

Making Common Fund data more findable: Catalyzing a Data Ecosystem

Amanda L Charbonneau^{1,C}, Arthur Brady², Saranya Canchi¹, Robert Carter², Kyle Chard³, Daniel J.B. Clarke⁴, Heather H Creasy², Mike D'Arcy⁵, Michelle Giglio², Alicia Gingrich⁶, Rayna M Harris¹, Theresa K Hodges², Minji Jeon⁴, Eryk Kropiwnicki⁷, Marisa C.W. Lim¹, R. Lee Liming³, Meisha Mandal⁸, James B. Munro², Suvarna Nadendla², Cia Romano^{5,9}, Philippe Rocca-Serra¹⁰, Robert E. Schuler⁵, Alex Waldrop⁸, Cris Williams⁵, Karen Word¹¹, Susanna-Assunta Sansone¹⁰, Avi Ma'ayan⁴, Rick Wagner¹², Ian Foster³, Carl Kesselman⁵, C. Titus Brown¹, Owen White²

¹ Population Health and Reproduction, UC Davis

² University of Maryland Institute for Genome Sciences, University of Maryland School of Medicine

³ University of Chicago and Argonne National Laboratory

⁴ Department of Pharmacological Sciences, Mount Sinai Center for Bioinformatics, Icahn School of Medicine at Mount Sinai

⁵ University of Southern California Information Sciences Institute

⁶ UC Davis Medical Center

⁷ IDG

⁸ RTI International

⁹ Interface Guru

¹⁰ Oxford e-Research Centre, Department of Engineering Science, University of Oxford, UK

¹¹ The Carpentries

¹² University of California San Diego

^C Corresponding Author achar@ucdavis.edu

Abstract

The Common Fund Data Ecosystem (CFDE) has created a flexible system of data federation that enables users to discover datasets from across the Common Fund without requiring the data owners to move, reformat, or rehost those data. The CFDEs federation system is centered on a catalog that ingests metadata from individual Common Fund Program's Data Coordination Centers (DCCs) into a uniform metadata model that can then be indexed and searched from a centralized portal. This uniform Crosscut Metadata Model (C2M2), supports the wide variety of data types and metadata terms used by the individual DCCs and is designed to enable easy expansion to accommodate new data types.

Introduction

Findability of existing data in biomedical research is important for data reuse. Reusing data can increase the speed of scientific discovery, as well as allow researchers to generate hypotheses (Thanos, 2017; Pronk, 2019; van de Sandt *et al.*, 2019). However, these benefits are highly dependent on the Findability, Accessibility, Interoperability, and Reusability (FAIRness) (Wilkinson *et al.*, 2016) of the individual datasets, and data reuse can only happen if the data is Findable by interested researchers. Improving reuse of datasets is increasingly a priority for both scientists and funding agencies (EU High-Level Group on Scientific Data, 2010; Federer *et al.*, 2015; *Biomedical Data Repositories and Knowledgebases*, 2021), and there are many data publishing and scientific data repositories that are designed to enable targeted search across a breadth of data types (Assante *et al.*, 2016). However, many researchers find it difficult to navigate the number of available data repositories or to predict which might hold data that is useful to them (Pampel *et al.*, 2013). Further, researchers browsing a data repository typically rely on topical information to determine if there are relevant data (Wu *et al.*, 2019; Liu *et al.*, 2020; Zhang *et al.*, 2021), and as most large data repositories typically do not impose constraints on terms used to describe datasets, they are not always well suited to browsing (Assante *et al.*, 2016). This may be especially true of biomedical researchers who show a strong interest in reusing data, but commonly lack expertise in using data repositories (Federer *et al.*, 2015).

The Common Fund (CF) was created in 2006 to fund biomedical research efforts that did not fit into the funding remit of any one NIH Institute or Center, with the intention of generating unique and catalytic datasets that would serve as resources to fuel future innovation. Nearly fifteen years later, dozens of CF programs have been funded and have created large, diverse collections of genomic, transcriptomic, proteomic, metagenomic, and imaging assets. These data are deep, derived from hundreds of studies, with samples collected from thousands of human subjects, cell lines and organoids, and animal models. Most datasets are hosted at program-specific data portals and curated by experts in their field to enable easy use by biomedical researchers. For example, the Genotype-Tissue Expression (GTEx) project (*GTEx Portal*, no date) offers sophisticated search and contextual display of processed gene expression, tissue characteristics, and quantitative trait loci (QTLs) at their web portal. The GTEx data via the data portal which sees over 15,000 users a month (Brown, Charbonneau and White, 2019a) has enabled hundreds of published studies where the GTEx data was reused by the broad research community.

Although the Common Fund was created to catalyze cross-cutting research and create reusable datasets for biomedical research (*About the NIH Common Fund*, 2013), the data is not accessible from a single repository. Rather, each of the fifty-five extant Programs has its own metadata, storage, and access solution. This is an artifact of the Common Fund funding model. As of 2019 individual CF projects were isolated, with few connections between active projects, and there were few incentives to integrate them (Brown, Charbonneau and White, 2020). To address this, the Common Fund Data Ecosystem (CFDE) was established in 2019 to build bridges between individual Common Fund programs and to encourage and enhance data

reuse. The CFDE continues to work to build a self-sustaining community that both harmonizes existing data and develops community standards that newly funded programs can use to interoperate with existing datasets. Governance of the CFDE includes extensive use of “cross-pollination” networking events, Requests for Comments (RFC) documents for community input, and documentation of use cases. In addition, working groups have been established to guide best practices in areas such as gene-centric knowledge representation, clinical metadata, ontologies, technical implementation strategies, and genetic variants.

Prior to the CFDE, Common Fund-funded programs typically had little or no contact, let alone active collaboration, and with no links between portals it was challenging for a researcher to navigate each of the Common Fund web resources. This independence and isolation of different Common Fund programs has benefits, in that it allows each program to tailor their data gathering, portals, and infrastructure to answer their domain-specific questions, and to respond nimbly to changes in program needs. However, this independence has also impeded data integration around common data types. Even the seemingly simple task of finding what data is available is hindered by differences in nomenclature.

Comparing data across programs is particularly challenging. Each program’s data portal provides a curated experience of analyzed data that usually does not support comparison with data from other CF programs. Moreover, many data can only be meaningfully compared to other data analyzed in the same way, and as each Common Fund program operates independently, data are stored, labeled, analyzed, curated and maintained in incompatible ways. Thus, a researcher interested in combining data across CF programs is faced with not only a huge volume, richness, and complexity of data, but also a wide variety, richness and complexity of data access systems with their own vocabularies, file types, and data structures. Reusing Common Fund data for new cross-cutting analysis requires expertise in working with large files, accessing data in the cloud, harmonization, and data transformation -- all before any scientific analysis can begin. Each stage presents an *individually* large challenge for a typical biomedical researcher or clinician which motivates labs to hire dedicated bioinformaticians (at considerable cost to NIH); confronting all of these challenges together is prohibitive for integrative analysis.

To make Common Fund data more findable, the CFDE has created a flexible system of data federation that enables users to discover datasets from across the CF at a centralized portal without requiring Common Fund programs to move, reformat, or rehost their data, similar to the federation strategy of the Research Data Alliance (Plante *et al.*, 2021) and The Australian Research Data Commons (Barker, Wilkinson and Treloar, 2019). The CFDE uses a sociotechnical federation system that combines proven, explicitly community driven approaches (Cruz *et al.*, 2019; DeBarry *et al.*, 2020; Plante *et al.*, 2021) with a model driven catalog that incorporates metadata submitted by individual CF Program Data Coordination Centers (DCCs) into a uniform metadata model that can then be indexed and searched from a centralized portal. This uniform *Crosscut Metadata Model* (C2M2), supports the wide variety of dataset types, vocabularies, and metadata terms used by the individual CF DCCs. This C2M2 is designed to enable easy expansion to accommodate new data types that may emerge from existing DCCs or as new DCCs join the CFDE, and as new use cases are embraced.

