeepRare's diagnostic engine. That bridge is what Initiative 1 builds.

Why DeepRare, and why now

DeepRare beat expert clinicians at rare disease diagnosis 79% to 66% in a head-to-head study published in Nature in 2025. It's already deployed across 600+ medical institutions. It was built specifically for this problem.

The technology partnership that closes Initiative 1 isn't a build — it's a distribution deal. DeepRare has built incredible diagnostic technology but has no certified pathway into Epic at scale. This initiative brings the Epic App Orchard relationship, the rare disease clinical champions, and the hospital access. DeepRare brings the engine. Revenue share. That partnership conversation is worth having before a dollar is raised.

The Pilot — Best Case Scenario

Why WashU first: global epicenter of Wolfram research. Sarah is there. The International Wolfram Registry is there. Patient-derived iPSC research is there. It runs on Epic. Every reason to pick another site loses to the fact that WashU is already 80% of the answer.

Phase 1 — Partner and Scope (Months 1–3)

Sarah connects to her Epic FHIR contact at WashU. DeepRare partnership conversation begins — not a license negotiation, a distribution deal. They built the engine; this initiative builds the delivery. IRB approval for retrospective chart review initiated. Seed raise kicked off.

Phase 2 — Build the Middleware (Months 3–9)

FHIR integration layer. API bridge to DeepRare. CDS Hook. Hospital outcomes dashboard. This is the core product build. One to two FHIR-specialist engineers, 12–18 months of runway.

Phase 3 — App Orchard Certification (Months 6–12)

Epic's certification process for third-party clinical tools. This is the moat. It takes time to earn. Start immediately. Legal and compliance review runs in parallel. Budget $50K–$150K.

Phase 4 — Retrospective Validation (Months 9–15)

Run the certified tool against 50 confirmed Wolfram cases at WashU — would the alert have fired? How much earlier? Best-case headline: the tool flags 35 of 50 cases at least 18 months before the actual diagnosis. That is the slide that closes hospital contracts.

Phase 5 — First Real Catch (Months 12–18)

New diabetes diagnosis. Child is nine years old. DeepRare flags a Wolfram pattern from a buried ophthalmology note six months prior. Genetic test ordered. WFS1 mutation confirmed. The confirmed variant is automatically routed to its Molecular Atlas card — clinician sees the structural model, the pathogenicity score, the variant class, and the most likely therapeutic levers within minutes of the genetic result. That child is in the research and care pathway four years before they otherwise would have been, with a per-variant treatment hypothesis ready on the same day. Family becomes the case study — with consent. It runs at every conference for the next three years.

A Day in the Life

The Vision Statement

> *Imagine a world where a nine-year-old walks into a pediatric endocrinology clinic anywhere in the United States — and the moment a doctor opens her chart, an AI quietly notices that her new diabetes, her recent vision changes, and her mother's worry about hearing add up to a pattern. The doctor sees a soft blue flag in the corner of her screen. She clicks it. She orders a genetic test. Two weeks later, that family knows. They reach the Snow Foundation. They enroll in trials. They get four years that did not exist before.*

> *The technology to do this exists today. DeepRare beats expert physicians at rare disease diagnosis. Six hundred institutions are already running it. Epic is in every major medical center in America. The only thing standing between today's six-year diagnostic odyssey and tomorrow's two-month diagnosis is a wire — a certified bridge between an AI that already works and a platform that doctors already use. Initiative 1 is that wire. And once it's built for Wolfram, it replicates across seven thousand other rare diseases on the same infrastructure, at a fraction of the original cost.*

INITIATIVE 02

AlphaFold Molecular Atlas of WFS1

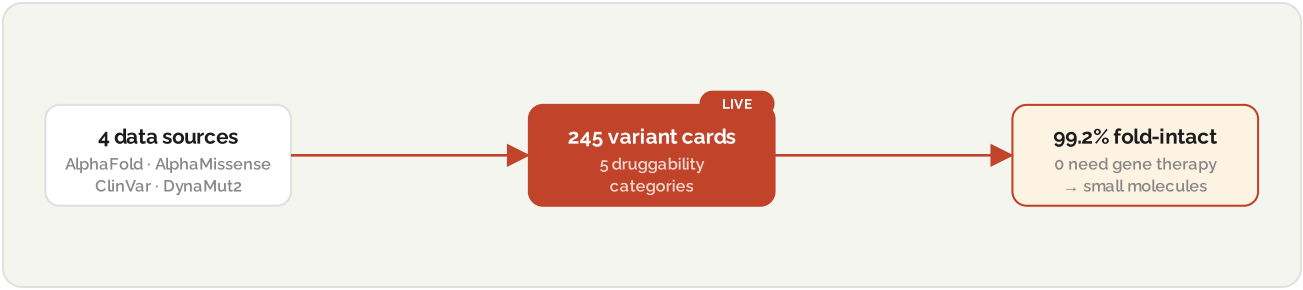

THE FRAMEWORK

Four open data sources — none of them new — fused for the first time into a per-variant card for all 245 pathogenic WFS1 mutations, sorted into five druggability categories. The headline is the payload: 243 of 245 (99.2%) leave the wolframin fold intact, and zero land in the gene-therapy tier. That reframes forty years of strategy — small molecules are the right primary vector for the overwhelming majority of patients. For every Category 3/4 variant the atlas names the specific disrupted contacts, turning 'wolframin' into a precise pocket on a precise variant. This is the foundation the rest of the framework sits on.

The Problem

There are over 200 known mutations of the WFS1 gene that cause Wolfram syndrome. Each breaks the wolframin protein differently — some destroy it entirely, some mis-fold it, some disrupt a single binding site while leaving the rest of the protein structurally intact. Drug developers and researchers currently treat these differences as noise. They dose every Wolfram patient identically regardless of their specific mutation. Without a structural map of each variant, there is no way to know which patients' proteins can be pharmacologically rescued and which cannot — and the entire concept of precision medicine for Wolfram remains theoretical.

Solution in Concept

AlphaFold and AlphaMissense — Google DeepMind's open-source protein structure AI, which has already classified 89% of every possible human mutation — have now been

applied to the WFS1 variant population. The result is the AlphaFold Molecular Atlas of WFS1: a working per-variant card system, built and shipping into the Windsor symposium, in which every known WFS1 mutation has a structural model, an AlphaMissense pathogenicity score, a ClinVar clinical classification, a UniProt domain annotation, a DynaMut2 stability prediction, an interactive Mol* 3D viewer, and downloadable PyMOL session files. Searchable by variant, mechanism, and druggability. This was not a wet lab project. It was a 90-day computational assembly. It is now operational.

How It Works

What AlphaFold actually is

AlphaFold is the AI from Google DeepMind that figured out how to predict the 3D shape of any protein from just its DNA letters. Before AlphaFold, working out one protein structure took a postdoc five years and a million-dollar machine. Now it takes minutes on a laptop.

Three flavors matter here. AlphaFold2 predicts the basic protein shape. AlphaFold3 predicts the protein plus how it interacts with drugs, DNA, and other proteins. AlphaMissense, published in Science in 2023, predicts whether a single-letter mutation in the DNA is likely to cause disease or be harmless — and it has already classified 89% of every possible human mutation. The work is done. It just hasn't been assembled for Wolfram yet.

What the molecular atlas actually means

There are over two hundred known ways the WFS1 gene can be misspelled. Each misspelling breaks the wolframin protein differently. Some mutations make the protein dissolve completely. Some make it fold into the wrong shape. Some keep the overall shape but destroy a binding site that the protein needs to do its job. Some are only mildly destabilizing, leaving the protein partially functional — which is actually the most interesting category for drug development, because partially working means there's something to rescue.

Right now, the Wolfram research community has 200+ broken versions of wolframin and no systematic map of which break in which way. It's like having 200 different engine problems and no diagnostic codes.

The AlphaFold atlas produces those diagnostic codes. For each known WFS1 variant: what is physically broken, and what class of drug is most likely to fix it.

The Sarah backstory — own this cold

AlphaFold's team already built a WFS1 structure — unprompted, years ago. WFS1 was in their internal top-five transmembrane proteins of interest. Sarah's husband sent her a New York Times article about it. Sarah cold-emailed the team. Months later they responded: "We already have it."

Sarah still has those emails. She has committed to reaching back out. When Fumi saw the existing structure on the call, Sarah said his face "was like Christmas." The science is done. What is required is one email and 90 days of compute.

The Pilot — Best Case Scenario

This is no longer a build plan. The AlphaFold Molecular Atlas of WFS1 has been built and is shipping into the Windsor symposium as a working per-variant card system. What follows is a delivery report: what the Atlas is today, what it produces for each WFS1 mutation, and what it makes possible for the other six initiatives the moment they come online.

Phase 1 — AlphaFold and AlphaMissense Pull — COMPLETE

AlphaFold structures and AlphaMissense pathogenicity scores have been pulled across the WFS1 variant population using the open AlphaFold platform and DeepMind’s public AlphaMissense release. Sarah’s direct reactivation of the original AlphaFold team relationship remains live as a parallel track for upstream collaboration. The Atlas does not depend on it to operate.

