Data First: The Case for FAIR-by-Design (Meta)data Standardization in New Approach Methodologies

Position Paper

Author: Damien Huzard
Affiliation: Neunonautix, Montpellier, France
Version: 1.0 (post Stage 4 revision, 2026-05-11)

Abstract

New Approach Methodologies (NAMs) are producing regulatory-grade evidence: the human Liver-Chip identifies drug-induced liver injury with 87% sensitivity and 100% specificity, with an estimated annual productivity gain exceeding three billion US dollars if adopted across small-molecule preclinical workflows (Ewart et al., 2022). Yet most NAM data sits in incompatible formats, cannot be reused across studies, and cannot be aggregated for the AI methods that would multiply its value. This paper argues that the bottleneck is no longer the science — it is the data. I advance three claims. First, with current technology, FAIR-by-design standardization is not a burden; it is a bounded upfront investment that pays compounding returns. Second, the infrastructure already exists: the New Approach Methodology Ontology and Schema (NAMO), maintained by the Monarch Initiative, gives NAM data a shared vocabulary and schema that makes it Findable, Accessible, Interoperable, and Reusable by construction. Third, the regulatory window is open now — the FDA's March 2026 draft guidance on NAMs makes structured, context-of-use-aware data a precondition for acceptance, and the organizations that standardize first will submit faster, collaborate more effectively, and support development of the AI models that will multiply every NAM platform's value. I close with a differentiated call to action for scientists, contract research organizations, sponsors, and regulators.

Keywords: new approach methodologies; FAIR data; data standardization; NAMO; regulatory toxicology; ontology

1. Introduction: The Pivot Moment

The human Liver-Chip study by Ewart and colleagues (2022) is the kind of result the NAM field has been waiting two decades to produce. Across a blinded panel of characterized drugs, the platform flagged 87% of compounds that cause drug-induced liver injury in humans and misclassified none of the safe ones. The authors estimate that broad adoption in preclinical small-molecule pipelines could return more than three billion dollars a year in avoided late-stage failures. That number alone justifies serious platform investment — with the caveat, addressed in Section 3, that the estimate was produced by the platform's developers and should be read as the best available projection rather than an independently verified figure.

But the value in that estimate is conditional. It assumes the data those platforms generate can be found, compared against historical results, pooled across laboratories, and submitted to regulators in a form they can evaluate. Right now, for most NAM data, none of that is true. Results are recorded in spreadsheets with idiosyncratic column names, in lab-specific schemas, in formats that the next study cannot read and that no machine can query. The asset is being created and discarded in the same motion.

Three things have converged. NAM science is delivering — not only the Liver-Chip, but a broader generation of microphysiological systems, high-content in vitro assays, and computational models maturing toward regulatory readiness (Basketter et al., 2012; OECD, 2005). Artificial intelligence is ready to consume this data at scale, provided the data is structured and harmonized. And the regulatory posture has shifted: in March 2026 the U.S. Food and Drug Administration issued draft guidance on the general use of NAMs in drug development, organizing acceptance around context of use, human biological relevance, technical characterization, and fit-for-purpose validation (FDA, 2026). That guidance remains in draft and may be revised before finalization; its direction is nonetheless clear, and each of its four pillars depends on metadata. You cannot establish a context of use for a dataset whose context was never recorded in a queryable form.

What has not kept pace is data infrastructure. No widely adopted framework currently connects NAM data generation to FAIR-by-design practice in routine laboratory settings. This paper argues that closing that gap is the highest-leverage investment a NAM stakeholder can make today, and that the means to do it — the NAMO schema, the FAIR principles, and a generation of AI-assisted curation tools — are already in hand. The argument runs from the problem (Section 2), through the regulatory stakes (Section 3), to the solution (Section 4) and its implementation (Section 5), and ends with what each part of the ecosystem needs to do (Section 6).

