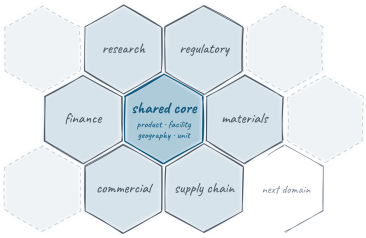
semantic arts

CASE STUDIES · PHARMA & LIFE SCIENCES · NARRATIVE EDITION

Knowledge graphs in pharma: six case studies

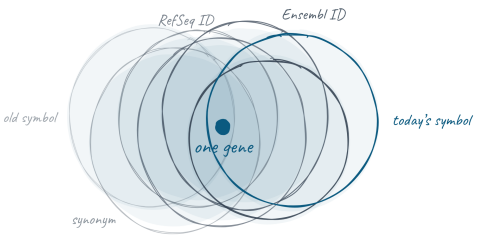
Each sketch below is a parti, the one-picture idea behind the project. Click one to read its case study.



PHARMA · ENTERPRISE DATA

Same thing, different words: one semantic layer for a biopharma leader's data

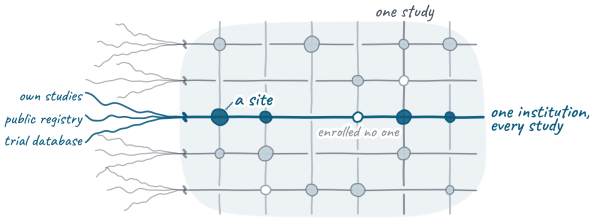
Read the case study › pages 2–3



PHARMA R&D · GENOMICS

One gene, many names: a release-aware gene knowledge graph for a US$35B+ pharma partner

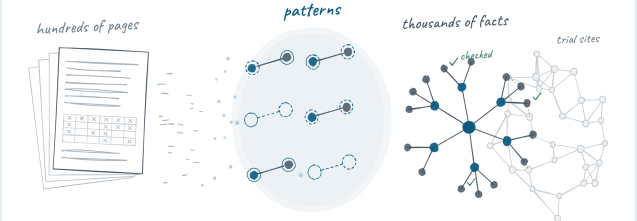
Read the case study › pages 4–5



PHARMA · CLINICAL OPERATIONS

A trial-site knowledge graph for a global clinical CRO

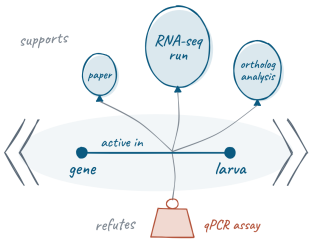
Read the case study › pages 6–7



PHARMA · CLINICAL TRIALS

Hundreds of pages, thousands of facts: digital clinical trial protocols for a global CRO

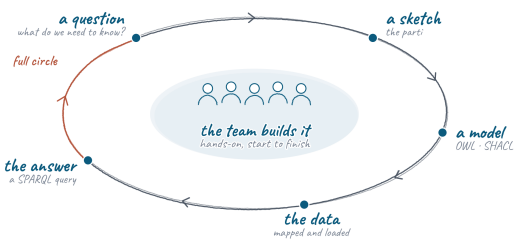
Read the case study › pages 8–10



LIFE SCIENCES · DISCOVERY RESEARCH

No facts, only claims: an evidence-graded discovery graph for a global agricultural-science leader

Read the case study › pages 11–13



PHARMA & LIFE SCIENCES · TRAINING

Full circle: ontology and knowledge-graph training for pharma teams

Read the case study › pages 14–16

semantic arts

CASE STUDY · PHARMA · ENTERPRISE DATA

Same thing, different words: one semantic layer for a biopharma leader's data

How Semantic Arts helped the company give research, operations, commercial and finance data one shared vocabulary, domain by domain.

CLIENT	Biopharmaceutical leader, 30K+ employees, US$35B+ total revenue
DOMAIN	Enterprise data across research, operations, supply chain, commercial and finance
WHAT WE DID	Ontology design, knowledge graph builds, method workshops and coaching
STANDARDS	OWL, SHACL, SPARQL; FAIR principles
WHEN	2021 to 2023
STATUS	Completed; follow-on work in regulatory, CMC and discovery biology

The challenge

A global biopharmaceutical company drew its data from a multitude of systems, each with its own data model and its own vocabulary. The same thing went by different names in research, manufacturing, supply chain and finance, so data from one function did not line up with data from another, and each new project made its own copy of the data it needed.

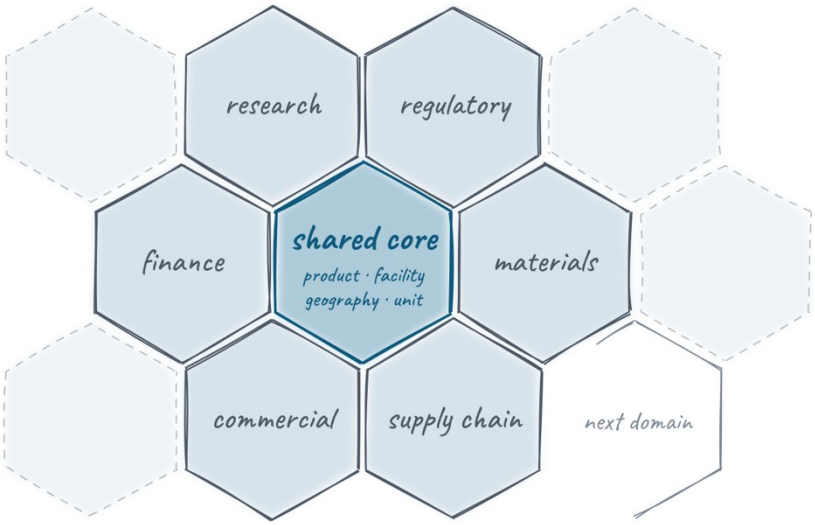


Fig 1: The idea. Each business domain is its own cell, modeled in turn; all build on one shared core of concepts, so the cells fit together.