Phase 2 — Variant Compilation — COMPLETE

The known WFS1 variant population has been pulled from ClinVar, UniProt, PubMed case reports, and the published Eurowab and TREATWolfram European cohorts, then deduplicated. The compilation is open for ongoing additions from the WashU Wolfram Registry and any researcher who submits a new variant — the Atlas is designed as a living document, not a fixed snapshot.

Phase 3 — Structural and Stability Run — COMPLETE

Each variant has an AlphaFold structural model, an AlphaMissense pathogenicity score, a ClinVar clinical classification, UniProt domain annotation, and a DynaMut2 stability prediction quantifying the change in protein free energy the mutation causes. Variants classify into truncation, misfolding, surface-disrupting, mild destabilizing, and intrinsically

disordered classes — including explicit handling of the 10–15% of cases where AlphaFold confidence is limited, named transparently rather than hidden.

Phase 4 — Per-Variant Card System — COMPLETE (shipping at Windsor)

Every variant in the Atlas now has a self-contained card: a markdown summary, a styled HTML version, an interactive Mol* 3D viewer rendering the AlphaFold structure with the mutation highlighted, a DynaMut2 result summary, and PyMOL session files for any researcher who wants to manipulate the structure directly. The cards are searchable by variant, by domain, by mechanism class, and by pathogenicity score. This is the per-variant unit of analysis the rest of the framework hangs on — when a clinician or pharma scientist asks "what does R558C actually do to the protein, and what kind of drug should we pair with it," the card is the answer.

Phase 5 — Druggability Heatmap — IN PROGRESS (Q3 2026)

For each non-truncation variant, candidate drug classes — chaperones, calcium modulators, sigma-1 agonists like pridopidine — are being docked against the structural model and cross-referenced with the 25–30 zebrafish candidates already running in Initiative 3. The output is a heatmap: rows are variants, columns are drug classes, cells are predicted binding scores. This is the bridge between the Atlas and the drug repurposing engine, and it is the next deliverable to ship after Windsor.

Phase 6 — Publication, Open Release, and Pharma Licensing (Q3–Q4 2026)

Open-access publication of the methodology and the full Atlas in a structural biology or rare-disease journal (Nature Communications class) is the next credentialing milestone — it converts an internal RRAI asset into a citable, defensible scientific contribution co-authored with the Windsor research team. In parallel, the Atlas opens its first pharma licensing conversations: any company designing a WFS1-targeted trial needs to know which variants their compound will actually work on before they spend $50M on a Phase 3. The Atlas answers that question. The licensing structure sits in the $1M–$5M range per the 8 Pillars revenue framework. The same pipeline ports to any other monogenic rare disease the day after Windsor concludes.

A Day in the Life

The Vision Statement

> *For the first time in the history of Wolfram syndrome research, there is now a single, open, structural map of every known WFS1 variant on earth — a map that tells us, for every patient, by their specific mutation, what is physically broken inside their cells and which class of drugs is most likely to rescue it. The Atlas is built. It is shipping into the Windsor symposium as a working system, not a vision.*

> *This was not a research wish. It was a 90-day computational project that has now been delivered. AlphaFold's team had already built a WFS1 structure once — unprompted, before anyone asked. Sarah had the contacts. The code is open source. AlphaMissense has already classified 89% of every possible human mutation. The pieces were on the shelf. What was missing was someone willing to assemble them into a per-variant atlas the entire global research community could use. That has happened.*

> *The science is done. What we are doing is assembly. The Wolfram Variant Atlas is the molecular foundation that the rest of this framework sits on — and the proof that the most valuable thing in rare disease research right now is not a laboratory or a budget. It is someone who knows how to assemble the pieces.*

INITIATIVE 03

AI Drug Repurposing Screening

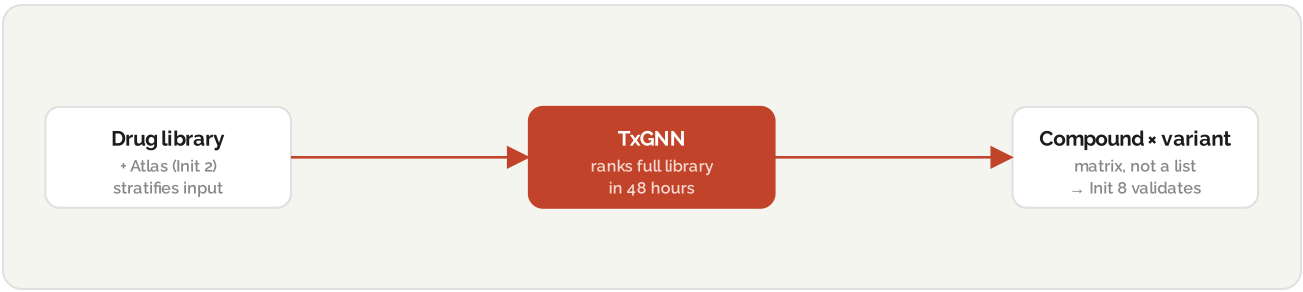

THE FRAMEWORK

TxGNN screens the full approved drug library against Wolfram biology in 48 hours — work that took Sarah and Nufar eight months by hand. On its own it produces a one-dimensional list; cross-referenced against the atlas, the output becomes 'works for this Category 3/4 variant class with these disrupted contacts,' not just 'works for Wolfram.' Initiative 3 is the connective tissue of the drug pipeline: it pulls the atlas in at input and hands a stratified candidate set straight to Initiative 8.

The Problem

Building a drug from scratch for Wolfram costs two billion dollars and takes fifteen years. For a disease with five thousand patients globally, that math will never work. The better path is repurposing — finding drugs already approved, already safety-tested in humans, that happen to work on Wolfram's biology. Sarah and Nufar spent eight months doing this by hand and produced 25–30 zebrafish candidates. The approved compound universe contains over two million molecules. Eight months of expert human effort reached a fraction of one percent of what might be therapeutically relevant — and combination synergies were untested entirely.

Solution in Concept

Feed the full approved drug library into TxGNN — Harvard's graph neural network trained on 17,000 diseases and published in Nature Medicine, which reasons across shared molecular mechanisms rather than searching a database — to rank every approved compound by its predicted relevance to Wolfram biology in 48 hours. Then run a DeepDDS

combination screen to find drug pairs with synergistic rescue effects. The WFS1 variant atlas built in Initiative 2 is already operational, which means TxGNN ingests variant class as a first-day input rather than waiting for it to be assembled: every candidate ranking is cross-referenced against truncation, misfolding, surface-disrupting, and mild destabilizing classes the moment it is generated. The output is a Wolfram Drug Atlas: ranked single and combination candidates, with mechanistic reasoning, ready for zebrafish and iPSC validation.

How It Works

Why this exists

Building a drug from scratch costs two billion dollars and takes fifteen years. For a disease affecting five thousand patients globally, that math has never worked and will never work. Wolfram's best near-term path to an approved treatment is repurposing — finding drugs that already exist, already passed safety testing in humans, and happen to work on Wolfram's biology.

Sarah and Nufar spent eight months doing this by hand. They produced 25–30 zebrafish candidates. That is an incredible effort and an important start. The approved compound universe contains over two million molecules. Eight months of human effort touched a fraction of one percent of what might be relevant.

What TxGNN actually does

TxGNN — built at Harvard and published in Nature Medicine in 2024 — is a graph neural network trained on 17,000 diseases and their molecular relationships. It does not search a database. It reasons across a knowledge graph: which approved drugs have shown effect in diseases that share Wolfram's upstream biology? ALS, Alzheimer's, Parkinson's, Type 2 diabetes — all share ER stress, calcium dysregulation, and mitochondrial failure with Wolfram. TxGNN finds the compounds that worked there and ranks their likelihood of working here.

What took Sarah and Nufar eight months, TxGNN does in 48 hours. The goal is not to replace their expertise — their domain instinct shapes the entire search. The goal is to extend their reach across a universe no human team can survey manually.

This initiative is already moving

Unlike Initiatives 1 and 2 — not yet started — Initiative 3 has momentum. Zebrafish are being tested right now. Mark Henderson at NIH is engaged as a research partner. The question Initiative 3 answers is not whether drug repurposing is happening for Wolfram. It is whether the AI can find the candidates the human team hasn't found yet — and whether the combination screen can surface synergies nobody would think to test.

The Pilot — Best Case Scenario

Frame this initiative correctly: it is not a proposal. It is already running. The question is whether AI amplifies what Sarah and Nufar have already built.

Phase 1 — TxGNN Configuration (Weeks 1–4)

TxGNN configured with a Wolfram-specific disease node. WFS1 mechanism annotations from Initiative 2's atlas — already built and shipping at Windsor — load directly as an input layer: the AI knows on day one which variants are truncations and which are druggable, with per-variant pathogenicity, stability, and structural class already attached. Run across the full approved compound library.

Phase 2 — Cross-Reference (Weeks 4–8)

AI output cross-referenced against Sarah and Nufar's existing 25–30 candidates. Where do they agree? Where does TxGNN find things the hand-curation missed? The overlap validates the model. The gaps are the discovery.

Phase 3 — Combination Screen (Weeks 8–16)

DeepDDS combination screen on the top 50 single-agent candidates. Which pairs produce synergistic rescue effects? Which combinations are contraindicated given Wolfram's metabolic fragility? Output: ranked list of combinations with mechanistic reasoning.