The primary user interface for the CFDE's metadata catalog is a Web-based portal (*CFDE Home*, no date) that supports multi-faceted search of metadata concepts such as anatomical location, species, and assay type, across a wide variety of datasets using controlled vocabularies. This style of search supports common researcher use cases (Wu *et al.*, 2019; Liu *et al.*, 2020) by giving them the ability to filter data based on their desired topical information to more easily discover datasets that would otherwise require idiosyncratic targeted search across multiple databases. Findability is more than just search, it is the user's experience of interacting with the search and understanding how to use it to find relevant results. To ensure that our portal actually meets the needs of users, we work closely with a professional useability testing team to assess how well our interface works for our intended audiences.

In this paper, we describe the motivation for the C2M2, detail the current C2M2, and discuss the portal that serves gathered metadata from across CF programs. We also describe the strategy guiding the development of the C2M2, as well as the processes by which the C2M2 evolves.

Results

Common Fund data cannot be found by uniform internet search terms

A hypothetical biomedical researcher interested in finding Common Fund RNAseq datasets created from human blood samples should, in theory, be able to find relevant data from at least five Common Fund programs: Genotype-Tissue Expression (GTEx), Gabriella Miller Kids First (GMKF), Human Microbiome Project (HMP), Extracellular RNA (ExRNA), and the Library of Integrated Network-Based Cellular Signatures (LINCS). Each of these programs hosts their data on a public website, and typically also have informational websites about their work, and so would be expected to be easily Findable. However, a Google search for 'human blood RNAseq "common fund"' returns 20,500 results, all but 55 of which are omitted by Google as they are "very similar to the 55 already displayed" (*human blood RNAseq 'common fund' - Google Search*, no date). These 55 results contain references to data from only three Common Fund programs: GTEx, Human BioMolecular Atlas Project (HuBMAP), and Illuminating the Druggable Genome (IDG). Of these three results, GTEx is in fact the only Common Fund program with RNAseq data from human blood samples. HuBMAP does not have data from blood samples, but Dr. Phil Blood is the director of HuBMAP, so his name matches the search. IDG also lacks RNAseq data from blood, but does feature a blog post that mentions both RNAseq and blood separately. The other four expected programs, GMKF, HMP, ExRNA and LINCS, do not appear in the results.

To illustrate why so few relevant data appear in these search results, we chose six concepts that are broadly applicable to biomedical data -- Sample Type, General Tissue, Specific Tissue, Anatomical Part, Analysis Pipeline, and Organism -- and used them to manually search the five Common Fund RNAseq datasets we know have RNAseq data from human blood. We then documented how each Common Fund program described these concepts in their respective data portals.

The example search in Table 1 highlights several common types of differences between Programs. For each of the six concepts, each sub-table lists the “key” used by each Program to refer to that idea, which is equivalent to the column name in a spreadsheet. The “value” is an example of the data you might find under that column, and here we display the values that best fit our “human blood RNAseq” search. These results show three general types of disagreement in term use: differing term values, differing keys in specific categories, and differing assumptions. Differing term values (“RNAseq” vs “RNA-seq” vs “RNA-Seq”) hinder search because the term of interest may not be matched by a search engine.

Differing keys rarely impact search, but introduce complexity in combining datasets, as each key must be manually harmonized. Table 1 shows one method of harmonization, but there are other valid interpretations. As each Program uses very different terms to describe these concepts, and those terms are often not intuitive, a researcher needs deep familiarity with each dataset to make decisions. For example, is HMP’s “Supersite” most analogous to GTEx’s “SMTS”? Or to GTEx’s “SMTSD”? Differing assumptions can be seen in the reporting for Organism. HMP, GTEx and GMKF only host human data, and thus do not specify species in their internal metadata, making them more difficult to find. Taken together, these differences make data discovery nearly impossible with a uniform set of search terms.

A listening tour identified obstacles to interoperability

We conducted in-depth interviews at nine different programs to better understand the obstacles that the Common Fund Data Coordination Centers (DCCs) face in making CF datasets more accessible to researchers and to learn what data they collect, how they model and store it, and their target user base (Brown, Charbonneau and White, 2019a, 2019b, 2020). We used these visits to identify CF program requirements and establish an initial working relationship with DCCs. A primary outcome of the listening tour in 2019 was a new draft of a common set of data elements that could describe data held at all current DCCs: the C2M2. Over the ensuing two years, we elaborated the C2M2 through a consensus-driven process, instantiated it in a rich relational database, ingested data from nine DCCs, and built a Web-based portal on top of the ingested data.

Entities and associations are key structural features of the C2M2

Entities are represented in C2M2 as **tables**: specifically, rectangular matrices, each **row** of which is a **metadata record** comprising a small collection of named columns (**fields**) containing metadata values. Each field has an agreed-upon meaning that helps to describe whatever entity the table represents. Entities (and the tables describing them) can refer to physical objects, like numbered biosamples; virtual objects, like digital files; or abstract concepts, like a project or a standard name for “salmon louse.” *Relationships between entities* are represented as **association tables**, whereby metadata records describing entities are linked to one another according to broad relationship definitions like “file describes biosample” or “biosample from subject.” See Figure 1 for a simplified **entity-relationship (ER) diagram** describing the model; see Figure supplement 1 for the full C2M2 ER diagram.

C2M2 currently supports three core entities describing fundamental types of experimental resources: **files** (digital bytestreams encoding experimental data); **biosamples** (living material collected and processed via experimental protocols); and **subjects** (organisms studied by experiment, both directly observed and as biosample donors).

C2M2 offers two “container” entities -- **project** and **collection** -- to allow DCCs to explicitly and flexibly group related entity records together into named sets. All valid C2M2 submissions must provide minimal information describing experimental project provenance, wherein every C2M2 core resource record (file, biosample or subject) is linked to exactly one C2M2 project record describing the particular research effort under which the resource was created or observed. Operations essential to discovery (sorting, searching, and binning) depend on this information, so that as C2M2 resource information is discovered by users, it can be more easily associated with its proper research context. The C2M2 collection entity is a conceptual generalization of “dataset” (a named, well-defined collection of data resources) to explicitly also include non-data resources (like biosamples and subjects). The aggregation of C2M2 records into collections is optional, with decisions defining scope and complexity of usage generally left to the submitting DCC. Collections can optionally be assigned persistent IDs (like DOIs) for stable citation, reference and retrieval.

Basic usability for any given C2M2 record will not depend on the presence of values for every possible metadata field -- especially as the C2M2 broadens to accept new data types and variants -- so most C2M2 field values are optional. Nearly every column, and most tables, can optionally be left blank, allowing each DCC to build their C2M2 submission with whatever level of richness best fits their capacity and presentation goals, while meeting conservatively minimal universal criteria designed to permit basic interoperability and discovery.

The C2M2 integrates standardized vocabularies

A key component of cross-collection metadata harmonization within the CFDE is support for the detailed description of C2M2 metadata with terms from standard scientific ontologies. C2M2 currently provides a variety of features by which controlled (standardized, curated) scientific vocabulary terms can be attached to core C2M2 entities. All C2M2 controlled vocabulary annotations are optional. Currently supported controlled vocabularies, listed in Table 2, include the Disease Ontology (*Disease ontology - institute for genome sciences @ university of Maryland*, no date); the Ontology for Biomedical Investigations (*OBI*, no date); the Uber-anatomy Ontology (UBERON; (Ontology Xref Service, no date); the NCBI Taxonomy (*Home - Taxonomy - NCBI*, no date); EDAM (*EDAM: Ontology of bioscientific data analysis and data management*, no date), an ontology for bioinformatics concepts including data types and formats; gene terms from Ensembl (Howe *et al.*, 2021), a database for researchers studying genomics in humans and other vertebrates and model organisms; and PubChem (Kim *et al.*, 2016) the world’s largest curated cheminformatics database.

If sufficiently specific terms cannot be found in the supported ontologies, we encourage DCC data managers to include, provisionally, more general parent terms (as available), and

simultaneously to contact the CFDE Ontology Working Group (WG) with descriptions of any needed additions to the supported controlled vocabularies. CFDE has established direct update channels with the curation authorities for each supported ontology, and the Ontology WG aims to expedite the addition of any missing terms on behalf of Common Fund DCCs. Our search portal directly supports both the official ontologies and the new provisional terms in its pipeline so that DCCs can use the best terms right away rather than resorting to inappropriate term usage.