The company's data leaders agreed that a common language would remove the ambiguities that slow the business. Their answer was a data-centric architecture: a data fabric with a semantic layer, in which data is organized into business domains and represented unambiguously, one model for every system and person. The company had the vision; it needed experienced ontologists to extend the semantic layer, domain by domain.

What Semantic Arts did

We started by blending methods. In joint workshops, the company's approach to ontology work was combined with ours, and its core ontology was compared with gist, Semantic Arts' open upper ontology. **Each domain then started from the business**, not from the systems: the team interviewed subject-matter experts, studied sample data, sketched the concepts and, with the company's experts, wrote competency questions, the business questions the graph must answer.

We modeled each domain as an OWL ontology written to the company's ontology standards, with SHACL shapes to check data loads and guide application builders. Over two years the work spanned research (assays, biologics production, animal studies, research taxonomies and genomics), regulatory eCMC, materials and sourcing, supply chain, shipping and logistics, commercial and finance, plus the upper-level properties shared across domains. **Every model was proved with data**: extracts from source systems were mapped to the model, loaded into the knowledge graph and queried with SPARQL to answer the competency questions. The mappings doubled as specifications for the company's developers building production loaders. Along the way, **we built the company's own capability** with learning sessions for subject-matter experts and coaching to take models into production.

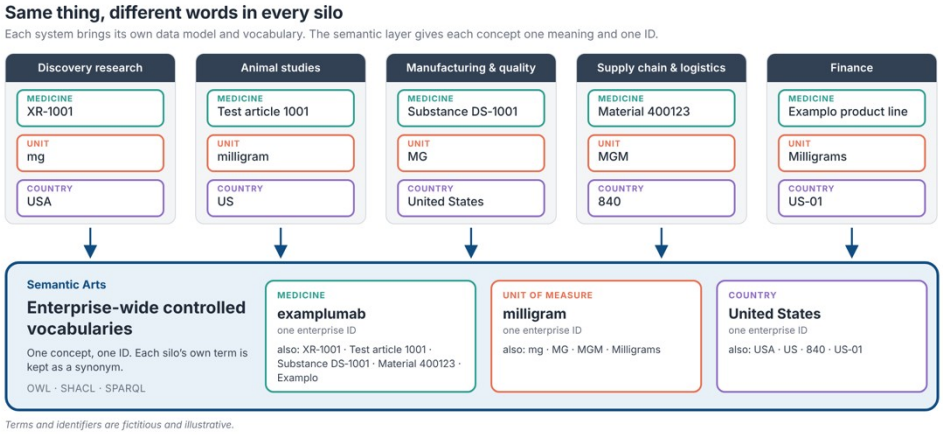


Fig 2: Each silo had its own word for the same thing; the semantic layer gives each concept one meaning and one ID.

What it makes possible

Research, operations, commercial and finance data now share one vocabulary in the semantic layer, organized as a map of business domains. The domains connect through shared concepts such as product, facility, geography and units of measure, so data from different functions fits together. New solutions start from the shared business model instead of a single application, so the same data can serve many applications. And scientific process documentation moves from human-only-readable to human- and machine-readable, ready for analytics and AI.

What's next

The foundation carried on. Semantic Arts has since worked with the same company on product identification aligned with the IDMP ontology and the Pistoia Alliance, end-to-end regulatory data, CMC, and a gene and protein knowledge graph for discovery biology.

Your next move

Bring Semantic Arts your most difficult enterprise data interoperability challenge, whether it spans research, manufacturing, supply chain or finance. We have likely solved something similar. We model your data semantically so you can store it, find it, use it all, scale it and own it completely. The data's meaning is self-contained and available for your APIs or LLM interactivity. Book an appointment with a data architect at info@semanticarts.com and let us know your vision.

semantic arts

CASE STUDY · PHARMA R&D · GENOMICS

One gene, many names: a release-aware gene knowledge graph for a US$35B+ pharma partner

How Semantic Arts enabled omics R&D teams to connect their legacy and current gene data across public databases, release by release.

CLIENT	Biopharmaceutical leader, 30K+ employees, US$35B+ total revenue
DOMAIN	Genomics and proteomics drug discovery
WHAT WE DID	Ontology design and knowledge graph build
SOURCES	RefSeq, Ensembl, GENCODE, UniProt, QIAGEN OmicSoft, and internal data
STANDARDS	RDF, OWL, SHACL, SPARQL, federated SPARQL
STATUS	In production; ongoing

The challenge

Genomics data identifies genes using public reference databases such as RefSeq, Ensembl and GENCODE, with new releases every few months. Gene records get renamed, merged by synonym linkage, split and retired between releases. Even a gene symbol that scientists use every day is not a stable identifier: the same symbol can point to different genes over time, and one gene can carry several symbols throughout the multi-decade lifespan of drug development and patient monitoring.