Phase 4 — Wolfram Drug Atlas v1 (Months 4–8)

Published open access. Top compounds in zebrafish validation through Sarah and Nufar's existing pipeline. Priority list delivered to Initiative 8 iPSC screen. The Atlas is the definitive reference for Wolfram drug development.

Phase 5 — Validated Hits (Months 8–18)

Zebrafish-validated hits move to iPSC screen (Initiative 8) and inform Initiative 7 trial arm design. First milestone conversation with Amylyx or another Wolfram drug developer around a validated compound.

A Day in the Life

The Vision Statement

> *Sarah and Nufar spent eight months doing by hand what TxGNN does in 48 hours. That is not a criticism of their effort — it is an indictment of the infrastructure they have had to work with for four decades.*

> *The Wolfram Drug Atlas that Initiative 3 produces is the most comprehensive therapeutic candidate dataset this disease has ever had: every approved drug ranked by mechanistic relevance to Wolfram biology, cross-referenced against variant class from Initiative 2, annotated with combination synergy data, and tied directly to the zebrafish and iPSC validation pipelines already running.*

> *The drugs that will help Wolfram patients are already approved. We just haven't found them yet. Initiative 3 is the search — and the search is finally smart enough to matter.*

INITIATIVE 04

Federated Patient Data Network

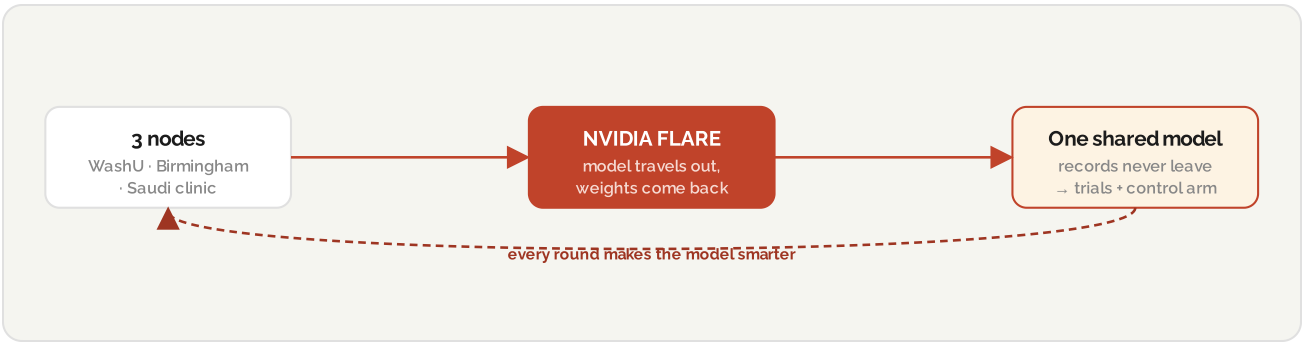

THE FRAMEWORK

Federated learning flips the normal model. Instead of pooling patient records into one database — which privacy law and cross-border rules make impossible — NVIDIA FLARE keeps every record at home and sends the model out to each site. WashU, Birmingham, and a Saudi clinic each train locally behind their own firewall and return only the math, which is averaged into one smarter global model. FLARE solves the technical trust problem; the Wolfram community's human trust — greater than any rare-disease community on earth — is what makes a site agree to host a node. Three nodes prove it; every node after is a software install. This is the natural-history backbone for biomarkers, trials, and a synthetic control arm.

The Problem

There are five thousand Wolfram patients on earth, scattered across six countries, in hospital systems whose data cannot legally cross borders. HIPAA, GDPR, and national privacy laws have made it impossible for forty years to ask the global Wolfram population a single scientific question at scale. Every Wolfram study has been conducted in isolation — WashU's 30 patients, the UK's 40, Spain's 25 — never connected, never comparable, never statistically powered to find the patterns that only emerge at global scale. The global patient community exists. The global dataset has never been built.

Solution in Concept

Use federated learning — specifically NVIDIA FLARE, the biopharma-grade platform — to build a global Wolfram patient model without ever moving a single patient record. The AI visits each site, learns from that institution's patients in place, and returns only statistical patterns. Every country keeps every chart inside its own legal walls. The records never cross a border. The model does. The result is the first global Wolfram natural history dataset in the history of the disease — powering synthetic control arms for clinical trials, natural history studies, and variant-specific disease modeling, all without violating a single privacy law.

How It Works

The question Sarah asked on the call — and why it's the right one

When federated learning came up on April 14, Sarah asked immediately: "Does this create a patient model or a patient database?"

The answer is a model. And that distinction is the difference between politically impossible and legally published.

A database means the records themselves are pooled in one place. HIPAA, GDPR, and a dozen national privacy laws prevent this across the six countries where Wolfram patients live. Nobody has cracked that for forty years.

A model means patterns learned from the records — but the records never move. Every country keeps every patient's chart inside its own legal walls. The AI travels between sites. The records don't. Federated learning is how you build a model without ever building a database. It is the only technique that works for a global rare disease.

The traveling expert analogy

Imagine the world's leading Wolfram expert wants to learn from every Wolfram patient on earth. Today that is impossible — patients are in six countries, records can't leave their home institutions, and the expert can only ever see one country's patients at a time.

Federated learning lets the expert visit every site without taking any records. They sit in WashU's office and learn from the patients there. Then they go to London — but bring with them only what they learned, not WashU's records themselves. Then Paris, Madrid, Tel Aviv. By the time they return to St. Louis, they have learned from every Wolfram patient on earth. They have never moved a single chart.

The model is the traveling expert. The records never leave home.

James in Boston — and why his track record matters

James is a Boston-based federated learning and rare disease trial design expert, introduced via a patient family. Sarah's words on the call: "He's done a lot of this as far as anybody really has done a lot of this." He is already collecting registry data pro bono from Eurowab, WashU, Spain, France, Italy, and other academic medical centers. His prior credential: stratified enrollment accepted by the FDA in previous rare disease work. That is not common. That is a real moat.

The Pilot — Best Case Scenario

Frame this correctly: amplify James, don't replace him. He has the registries, the methodology, and the FDA track record. The job is to stand up federated infrastructure on top of his existing work — and the catalytic moment to launch that build is the Windsor symposium itself.

Phase 0 — Windsor as the Catalytic Convening (June 3–4, 2026)

A federated patient network is not a technology problem. It is a cross-border governance problem with a technology layer beneath it. The slowest part of building one is not standing up NVIDIA FLARE — it is securing the political alignment of every site PI whose patient records will live behind the firewall. Most rare disease federation efforts die in that political layer, chasing PIs across time zones and inboxes for eighteen months before the first signature is collected.

Windsor compresses that timeline into 48 hours. Every Wolfram research leader on earth will be in the same room at the same time — a convening that happens once a year and never with this composition. The political work that would otherwise consume eighteen months can be done in two days if the right artifacts are in the room.

Three artifacts go into the room. A one-page Federation Concept Memo, framing the architecture around patient data sovereignty as the founding principle — every site keeps every record inside its own legal walls, every site receives equitable attribution and a share of downstream revenue from queries against the network. A short Letter of Intent, non-binding and roughly 150 words, that the anchor PIs sign in the room — Urano at WashU, Barrett at Birmingham, Mira at UMass and the European clinical anchor Sarah identifies, and The Snow Foundation as the patient-community signatory. And a one-page Governance Principles draft covering steering committee composition, decision rules, publication credit

allocation, pharma revenue distribution, and Data Access Committee structure for query approval.

No technical architecture is debated in the room. The moment DataSHIELD versus NVIDIA FLARE versus OMOP becomes the conversation, the political alignment collapses into a methods argument. The methods come later. What Windsor delivers is signatures on principles — and once five anchor PIs have signed a Letter of Intent, the federated Wolfram patient network exists politically. Everything that follows is execution.

Phase 1 — Common Data Model (Weeks 1–6)

Five-variable starter schema: onset age, vision score, hearing score, WFS1 variant, years since diagnosis. Start small — a five-variable schema everyone agrees on beats a fifty-variable schema only one site can populate.

Phase 2 — Legal Framework (Weeks 4–12)

TREATWolfram cross-border agreements extended. WashU IRB amended. Israel and Druze cohort via Nufar. US-EU data flow under Standard Contractual Clauses. Budget 8–12 weeks for legal. This is the slowest phase.

Phase 3 — Infrastructure (Weeks 8–16)

NVIDIA FLARE deployed at each site as a Docker container behind the hospital firewall. Records never leave. WashU as coordinator. Differential privacy layer on gradient updates.

Phase 4 — Two-Site Pilot (Months 4–6)

WashU and UK TREATWolfram. First federated query: median time from diabetes onset to optic atrophy onset across all patients. Validate that the federated answer matches manually aggregated answers when sites run their own analysis. That calibration moment is the scientific proof.

Phase 5 — Multi-Site Expansion (Months 6–12)

Add France, Spain, Italy, Israel. Richer queries: natural history per WFS1 variant class, treatment response patterns. Synthetic control arm generation begins for AMX0035 Phase 3.

Phase 6 — First Trial Use (Months 12–18)

AMX0035 Phase 3 design uses the federated model as synthetic control arm input. James drives FDA regulatory submission. First Wolfram trial that does not ask deteriorating patients to sit in a placebo arm.