The C2M2 supports a flexible system of internal and global identifiers

The C2M2 is designed to be a framework for sharing information with the global research community about data deriving from experimental resources. Experimental metadata is created at different times by different DCCs working independently, so any system trying to federate such information must establish a standard way for DCCs to generate stable identifiers (IDs) without requiring DCCs to coordinate ID usage directly with each another. At the same time, any integrated system must guarantee unambiguous IDs, so a minimally effective ID scheme must also allow DCCs to create IDs for their C2M2 metadata that do not (and will never) clash with C2M2 IDs created by other (unknown and possibly inaccessible) DCCs.

The C2M2 provides two types of IDs: a mandatory “C2M2 ID” and an optional “persistent ID.” Every entity representing an experimental resource (file, biosample, project, subject, collection) must be labeled with a C2M2 ID. C2M2 IDs comprise are two parts: a prefix (`id_namespace`) representing the DCC (i.e., point of origin), and a suffix (`local_id`) representing the specific entity (file, project, etc.) being identified. The two parts of each C2M2 ID, concatenated, serve as an ID for each resource that is unambiguous across the entire CFDE ecosystem. This scheme allows DCCs to import their preferred intramural ID scheme directly into the `local_id` component, generally without modification. Persistent IDs encode actionable information that users or automated software can follow to further access the resource named by the ID. The CFDE system will accept a wide variety of particular persistent ID schemes such as minids, Data Repository Service (DRS) IDs and digital object identifiers (DOIs); see the Identifiers Supplement for details.

Independent “data packages” are submitted to the CFDE by each program

The C2M2 is designed to integrate submissions from multiple DCCs, operating independently. Each submission comes as a “data package”, a collection of data tables encoded as tab-separated value files (TSVs). Each DCC collects metadata for data resources within their purview in data packages that it then submits to CFDE.

A C2M2 data package consists of 26 TSV files (as of 9/29/21) populated with interrelated metadata about DCC data assets. Precise formatting requirements for data packages are specified by a JSON Schema document. This schema is an instance of the Data Package meta-specification published by the Frictionless Data group (*Frictionless Data*, no date), a

platform-agnostic toolkit for defining format and content requirements for files on which automatic validation can then be performed. Using this toolkit, the C2M2 JSON Schema specification defines foreign-key relationships between metadata fields (TSV columns), rules governing missing data, required content types and formats for particular fields, and other constraints. These architectural rules provide guarantees for the internal structural integrity of each C2M2 submission, while also serving as a baseline standard to create compatibility across multiple submissions received from different DCCs.

A data package can be created at several levels of complexity, as many of the columns and several entire tables can be left empty and still produce a valid package for submission. Only three metadata records (three rows, across three C2M2 tables) are strictly required, so most tables can optionally be left empty in a minimally compliant submission. The three required records are:

1. a short contact table (name, email address, and other contact details) referencing the DCC technical contact responsible for the submission;
2. a single project record representing the submitting DCC (for resource attribution); and
3. at least one identifier (ID) namespace, registered in advance with the CFDE, that protects IDs from conflicts with IDs generated by other DCCs.

The simplest *usable* submission configuration will also contain at least one non-empty data table representing a flat inventory of experimental resources (e.g., data files or biosamples). A more complex submission might inventory a few different resources like biosamples, files, and subjects, and then also encode basic associative relationships among those resources. For example, a submission might denote which biosamples were materially descended from which subjects, or list which files have been derived from which biosamples. Beyond the single mandatory resource attribution record, DCCs can also attach a hierarchy of project records to their experimental metadata, in order to group resources by funding source. In future work, we anticipate allowing submitters to model events and timing (both for provenance and to describe time-indexed data), among other anticipated extensions such as genes and chemical substances.

Preparing a project submission for ingest into the catalog

Each TSV file in a C2M2 submission is a plain-text file representing a tabular data matrix, with rows delimited by newlines and fields (columns) delimited by tab characters. Field values in TSV files must conform to all formatting and relational constraints specified in the C2M2 schema document (Brady, 2021). Any blank table will be represented by a TSV file containing just one tab-separated header line which lists the (empty) table's field names. Requiring that files exist even for empty tables differentiates intentional data omission from accidental file omission.

For each controlled vocabulary supported by C2M2, a term table must be included as part of any valid submission (see Figure 1, Supplemental Figure 1, green tables). Each such table will contain one row for each (unique) controlled vocabulary term used anywhere in the containing

C2M2 submission, along with basic descriptive information for each term that empowers both downstream user searches and automated display interfaces. All term metadata are loaded directly from the ontology reference data: once the C2M2 entity and association tables are prepared, a CFDE-provided script (Brady *et al.*, 2021) is used to automatically scan the prepared tables for controlled vocabulary terms. These controlled vocabulary terms are validated against externally-provided ontology reference files, and then combined with descriptive data drawn directly from the reference files. The resulting information is then used to automatically build all necessary term tables. These automatically-generated term tables (organized as TSVs) are then bundled along with the rest of the C2M2 submission.

The CFDE provides a robust data ingest and validation process for data packages

The CFDE Coordination Center provides a full service submission system for data packages. DCCs submit candidate data packages to the CFDE Data Submission System by using the *cfde-submit* tool, a lightweight command-line Python package (*cfde-submit*, no date) that enables authenticated upload to the CFDE Portal via Globus Flows (*Globus Flows*, no date). This tool takes a directory of TSVs (the tables in Figure 1) as input, performs initial validation against the C2M2 model, then builds the directory into a *bdbag* (*bdbag: Big Data Bag Utilities*, no date) data package that it securely uploads to a Globus (*Research data management simplified*, no date) endpoint. Each DCC has a separate, secure Globus endpoint location that is created by the CFDE Coordination Center as part of DCC onboarding, and only authorized DCC users can submit to each DCC's location.

Once a data package is uploaded to the Globus endpoint, Deriva (Bugacov *et al.*, 2017) automatically begins ingesting it, performing further validation using the frictionless validator (*frictionless*, no date). Depending on the size and complexity of the data package, this process can take up to an hour to complete. Users are notified by email when the submission process has completed, and are provided with a link to view the data or a description of any errors encountered. Processed data packages are viewable by the submitting DCC in secure set of pages of the CFDE Portal. After submission, DCCs can log in to view, search, and verify their data package. Sandboxed instances offer both an instance of the exact search pages that will be available to the public as well as a high-level overview of the dataset and summary statistics (Figure 3).

DCCs can have any number of submitted data packages, and use the portal to view each submission in multiple ways, for example to ensure that it is structured as intended and/or to compare different ways of modeling their data. No DCC submissions are viewable or searchable by the public until they are approved by the submitting DCC for inclusion in the public release. Although DCCs can submit any number of submissions for testing and review, each public release includes, at most, a single submission from any DCC. At each public release date, the most recently approved data package for each DCC is merged into the public catalog and becomes searchable in the portal. If a DCC does not submit a new data package between

releases, their current public data package remains in the portal. If a DCC has submitted a new data package, it completely replaces any previous data that was available; submissions cannot currently be incrementally updated.

The combined C2M2 catalog can be queried via the CFDE portal Web site.

Fully processed and approved data packages are merged into the CFDE Portal (<https://app.nih-cfde.org/>) catalog on a quarterly basis. This catalog is a customized instance of [Deriva](#) (Bugacov *et al.*, 2017), an asset management platform that provides both a model-driven web interface and a command line interface (deriva-py), with a relational database store that conforms to any given model.

The CFDE web interface supports three basic types of search:

1. Search for a specific core entity type (files, biosamples, or subjects) faceted by any number of controlled vocabulary terms
2. Free text search of a single core entity's controlled vocabulary terms, descriptions, and synonyms
3. Search for core entities associated with a single controlled vocabulary term

The first two search types support use cases where users are interested in finding instances of a specific type of asset (a set of files, a set of biosamples, or a set of subjects). Here, a researcher chooses the type of core entity they are interested in, and filters from there using a simple faceted search, faceted search with Boolean operators, or free text search to create a collection of files, biosamples, or subjects. Such searches are especially useful for researchers looking for data similar to their own experiments.

Users who are interested in finding data from a given assay, tissue type, disease, or other controlled vocabulary term can use the third search type to filter the CFDE catalog to show all files, biosamples, *and* subjects that are associated with that term. This use case allows researchers to quickly assess whether specific data (e.g., mass spectrometry data or psoriasis data) exists in any Common Fund dataset, without needing to specify where the data comes from.