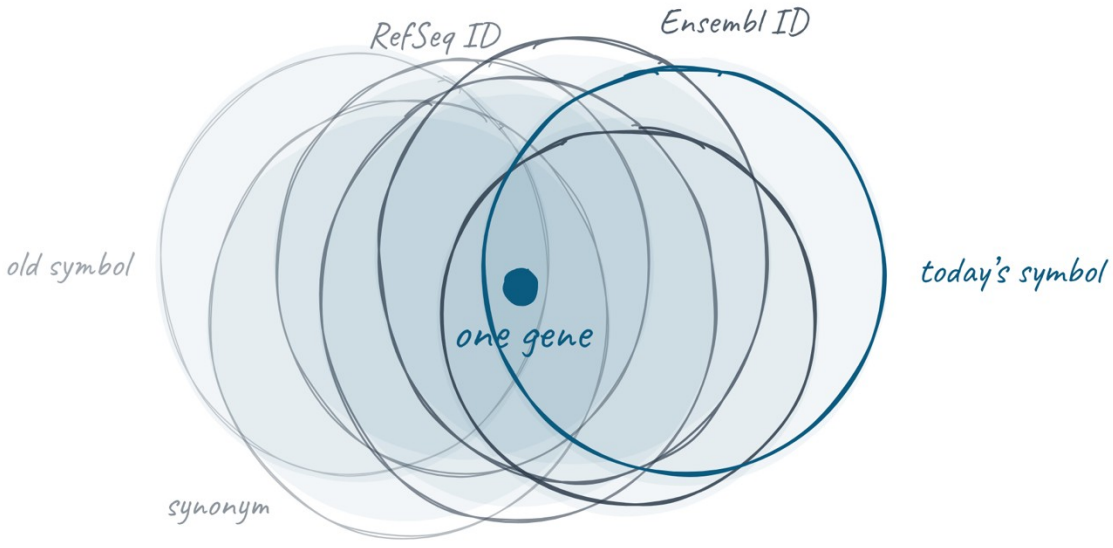


Fig 1: Early concept. The names overlap but never line up exactly. Old symbols, synonyms and database IDs drift from release to release. The graph should anchor each onto one gene concept.

A top-10 biopharmaceutical leader decided to combine years of legacy data with modern data streams, both from its internal laboratories and from outside partners, each labeled its own way. Without a reliable map between old and

new, gene data was being left behind or mis-mapped when data sets were merged, preventing comprehensive meta-analysis.

What Semantic Arts did

We modeled the domain, delivering an RDF/OWL solution to separate the concept of a gene or protein from each release-specific record of it, so a gene's history becomes data that can be queried. Historical and ongoing releases of RefSeq, Ensembl, UniProt and GENCODE were **loaded into the knowledge graph,** with each gene record linked to its successors, transcripts and proteins as well as the cross-species homologs. The result was a **dynamic bridge** across reference sources with seamless **interoperability.** The graph includes UniProt's identifiers and can query UniProt's public RDF endpoint directly, which allows queries for additional information on cross-species gene homologs. We also aligned to the company's broader enterprise ontology and its internal *in vitro* assay and *in vivo* study graph, so all pertinent data can be queried collectively.

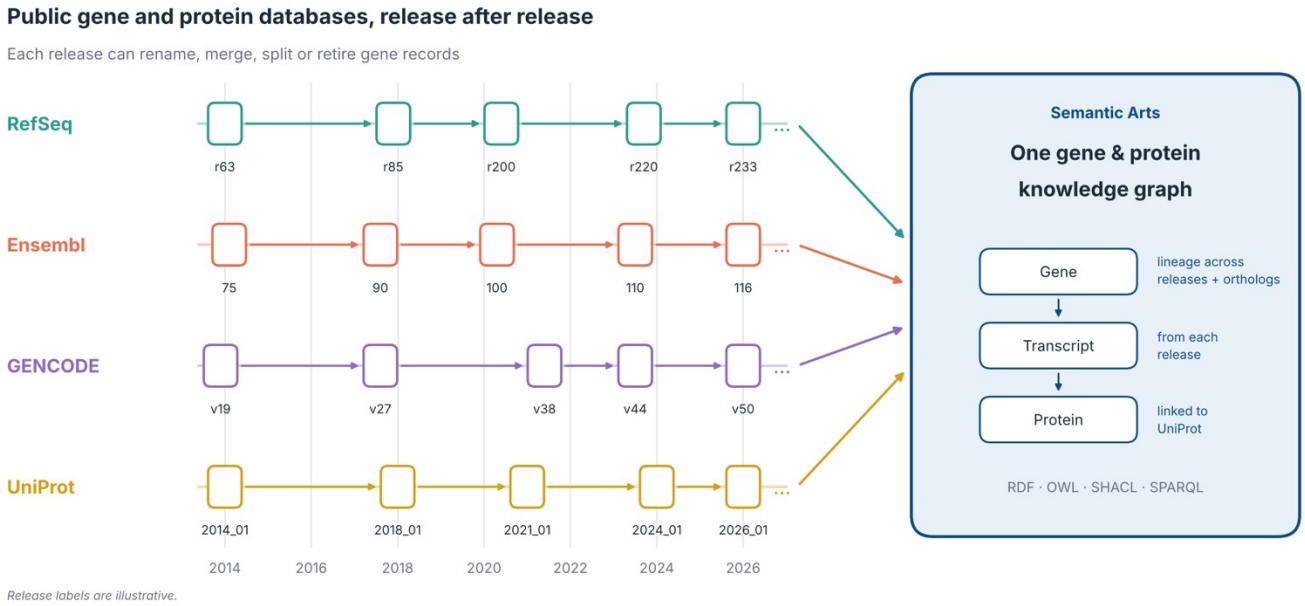


Fig 2: Public gene/protein databases publish new releases year after year. Semantic Arts brought their history into one knowledge graph.

What's now possible

With the production release of the gene knowledge graph, a scientist can now follow a gene record's edits as a lineage over time and across curated reference sources. If an old dataset contains a list of gene symbols but lacks bioinformatics metadata, the reference database release it most likely came from can be determined. To increase statistical power, an analyst can now pool samples from CROs and internal labs, labeled with different identifiers, for comprehensive biomarker meta-analyses. Lastly, because our graph links human genes to their counterparts in multiple pertinent animal species, the *in vivo* team can select animal models most likely to mimic the molecular human disease state for pre-clinical studies.

What's next

Bring Semantic Arts your most difficult data interoperability challenge in genomics, assay documentation or pre-clinical designs. We have likely solved something similar. We model your data semantically so you can store it, find it, use it all, scale it and own it completely. The data's meaning is self-contained and available for your APIs or LLM interactivity. Book an appointment with a data architect at info@semanticarts.com and let us know your vision.

semantic arts

CASE STUDY · PHARMA · CLINICAL OPERATIONS

A trial-site knowledge graph for a global clinical CRO

How Semantic Arts helped a global contract research organization bring its own trial history, public registries and commercial data into one knowledge graph of sites, institutions and investigators.