---

2. The Data Crisis in NAMs

Ask a NAM scientist what slows their work and the answer is rarely the assay. It is the data around the assay: the inability to reuse it, the inability to combine it with anyone else's, and the dead weight of historical results that no longer mean anything because the context that gave them meaning was never written down in a usable way. These are not three separate problems. They are three faces of one absence — there is no common data model for NAM studies.

Data cannot be reused. A result generated for one submission cannot be repurposed for the next without manual reconstruction of what the variables meant, how the assay was run, and which controls applied. Each study starts from a blank page even when the laboratory has run the same protocol fifty times. The cost is paid again and again.

Data is not interoperable. Two laboratories running the same organ-on-a-chip assay record their endpoints under different names, in different units, against different implicit ontologies of cell type and tissue. Pooling their data — the basic move behind any meta-analysis, any cross-platform comparison, any weight-of-evidence argument — requires a bespoke harmonization project each time. At the scale the field is now operating, that does not happen. The data stays siloed because integrating it is too expensive to be worth it case by case.

Historical data cannot be leveraged. Years of NAM results exist as unstructured records: free-text protocols, PDFs, spreadsheets with no schema. For the AI methods that need large, harmonized training sets, this corpus is effectively invisible. The longer it accumulates without structure, the more expensive the eventual remediation becomes, and the more of it ages past the point where reconstruction is feasible at all.

These failures matter more now than they did five years ago for a simple reason: there is more NAM data than ever, the rate is increasing, and AI is finally able to do something useful with it — but only if the data is FAIR (Wilkinson et al., 2016). A field that generates data faster than it can structure it is digging a hole, and the slope is getting steeper.

Why has this not already been fixed? The question is worth asking directly, because it demands a straight answer from anyone making the case for standardization. The reasons are real. The tooling for ontology-based data capture has historically required specialised informatics expertise that most experimental laboratories do not have in-house. Vocabulary mapping is time-consuming when done manually. LIMS vendors have had little commercial incentive to build NAMO-compliant export modules before NAMO had a critical mass of adopters — a classic coordination problem. And until the FDA formalized NAM data expectations in 2026, the regulatory forcing function that drove SEND adoption in in vivo nonclinical work was absent from the NAM side. None of these reasons make the status quo acceptable, but they do explain it, and any realistic implementation path has to address them. Section 5 argues that the tooling has now moved past these barriers for the forward-looking problem, even if legacy remediation remains harder.

There is an emotional truth here that the rest of this paper builds on. Scientists in this field already feel the problem. They have lived the harmonization project that ate a month. They have opened a five-year-old dataset and found it unreadable. The argument I want to make is not that this pain is real — they know it is real — but that the solution already exists, and that it costs far less than the pain does.

---

3. The Regulatory Imperative: Data Standard Before Artifact

Regulators are not, in the first instance, asking for better assays. They are asking for better data about assays. The FDA's March 2026 draft guidance frames NAM acceptance around four features of any validation approach: context of use, human biological relevance, technical characterization, and fit-for-purpose (FDA, 2026). Read those again as data requirements. A context of use is a structured statement about what question a method answers, in what population, under what conditions — it is metadata, and an unstructured dataset cannot carry it. Technical characterization is a record of assay performance parameters in a form that can be compared across runs and laboratories — it is a schema. Fit-for-purpose is a judgment that can only be made if the purpose, and the evidence bearing on it, are both legible to the evaluator. The guidance does not specify NAMO by name; it specifies a set of data properties that NAMO is designed to provide. The case for NAMO in this paper is that it is the appropriate vehicle to satisfy these regulatory data requirements — that argument should be weighed on its merits, not taken as an official endorsement.

The infrastructure for nonclinical data submission already exists in a partial form. The FDA's data standards catalog makes the CDISC Standard for Exchange of Nonclinical Data (SEND) the required format for nonclinical study data in electronic submissions (CDISC, 2023; FDA, 2022). SEND is a real, mandatory, working standard — and it was built primarily for in vivo studies. Some specialized implementation guides have followed: SENDIG-Genetox v1.0, which covers genetic toxicology studies including certain in vitro genotoxicity assays, became mandatory for relevant studies starting after March 2025. But organ-on-a-chip readouts, microphysiological system endpoints, and in silico model outputs remain outside any validated, published SEND extension. The field therefore has a regulatory data standard that covers an important but narrow slice of the NAM data it is now producing, and a regulatory expectation — increasingly explicit — that NAM data will arrive in evaluable form anyway.