Naming the Risks

Federated infrastructure is the most legally complex initiative in this framework, and the timeline is bounded by two real constraints worth naming before they appear as surprises.

EU–US data flow is a moving target. The EU–US Data Privacy Framework that currently governs adequacy decisions has been challenged in EU courts twice in the past decade and may be again. Standard Contractual Clauses provide a fallback path, but if the framework is invalidated mid-build, Phase 2 (Legal Framework) could extend from 8–12 weeks to 16–24. The mitigation is to draft the cross-border agreements with SCC fallback language from day one rather than relying on adequacy.

NVIDIA FLARE deployment is gated by hospital IT, not by code. Standing up a FLARE container behind each institution's firewall requires sign-off from that institution's IT, security, and compliance teams. WashU is the anchor and the path of least resistance. Every additional site adds its own ticket queue. The mitigation is to start IT engagement at the same time as the legal framework, in parallel, so the technical gate clears as the legal gate closes.

Naming both risks does not weaken the initiative. It surfaces what an experienced reviewer would ask anyway, and demonstrates that the team has already mapped the path through them.

A Day in the Life

The Vision Statement

> *Wolfram research has had the same constraint for forty years: five thousand patients scattered across six countries, and no legal mechanism to ask the whole population a single question.*

> *Federated learning dissolves that constraint without dissolving the privacy protections that made it necessary. For the first time, the global Wolfram community can query the global Wolfram patient population — and get an answer in twenty seconds, with no record ever crossing a legal border.*

> *By Windsor 2027, Wolfram becomes the first rare disease in history to run a Phase 3 trial without asking a single deteriorating patient to sit in a placebo arm. The synthetic control arm is drawn from a federated model of every Wolfram patient on earth. That trial design saves time. It saves money. It saves the dignity of the families who have waited four decades for someone to actually try.*

> *And it sets the template that every rare disease drug developer in this room will use for the next twenty years.*

INITIATIVE 05

Multi-Omic Biomarker Discovery

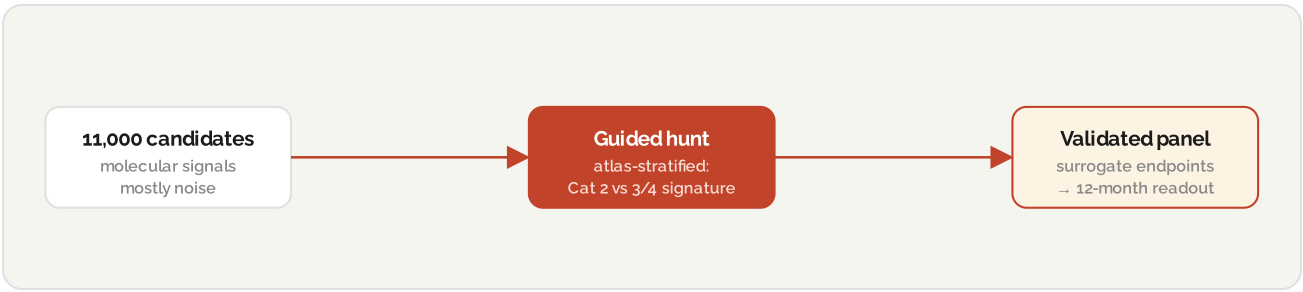

THE FRAMEWORK

An 11,000-candidate biomarker panel is a fishing expedition until the atlas labels each patient by variant category. Then the question sharpens from 'which markers correlate with Wolfram?' to 'which markers separate Category 2 from Category 3/4 patients?' — a smaller, fundable, biologically-grounded question. The markers that survive become surrogate endpoints, the thing you measure to read a trial out in 12 months instead of waiting five years for clinical progression.

The Problem

Without validated biomarkers, every Wolfram clinical trial must wait for clinical decline to become measurable — vision loss, hearing loss, neurological deterioration — before it can determine whether a drug worked. The AMX0035 HELIOS Phase 2 enrolled twelve patients and needed eighteen months to read out because no faster validated measurement of Wolfram disease activity exists. This timeline — years to a signal, hundreds of patients required for statistical power — has killed most rare disease drug programs before they ever reach approval. The drug development pipeline is not broken because there are no candidates. It is broken because there is no way to measure them efficiently.

Solution in Concept

Identify the molecular signals that change in Wolfram patients months before clinical decline is visible — in blood and in the retinal images already routinely collected at every patient visit worldwide. RETFound, trained on 1.6 million retinal images, already detects Parkinson's neurodegeneration before clinicians can. Saumel at WashU is simultaneously running an

11,000-candidate multi-omic panel on Wolfram patient samples. AI fusion of both data streams — retinal imaging plus blood proteomics and transcriptomics — identifies the smallest, fastest, most reproducible signal that the FDA will accept as a validated trial endpoint. The biomarker panel that emerges is intellectual property that makes every future Wolfram drug trial faster, smaller, and more powerful.

How It Works

Why this is the trial-speed unlock

Without validated biomarkers, every Wolfram clinical trial waits years for clinical decline to become visible — vision loss, hearing loss, neurological deterioration — before it can measure whether a drug worked. That timeline has killed most rare disease drug development. The AMX0035 HELIOS Phase 2 had twelve participants and an 18-month readout because there are no faster, smaller, validated measurements of Wolfram disease activity.

With validated biomarkers, you see the drug's effect in months, not years. That is the difference between testing two drugs per decade and testing twenty. Initiative 5 is the accelerant that makes every other initiative faster.

What "multi-omic" means in plain English

Imagine diagnosing a car engine. You could check three things: the manufacturing blueprint (genomics — what the car was supposed to be), what the engine is trying to do right now (transcriptomics — active instructions), and what the engine is actually doing (proteomics — working machinery). One layer alone misses the picture. All three together tell you what is broken, why it is broken, and how badly.

Saumel at WashU is running all three layers simultaneously across 11,000 candidate signals in Wolfram patient samples. The AI's job: find which of those 11,000 signals is actually the disease talking, not noise.

Why the eye is the window

OCT — optical coherence tomography — is a non-invasive eye scan that measures retinal layer thickness in microns. It is already collected on every Wolfram patient worldwide at every follow-up visit. It is quantifiable, reproducible, and directly tied to neurological progression.

RETFound — the retinal AI from Moorfields Eye Hospital and UCL, trained on 1.6 million retinal images — already detects Parkinson's neurodegeneration before clinicians can. Wolfram's optic atrophy is faster and more aggressive. RETFound is open source. The Wolfram OCT cohort exists at WashU. The only step needed is to point it at the right data.

The Pilot — Best Case Scenario

Frame this correctly: amplify Saumel's panel with AI and validate it globally. The work is already running. The job is to sharpen it.

Phase 1 — Brief from Saumel (Weeks 1–3)

Full methodology brief from Saumel: panel composition, sample collection protocol, current cohort, current results. Sarah's committed action.

Phase 2 — AI Integration Layer (Weeks 4–10)

Multi-modal transformer fusion on Saumel's 11,000-candidate panel. Cross-reference against the WFS1 variant atlas — operational since Windsor — so every candidate signal is filtered through known variant mechanism, pathogenicity, and structural class. Output: ranked subset of 50–200 candidates with mechanistic explanations grounded in the Atlas, not in generic protein assumptions.

Phase 3 — Retinal AI Layer (Weeks 8–14)

RETFound fine-tuned on WashU's existing Wolfram OCT cohort. Output: AI-derived OCT signatures predicting progression rate and preclinical changes. Cross-validate against the blood panel — convergent signals across imaging and blood are the strongest candidates.

Phase 4 — Federated Validation (Months 4–8)

Candidate panel validated across the six-country network. Adjusted for ethnic and ancestral variation including the Druze cohort. Panel with documented sensitivity and specificity across the global Wolfram population.

Phase 5 — Regulatory Submission (Months 8–14)

FDA Biomarker Qualification Program submission. Parallel EMA submission for European trials. Goal: regulatory acceptance as a trial endpoint.

Phase 6 — First Trial Use (Months 14–24)

AMX0035 Phase 3 uses the panel as primary or secondary endpoint. Every future Wolfram drug trial uses it. The biomarker panel becomes the standard.