CLIENT	Global full-service contract research organization (CRO), 3,000+ employees in 50+ countries
DOMAIN	Clinical-trial operations: identifying and selecting trial sites
WHAT WE DID	Enterprise ontology design and knowledge graph build
SOURCES	The CRO's own study history, ClinicalTrials.gov and a commercial trial database
STANDARDS	RDF, OWL, SHACL, SPARQL; gist and gistPharma; MONDO disease ontology
WHEN	2025, with data refreshes in 2026
STATUS	Completed; follow-on work on digital trial protocols

The challenge

When a drug sponsor puts a clinical trial out to bid, a contract research organization (CRO) must estimate what it will cost to activate potential trial sites, to attract and screen potential patients, and then enroll a subset. The choice is a gamble. Industry research finds that more than a third of trial sites enroll poorly and about one in ten never enrolls a single patient (Tufts CSDD, 2012). A site that enrolls no one still costs time and start-up money.

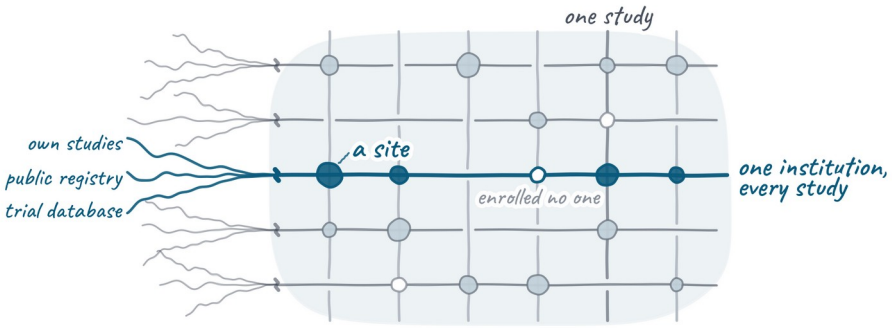


Fig 1: Early concept. An intersection of an institution and a study is a clinical site, even if no subjects were enrolled.

A global full-service CRO wanted to learn from experience: which sites, institutions and investigators have recruited well, and for which conditions. The evidence existed but was scattered. Its own study history, the public ClinicalTrials.gov registry and a commercial trial database each described the same hospitals, sites and investigators their own way, with their own identifiers, across millions of records.

What Semantic Arts did

Semantic Arts **thought big and started small**. We modeled the CRO's enterprise ontology from gistPharma, our life-sciences extension of the gist upper ontology, then populated it one dataset at a time. The pivotal decision was **what a site is**: in the model, a site is an institution's contribution to one study, at one location, with its investigators and

staff. The same institution can then be followed across every study it has taken part in, whichever source reported it, and recruiting measures such as patients screened and enrolled, screen failures, start-up timelines and protocol deviations can be compared across studies and countries.

Records on institutions, study sites, investigators, studies and therapeutic areas were **converted and loaded**, with SHACL shapes and quality checks at each step, into a knowledge graph of more than 20 million triples in the first phase, close to 150 million with the 2026 work. Identifiers were designed so the same entity from different sources **aligns**. Built-in vector search proposed links from the CRO's disease terms to the MONDO disease ontology, with no separate vector database. Competency questions, the business questions the graph must answer, were defined with the CRO's experts; in the first phase, the graph answered most of them using only SPARQL.

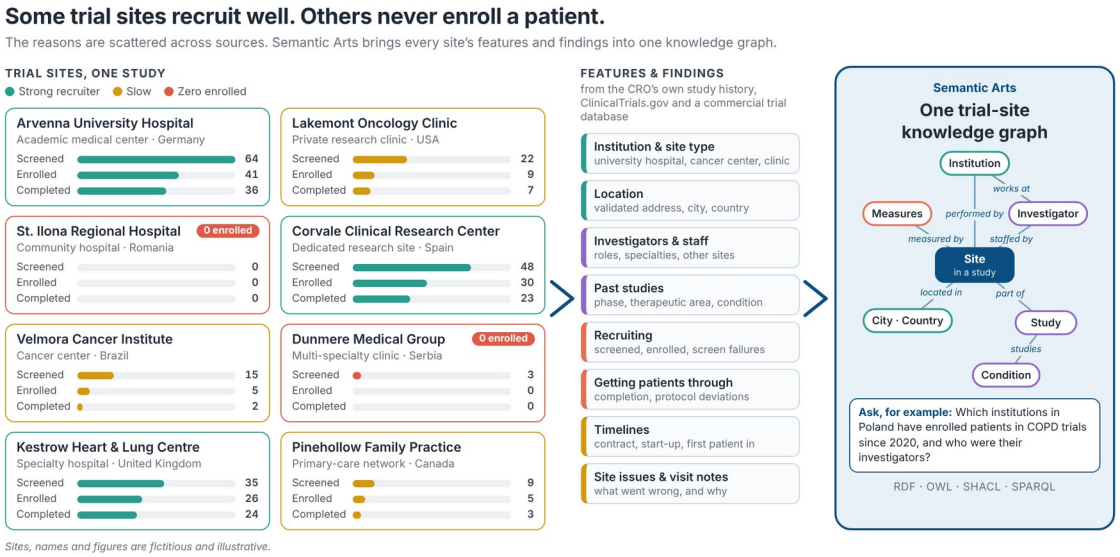


Fig 2: Sites with very different recruiting records, and everything known about them, brought into one knowledge graph.

What it makes possible

The graph puts a site's track record across the CRO's own studies and everyone else's in one place. It can find institutions with experience in a given condition and country, and those with none; follow an investigator across institutions and studies; and compare sites on recruiting speed, screen failures, start-up time and protocol deviations. With conditions linked to a standard, studies can be grouped by disease whatever words the sources use.

What's next

The graph became a foundation. Semantic Arts went on to design the ontology, patterns and taxonomy that turn ubiquitous trial protocol documents into knowledge-graph data, organized by the ICH M11 structured-protocol template. As protocols flow in, the CRO can lean on the historical graph to choose sites. Semantic Arts is scoping how to add site issue reports and meeting notes.