The gap has downstream consequences. A weight-of-evidence assessment, the structure regulators use to integrate multiple lines of NAM evidence, is a graph: nodes are studies and endpoints, edges are the relationships between them. Building that graph computationally — which is the only way to build it at scale — requires that the underlying data share a vocabulary. The OECD's Adverse Outcome Pathway Knowledge Base has shown what becomes possible when toxicological knowledge is given a controlled structure: Martens and colleagues (2022) converted the entire AOP-Wiki content into semantic-web RDF triples using more than twenty ontologies, making its content queryable and linkable in ways free text never was. NAM study data deserves the same treatment, and for the same reason. Without it, every weight-of-evidence argument is reassembled by hand.

The regulatory window extends beyond the United States. The European Medicines Agency has published reflection papers on the use of new approach methodologies in regulatory submissions, and the ICH S17 guideline covers NAM use in nonclinical safety evaluation. The direction of travel is consistent globally: regulators are formalizing how NAM evidence will be evaluated, and the data-quality requirements implicit in that formalization will eventually apply to all major markets. This is not a US-only opportunity.

Here is the position, stated plainly: the data standard matters more than the artifact the regulator is currently foregrounding. A perfectly validated assay whose data cannot be submitted in the required structure is a blocked asset. And the forcing function is already visible — regulatory platforms that require data in a specific format push that requirement upstream, into how laboratories collect data in the first place. The cleanest way to satisfy a downstream format requirement is not to convert at the end. It is to be FAIR by design from the start. Organizations that internalize this now will be submitting structured NAM data while their competitors are still building converters.

---

4. NAMO and FAIR-by-Design: The Solution Already Exists

The phrase "FAIR-by-design" does real work, and it is worth being precise about it. FAIRification can happen at two moments. It can happen after the fact — taking a finished, messy dataset and retrofitting identifiers, vocabularies, and provenance onto it. Or it can happen at design time — making the standardization choices when the assay, the data capture system, or the laboratory information system is being set up, so that every dataset is born structured. The first is a project. The second is a bounded upfront investment. Rocca-Serra and colleagues (2023a) lay out a flexible, multi-level framework for FAIRification, and the companion FAIR Cookbook (Rocca-Serra et al., 2023b) collects the practitioner recipes for doing it; the consistent lesson across both is that the earlier in the data lifecycle the FAIR decisions are made, the cheaper and more durable they are (see also Jacobsen et al., 2020, for a detailed operationalization of what each FAIR principle entails in practice). Post-hoc FAIRification is sometimes unavoidable, and Section 5 addresses it. But it is the expensive path, and choosing it by default — by simply never deciding — is the most expensive path of all.

What makes FAIR-by-design tractable for NAMs specifically is that the domain vocabulary already exists. The New Approach Methodology Ontology and Schema, NAMO, is maintained by the Monarch Initiative and is openly available (Monarch Initiative, n.d.). NAMO is built on LinkML; it provides a shared schema and controlled vocabulary for describing the full range of NAM model systems — cellular cultures, organ-on-a-chip devices, computational models, digital twins — together with their specifications, concordance metrics, and validation data. It integrates with established biomedical ontologies, including the Cell Ontology, the anatomy ontology UBERON, and the disease ontology MONDO, so that a NAM dataset described with NAMO inherits semantic links to the wider biomedical knowledge graph. And it exports to OWL, JSON Schema, and Python models, so it slots into whatever technical stack a laboratory or platform already runs. (NAMO is best understood as a semantic data model with ontology-like structure that is exported as an ontology for browsing and querying — not as a hand-built description-logic ontology. The distinction matters for how it is adopted: a laboratory maps its data to a schema, not to a logical theory.)