A Day in the Life

The Vision Statement

> *Without validated biomarkers, the field tests two Wolfram drugs per decade. With them, it tests twenty. Initiative 5 is the trial-speed unlock for the entire framework.*

> *The technology converges from two directions simultaneously: Saumel's 11,000-candidate multi-omic screen, already running at WashU, and RETFound, the retinal AI that detects Parkinson's neurodegeneration before clinicians can. Between them — AI-integrated, federally validated, regulatorily submitted — Initiative 5 produces the first regulator-accepted biomarker panel in the history of Wolfram syndrome research.*

> *Every future drug trial is faster, smaller, and more definitive because of it. That is not a research milestone. That is a commercial inflection point for every company that has ever considered developing a Wolfram drug and decided the clinical timeline was too long.*

INITIATIVE 06

Wolfram Research Intelligence (LLM + Chatbot + Digest)

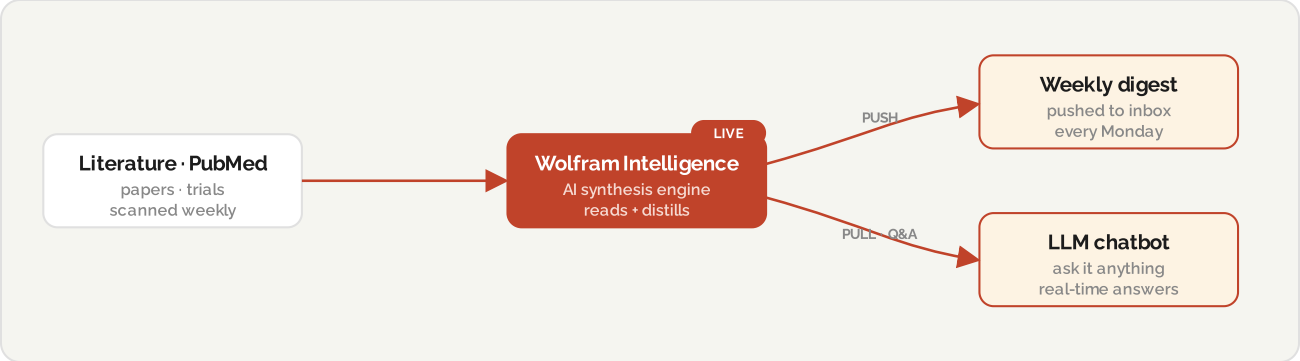

THE FRAMEWORK

The only unambiguously live initiative, and the one with the cleanest origin story: Sarah said 'Can I have that?' on April 14, and it was in her inbox six days later — delivered every Monday since April 20. One engine reads the world's Wolfram literature off PubMed weekly and produces two things: a push and a pull. The push is the weekly digest that lands in leadership's inbox; the pull is an LLM chatbot the community can query in real time — ask it anything about a variant, a paper, or a mechanism and get a sourced answer. Its real role is structural: it is the connective tissue that keeps all eight initiatives current and connected to one another.

The Problem

PubMed publishes over 50,000 new biomedical papers every month. The paper that could unblock a Wolfram lab is frequently published in a field that lab never monitors — ALS, Parkinson's, Type 2 diabetes, macular degeneration — all sharing upstream mechanisms with Wolfram but tracked by different research communities. Every week, critical insights relevant to Wolfram disease biology are published and go unread for months or years. There is no systematic way for any researcher, no matter how dedicated, to maintain awareness across the adjacent literature that matters. The knowledge exists. No one has time to find it.

Solution in Concept

A disease-specific AI research intelligence system is now operating: a weekly digest that automatically scans every Wolfram-relevant paper published in the prior two weeks, rates each paper for significance using a Snow Foundation research-priorities system prompt, and delivers a plain-English summary with a recommended action to each researcher's inbox — every Monday, forever, with zero ongoing labor. A natural-language chatbot allows any researcher to query the full literature in twenty seconds. This initiative is not a proposal. It is live. It shipped in three weeks from concept. It has been running every Monday since April 20, 2026, into The Snow Foundation's leadership inboxes — and it expands into the broader Wolfram research community at Windsor.

How It Works

The problem it solves

PubMed publishes over 50,000 new biomedical papers every month. No Wolfram researcher has time to read them all — and the paper that would unblock a Wolfram lab is almost always published in a field that lab doesn't watch. Parkinson's, Type 2 diabetes, ALS, Alzheimer's, macular degeneration — all share upstream mechanisms with Wolfram, and almost none of them are systematically tracked by the Wolfram research community. Critical insights are being published every week and going unread for months or years.

What is actually shipped today

The Wolfram Research Intelligence Digest runs every Monday at 8 AM Eastern. It pulls every Wolfram-relevant paper published in the previous fourteen days, sends each to Claude with a Snow Foundation research-priorities system prompt, and receives back a significance rating of 1–5, a plain-English summary, an explanation of why it matters to Wolfram research, and a proposed action for the research team. An HTML email goes out. Sarah reads it with morning coffee.

Zero ongoing labor per run. The Mac wakes itself, runs the cron job, the email goes out. Every week. Forever.

The chatbot — already deployed — allows any researcher to ask a natural-language question and get a precise, cited answer in twenty seconds. Sarah said "Can I have that?" on April 14. It was in her inbox six days later.

Sarah's framing — use it on stage

Sarah called this "social media for researchers — they don't care about body-slimming Spanx, but they care about the data that influences their work. This is putting the content they need into accessible units so they can see what's out there."

Wicked high-level data content creation. Quote her. The framing is hers.

What is about to ship

The researcher-facing web app deploying before Windsor adds magic-link login, inline paper ratings and feedback, and a feedback loop that changes the entire game: the AI reads the last twenty researcher ratings before generating the next digest and learns what each person actually cares about. Static newsletter becomes living, calibrating intelligence. A searchable archive of every digest. Admin dashboard. The system gets smarter every week without any human intervention.

The Pilot — Best Case Scenario

The pilot is already running. The phases ahead are about scaling what works.

Phase A — Internal Beta (April–May 2026, LIVE)

Sarah, Stephanie, and Fumi on the recipient list. Weekly digest running since April 20. Informal feedback collection. Sarah testimonial in progress.

Phase B — Web App (May–June 2026)

Railway deployment. Magic-link auth. Inline paper rating and feedback loop. Searchable archive. The static digest becomes a living intelligence platform.

Phase C — Windsor Launch (June 3–4, 2026)

Live demo of the digest and chatbot during the framework presentation. Open sign-up for all 14 attending researchers. Sarah or Fumi speaks to what it actually did for their research in the prior six weeks. The conference becomes a product launch.

Phase D — Community Scaling (July 2026+)

Open to the wider Wolfram research community through WolframNow.org integration. Researcher upload pathway: any researcher submits papers, lab notes, or transcripts to

enrich the knowledge base. Self-enriching system — every new researcher makes it smarter for everyone else.

A Day in the Life

The Vision Statement

> *There are five thousand Wolfram patients on earth and roughly fifty thousand new biomedical papers published every month. The Wolfram Research Intelligence Digest is the system that closes that gap.*

> *This initiative matters for two reasons. First: it gives every Wolfram researcher in the world the equivalent of a tireless graduate student who reads every paper in every adjacent field and surfaces only what matters — every Monday, forever, for free. Second, and more importantly: this is the existence proof for the rest of the framework.*

> *Eight initiatives. One built and shipped in three weeks by one operator with one open AI platform. The other seven follow the same pattern. Every Monday at 8 AM Eastern, the realistic version of this framework is in your inbox. The question was never whether this could be built. The question was whether someone would build it.*

INITIATIVE 07

AI-Optimized Clinical Trial Design

THE FRAMEWORK

Where the whole stack converges. The atlas stratifies enrollment, the biomarkers become endpoints, the federated network finds and consents patients, and the drug pipeline supplies the candidate. Together they make a Phase 3 that enrolls 70 patients instead of 700, reads out in 12 months instead of five years, and gives every patient a drug instead of a coin flip — built around the FDA pathway and a synthetic control arm drawn from the registry.

The Problem

Traditional Phase 3 drug trials need hundreds of patients to achieve the statistical power required for FDA approval. There are five thousand Wolfram patients on earth. A traditionally designed Wolfram Phase 3 is mathematically impossible at scale. The AMX0035 HELIOS Phase 2 enrolled twelve participants — the practical ceiling the current model allows. Every Wolfram trial before that has hit the same wall: too few patients, too long a timeline, too much asked of families watching their children deteriorate in a placebo arm, waiting for a statistical readout that takes years and still cannot confirm efficacy with confidence.

Solution in Concept

Apply three AI techniques that collectively change the trial math. Synthetic control arms replace the placebo group with a historical cohort built from the global federated patient registry — every enrolled patient gets the drug, no one waits in placebo. Bayesian adaptive design lets the trial adjust in real time based on interim data — shortening when the signal is

clear, extending when it is not. Variant-stratified enrollment, guided by the AlphaFold atlas, places patients in arms matched to their specific biology rather than treating all 5,000 Wolfram patients as interchangeable. The result: a Phase 3 that enrolls 70 patients instead of 700, reads out at 12 months instead of 5 years, and produces every patient a drug instead of a coin flip.

How It Works

The math problem Wolfram has always had

Traditional Phase 3 drug trials need hundreds of patients to achieve statistical power. There are five thousand Wolfram patients on earth. A traditionally designed Wolfram Phase 3 is mathematically impossible. The AMX0035 HELIOS Phase 2 enrolled twelve participants — the ceiling that the current model allows. Every trial before that was similarly constrained: too few patients, too long a timeline, too much asked of families who are watching their children deteriorate.

Three AI techniques that change the math

Synthetic control arms: instead of asking 35 patients to sit in placebo for 18 months, an AI builds a comparison cohort from existing Wolfram patient registry data. No patient gets placebo. Every patient gets the drug.

Adaptive design: instead of locking the protocol on day one and running it rigidly for four years, the trial adjusts in real time. Dosing, enrollment, and endpoints can be tuned by Bayesian statistics while the trial is running. I-SPY 2 in breast cancer accelerated sixteen drugs through Phase 2 using this exact architecture.