Your next move

Bring Semantic Arts your most difficult ClinOps data challenge, whether it's site selection, feasibility or protocol complexity. We have likely solved something similar. We model your data semantically so you can store it, find it, use it in its entirety, scale it and own it completely. The data's meaning is self-contained and available for your APIs or LLM interactivity. Book an appointment with a data architect at info@semanticarts.com and let us know your vision.

semantic arts

CASE STUDY · PHARMA · CLINICAL TRIALS

Hundreds of pages, thousands of facts: digital clinical trial protocols for a global CRO

How Semantic Arts helped a global contract research organization turn clinical-trial protocols from PDF documents into connected, standards-aligned data, with a large language model guided by an ontology, data patterns and a human-in-the-loop (HITL) checking the result.

CLIENT	Global full-service contract research organization (CRO), 3,000+ employees in 50+ countries
DOMAIN	Clinical-trial protocols: bidding, planning and site selection
WHAT WE DID	Ontology, triple patterns and taxonomy that guide AI extraction; human-review and audit-trail design
SOURCES	Sponsors' clinical-trial protocols, as PDF documents
STANDARDS	RDF, OWL, SHACL, SPARQL, JSON-LD; gist and gistPharma; ICH M11 structured protocol
WHEN	Late 2025 to 2026
STATUS	In final testing (Oct 2026); work continues

The challenge

Every clinical trial starts with a protocol: the sponsor's plan for who can take part, what treatment they receive and which tests happen at each visit. A contract research organization (CRO) bidding to run the trial must read it closely, because it drives cost, site choice and timeline, and the CRO needs to decide, within days, how to bid.

Protocols arrive as PDFs of 50 to 150 pages or more, with the facts spread through narrative text and large tables such as the schedule of activities. Pulling out keywords and passages of text is not enough: the facts must connect to each other and to a reference database at the CRO to allow decision support.

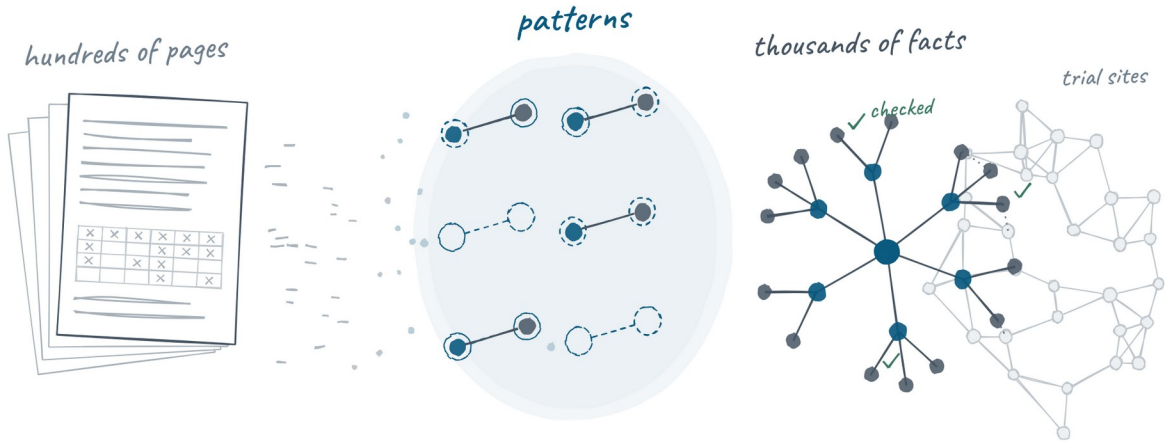


Fig 1: Concept. The AI reads every protocol through a lens of patterns, so hundreds of pages come out as facts of a known shape, checked by a person and joined to the graph of the CRO's trial sites.

This CRO already had its trial-site history in a knowledge graph built with Semantic Arts. Next, it wanted every protocol digested (ETL) into that same graph.

What Semantic Arts did

We wrote the rules the AI follows. An ontology defines the terms and their valid relationships. About 350 **triple patterns** then specify exactly which fact types must be identified from each part of a protocol, organized by the sections of **ICH M11**, the new international template for structured protocols. The patterns draw on **controlled vocabularies** and on the CRO's own ontology, built on gist (Semantic Arts' open upper ontology) and gistPharma, so the LLM picks agreed terms instead of inventing its own.

We proved the method, then handed it over. The Semantic Arts team produced output that convinced the client. A **reference pipeline**, tested on public protocols, splits each protocol into M11 sections, extracts and merges them with one identifier per entity, and checks the result. The CRO runs it on a large language model inside its own firewall. Quality checks are built in: more than **800 automated tests** guard the patterns, and every extraction is checked for structure and vocabulary, with automatic fixes, before review.

We kept a person in the loop. In the review design, people confirm or correct one fact at a time, beside the PDF page it came from, and each decision records its status, page, reviewer and time. Original, edited and final versions sit in separate named graphs. Extra patterns tame the **schedule of activities**: the graph holds planned visits and the assessments due at each, and keeps the table's rows and columns so the review screen can show the familiar grid. Each protocol then loads, in its own named graph, into **the same knowledge graph** as the CRO's studies, sites, institutions and investigators.