A note on NAMO's current status is warranted. NAMO is under active development at the time of writing. Coverage is most complete for cellular and microphysiological model systems; schema modules for computational models and digital twins are less fully specified. This paper advocates NAMO as the schema toward which the field should converge and invest — not as a fully mature, field-wide standard already deployed at scale. The governance model for schema evolution (how breaking changes are managed, how the community contributes new NAM types) is an open question that the Monarch Initiative and the NAM community will need to address as adoption grows. These are real constraints on the implementation timeline, but they are engineering and governance problems, not reasons to build competing silos while waiting for a perfect standard.

NAMO is also not the only structured data option. Other biomedical data models — OBI (Ontology for Biomedical Investigations), ISA-Tab/ISA-JSON, MIFlowCyt for flow cytometry, and the broader EBI Ontology Lookup Service ecosystem — each address parts of the problem. NAMO's advantage for the NAM domain specifically is its scope: it is designed around the full range of NAM system types, it is maintained by the same Monarch Initiative group that developed SSSOM and OntoGPT (making tooling interoperability likely), and it integrates the specific biomedical ontologies most relevant to NAM study design. The recommendation here is to invest in converging on NAMO, not to treat it as the only possible vocabulary — but convergence on a single schema is the precondition for interoperability, and deferring that choice indefinitely is a cost, not a hedge.

The payoff is the FAIR acronym, satisfied by construction rather than by retrofit. A dataset built on NAMO is **Findable**, because its entities carry persistent identifiers and rich metadata. It is **Accessible**, because it is expressed in open, standard serializations. It is **Interoperable**, because it speaks a shared vocabulary that other NAMO datasets — and the biomedical ontologies NAMO links to — also speak. And it is **Reusable**, because the provenance, the assay conditions, and the validation context travel with the data instead of living in someone's memory. The four problems of Section 2 are not patched one at a time. They dissolve, because they were all symptoms of the missing common model, and NAMO is that model.

This is not an experiment with an uncertain payoff. The precedent is one of the best-documented successes in the history of biological data infrastructure. The Gene Ontology, launched in 2000 as a shared controlled vocabulary across model-organism databases, transformed how biological function is recorded, queried, and computed on (Ashburner et al., 2000). The critical caveat is that GO's success rested partly on a forcing function that does not yet exist for NAMO: major journals and model-organism databases made GO annotation a condition of publication and data submission. The NAM field's analogue to that forcing function — a regulatory requirement, a consortium data-sharing agreement, or a funder mandate — is the missing accelerant. The Roadmap call to action in Section 6 is partly about creating that forcing function. The Gene Ontology analogy is instructive precisely because it shows what becomes possible once the vocabulary exists and the community commits: biomedical ontologies have since proven their value across knowledge management, data integration, and decision support in setting after setting (Bodenreider, 2008). The Monarch Initiative itself has spent more than a decade integrating phenotype and genotype data across dozens of species databases through exactly this ontology-driven approach (Mungall et al., 2017), and has contributed standards such as the Simple Standard for Sharing Ontological Mappings (SSSOM) that make the cross-vocabulary mapping step itself reproducible (Matentzoglu et al., 2022). NAMO follows a model that works. The community context reinforces it: domain-specific reporting standards have repeatedly proven more actionable than generic ones (Taylor et al., 2008), and NAMO development is happening in collaboration across institutions. The path is well trodden and it is not lonely.

---

5. Implementation: Lower Barrier Than You Think

The honest summary of what adoption costs is short: it is not that difficult for the forward-looking problem, and it is harder but not hopeless for the legacy problem. With current tooling, FAIR-by-design is a bounded upfront investment — primarily vocabulary mapping and LIMS or data-pipeline integration — that does not require per-study effort once in place.