Researchers can search the CFDE portal without registration, or can register to access a personalized dashboard page where they can view interactive summary plots, save searches, and build personalized collections of 'favorite' items. Registered users can authenticate to the CFDE portal by using a number of identity providers, including eRA Commons.

Content searchable at the CFDE Portal continues to expand

Following a series of internal prototypes, our first portal release with live data, on March 30, 2021, included submissions from seven Common Fund programs. This first release allowed researchers to search across a combined 594,507 Files, 425,341 Biosamples, and 6,689 Subjects (Figure 4). Two new Programs added data for our second release in June for a total of 2,558,248 Files, 1,451,443 Biosamples, and 27,951 Subjects from nine Common Fund programs. We expect at least two more programs to begin submitting data in 2022.

The C2M2 is a living standard, and is constantly being expanded to allow new datatypes, new associations, and better ways of describing the underlying data. Over time, DCCs also get better at using the C2M2 to describe their data, a phenomenon clearly visible in the changing number of Biosamples reported over time in Figure 4. In early versions of the portal, DCCs treated cell line replicates as unique Biosamples; as of October, replicate names are collapsed so that the search interface will return all uses of a given named cell line.

The C2M2 and the CFDE portal are evolving over time to better support search

In the first iteration of the CFDE portal, metadata search was a direct extension of the C2M2 model. However, once DCCs began submitting data packages, it quickly became evident that we needed to provide an extra layer of abstraction between the C2M2 model and users to support intuitive search. The reason was that the terms that DCCs used to describe data often did not correspond to the terms employed by users to search for data.

We use the example of controlled vocabulary anatomy terms to illustrate the difficulty. C2M2 uses UBERON for these terms, and each DCC then maps their local anatomy specifications to UBERON for inclusion in their data package. However, while we expected this mapping process to harmonize anatomy terms across DCCs for easy search, we found in practice that there was no more overlap in term use when the DCCs all used the same vocabulary than when they each used their own idiosyncratic vocabulary. As UBERON has, at present, 21,911 unique terms, many describing subtle shades of difference in anatomical structure; in choosing the UBERON term most like their local terminology, DCCs often chose terms at different levels of specificity. We encountered similar results for all controlled vocabularies in our model.

To overcome this barrier to searchability, we have instituted two new practices, one social and one technological. The social solution was to create a working group specifically for dealing with ontology issues, where DCC members can discuss and agree on best practices for choosing ontology terms that are simultaneously a good fit for their data and meaningful to end users. However, even with best practices, there will always be some disagreement on usage. Therefore, we have also created a layer of abstraction in the search portal that allows users to search on higher-level, more general terms under which more specific terms can be grouped through the use of ontology “slims”. For the UBERON anatomy, these “slimmed” search terms are mostly system level names such as nervous system (UBERON:0001016) and connective tissue (UBERON:0002384).

In the CFDE portal, end users can now choose to search by all anatomy terms, only slimmed terms, or both. We are working to add similar slim search capabilities to the portal for all controlled vocabularies available in C2M2.

The C2M2 and the CFDE portal are evolving over time to improve user experience

A known challenge for the design of user interfaces to complex data repositories is to ensure the user intuitively knows how that system should be used. To ensure our portal functionally improves Findability, we have done two rounds of usability testing for the CFDE portal interface since launch. These hour-long in-depth interviews were conducted by a professional user experience team (*Interface Guru - Expert user experience design*, 2013) to determine both how users currently interact with the portal, as well as how they would like to interact with it. This process revealed a number of assumptions and preferences that we then used to refine the interface to best support user search. In keeping with previous findings, key elements users sought on the user interface and portal included: 1) Highly specific terminology or controlled vocabulary; 2) Types and extent of data available within the portal; 3) The ability to easily find and complete key tasks; 4) Consistent, contextual navigational elements; 5) Easily comprehensible data visualizations; 6) The ability to compare tabular data against data visualizations; and 7) Dates of data submission. All of these features, as well as many others suggested during these interviews, have been implemented in the existing portal, and will be subject to further refinement from our next testing cohort.

Discussion

The CFDE portal is a central search solution for locating CF datasets. While a researcher conducting a search sees only a relatively simple user interface, the portal has a complex underlying architecture that relies on the C2M2. The search capabilities of the portal are a visual manifestation of the underlying C2M2, which itself is only useful when populated by data packages. While the CFDE Coordination Center manages and maintains the CFDE infrastructure, we rely on the DCCs to populate the portal with information. This recognition has driven a number of our design decisions around incremental adoption of the C2M2, inclusion of standards, our approach to evolving the C2M2, and our overall approach to federation.

The C2M2 is designed to support incremental adoption

One specific example of an important evolution that supports improved human engagement on both the data submitter and end user is the addition of slimmed ontology terms; see Results section, above. In brief, after soliciting precise ontology terms from DCCs, we realized that a highly precise terminology hindered discoverability of datasets by end users. We then developed a “slim” ontology that we imposed centrally in order to meet user needs, but did not require DCC engagement. We also created a CFDE-CC working group to support DCC engagement where it was valuable to them to guide future evolution of the slimmed ontologies.

This working group serves as an ongoing interface between the CFDE-CC infrastructure, DCCs, and end users, and supports ongoing engagement around this particular feature.

C2M2 is intentionally designed with the goal of supporting incremental adoption and use by program participants. In particular, the C2M2 design facilitates the graded introduction of metadata from CF programs into the CFDE system, through submission of data packages of gradually increasing with additional metadata and more detailed metadata modeling added to support more sophisticated searches and data harmonization. These layers range from basic flat asset inventories to well-decorated networks of relationships among resources, described in finer detail.

The C2M2 requires that DCCs meet a fairly sparse set of *minimum* structural benchmarks when building a submission. The general idea is that DCC resource collections can initially be represented quickly (to enable rapid downstream use) via metadata that meets minimal richness requirements -- enough to provide a basic level of harmonization with biomedical experimental metadata coming from other C2M2 sources (DCCs). Over time, DCC data managers can upgrade their C2M2 metadata submissions by adding more detailed descriptive information to their resource records; by elaborating on provenance, timing and other relationships between resources; and by working with CFDE to expand C2M2 itself to better fit models and automation requirements already in production elsewhere.

A minimally compliant submission -- containing just the three required records (Results, above) and no more -- would clearly be of little use. Search capability at the CFDE portal is highly correlated with table completeness, and submitters encouraged to improve the completeness of their data package over time. This encouragement is built into the submission system, which displays summary statistics of draft submissions before integration with other program submissions for quarterly release.

The ability to submit C2M2 metadata in managed stages of sophistication serves three important purposes. First, it flattens the learning curve for onboarding of DCC data managers by making it possible to create immediately useful submissions with little effort, while encouraging incremental additions over time. Second, it lets DCC data managers test how downstream functionality (e.g., overlapping metadata terms across CFDE) interacts with their C2M2 metadata before investing more heavily in creating more complex C2M2 submissions. Third, it allows *submitters* to provide feedback to CFDE to modify C2M2 in response to submission needs, albeit over longer timelines.

The C2M2 interfaces seamlessly with existing standards

The world has no shortage of standards, and we have specifically designed our model to leverage mature scientific and technical standards wherever possible. Many metadata models have been developed to find and access datasets in different repositories: e.g (Perez-Riverol *et al.*, 2017; Albertoni *et al.*, 2020). Ultimately a successful metadata model is one that fulfills community needs. In keeping with this philosophy, our initial version of the C2M2 was an

evolved version of the DATS model (Sansone *et al.*, 2017), where each data contributor used somewhat different encodings to describe their data, in an ultra flexible system (*README.md at master · dcppc/crosscut-metadata*, 2018). However, during the in depth interviews conducted during our listening tour, we were able to determine the specific needs of each DCC for modeling their data, as well as learning what metadata was most important to their users. As a result, we completely reimplemented the C2M2 to rely on controlled vocabularies (Table 2) for harmonization, and to require a relatively strict set of metadata tables. This resulted in a somewhat less flexible model, but one that is still more than nimble enough to meet the needs of the CFDE community, while also supporting the types of faceted search that biomedical researchers prefer.