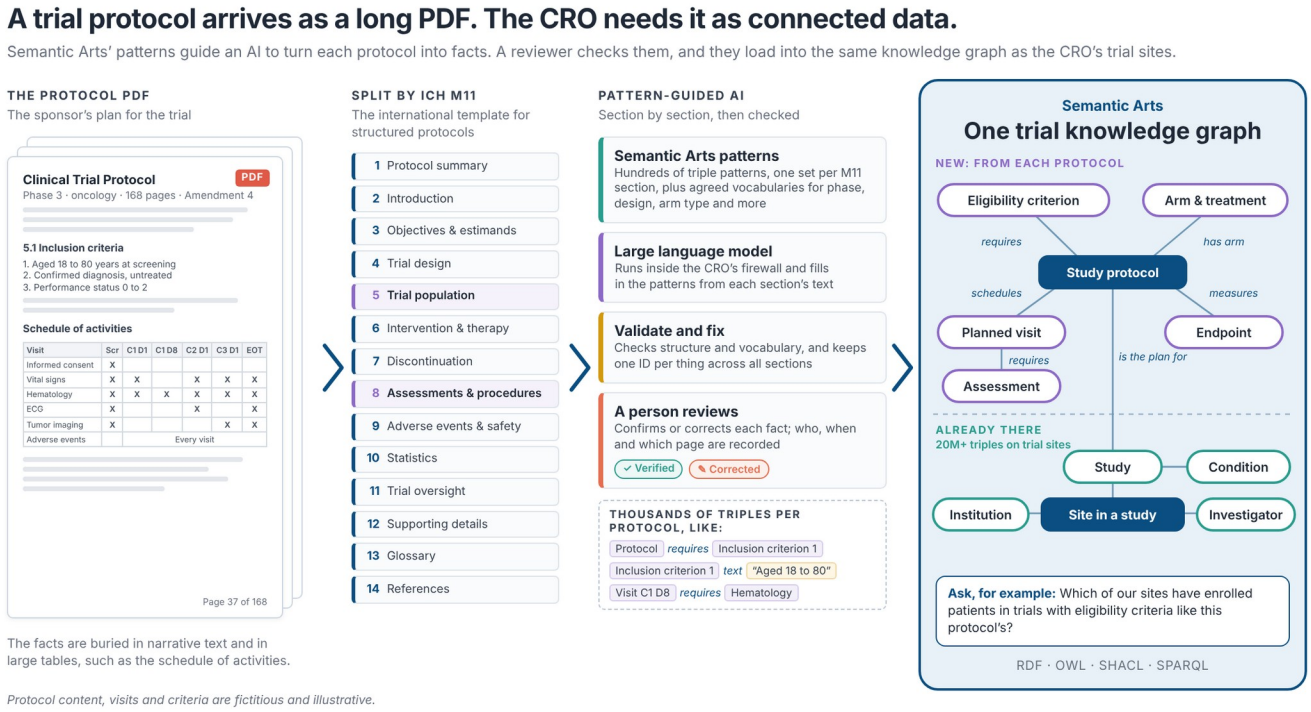


Fig 2: A protocol PDF becomes thousands of connected facts, reviewed and loaded into the same knowledge graph as the historical data.

What it makes possible

The approach lets the CRO turn a new protocol into structured, queryable data in one automated run, ready for a person to check. Analysts can ask questions across protocols, such as which trials require a given assessment or share eligibility criteria, and put a protocol's requirements next to the track records of the sites that might run it. Every fact traces back to its place in the document and to the person who confirmed it. And because the work is organized by ICH M11, it stays in step as sponsors and regulators move to structured protocols.

What's next

The parser and review flow are in final testing. Next, the CRO plans to use protocol requirements to choose sites, and Semantic Arts is scoping the same approach for site issue reports and meeting notes.

Your next move

Bring Semantic Arts your most difficult clinical-trial document challenge, whether it is protocols, amendments or other trial documents. We have likely solved something similar. We model your data semantically so you can store it, find it, use it all, scale it and own it completely. The data's meaning is self-contained and available for your APIs or LLM interactivity. Your next step: Book an appointment with a data architect at info@semanticarts.com and let us know your vision.

semantic arts

CASE STUDY · LIFE SCIENCES · DISCOVERY RESEARCH

No facts, only claims: an evidence-graded discovery graph for a global agricultural-science leader

How Semantic Arts helped an international crop-protection R&D team build a knowledge graph in which what isn't a fact is a claim, traced to the analysis and method that produced it.

CLIENT	Global agricultural-science leader; bioinformatics in crop-protection R&D
DOMAIN	Discovery research on crop-protection targets; a close analog of drug-target discovery
WHAT WE DID	Technology choice; ontologies on gist; a claims model with provenance; automated load pipelines
SOURCES	Ortholog analyses, RNA-seq expression data and public biology (Gene Ontology, FlyBase, UniProt)
STANDARDS	RDF and RDF-star (ready for RDF 1.2), OWL, SKOS, SPARQL; gist; public ontology identifiers
WHEN	2026
STATUS	Proof of concept complete; scaling to about 200 species (Oct 2026)

The challenge

A global agricultural-science leader develops crop-protection products that must stop the target pest but spare bees, birds, fish, pets and people. That means joining genome comparisons across dozens of species with gene expression data from internal lab experiments and from public biology. Its pipelines wrote results to unlinked tables, which each scientist analyzed using their own scripts.

The company's bioinformatics team started from a principle: in biology there are almost no facts. Two analyses can disagree about the same gene, each for good reasons, and before choosing the next experiment a scientist needs to know what stands behind each result.

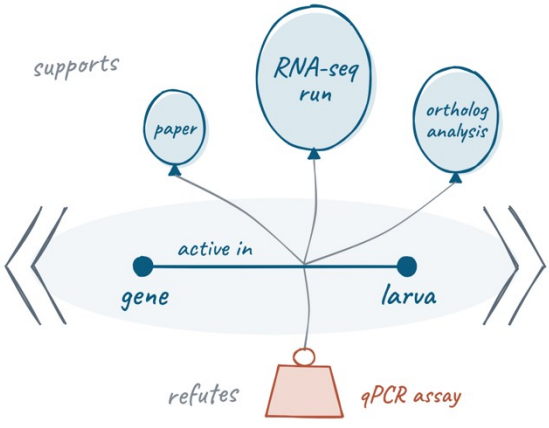


Fig 1: Early concept. A result is not stored as a bare fact but held as a claim. The analyses that buoy or sink it stay attached, each with its respective confidence.

What Semantic Arts did

Semantic Arts first **helped choose the technology.** We laid out the options, and the R&D team chose an RDF knowledge graph, hosted on its own infrastructure, for model expressivity, total data provenance, speed, auditable security and compatibility with natural-language querying. We then **modeled the biology and the science** as two models on gist, Semantic Arts' open upper ontology: one for genes, proteins, pathways and mode-of-action targets, and one for the scientific process, from hypothesis to experiment, equipment and finding. Together they are designed to cover all of crop-protection R&D.