Stratified enrollment: the AlphaFold variant atlas — already built and shipping into Windsor — tells you which arm a patient joins. Truncation variants go to the gene therapy track. Surface-disrupting variants go to the calcium modulator arm. Mild destabilizing variants go to the chaperone arm. The trial is built around the biology, not around the assumption that all Wolfram patients are the same. The Atlas is the precondition that makes this kind of stratification possible at Phase 3 today rather than five years from now.

Why this is the most time-sensitive initiative in the framework

Amylyx is finalizing AMX0035 Phase 3 design with the FDA right now. The Windsor symposium in June sits inside that alignment window. Sarah said on the call: "This is what we would most want to do." James in Boston has prior FDA acceptance of stratified enrollment

in rare disease work. Sarah has committed to connecting James to Amylyx's clinical development team.

This is not an initiative to do next year. The trial is launching this year. The window to shape its design is measured in weeks.

The Pilot — Best Case Scenario

Unlike Initiatives 1 through 6, Initiative 7 is not a new build. AMX0035 Phase 3 is a real trial happening this year regardless. The question is whether AI-optimized design gets into the protocol before it is finalized.

Phase 1 — Sarah Connects James to Amylyx (Weeks 1–4)

30-minute call: stratified enrollment plus synthetic control arm proposal. Goal: open door for joint FDA pre-IND discussion. This is the highest-leverage call in the framework.

Phase 2 — Joint FDA Pre-IND Meeting (Months 1–3)

Amylyx plus James plus Sarah. Proposal: 70-patient Phase 3, stratified by WFS1 variant class, synthetic control from federated network, crossover at month 6. Best case: FDA accepts. Realistic: hybrid design.

Phase 3 — Protocol Finalization (Months 3–6)

Adaptive rules pre-specified. Bayesian framework with FDA input. Initiative 5 biomarker panel as endpoints. Synthetic control cohort specifications finalized. Sites: WashU, UK TREATWolfram, EU partners.

Phase 4 — Trial Launch (H2 2026)

First patient enrolled. Real-time monitoring via federated infrastructure. Six-month interim analysis.

Phase 5 — Readout and Submission (2027–2028)

Faster readout than traditional Phase 3 because of biomarker endpoints and adaptive interim analyses. BLA or NDA submitted with synthetic control evidence package. AMX0035 becomes the first approved Wolfram drug.

Phase 6 — Template (2028+)

Tirzepatide, pridopidine, and AUDIOWOLF trials all use the same AI-optimized architecture. Every future Wolfram trial — and a growing number of rare disease trials in other conditions — cites this design.

A Day in the Life

The Vision Statement

> *AMX0035 Phase 3 begins in the second half of this year. The design is being finalized with the FDA right now. The window to shape that design — with AI-optimized stratification, synthetic controls, and adaptive endpoints — is open for weeks, not months.*

> *If a Phase 3 needs 700 patients and there are five thousand in the world, the traditional math has already failed. Stratified enrollment plus synthetic controls drawn from a federated global registry plus adaptive Bayesian design is how a Wolfram Phase 3 runs with 70 patients instead of 700. Once it works — once the first rare disease in history completes a Phase 3 without a placebo arm and with half the usual enrollment — every rare disease drug developer in the world adopts the playbook.*

> *Initiative 7 is where the entire framework lands in a patient's hand as approved medicine. Everything else we are building feeds this trial. This trial proves the framework works.*

INITIATIVE 08

iPSC + AI High-Content Drug Screening

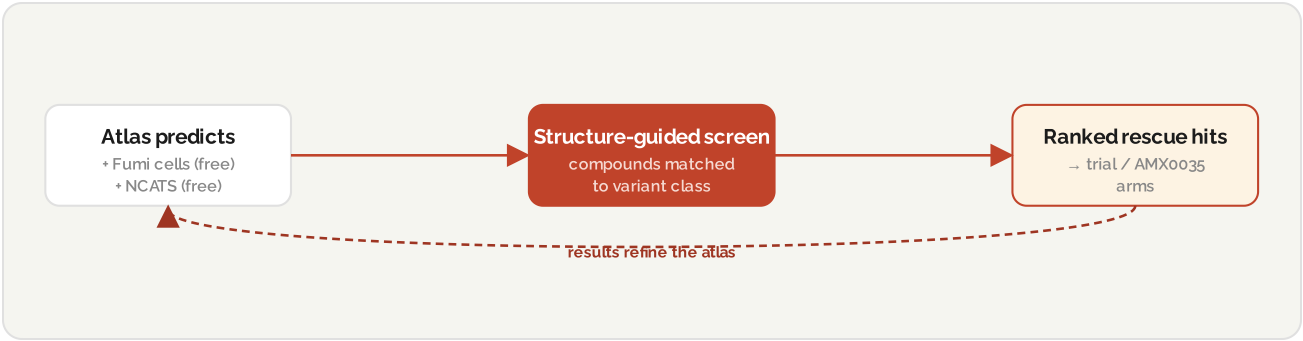

THE FRAMEWORK

The only initiative shaped as a closed loop, and that shape is what makes it fundable. The atlas tells the screen which compounds to test on which variant class; Fumi's patient-derived iPSC cells and NCATS Cell Painting validate them in living cells; and the results loop back to sharpen the atlas. It is exactly what NCATS funds — predict, validate, refine, deliver. The atlas closes the 'what are you screening for?' gap; the screen closes the atlas's 'where's your wet-lab proof?' gap. Top leads feed Amylyx's AMX0035 Phase 3 arms.

The Problem

Even with a ranked list of promising drug candidates, drug development hits a fundamental biological bottleneck: someone has to test each compound in actual patient cells to see if it physically rescues the disease. This normally requires five to ten years to build patient-derived cell models and access to pharmaceutical-grade robotic screening infrastructure that only the largest drug companies own. For a rare disease with five thousand patients worldwide, neither resource has ever been fully accessible. Drug candidates have been proposed. The infrastructure to validate them efficiently has not existed for Wolfram.

Solution in Concept

Combine two already-built resources that have never been pointed at Wolfram together. First: Fumi's patient-derived iPSC beta cells, retinal organoids, and brain organoids — built and quality-controlled at WashU over ten years, covering multiple WFS1 variant classes,

available at no cost. Second: NIH NCATS's industrial-grade Cell Painting screening infrastructure — the same methodology Recursion Pharmaceuticals used to earn thirty-million-dollar milestone payments from Genentech, now available to rare disease researchers through a collaboration program that is actively seeking stem cell partners. AI image analysis reads hundreds of thousands of cellular images per run, extracts 1,500 morphological features per cell, and ranks every compound by how completely it rescues diseased cells back to healthy morphology. The NIH proposal was submitted April 14, 2026. The cells are ready.

How It Works

What "high-content screening" actually means

Drug discovery dream: take a thousand candidate drugs, expose them to actual diseased cells, watch what happens, pick the ones that fix the disease.

Drug discovery problem: a human pathologist examines maybe fifty cells per day. A drug screen needs to look at hundreds of millions.

High-content screening solves this. Robotic systems load cells in 384-well plates. Robotic liquid handling exposes each well to a different drug. Automated microscopes photograph everything. AI image analysis processes hundreds of thousands of images per run, extracts approximately 1,500 morphological features per cell — things no human eye can quantify — and ranks compounds by how well they rescue the diseased cell's appearance back to a healthy state.

The gold standard method is Cell Painting, developed by Anne Carpenter at the Broad Institute. Six fluorescent dyes, five channels, eight cellular components. Already at 100,000-compound scale. Open source.

Why Wolfram is uniquely positioned for this

The two hardest parts of any high-content screen are building patient-derived cell models — usually five to ten years — and accessing screening infrastructure — usually requires major pharma.

Both bottlenecks are already solved for Wolfram.

Fumi's lab at WashU has patient-derived iPSC beta cells, retinal organoids, and brain organoids from multiple WFS1 variant classes. Already published. Already quality controlled. Fumi provides them free.

NIH NCATS — the National Center for Advancing Translational Sciences — has the infrastructure, the compound library, the robotics, and the imaging. They are actively recruiting stem cell research collaborators. The proposal was submitted to Mark Henderson at NCATS on April 14, the same day Sarah said "I think you just found your new company."

How Recursion proved the stack works

Recursion Pharmaceuticals built this exact stack at industrial scale. They have one trillion iPSC-derived neurons in their database. They have earned thirty-million-dollar milestone payments from Genentech twice. Their partnership with Roche and Sanofi is built entirely on this methodology.

We are not asking NIH to invent something new. We are asking them to point an existing industrial-grade methodology at Wolfram. Fumi provides cells for free. NIH provides infrastructure for free. The total cost to Snow Foundation: close to zero.

The Pilot — Best Case Scenario

The proposal is already submitted. The phases ahead assume NIH says yes — and there is no structural reason they would say no.

Phase 0 — Proposal Submitted (April 14, 2026 — DONE)

NIH collaboration proposal submitted to Mark Henderson. Sarah committed to independent follow-up and will copy Mike on all correspondence.

Phase 1 — NIH Acceptance (Weeks 1–6)

NIH evaluates the proposal. Best case: SCTL collaboration, assays designed jointly with Fumi. Realistic: pilot scope of 1,200 compounds across five cell types with expansion option.