The C2M2 mission is to operationalize and anchor our guiding principles of data stewardship -- findability, accessibility, interoperability, and reusability (FAIR) -- all of which are enhanced by the integration of established standards directly into the model framework. Findability is streamlined by the use of common terms to describe scientific concepts: aggregating data according to harmonized and universal descriptive metadata helps users find information and enhances discoverability of relevant related data. Accessibility benefits from integration of technical standards allowing users uniform access to heterogeneous data sources without needing to use multiple bespoke access interfaces. Interoperability is defined by how well data flows interact with other information systems: adoption of technical interface standards directly determines how interoperable any system will be. Reusability depends both on persistence of data over time (encouraged directly in C2M2 by rules governing persistent identifiers) and on the implementation of conceptual standards defining meaning and context (so future users can properly explore the data for their own purposes).

By design, the C2M2 will be amended and extended over time to include additional metadata and relationships, including new community standards, so it can flexibly grow to support any biomedical metadata type.

Future work to improve global discoverability of CFDE Portal resources may for example include integrations with harmonizing efforts like schema.org (Patel-Schneider, 2014) or bioschemas (Goble *et al.*, 2017). Technologies like these aim both to increase data accessibility via common query interfaces and to improve search engine visibility for indexed resources and datasets.

The C2M2 supports two distinct identifiers for each entity, a (required) C2M2-specific ID and also an (optional) persistent ID which is globally resolvable. Required C2M2 IDs can be automatically generated from local DCC identifiers, avoiding the need to mint new IDs before submitting to the CFDE portal and thereby reducing cost and complexity for DCCs preparing C2M2 metadata submissions. Optional persistent IDs can be used to give users direct access (via extramural protocols and APIs) to further metadata describing experimental resources of interest, including direct or programmatic download access to data files indexed by CFDE. Persistent IDs constitute a critical element of the C2M2 framework: they facilitate structured, stable and reliable access to research information housed outside the CFDE portal.

Use of the C2M2 is supported by extensive documentation

Participating DCCs create their submissions by mapping their internal data model to the C2M2; depending on the complexity of their data, this can be a difficult task. To support DCCs in creating their data packages, we provide full technical documentation (*C2M2 documentation - common fund data ecosystem documentation*, no date), a more novice friendly wiki (*published-documentation Wiki*, no date), a bug and request tracker (*published-documentation*, no date), and a full service helpdesk. Our submission system and portal review system also provide detailed error messages for invalid data packages.

The C2M2 is constantly evolving and expanding.

The purpose of the C2M2 is to facilitate metadata harmonization: wherever possible, it should represent legitimately comparable entities in standard ways without compromising meaning, context, or accuracy. Where it may be useful to weaken precision to preserve search recall, slims ensure that the underlying metadata remain accessible. Importantly, the evolution of the C2M2 is an ongoing process.

Most DCCs already use some form of internal metadata model for their own curation operations. C2M2 representation of similar but distinct packages of important information, taken from multiple independently-developed custom DCC metadata systems (e.g., metadata describing people and organizations, data provenance, experimental protocols, or detailed event sequences) requires ongoing, iterative, case-based design and consensus-driven decision-making, coordinated across multiple research groups. Design and decision-making in such contexts requires long-term planning, testing, and execution. CFDE is committed to handling new metadata that is difficult to integrate and harmonize by the creation of generalizable, well-defined extensions to C2M2 if possible, and by pruning (at least in the short term) if not.

We aim with the flexible C2M2 design to split the difference between the ease of evolution inherent in a simple model and the operational power provided to downstream applications by more detailed and difficult-to-maintain extended frameworks.

This flexibility is also intended to address the needs of different DCCs that inevitably operate at widely different scales of data complexity or funding level as well as organization life-cycle phases, research scope, etc. DCCs with advanced, operationalized metadata modeling systems of their own should not encounter arbitrary barriers to C2M2 support for more extensive relational modeling of their metadata if they want it; newer or smaller DCCs, by contrast, may not have enough readily-available information to feasibly describe their experimental resources beyond giving basic asset lists and project attributions. By committing to developing modular C2M2 extensions (Figure 5) for the most advanced DCC metadata, while also offering simpler but well-structured model options for simpler data (that are, furthermore, already harmonized across C2M2 metadata from other DCCs) we aim to minimize barriers to rapid entry into the C2M2 ecosystem and its downstream applications. This approach both allows us to meet the

needs of an ever expanding group of stakeholders, and makes C2M2 an ideal framework for other consortia to adopt for their own data curation needs.

We recognized in formulating the CFDE approach that both the human and technical elements of the cross-DCC **sociotechnical system** (Ropohl, 1999; Baxter and Sommerville, 2011) are vital to achieving greater data integration across diverse CF projects. Our federation of disparate sources using a single model was only possible due to our carefully designed combination of appropriate human engagements (e.g., for understanding requirements and building consensus) that was then supported with mechanisms for data ingest, evaluation, indexing, linking, and access. Also important was a strategy of forming agreement around universal and uncontroversial metadata concepts for streamlined submission construction and usable deployment of DCC metadata. Use of widely recognized controlled vocabularies and ontologies encouraged stable and easily referenced data attributes, and has the added benefit of working in accordance to commonly accepted FAIR practices. (Wilkinson *et al.*, 2016). Our long-term goal with C2M2 is to promote robust support for new CF projects such as (*RFA-RM-21-008: Cellular Senescence Network: Tissue Mapping Centers (U54 Clinical Trial Not Allowed)*, no date, *RFA-RM-21-023: Integration, Dissemination, and Evaluation (BRIDGE) Center for the NIH Bridge to Artificial Intelligence (Bridge2AI) Program (U54 Clinical Trial Not Allowed)*, no date) and other projects that have yet to be announced. We also note that our approach is not prohibitively incompatible with other proposed federation models that use schemas based on relational or RDF systems (Bug *et al.*, 2008; Zhang *et al.*, 2011; Hasnain *et al.*, 2017; Sima *et al.*, 2019). We also anticipate the C2M2 will be fully compatible with other NIH ecosystem projects using data platforms and GA4GH standards (Schatz *et al.*, 2021).

Funding

NIH Common Fund OT3OD025459-01 for the CFDE Coordinating Center

Data Availability

Full C2M2 Documentation: https://docs.nih-cfde.org/en/latest/c2m2/draft-C2M2_specification/

C2M2 Supplemental Wiki: <https://github.com/nih-cfde/published-documentation/wiki>

C2M2 Submission System Documentation: <https://docs.nih-cfde.org/en/latest/cfde-submit/docs/>

Portal code:

CFDE specific Deriva fork: <https://github.com/nih-cfde/cfde-deriva>

Portal UI customization: <https://github.com/nih-cfde/dashboard>

Submission code:

CFDE specific Globus flow fork: <https://github.com/nih-cfde/cfde-submit>

CFDE specific Globus action provider fork: <https://github.com/nih-cfde/deriva-action-provider>