The FAIR-by-design workflow. The standardization investment is made once, when the data capture system or analysis pipeline is configured. After that, the standardization runs in the background. Scientists work the way they always have; the structure is applied as data is recorded, not bolted on afterward. There are two practical patterns. In the first, NAMO vocabulary is integrated at the data-collection layer — the LIMS, the assay export template, the electronic lab notebook — so that data is born NAMO-compliant. In the second, the laboratory works through an application that is connected to NAMO and to the relevant third-party systems and document repositories, and the application handles the mapping. Either way, the scientist's daily workflow does not change. The investment is structural — it involves LIMS integration work, controlled-vocabulary governance, and a mapping exercise described below — and it is made by whoever stands up the system, not by each experimenter anew.

The mapping step. Adoption does involve one real piece of work: a laboratory has to map its own entity names — its terms for cell types, endpoints, conditions, model systems — onto the NAMO vocabulary. Done by hand for a large, multi-assay operation, that is non-trivial. But it does not have to be done by hand. The mapping problem is exactly what AI-assisted curation tools are now good at: a large language model proposes the mapping, a curator reviews it, and the human role shifts from generation to verification. Tools in this space — SPIRES and the OntoGPT family, developed by the same Monarch-affiliated group that maintains NAMO — extract schema-conformant, ontology-grounded structured data from text using zero-shot prompting and existing vocabularies for identifier resolution (Caufield et al., 2024). For complex operations, multi-agent systems can carry the mapping across many assays and data types at once. And the mapping is largely a one-time effort, refreshed only when the laboratory adds genuinely new kinds of data.

Legacy data is not hopeless. The hardest version of the problem is a back catalogue of unstructured results — free-text protocols, PDFs, schemaless spreadsheets — that someone now wants to make usable. This is the case where people give up. They should not. The same AI-assisted extraction approaches that handle forward mapping can mine structured metadata out of historical records (Caufield et al., 2024). It is not full automation, and it should not be sold as such — the extracted mappings need human validation, language models hallucinate, and provenance has to be traced. But the relevant comparison is not "perfect" versus "imperfect." It is "AI-assisted review" versus "fully manual reconstruction," and against that baseline the AI-assisted path is dramatically faster. Even partial remediation of a legacy corpus creates value, and the value compounds as more of it comes online.

Addressing the coordination problem. The section above explained why FAIR-by-design has not already happened: tooling immaturity, LIMS vendor inertia, and a missing regulatory mandate. The tooling barrier has materially lowered with AI-assisted curation. The vendor-inertia problem has a solution: CROs and sponsors who write NAMO-compliant deliverables into contracts create the market signal that moves vendors. The regulatory mandate is forming with the FDA's 2026 guidance and the global regulatory direction described in Section 3. None of these forces are yet aligned into a single forcing function. The call to action in Section 6 is, in part, about how different actors can contribute to creating that alignment.

The compounding return. Once NAM data is NAMO-compliant and FAIR, a series of things that were previously bespoke projects become routine queries. Cross-study meta-analysis becomes possible because the studies share a vocabulary. Weight-of-evidence graphs can be assembled computationally instead of by hand. Context-of-use statements become precise and auditable, which is exactly what the FDA's 2026 framing asks for. Regulatory submission packages can be generated from the structured data rather than re-keyed. And AI models — the ones that would multiply the value of every NAM platform — can be trained on a curated, harmonized corpus instead of a pile of PDFs. Each of these is a return on the same bounded upfront investment. That is what "by design" buys: not a single fix, but a foundation that keeps paying.

---

6. Call to Action: What Each Stakeholder Must Do

The responsibility here is distributed, and that is the point. No single actor can fix this alone, and inaction by any one of them slows the whole ecosystem. The good news is that each actor's part is concrete and within reach.

NAM scientists should treat data as a first-class research output, not a byproduct of the experiment. In practice that means learning to handle data better — adopting NAMO vocabulary at the point of collection, recording context and provenance as the experiment runs, and resisting the habit of treating the spreadsheet as a private artifact. The shift is partly cultural: the most important part of the process is the data, and that is where the care should go. Scientists who contribute to NAMO by submitting new NAM-type definitions and schema modules accelerate the standard for everyone.