Phase 2 — Cell Line Transfer (Weeks 6–12)

Fumi's iPSCs transferred to NCATS: Wolfram patient-derived beta cells across multiple variant classes, retinal organoids, brain organoids, and CRISPR-corrected isogenic controls. Quality control and protocol harmonization.

Phase 3 — Pilot Screen Design (Weeks 8–16)

Top 1,200 candidates from NCATS pharmaceutical collection, Initiative 3 TxGNN ranking, and Initiative 2 atlas mechanism predictions. Cell Painting plus Wolfram-specific assay panel:

ER stress markers, calcium dynamics, mitochondrial membrane potential, cell viability.

Phase 4 — Screen Execution (Months 4–8)

384-well plate screening at NCATS. Hundreds of thousands of images. CellProfiler AI analysis. Ranked compound list with phenotypic rescue scores per cell type per variant class.

Phase 5 — Hit Validation (Months 8–14)

Top 50 hits through dose-response, cross-cell-type confirmation, and mechanism profiling. Sovan Sarkar's Birmingham brain organoids validate top neurological hits. Sarah and Nufar's zebrafish validate in-vivo signals.

Phase 6 — Hand-Off to Initiative 7 (Months 14–18)

Top five to ten validated leads with full mechanistic profiles delivered to Amylyx and Wolfram drug developers. Feeds AMX0035 Phase 3 expansion arms and informs every subsequent Wolfram trial.

Phase 7 — Full Library Screen (Months 18–24+)

Expand to the full NCATS compound library if the pilot succeeds. Most comprehensive therapeutic candidate dataset in Wolfram history. Open-access publication.

If NIH Says No

There is no structural reason NIH NCATS declines this proposal — the cells are ready, the infrastructure is already pointed at exactly this kind of collaboration, and Recursion Pharmaceuticals has already proven the stack works at industrial scale. But proposals slip, review cycles stretch, and a contingency path matters.

Three fallback paths exist. The first is Recursion Pharmaceuticals directly — their entire Cell Painting platform is built for exactly this kind of partnership, and their established Roche and Sanofi relationships demonstrate they license screening capacity. The second is the Broad Institute, where Cell Painting was developed; they run academic collaboration programs and accept patient-derived iPSCs from outside labs. The third is a CRO-funded path — a Phase 2 grant from NORD, Global Genes, or a rare disease impact investor pays for screening through a contract research organization at roughly $200K–$400K, an order of magnitude smaller than building the infrastructure from scratch.

The proposal sits on Mark Henderson's desk. None of the fallback paths are activated. They exist because Initiative 8 is the validation gate for Initiatives 3 and 7 — the drug pipeline depends on it, and the program does not stall waiting for a single inbox.

A Day in the Life

The Vision Statement

> *We started this framework with a child walking into a pediatric endocrinology clinic and a soft blue flag appearing on the doctor's screen. We end with the drugs that flag will lead to.*

> *Every other initiative in this framework predicts. Initiative 8 validates. The infrastructure for all of this is already built. Fumi spent ten years making patient-derived iPSCs. NIH spent twenty years building NCATS. Anne Carpenter spent fifteen years building Cell Painting. Recursion Pharmaceuticals spent ten years and a billion dollars proving the stack works at industrial scale — with thirty-million-dollar milestone payments from Genentech to show for it. None of this needs to be invented. None of this needs to be funded by Snow. Fumi provides cells free. NIH provides infrastructure free. The methodology is open source. The proposal is on Mark Henderson's desk.*

> *Within eighteen months, we will have the most comprehensive therapeutic candidate dataset in the history of Wolfram syndrome. Eight initiatives. Seven months from concept to framework. One operator. One open AI platform. One Snow Foundation. The other rare diseases watching this room — there are seven thousand of them — will follow what they see. This isn't the framework's closer. It's the framework's beginning.*

The Ecosystem

> *How the Eight Initiatives Connect, Compound, and Converge*

Initiative 6 is the connective tissue

Start here, because Initiative 6 — the Research Intelligence Digest and AI chatbot — is the only initiative already live. Every Monday since April 20, 2026, an automated synthesis of every Wolfram-relevant paper published in the prior two weeks lands in the inboxes of the Snow Foundation's research leadership. That is not a convenience tool. It is the nervous system of the ecosystem.

Every other initiative generates findings. Publications. Structural models. Drug candidates. Biomarker signatures. Trial results. Without a system that reads all of it, integrates it, and surfaces the cross-connections that no individual researcher has time to find, the ecosystem's outputs pile up in silos — the same problem that has plagued rare disease research for forty years.

Initiative 6 is what keeps the other seven talking to each other. When Initiative 2 publishes a new variant structural model, Initiative 6 finds it, digests it, and flags its relevance to the drug screen running in Initiative 3. When Initiative 5 identifies a new biomarker signal, Initiative 6 connects it to the trial design conversation in Initiative 7. The chatbot is the shared memory of the entire framework.

Initiative 2 is the molecular foundation everything else is built on

There are over 200 known ways the WFS1 gene can be misspelled. Each spells a different version of Wolfram — a different rate of progression, a different set of organs affected first, and critically, a different class of drug that might work.

The AlphaFold Molecular Atlas maps all 200. For every known variant: what is physically broken inside the cell, and what type of intervention is most likely to rescue it. That map is not just a research tool — it is the molecular index that every other initiative references.

Initiative 3 (drug repurposing) uses the atlas to rank candidates not just by disease relevance but by variant-specific mechanism — so the sigma-1 agonist goes to the patients

whose protein can still be partially rescued, and the gene therapy candidates go to the patients whose protein is simply gone. Initiative 5 (biomarkers) uses the atlas to look for variant-specific molecular signatures in blood and imaging — because a truncation variant and a misfolding variant may need different biomarkers to track disease progression. Initiative 7 (trial design) uses the atlas to stratify enrollment — placing patients in treatment arms matched to their actual biology rather than treating all Wolfram patients as interchangeable.

Without the atlas, the rest of the framework operates in the dark. With it, every downstream initiative gets faster, sharper, and more statistically powerful.

Initiatives 1 and 4 are the patient infrastructure

Two initiatives exist specifically to solve Wolfram's most fundamental research problem: there are not enough patients in any one place to see the patterns that only emerge at global scale.

Initiative 1 — AI-powered early diagnosis — finds patients earlier. A child diagnosed at nine instead of thirteen represents four additional years in the natural history study, four additional years of pre-treatment imaging and biomarker data, and four additional years of trial eligibility before irreversible neurological decline sets in. Every patient found early strengthens every other initiative that depends on patient data.

Initiative 4 — the federated patient data network — connects the patients who are already diagnosed. Six countries. Multiple academic medical centers. Patient records that cannot legally cross borders. NVIDIA FLARE visits each institution, learns from the data in place, and returns only statistical patterns — never moving a single patient record. The result is the first global Wolfram natural history dataset in the history of the disease.

Together, Initiatives 1 and 4 solve the patient infrastructure problem from both ends. Initiative 1 grows the population. Initiative 4 connects the population that already exists. Both feed every initiative that requires patient data — which is all of them.

Initiatives 3 and 8 are the drug pipeline

Wolfram has no approved treatment. The path to one runs through two initiatives that are designed to operate as a sequential pipeline.

Initiative 3 starts broad. TxGNN — Harvard's graph neural network trained on 17,000 diseases — screens the full approved drug library against Wolfram's molecular biology in 48 hours. The output is a ranked list: single agents, combination candidates, and

contraindications. Cross-referenced against the variant atlas from Initiative 2, so the ranking is not just "works for Wolfram" but "works for this class of variant."

Initiative 8 validates what Initiative 3 identifies. Fumi's patient-derived iPSC beta cells, retinal organoids, and brain organoids — built over ten years at WashU, covering multiple WFS1 variant classes — are the biological testing ground. NIH NCATS's Cell Painting infrastructure reads 1,500 morphological features per cell and ranks every candidate by how completely it rescues diseased cells back to healthy morphology. The top candidates from Initiative 3's computational screen enter Initiative 8's biological screen. The ones that survive both move into clinical trial design.

This is a pipeline, not two separate projects. Initiative 3 generates the candidates. Initiative 8 confirms them. Together they compress what would be a decade of sequential wet lab work into a parallel, AI-accelerated screening process.

Initiative 5 is what makes clinical trials possible

The AMX0035 HELIOS Phase 2 enrolled twelve patients and needed eighteen months to read out because there was no faster way to measure whether the drug was working. Clinical decline in Wolfram — vision loss, hearing loss, neurological deterioration — takes years to become measurable. Without a faster signal, every Wolfram trial operates on a timeline that exhausts families, depletes budgets, and kills drug programs before they reach approval.

Initiative 5 finds that faster signal. RETFound, trained on 1.6 million retinal images, can detect neurodegeneration months before it becomes clinically visible. Saumel's 11,000-candidate multi-omic panel is running on Wolfram patient samples at WashU right now, looking for the molecular change in blood that precedes the clinical change in the clinic. AI fusion of retinal imaging and blood proteomics identifies the smallest, fastest, most reproducible signal that the FDA will accept as a validated trial endpoint.

That validated biomarker panel changes the trial math for every future Wolfram study. Smaller enrollment. Shorter timeline. Faster readout. Initiative 7 uses it directly — the synthetic control arm design and the adaptive protocol depend on having a measurable endpoint that moves faster than clinical decline. Initiative 5 is the unlock that makes the trial framework in Initiative 7 actually executable.