References

- About the NIH Common Fund* (2013). Available at: <https://commonfund.nih.gov/about> (Accessed: 5 November 2021).
- Albertoni, R. *et al.* (2020) 'Data Catalog Vocabulary (DCAT)-Version 2', *World Wide Web Consortium*. Available at: <https://www.w3.org/TR/vocab-dcat-3/>.
- Assante, M. *et al.* (2016) 'Are scientific data repositories coping with research data publishing?', *Data science journal*, 15. doi: 10.5334/dsj-2016-006.
- Barker, M., Wilkinson, R. and Treloar, A. (2019) 'The Australian research data commons', *Data science journal*, 18. doi: 10.5334/dsj-2019-044.
- Baxter, G. and Sommerville, I. (2011) 'Socio-technical systems: From design methods to systems engineering', *Interacting with computers*, 23(1), pp. 4–17. doi: 10.1016/j.intcom.2010.07.003.
- bdbag: Big Data Bag Utilities* (no date). Github. Available at: <https://github.com/fair-research/bdbag> (Accessed: 1 November 2021).
- Biomedical Data Repositories and Knowledgebases* (2021). Available at: <https://datascience.nih.gov/biomedical-data-repositories-and-knowledgebases> (Accessed: 5 November 2021).
- Brady, A. (2021) 'C2M2 master JSON Schema'. Open Science Framework. doi: 10.17605/OSF.IO/C63AW.
- Brady, A. *et al.* (2021) 'C2M2 term-table builder script + ontology support files'. Available at: <https://osf.io/bq6k9/> (Accessed: 1 November 2021).
- Brown, C. T., Charbonneau, A. and White, O. (2019a) '2019-July_CFDE_AssessmentReport.pdf'. doi: 10.6084/m9.figshare.9588374.v1.
- Brown, C. T., Charbonneau, A. and White, O. (2019b) '2019-October_CFDE_AssessmentReport.pdf'. doi: 10.6084/m9.figshare.10261055.v1.
- Brown, C. T., Charbonneau, A. and White, O. (2020) '2019-December_CFDE_AssessmentReport.pdf'. doi: 10.6084/m9.figshare.11592936.v1.
- Bugacov, A. *et al.* (2017) 'Experiences with DERIVA: An Asset Management Platform for Accelerating eScience', in *2017 IEEE 13th International Conference on e-Science (e-Science)*, pp. 79–88. doi: 10.1109/eScience.2017.20.
- Bug, W. *et al.* (2008) 'Data federation in the Biomedical Informatics Research Network: tools for semantic annotation and query of distributed multiscale brain data', *AMIA ... Annual Symposium proceedings / AMIA Symposium. AMIA Symposium*, p. 1220. Available at: <https://www.ncbi.nlm.nih.gov/pubmed/18999211>.
- C2M2 documentation - common fund data ecosystem documentation* (no date). Available at: https://docs.nih-cfde.org/en/latest/c2m2/draft-C2M2_specification/ (Accessed: 1 November 2021).

CFDE Home (no date). Available at: <https://app.nih-cfde.org> (Accessed: 1 November 2021).

cfde-submit (no date). Available at: <https://pypi.org/project/cfde-submit/> (Accessed: 1 November 2021).

Cruz, M. *et al.* (2019) 'Policy needs to go hand in hand with practice: The learning and listening approach to data management', *Data science journal*, 18. doi: 10.5334/dsj-2019-045.

DeBarry, J. D. *et al.* (2020) 'Practical recommendations for supporting a systems biology cyberinfrastructure', *Data science journal*, 19. doi: 10.5334/dsj-2020-024.

Disease ontology - institute for genome sciences @ university of Maryland (no date). Available at: <https://disease-ontology.org/> (Accessed: 1 November 2021).

EDAM: Ontology of bioscientific data analysis and data management (no date). Available at: <http://edamontology.org> (Accessed: 1 November 2021).

EU High-Level Group on Scientific Data (2010) *Riding the Wave: How Europe can gain from the rising tide of scientific data*. Available at: <https://www.fosteropenscience.eu/sites/default/files/original/831.pdf>.

Federer, L. M. *et al.* (2015) 'Biomedical Data Sharing and Reuse: Attitudes and Practices of Clinical and Scientific Research Staff', *PloS one*, 10(6), p. e0129506. doi: 10.1371/journal.pone.0129506.

frictionless (no date). Available at: <https://pypi.org/project/frictionless/> (Accessed: 1 November 2021).

Frictionless Data (no date). Available at: <https://frictionlessdata.io/> (Accessed: 1 November 2021).

Globus Flows (no date). Available at: <https://www.globus.org/platform/services/flows> (Accessed: 1 November 2021).

Goble, C. *et al.* (2017) 'Bioschemas. org', in *Joint 25th Annual International Conference on Intelligent Systems for Molecular Biology (ISMB) and 16th European Conference on Computational Biology (ECCB) 2017*. Available at: <http://bioschemas.org/>.

GTEX Portal (no date). Available at: <https://gtexportal.org/> (Accessed: 1 November 2021).

Hasnain, A. *et al.* (2017) 'BioFed: federated query processing over life sciences linked open data', *Journal of biomedical semantics*, 8(1), p. 13. doi: 10.1186/s13326-017-0118-0.

Home - Taxonomy - NCBI (no date). Available at: <https://www.ncbi.nlm.nih.gov/taxonomy> (Accessed: 1 November 2021).

Howe, K. L. *et al.* (2021) 'Ensembl 2021', *Nucleic acids research*, 49(D1), pp. D884–D891. doi: 10.1093/nar/gkaa942.

human blood RNAseq 'common fund' - Google Search (no date). Available at: <https://www.google.com/search?q=human+blood+RNAseq+%22common+fund%22&ei=yKaCYa fPGLOF9PwPyNO8iA4&start=0&sa=N&ved=2ahUKEwin052rvfzzAhWzAp0JHcgpD-E4PBDy0wN6BAgBEDc&biw=1461&bih=896&dpr=1> (Accessed: 3 November 2021).

Interface Guru - Expert user experience design (2013). Available at: <https://www.interfaceguru.com/> (Accessed: 5 November 2021).

Kim, S. *et al.* (2016) 'PubChem Substance and Compound databases', *Nucleic acids research*, 44(D1), pp. D1202–13. doi: 10.1093/nar/gkv951.

Liu, J. *et al.* (2020) 'How do people make relevance judgment of scientific data?', *Data science journal*, 19(1), p. 9. doi: 10.5334/dsj-2020-009.

OBI (no date). Available at: <http://obi-ontology.org/> (Accessed: 1 November 2021).

Ontology Xref Service (no date) *Uber-anatomy ontology < Ontology Lookup Service < EMBL-EBI*. Available at: <https://www.ebi.ac.uk/ols/ontologies/uberon> (Accessed: 1 November 2021).

Pampel, H. *et al.* (2013) 'Making research data repositories visible: the re3data.org Registry', *PloS one*, 8(11), p. e78080. doi: 10.1371/journal.pone.0078080.

Patel-Schneider, P. F. (2014) 'Analyzing Schema.org', in *The Semantic Web – ISWC 2014*. Springer International Publishing, pp. 261–276. doi: 10.1007/978-3-319-11964-9_17.

Perez-Riverol, Y. *et al.* (2017) 'Discovering and linking public omics data sets using the Omics Discovery Index', *Nature biotechnology*, 35(5), pp. 406–409. doi: 10.1038/nbt.3790.

Plante, R. L. *et al.* (2021) 'Implementing a registry federation for materials science data discovery', *Data science journal*, 20. doi: 10.5334/dsj-2021-015.

Pronk, T. E. (2019) 'The time efficiency gain in sharing and reuse of research data', *Data science journal*, 18. doi: 10.5334/dsj-2019-010.

published-documentation (no date). Github. Available at: <https://github.com/nih-cfde/published-documentation> (Accessed: 1 November 2021).

published-documentation Wiki (no date). Github. Available at: <https://github.com/nih-cfde/published-documentation> (Accessed: 1 November 2021).

README.md at master · dcppc/crosscut-metadata (2018). Github. Available at: <https://github.com/dcppc/crosscut-metadata> (Accessed: 5 November 2021).

Research data management simplified (no date). Available at: <https://www.globus.org/> (Accessed: 1 November 2021).

RFA-RM-21-008: Cellular Senescence Network: Tissue Mapping Centers (U54 Clinical Trial Not Allowed) (no date). Available at: <https://grants.nih.gov/grants/guide/rfa-files/RFA-RM-21-008.html> (Accessed: 1 November 2021).

RFA-RM-21-023: Integration, Dissemination, and Evaluation (BRIDGE) Center for the NIH Bridge to Artificial Intelligence (Bridge2AI) Program (U54 Clinical Trial Not Allowed) (no date). Available at: <https://grants.nih.gov/grants/guide/rfa-files/RFA-RM-21-023.html> (Accessed: 1 November 2021).

Ropohl, G. (1999) 'Philosophy of Socio-Technical Systems', *Society for Philosophy and Technology Quarterly Electronic Journal*, 4(3), pp. 186–194. doi: 10.5840/techne19994311.

van de Sandt, S. *et al.* (2019) 'The definition of reuse', *Data science journal*, 18. doi: 10.5334/dsj-2019-022.

Sansone, S.-A. *et al.* (2017) 'DATS, the data tag suite to enable discoverability of datasets', *Scientific data*, 4, p. 170059. doi: 10.1038/sdata.2017.59.