Contract research organizations, biotechs, and pharmaceutical sponsors should require better data handling, both internally and from their partners. A CRO contract can specify NAMO-compliant deliverables. A sponsor's data governance can mandate FAIR-by-design infrastructure for the assays it commissions. CROs and sponsors who do this collectively are the market signal that moves LIMS vendors. The investment is bounded — a configuration decision and a mapping exercise — and the return is the compounding one described in Section 5, plus a faster path through regulatory review.

Regulators should make the implicit requirement explicit. The FDA's 2026 draft guidance already organizes NAM acceptance around context of use and fit-for-purpose; the natural next step is to name NAMO or an equivalent structured format as the expected vehicle for NAM data submission, and to extend or supplement SEND so that NAM endpoints have a defined submission format. The EMA and ICH S17 working groups face the same problem and an aligned global position — "NAM data should be FAIR-by-design, with NAMO as the recommended schema" — would be the most powerful coordination signal the field has ever received. A regulatory platform that expects data in a specific structure is the most effective forcing function there is, because it makes FAIR-by-design the path of least resistance for everyone upstream.

The register of this call is not accusatory. It is the opposite. The message to every one of these actors is that they have more leverage than they think and that the tools already exist. The scientist who works through this should come away thinking two things: it is not that difficult, and — now that the data is in order — I can do things with it I could not do before.

---

7. Conclusion: Act Now or Compound Later

The NAM field is at an inflection point. The science is delivering regulatory-grade evidence. AI is ready to consume that evidence at scale. Regulators are formalizing how NAM evidence will be evaluated. The one piece that has not caught up is the data infrastructure, and the means to fix it — the NAMO schema, the FAIR principles, AI-assisted curation — are already available and already proven in adjacent domains.

The three claims advanced at the outset have been supported. First, FAIR-by-design standardization via NAMO is not a per-study burden; it is a bounded upfront investment in vocabulary mapping and LIMS integration that does not recur, and whose returns compound across every future study, submission, and model. Second, the infrastructure exists: NAMO provides the NAM-specific schema that the FAIR principles require, with the caveats that NAMO is still maturing and that adoption requires a community commitment the field has not yet made. Third, the regulatory window is open: the FDA's 2026 draft guidance makes structured, context-of-use-aware data a functional precondition for NAM acceptance, and the same direction is visible across the EMA and ICH; the window will not stay open indefinitely.

The logic of compounding runs in both directions. An organization that adopts FAIR-by-design standardization now starts accumulating a structured, reusable, AI-ready data asset, and every future study, submission, and model benefits from it. An organization that defers keeps accumulating an unstructured liability, and the eventual remediation gets more expensive every year, with more of the back catalogue aging past the point of recovery. The question is not whether NAM data standardization will happen. The regulatory direction and the economics both point one way. The question is whether you will be ahead of it or behind it.

Limitations

This is a position paper, not an empirical study, and its argument should be read accordingly. I have not measured the cost of NAMO adoption in a working laboratory, the time required for the mapping step, or current NAMO uptake across the field — to my knowledge those figures have not been published, and establishing them is itself a research gap worth filling. The economic case rests substantially on a single study (Ewart et al., 2022) in one therapeutic context (small-molecule hepatotoxicity); the generalization to other NAMs and endpoints is plausible but not demonstrated, and the $3B figure is an estimate by the same authors who built the Liver-Chip platform. The claim that FAIR-by-design outperforms post-hoc FAIRification is argued from first principles and from practitioner experience codified in the FAIR Cookbook (Rocca-Serra et al., 2023b), not from a controlled head-to-head comparison, which has not been done. The FDA 2026 guidance is a draft document that may be revised before finalization, and the regulatory stakes argued in Section 3 are contingent on the final form. And the AI-assisted curation path for legacy data, while real, carries known risks — hallucination, provenance loss — that mandate human validation and that I do not want to minimize. None of these caveats undercuts the direction of the argument, but they bound its strength, and a sponsor weighing investment should weigh them.