Initiative 7 is where everything converges

Every other initiative feeds Initiative 7 — because Initiative 7 is the one that gets a drug approved.

The trial design problem for Wolfram is mathematical. Five thousand patients globally. A traditional Phase 3 needs hundreds. That gap has blocked every serious drug program from reaching approval. Initiative 7 closes the gap with three AI techniques that collectively change the trial math.

Synthetic control arms — built from the federated patient registry in Initiative 4 — replace the placebo group with a historical cohort. Every enrolled patient gets the drug. No family watches their child deteriorate in a placebo arm waiting for a statistical readout. Bayesian adaptive design lets the trial adjust in real time, shortening when the signal is clear. Variant-stratified enrollment, guided by the atlas from Initiative 2, places patients in arms matched to their specific biology.

The biomarker panel from Initiative 5 is the endpoint. The drug candidates from Initiatives 3 and 8 are what's being tested. The patients from Initiatives 1 and 4 are the enrolled population. Initiative 7 is not a separate project — it is the destination that every other initiative is pointed toward.

A Phase 3 that enrolls 70 patients instead of 700. A readout at 12 months instead of 5 years. Every patient receiving the drug. That is what convergence looks like.

The full picture: a research acceleration engine

Laid flat, the eight initiatives look like a list. Laid in motion, they look like this:

Initiative 6 reads everything and keeps the ecosystem connected. Initiative 2 maps the molecular landscape every other initiative navigates. Initiatives 1 and 4 build and connect the patient infrastructure. Initiatives 3 and 8 run the drug pipeline. Initiative 5 creates the measurement tools that make trials fast. Initiative 7 brings the drug through approval.

Each initiative accelerates the others. Earlier diagnosis feeds richer natural history data into the federated network. The federated network powers stronger synthetic control arms for trials. The variant atlas sharpens the drug screen. The drug screen feeds the iPSC validation. The validated candidates enter a trial designed by Initiative 7 using the biomarkers from Initiative 5.

The result is not eight parallel projects moving independently. It is a single compounding system — one where every advance in any initiative shortens the timeline in every other. That compounding is what makes this framework different from forty years of fragmented rare disease research. And that compounding is why, for the first time, the phrase "treatment for Wolfram" is not a hope. It is an engineering problem with a plan.

Sequence of Build

> *What Ships When, and Why the Order Matters*

Eight initiatives are not equally ready to begin. Two are already built. Six follow in a sequence dictated by which initiative makes the next initiative possible. What follows is not a wishlist. It is a build order with internal dependencies, capital implications, and a political timeline anchored to the Windsor symposium.

Phase A — Built and Shipping at Windsor (Initiatives 6 + 2)

The Wolfram Research Intelligence Digest and chatbot (Initiative 6) has been delivering weekly automated synthesis to The Snow Foundation's leadership since April 20, 2026. The AlphaFold Molecular Atlas of WFS1 (Initiative 2) ships at Windsor as a working per-variant card system. Together they form the substrate the other six initiatives run on. Every downstream initiative either ingests from them, sells through them, or stratifies because of them. Walking into Windsor with these two operational changes what the framework is — from a roadmap of eight things to build into an operating platform with six modules in queue.

Phase B — Patient Layer (Initiatives 4 + 1) — 90 to 180 Days Post-Windsor

Initiative 4 — the federated patient data network — leads this phase because Windsor is the once-a-year convening that makes it politically possible. Anchor PIs from every site that holds Wolfram patient data will be in one room for 48 hours. The Letter of Intent signed in that room is the launch event. Standing up the federated infrastructure takes the next 90 to 180 days: Common Data Model, IRB amendments at each site, NVIDIA FLARE deployment behind each hospital firewall. Initiative 1 — AI-powered early diagnosis through DeepRare and Epic integration — runs in tandem. One initiative grows the patient population; the other connects the population that already exists. Both anchor to the Molecular Atlas, so every flagged or enrolled patient inherits a variant card the moment they enter the system. Without the patient layer, nothing downstream validates — and without Windsor, the patient layer takes eighteen months instead of six.

Phase C — Signal Layer (Initiative 5) — Months 6 to 18

Multi-omic biomarker discovery cannot begin in earnest until the patient layer is operating. Saumel's 11,000-candidate panel at WashU has the methodology. The federated network provides the longitudinal samples and outcomes across the global population. The Atlas provides the mechanism priors that say which of those 11,000 signals are most likely to be the disease talking. RETFound runs in parallel on the OCT imagery already collected at every Wolfram visit worldwide. This is the trial-speed unlock — the layer that turns the platform from R&D intelligence into trial enablement and qualifies endpoints with the FDA Biomarker Qualification Program.

Phase D — Drug Pipeline (Initiatives 3 + 8) — Parallel Track from Month 4

Drug repurposing (Initiative 3) and iPSC high-content screening (Initiative 8) form a pipeline rather than two separate projects. Initiative 3 generates ranked candidates: TxGNN screens the full approved drug library against Wolfram's molecular biology, informed by the variant atlas, in 48 hours. Initiative 8 validates the candidates: Fumi's patient-derived iPSC beta cells, retinal organoids, and brain organoids meet NIH NCATS's screening infrastructure at WashU. Initiative 8 can begin on a parallel track from the moment NIH accepts the proposal — it depends on cells and infrastructure, not on patient flow — which is why it appears here as a parallel rather than sequential build. Together they compress what would be a decade of sequential wet-lab work into a parallel 18-month engine.

Phase E — Trial Design (Initiative 7) — Year 2 and the AMX0035 Window

AI-optimized trial design is the assembly layer. It needs everything before it: the federated patient population to draw a synthetic control arm from, the validated biomarker panel to use as endpoints, the variant atlas to stratify enrollment, and the drug candidates from Initiatives 3 and 8 to test. The one timeline override is AMX0035: Amylyx is finalizing the Phase 3 design with the FDA inside the Windsor window, which means Initiative 7 has a forcing function that runs in parallel with Phases B and C, not strictly after them. Sarah's introduction of James to Amylyx is the highest-leverage call in the framework — it gets AI-optimized design into the AMX0035 protocol before the protocol locks.

Why the Order Matters

Built in the wrong order, this framework collapses. Trial design before biomarkers means measuring the wrong endpoint. Drug repurposing before the variant atlas means screening against a generic protein. A federated network attempted without the Windsor convening means eighteen months of cross-border negotiation that may never close. The order is not

arbitrary. It is what makes each subsequent initiative cheaper, faster, and more defensible than the one before it — and what turns eight parallel research projects into a single compounding system that walks into FDA review with the receipts already in hand.

About the Author

Mike Wallace II is a speaker, entrepreneur, and culture architect based in Los Angeles. He played D1 football and ran Junior Olympic Track at Penn State, and went on to earn an Executive MBA from the same institution. His career has taken him from the Jay-Z 4:44 Tour — where he served as Content Director — to Nicki Minaj's world tour, to a songwriting credit on Katy Perry's Lifetimes Tour. He is the Founding Creative Director of 1500 Sound Academy, which earned the Key to the City of Los Angeles, and served as Chief Business Officer at Socanomics, a music data company operating in 126 countries with partners including Meta, Lululemon, and Visa. He co-created Real Thing Records, a boutique joint venture between Universal Music Group and Coca-Cola. He delivered a TEDx talk titled "Living the Free Life: Turning Setbacks into Superpowers" and hosted the Young Investigator Draft 2026 at Lincoln Financial Field.

In 2024, Mike was invited to support The Snow Foundation through the Uplifting Athletes Rare Disease Champions program — a network that connects athlete-advocates to rare disease organizations to amplify awareness and mobilize resources. He was matched with Tara Zagoni, then Chief Strategy Officer at The Snow Foundation and a patient herself, who introduced him to the Foundation's work on Wolfram syndrome. That introduction led to a deeper relationship with the Foundation, including the connection to Brett Brackett, and to Stephanie and her research on the natural history of Wolfram in a family cohort.

The conversation that crystallized this framework happened in April 2026, when Mike sat down with Sarah Gladstone, the rare disease researcher and clinician at Washington University whose network, clinical instincts, and institutional relationships are the scientific backbone of every initiative described here. Mike had been mapping the AI landscape — what was actually built, what was deployable, what was waiting to be assembled — and walked Sarah through the eight-initiative framework one piece at a time.

When he finished, Sarah looked up and said: "I think you just found your new company."

That is what this document is. A company — or more precisely, a framework that becomes one. Mike’s background is not in science. It is in assembling rooms, building bridges between people who couldn’t otherwise reach each other, and deploying cultural and operational fluency to move things that were otherwise stuck. In the seven weeks between The Snow Foundation conversation and the Windsor symposium, that assembly fluency produced two

operational AI products — Initiative 6 shipping a weekly automated research digest and chatbot to The Snow Foundation's leadership, and Initiative 2 shipping a working per-variant Molecular Atlas of WFS1 — built using open-source AI infrastructure and the same agentic systems any operator can now deploy. The eight initiatives in this volume are not research proposals. They are deployment plans for technology that already exists, assembled by someone whose job is assembly, and two of them are already running.