Schatz, M. C. *et al.* (2021) 'Inverting the model of genomics data sharing with the NHGRI Genomic Data Science Analysis, Visualization, and Informatics Lab-space (AnVIL)', *bioRxiv*. doi: 10.1101/2021.04.22.436044.

Sima, A. C. *et al.* (2019) 'Enabling semantic queries across federated bioinformatics databases', *Database: the journal of biological databases and curation*, 2019. doi: 10.1093/database/baz106.

Thanos, C. (2017) 'Research Data Reusability: Conceptual Foundations, Barriers and Enabling Technologies', *Publications*, 5(1), p. 2. doi: 10.3390/publications5010002.

Wilkinson, M. D. *et al.* (2016) 'The FAIR Guiding Principles for scientific data management and stewardship', *Scientific data*, 3, p. 160018. doi: 10.1038/sdata.2016.18.

Wu, M. *et al.* (2019) 'Data Discovery Paradigms: User Requirements and Recommendations for Data Repositories', *Data Science Journal*, 18(1), p. 3. doi: 10.5334/dsj-2019-003.

Zhang, G. *et al.* (2021) 'Relationship between the metadata and relevance criteria of scientific data', *Data science journal*, 20(1), p. 5. doi: 10.5334/dsj-2021-005.

Zhang, J. *et al.* (2011) 'BioMart: a data federation framework for large collaborative projects', *Database: the journal of biological databases and curation*, 2011, p. bar038. doi: 10.1093/database/bar038.

Figures

Sample Type		
CF Program	Key	Value(s)
ExRNA	RNA Isolation Kit	Serum and Plasma Kit (Quigen)
GTEEx	SMNABTCHT	RNA isolation_PAXgene Blood RNA (Manual)
GMKF	<i>Analyte Type</i>	<i>RNA</i>
HMP	Body Product	blood cell
LINCS	Cell line	THP1

General Tissue		
CF Program	Key	Value(s)
ExRNA	Anatomical Location	Entire cardiovascular system
GTEEx	SMTS	Blood
GMKF	<i>Anatomical Site (Source Text)</i>	<i>Peripheral Whole Blood, Peripheral blood, Blood, Blood Ficolled Wbc</i>
HMP	Supersite	blood cell
LINCS	Tissue	Blood

Specific Tissue		
CF Program	Key	Value(s)
ExRNA	Biofluid	plasma
GTEEx	SMTSD	Whole Blood
GMKF	<i>Composition</i>	<i>Peripheral Whole Blood , Blood, Primary Blood Derived Cancer - Bone Marrow, Blood-Lymphocyte, Primary Blood Derived Cancer - Peripheral Blood, Blood Derived Cancer - Bone Marrow, Post-treatment, Blood or Saliva, Recurrent Blood Derived Cancer - Peripheral Blood, Blood Derived Cancer - Peripheral Blood, Post-treatment</i>
HMP	Body Site	blood cell
LINCS	NA	NA

Anatomical Part		
CF Program	Key	Value(s)
ExRNA	NA	NA
GTEEx	SMUBRID	13756
GMKF	<i>Tissue Type (Source Text)</i>	<i>Primary Blood Derived Cancer - Bone Marrow, Primary Blood Derived Cancer - Peripheral Blood, Blood Derived Cancer - Bone Marrow, Post-treatment, Recurrent Blood Derived Cancer - Peripheral Blood, Blood Derived Cancer - Peripheral Blood, Post-treatment</i>
HMP	Material	blood(ENVO:02000020)
LINCS	NA	NA

Analysis Pipeline		
CF Program	Key	Value(s)
ExRNA	Profiling Assay	small RNA-Seq
GTEEx	SMAFRZE	RNASEQ
GMKF	<i>Experiment</i>	<i>RNA-Seq</i>
HMP	Type	16s_raw_seq_set, 16s_trimmed_seq_set, host_transcriptomics_raw_seq_set, microb_transcriptomics_raw_seq_set
LINCS	Assay	RNA-seq

Organism		
CF Program	Key	Value(s)
ExRNA	Species	Human
GTEEx	NA	NA
GMKF	NA	NA
HMP	NA	NA
LINCS	Organism:	human

Table 1: For each of the six concepts (Sample Type, General Tissue, Specific Tissue, Anatomical Part, Analysis Pipeline, Organism) that would be relevant for finding existing datasets for “human blood RNAseq”, we list the Key and Value used by each of five Common Fund programs that host this type of data in their search portals. Keys are analogous to column headers in a metadata file, and the values shown are the specific values used at that program that are good matches for this search. NAs indicate that information for that concept is not an available search term at a given portal. GMKF Keys and Values shown as italics denote that while those terms are publicly available, they can only be searched while logged into the GMKF portal, and so do not appear in Google searches.

CV field or association	ontology	description
file.assay_type	OBI	the type of experiment that produced a file
file.file_format	EDAM	the digital format or encoding of a file (e.g. "FASTQ")
file.compression_format	EDAM	the compression format of a file (e.g. "bzip2", "gzip")
file.data_type	EDAM	the type of information contained in a `file` (e.g. "sequence data")
biosample.assay_type	OBI	the type of experiment that produced a biosample
biosample.anatomy	UBERON	the physiological source location in or on the `subject` from which a `biosample` was derived
biosample_disease, subject_disease	Disease Ontology	link biosamples and subjects to observations about diseases
biosample_gene	Ensembl	link biosamples to individually relevant genes (e.g. knockdown targets)
biosample_substance, subject_substance	PubChem	link biosamples and subjects to drugs, reagents, other small molecules
ncbi_taxonomy.id	NCBI Taxonomy	a taxonomic name associated with a subject record

Table 2. Controlled vocabularies currently supported in C2M2. Entity term fields are listed as C2M2_entity_table.field_name; term association tables (one-to-many relationships between entities and vocabulary terms) are listed by table name. We give the source ontology for each vocabulary, along with a general description of its annotation role within C2M2.

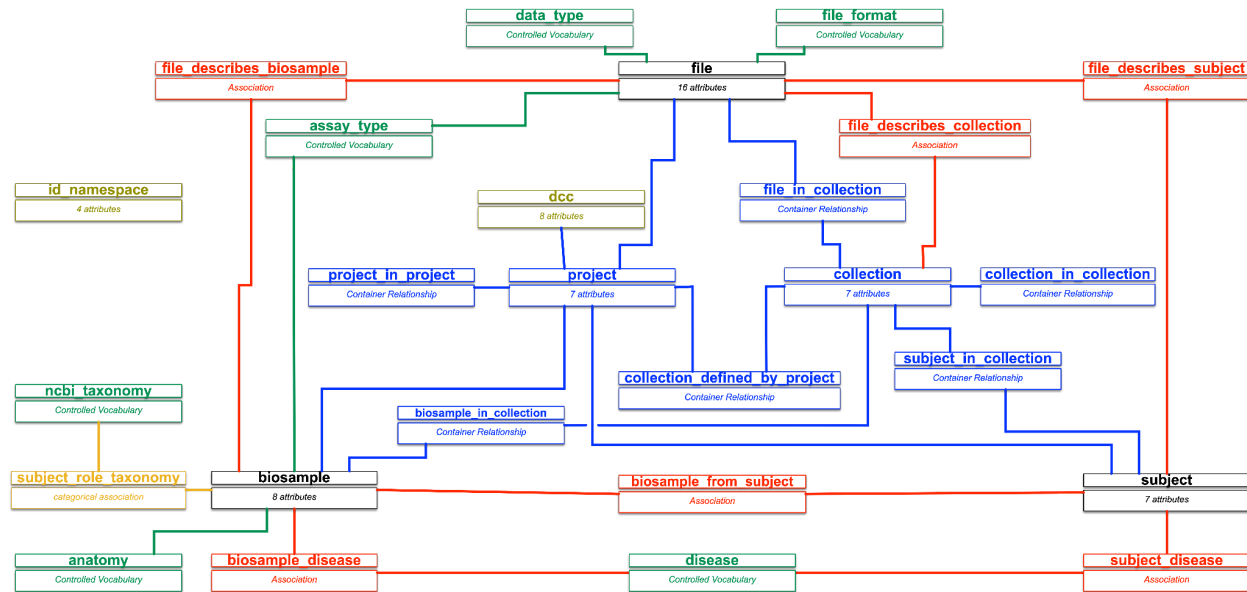


Figure 1: A simplified entity relationship diagram for the C2M2. Entities and associations (relationships between entities) are named inside small boxes: arrows are drawn connecting each association with the entities that participate in the relationship that the association represents.