Declarations

Author Contributions (CRediT). Damien Huzard: Conceptualization, Methodology, Writing – original draft, Writing – review & editing.

Funding. This work received no external funding. The research was conducted independently by the author.

Institutional Review Board Statement. Not applicable.

Informed Consent Statement. Not applicable.

Data Availability Statement. This position paper does not report new primary data. All sources cited are publicly available; the NAMO schema is openly available from the Monarch Initiative (https://monarch-initiative.github.io/namo/).

Conflicts of Interest. The author declares no conflict of interest.

AI Usage Disclosure. This paper was developed with the assistance of the Academic Research Skills plugin v3.7.0 (ARS, Claude Code, Anthropic). AI assistance was used for literature synthesis, argument structuring, and drafting. All conceptual framing, editorial decisions, and factual claims were directed and reviewed by the author. AI-generated citations were independently verified prior to drafting; seven fabricated citations identified in an initial bibliography were removed or replaced with verified sources.

Affiliation. Damien Huzard, Neunonautix, Montpellier, France.

References

Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., Harris, M. A., Hill, D. P., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J. C., Richardson, J. E., Ringwald, M., Rubin, G. M., & Sherlock, G. (2000). Gene Ontology: Tool for the unification of biology. *Nature Genetics, 25*(1), 25–29. https://doi.org/10.1038/75556

Basketter, D. A., Clewell, H., Kimber, I., Macmillan, A., Blaauboer, B., Boverhof, D., Burton, A., Cockshott, A., Cronin, M., Doe, J., Embry, M., Foster, J., Hartung, T., Jacobs, M., Knight, D., MacKay, C., Mahony, C., Margiotta-Casaluci, L., Marshall, P., & the systemic toxicity working group. (2012). A roadmap for the development of alternative (non-animal) methods for systemic toxicity testing. *ALTEX, 29*(1), 3–91. https://doi.org/10.14573/altex.2012.1.003

Bodenreider, O. (2008). Biomedical ontologies in action: Role in knowledge management, data integration and decision support. *Yearbook of Medical Informatics, 17*(1), 67–79. https://doi.org/10.1055/s-0038-1638585

Caufield, J. H., Hegde, H., Emonet, V., Harris, N. L., Joachimiak, M. P., Matentzoglu, N., Kim, H., Moxon, S. A. T., Reese, J. T., Haendel, M. A., Robinson, P. N., & Mungall, C. J. (2024). Structured Prompt Interrogation and Recursive Extraction of Semantics (SPIRES): A method for populating knowledge bases using zero-shot learning. *Bioinformatics, 40*(3), btae104. https://doi.org/10.1093/bioinformatics/btae104

Clinical Data Interchange Standards Consortium. (2023). *Standard for Exchange of Nonclinical Data Implementation Guide, version 3.1.1* (SENDIG v3.1.1). CDISC. https://www.cdisc.org/standards/foundational/send

Ewart, L., Apostolou, A., Briggs, S. A., Carman, C. V., Chaff, J. T., Heng, A. R., Jadalannagari, S., Janardhanan, J., Jang, K.-J., Joshipura, S. R., Kadam, M. M., Kanellias, M., Kujala, V. J., Kulkarni, G., Le, C. Y., Lucchesi, C., Manatakis, D. V., Maniar, K. K., Quinn, M. E., & Hamilton, G. A. (2022). Performance assessment and economic analysis of a human Liver-Chip for predictive toxicology. *Communications Medicine, 2*, 154. https://doi.org/10.1038/s43856-022-00209-1

Jacobsen, A., de Miranda Azevedo, R., Juty, N., Batista, D., Coles, S., Cornet, R., Courtot, M., Crosas, M., Dumontier, M., Evelo, C. T., Goble, C., Guizzardi, G., Hansen, K. K., Hasnain, A., Hettne, K., Heringa, J., Hooft, R. W. W., Imming, M., Jeffery, K. G., & Schultes, E. (2020). FAIR Principles: Interpretations and implementation considerations. *Data Intelligence, 2*(1–2), 10–29. https://doi.org/10.1162/dint_r_00024