The central design choice was to **create relationships as claims.** With RDF-star, most relationships are stored as claims that point to the analysis behind them, whether a wet-lab experiment, a published paper or a bioinformatics run, and to its method and method version. Contradictory findings live side by side, available to drive decision support. Public biology is **reused on its own terms:** identifiers from Gene Ontology, FlyBase and NCBI Taxonomy are valued so the data contains connections to these outside sources, while their terms are positioned according to standards within the overarching model. Each source has its own named graph, so an update is a quick reload.

Semantic Arts also **automated the loading,** with pipelines that turn the outputs of the team's genome-comparison and expression workflows into graph data with no hand-editing, and **proved it with real questions.** The proof of concept, about 36 million triples, joined two insect species (one a pest with no public gene annotations) with an expression experiment of about 100 samples. It answered the agreed questions, such as which genes active in pest larvae have well-annotated counterparts in the reference fruit fly genome.

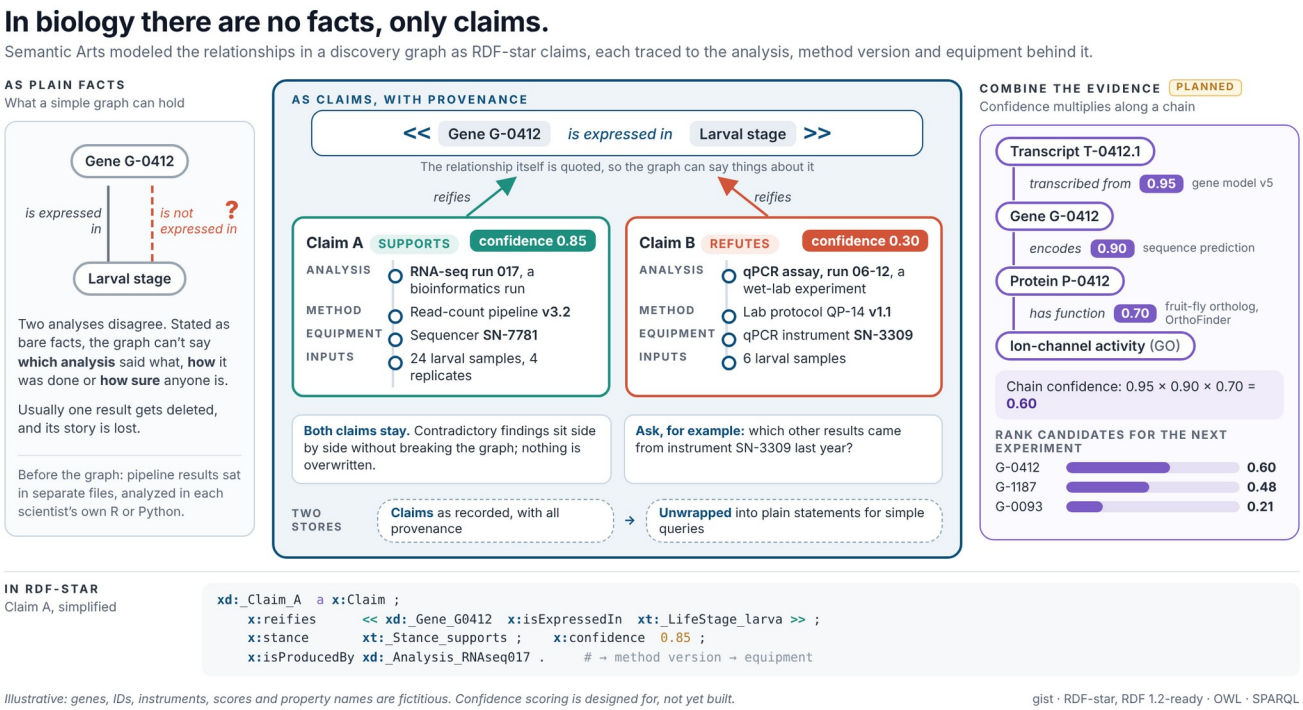


Fig 2: One relationship, two analyses that disagree. Each claim keeps its analysis, method version, equipment and confidence, so the evidence is weighed, not overwritten. Illustrative data.

What it makes possible

Scientists can now see every result about a gene or target together with the analysis and method behind it. Conflicting findings stay side by side, so decisions are made with all the evidence in view. And knowledge can be carried from well-studied species to poorly studied pests through gene orthologs, while scientists also check for counterparts in humans and other species that must remain unaffected.

Why it matters in pharma

Drug discovery has the same problem. Target validation weighs assay results, literature and predictions that often disagree; translational teams move between animal models and human biology; safety teams look for off-target effects through human homologs. A claims graph with provenance keeps all of that evidence in one place, traceable and ready to weigh, rather than flattened into one answer.

What's next

Semantic Arts is now scaling the graph from 2 species to 200, with a billion triples. The claims are built with the downstream step in focus: a confidence score on each one, combined from transcript to gene to protein to function, so scientists can rank targets for their next experiments.

Your next move

Bring Semantic Arts your most difficult discovery-data challenge, whether it is conflicting assay results, cross-species evidence or experiment provenance. We have likely solved something similar. We model your data semantically so you can store it, find it, use it all, scale it and own it completely. The data's meaning is self-contained and available for your APIs or LLM interactivity. Book an appointment with a data architect at info@semanticarts.com and let us know your vision.

semantic arts

CASE STUDY · PHARMA & LIFE SCIENCES · TRAINING & ENABLEMENT

Full circle: ontology and knowledge-graph training for pharma teams

How Semantic Arts teaches pharma, biotech and healthcare teams to model, build and query knowledge graphs themselves, from a two-day leadership workshop to a hands-on course for data engineers.