- Black: **Core entities** (basic experimental resources)
- Dark red: **Association relationships** between entities
- Blue: **Container entities** (project and collection) and their containment relationships
- Green: **Term entities** recording all standardized controlled-vocabulary terms
- Gold: **Administrative entities** giving basic contact information for DCC creators of C2M2 submissions and describing CFDE-registered, DCC-controlled identifier namespaces
- Yellow: **Subject Role Taxonomy**: a special association relationship optionally linking each subject record with NCBI Taxonomy IDs and roles

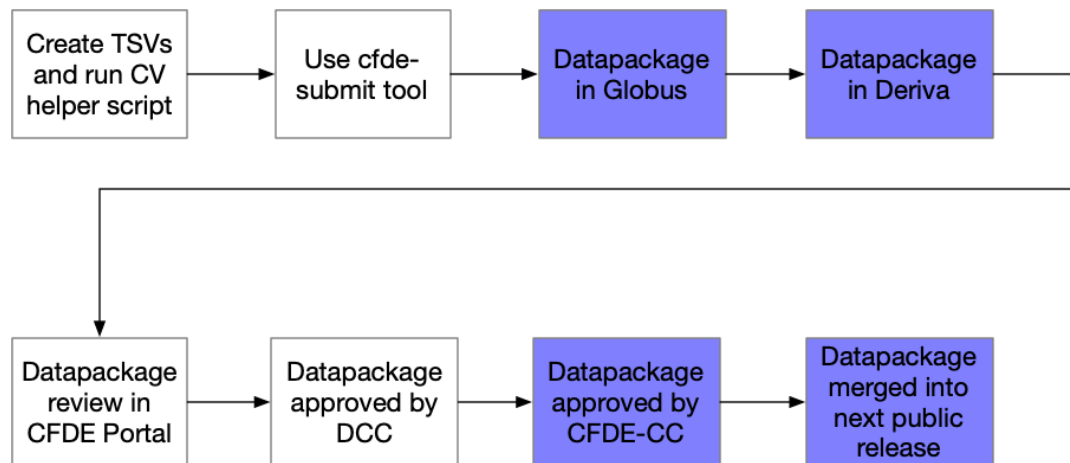


Figure 2: Graphic overview of the steps for data submission. White boxes are user steps, blue boxes are automated.

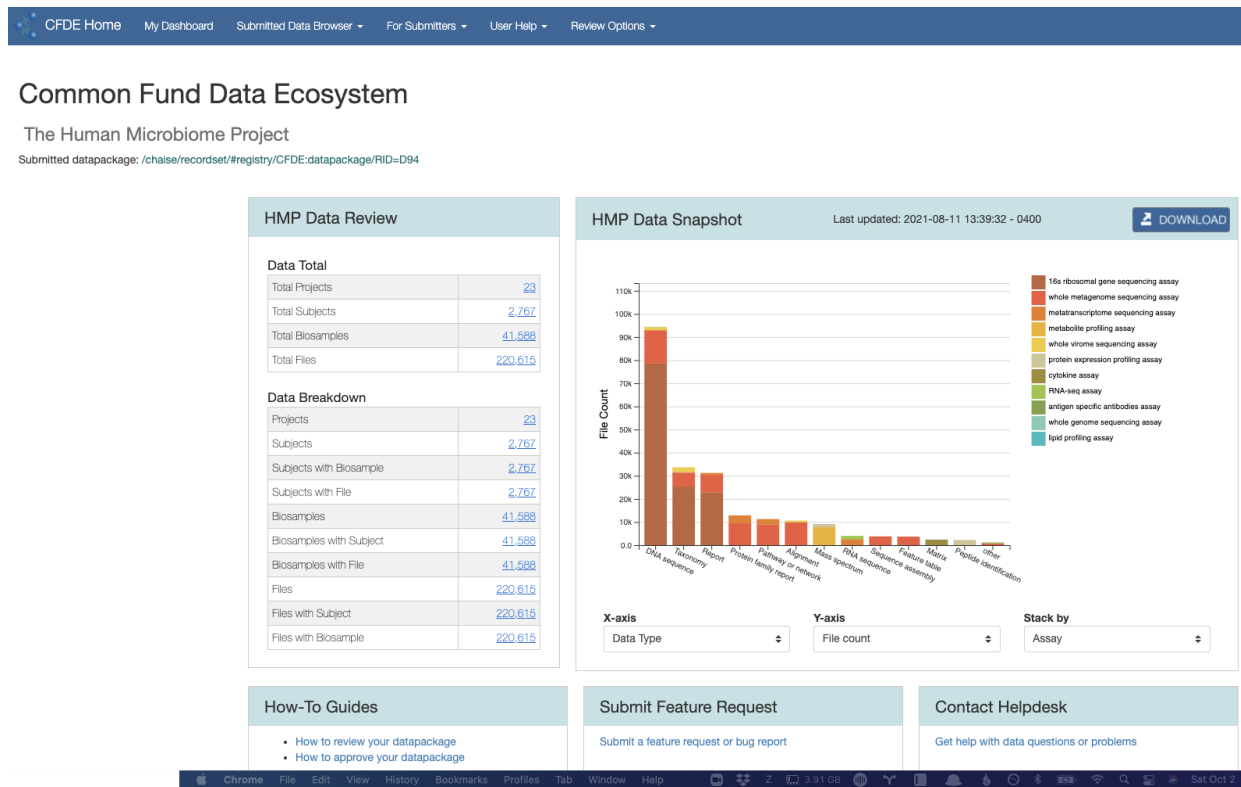


Figure 3: Summary page of a submitted data package with interactive chart, and summary statistics




Figure 4: Core data available for search at the CFDE portal over time. The sharp decrease in biosamples in October 2020 is due to replicate cell line data being more appropriately modeled as from a single biosample.

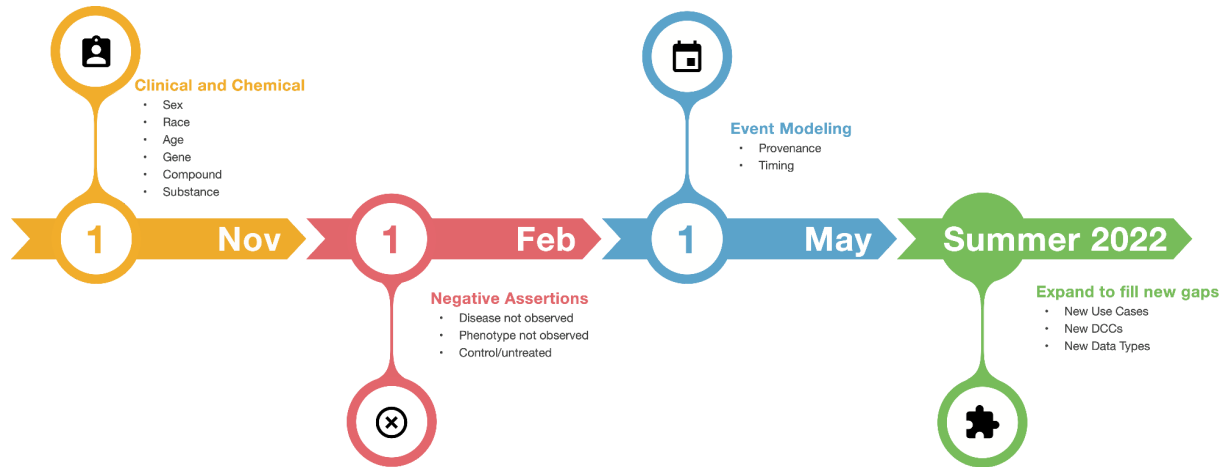


Figure 5: C2M2 Roadmap for 2021-2022. Each list is formally subject to change until its release date, but can be viewed as well established. Model updates include comprehensive client-side data package validation software. In the coming months, we will also introduce a formal versioning scheme for C2M2. The portal only accepts C2M2 data packages from Common Fund programs, and does not currently store any protected (access-controlled) metadata.