Martens, M., Evelo, C. T., & Willighagen, E. L. (2022). Providing adverse outcome pathways from the AOP-Wiki in a semantic web format to increase usability and accessibility of the content. *Applied in Vitro Toxicology, 8*(1), 2–13. https://doi.org/10.1089/aivt.2021.0010

Matentzoglu, N., Balhoff, J. P., Bello, S. M., Bizon, C., Brush, M., Caron, A. R., Chan, L. E., Bradford, Y. M., Carmody, L. C., Cuzick, A., D'Eustachio, P., Dröge, J., Engel, S. R., Fey, P., Garde, S., Gerken, J., Gourdine, J.-P., Harris, N. L., Harris, T. W., & Mungall, C. J. (2022). A Simple Standard for Sharing Ontological Mappings (SSSOM). *Database, 2022*, baac035. https://doi.org/10.1093/database/baac035

Monarch Initiative. (n.d.). *NAMO: New Approach Methodology Ontology and Schema*. Retrieved May 2026, from https://monarch-initiative.github.io/namo/

Mungall, C. J., McMurry, J. A., Köhler, S., Balhoff, J. P., Borromeo, C., Brush, M., Carbon, S., Conlin, T., Dunn, N., Engelstad, M., Foster, E., Gourdine, J.-P., Jacobsen, J. O. B., Keith, D., Laraway, B., Lewis, S. E., Nguyen Xuan, J., Shefchek, K., Vasilevsky, N., & Haendel, M. A. (2017). The Monarch Initiative: An integrative data and analytic platform connecting phenotypes to genotypes across species. *Nucleic Acids Research, 45*(D1), D712–D722. https://doi.org/10.1093/nar/gkw1128

Organisation for Economic Co-operation and Development. (2005). *Guidance Document No. 34 on the validation and international acceptance of new or updated test methods for hazard assessment* (ENV/JM/MONO(2005)14). OECD Publishing.

Rocca-Serra, P., Welter, D., Juty, N., Abbassi-Daloii, T., Capella-Gutierrez, S., Chandramouliswaran, I., Splendiani, A., Burdett, T., Giessmann, R. T., Henderson, D., Batista, D., Emam, I., Gadiya, Y., Giovanni, P. D., Willighagen, E., Evelo, C., Gormanns, P., & Sansone, S.-A. (2023a). FAIR in action — A flexible framework to guide FAIRification. *Scientific Data, 10*, 291. https://doi.org/10.1038/s41597-023-02167-2

Rocca-Serra, P., Gu, W., Ioannidis, V., Abbassi-Daloii, T., Capella-Gutierrez, S., & Sansone, S.-A. (2023b). The FAIR Cookbook — The essential resource for and by FAIR doers. *Scientific Data, 10*, 292. https://doi.org/10.1038/s41597-023-02166-3

Taylor, C. F., Field, D., Sansone, S.-A., Aerts, J., Apweiler, R., Ashburner, M., Ball, C. A., Binz, P.-A., Bogue, M., Booth, T., Brazma, A., Brinkman, R. R., Clark, A. M., Deutsch, E. W., Fiehn, O., Fostel, J., Ghazal, P., Gibson, F., Gray, T., & Wiemann, S. (2008). Promoting coherent minimum reporting guidelines for biological and biomedical investigations: The MIBBI project. *Nature Biotechnology, 26*(8), 889–896. https://doi.org/10.1038/nbt.1411

U.S. Food and Drug Administration. (2022). *Data standards catalog*. FDA. https://www.fda.gov/industry/fda-data-standards-advisory-board/data-standards-catalog

U.S. Food and Drug Administration. (2026). *General considerations for the use of new approach methodologies in drug development: Draft guidance for industry*. Center for Drug Evaluation and Research, FDA. *(Issued March 18, 2026; public comment period closed May 18, 2026.)*

Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L. B., Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data, 3*, 160018. https://doi.org/10.1038/sdata.2016.18