CLIENTS	Data-first organizations in the US and Europe from the pharmaceutical, healthcare, supply chain and financial sectors
DOMAIN	R&D, operations, customer delivery
WHAT WE DID	Custom ontology courses, hands-on knowledge-graph training, strategy workshops, coaching on project planning and oversight
FORMATS	Two-day workshops; six-session courses, in person or virtual; modules of Semantic Arts' ontology design course
STANDARDS	RDF, OWL, SHACL, SPARQL; the gist upper ontology; tools such as Protégé, TopBraid EDG and TARQL
WHEN	2015 to 2026

The challenge

Pharmaceutical corporations and other complex organizations are building knowledge graphs across research, manufacturing, regulatory and clinical data, and they keep running into the same limit: too few people who can model. Most data engineers are strong on tables and SQL, but standards such as RDF, OWL, SPARQL and semantic tools are new to them. Architects need judgment a TikTok doesn't give: which upper ontology to start from, how a taxonomy sits beside an ontology, and when a rule belongs in inference rather than in a data-quality check. Leaders want to know where semantics will pay off before they steer the ship toward the future.

Underneath is a simple wish. A company that brings in consultants to build its first models wants its own people to leverage and extend those models, consistently, once the consultants step back.

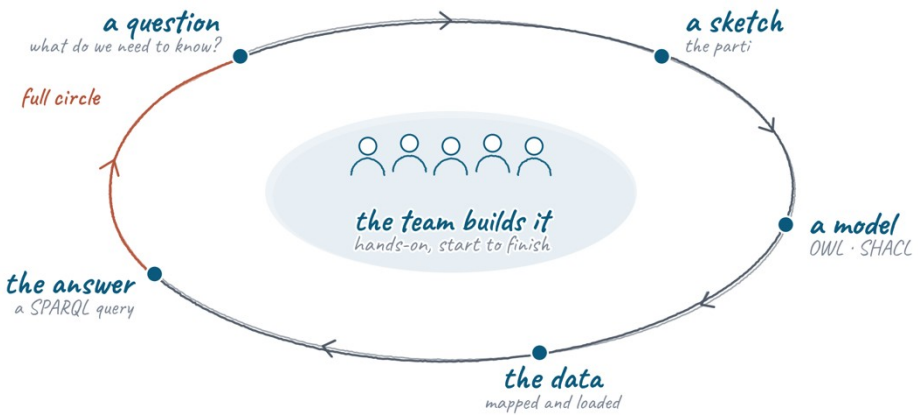


Fig 1: The idea behind the hands-on courses. The class runs one full loop itself: it writes a question, sketches and models an answer, loads the data, and finishes when the graph it built answers the question it started with.

What Semantic Arts can provide

Semantic Arts shapes each program around who is in the room and what they will build. For instance, for R&D informatics and knowledge-management staff, it ran a bespoke two-day workshop. Day one was a visioning exercise and a map of key capabilities onto the company's own use cases; day two worked toward an economic-impact model and asked where to start and which capabilities are foundational. Semantic Arts brought in an outside expert for it, a former DARPA program manager for the research that led to OWL.

For a semantic-layer team (its ontologists, data architects and data engineers), Semantic Arts taught a six-session course: OWL building blocks, properties and classes, the gist upper ontology, taxonomies alongside ontologies, SPARQL and SHACL. The client shared its own ontology modules and modeling guidelines with the instructors, and during "office hours" they sketched models for questions the team brought, such as clinical-trial milestones and a bill of materials. The course took on the team's live design questions, including which upper ontology to build on and how an OWL inference differs from a SHACL constraint.

For onboarding data engineers, data scientists and subject-matter experts to participate in knowledge-graph projects, Semantic Arts ran the full loop shown in Fig 1 in six two-hour sessions, capped at 20 learners so the instructors could offer individualized assistance. Learners queried a public knowledge graph to see what is possible, wrote competency questions, sketched a parti, turned it into an ontology and SHACL shapes, and mapped public COVID-19 data sets to triples with TARQL, minting identifiers and using named graphs so a bad load could be backed out. In the last session they wrote the SPARQL queries that answered the questions they had started with. An engineer from the client's staff taught one segment and gave a tour of the client's internal modeling tool.

The same approach, adjusted, has served other clients. A specialty-materials company took a tailored version of Semantic Arts' ontology design course, with hands-on sessions in the Protégé editor, plus a review of a product taxonomy it had built. A European pharmaceutical company heard Semantic Arts' shortcut to FAIR data: build a core enterprise ontology, populate it with real data, then catalog data sets against it. Other pharma and healthcare clients have taken customized training and advisory on semantics and knowledge graphs. The instructors are practicing Semantic Arts ontologists, the same people who build client graphs.

What it makes possible

Each program is built so that learners leave having done the work themselves: writing the questions a graph must answer, sketching a model, expressing it in OWL and SHACL, loading data and proving the result with SPARQL. Modelers and data engineers come away with one vocabulary and one method, so a model can keep growing without its first authors in the room. Leaders get a concrete picture of where semantics fits, what it could be worth and where to start. And because the instructors also build graphs, teaching continues on the job: at one biopharmaceutical leader, Semantic Arts consultants went on to coach client staff in semantic modeling during later projects.

Why it matters in pharma

Pharma knowledge graphs reach across discovery, CMC, regulatory and clinical data, and they lean on standards such as BFO, IDMP-O, CDISC and HL7 FHIR that must be modeled the same way every time. A graph like that keeps its value only if the company's own people can extend it, and AI assistants that read the graph are only as reliable as its model. Training the team is how a semantic layer outlives the engagement that built it.

What's next

Semantic Arts tailors each program to the team and its domain, from a leadership workshop to multi-session hands-on courses and coaching, in person or virtually. Training is often paired with a first modeling project, so the team practices on the domain it will build first.

Your next move

Bring Semantic Arts the team you want to grow, whether it is data engineers new to graphs, modelers ready for OWL and SHACL, or leaders deciding where semantics pays off. We will shape the course around your domain, your tools and your questions. We model your data semantically so you can store it, find it, use it all, scale it and own it completely. Book an appointment with an ontologist at info@semanticarts.com and let us know your training vision.