

# UC San Diego

## UC San Diego Previously Published Works

### Title

DATS, the data tag suite to enable discoverability of datasets.

### Permalink

<https://escholarship.org/uc/item/0007g5qq>

### Journal

Scientific data, 4(1)

### ISSN

2052-4463

### Authors

Sansone, Susanna-Assunta  
Gonzalez-Beltran, Alejandra  
Rocca-Serra, Philippe  
et al.

### Publication Date

2017-06-01

### DOI

10.1038/sdata.2017.59

Peer reviewed

# SCIENTIFIC DATA

## OPEN DATS, the data tag suite to enable discoverability of datasets

Susanna-Assunta Sansone<sup>1,\*</sup>, Alejandra Gonzalez-Beltran<sup>1,\*</sup>, Philippe Rocca-Serra<sup>1,\*</sup>, George Alter<sup>2</sup>, Jeffrey S. Grethe<sup>3</sup>, Hua Xu<sup>4</sup>, Ian M. Fore<sup>5</sup>, Jared Lyle<sup>2</sup>, Anupama E. Gururaj<sup>4</sup>, Xiaoling Chen<sup>4</sup>, Hyeon-eui Kim<sup>3</sup>, Nansu Zong<sup>3</sup>, Yueling Li<sup>3</sup>, Ruiling Liu<sup>4</sup>, I. Burak Ozyurt<sup>3</sup> & Lucila Ohno-Machado<sup>3</sup>

Received: 27 January 2017

Accepted: 30 March 2017

Published: 6 June 2017

Today's science increasingly requires effective ways to find and access existing datasets that are distributed across a range of repositories. For researchers in the life sciences, discoverability of datasets may soon become as essential as identifying the latest publications via PubMed. Through an international collaborative effort funded by the National Institutes of Health (NIH)'s Big Data to Knowledge (BD2K) initiative, we have designed and implemented the Data Tag Suite (DATS) model to support the DataMed data discovery index. DataMed's goal is to be for data what PubMed has been for the scientific literature. Akin to the Journal Article Tag Suite (JATS) used in PubMed, the DATS model enables submission of metadata on datasets to DataMed. DATS has a core set of elements, which are generic and applicable to any type of dataset, and an extended set that can accommodate more specialized data types. DATS is a platform-independent model also available as an annotated serialization in schema.org, which in turn is widely used by major search engines like Google, Microsoft, Yahoo and Yandex.

<sup>1</sup>University of Oxford, Oxford e-Research Centre, 7 Keble Road, Oxford OX1 3QG, UK. <sup>2</sup>Inter-university Consortium for Political and Social Research, University of Michigan, PO Box 1248, Ann Arbor, Michigan 48106-1248, USA. <sup>3</sup>University of California San Diego, 9500 Gilman Dr, La Jolla, California 92093, USA. <sup>4</sup>The University of Texas Health Science Center at Houston, 7000 Fannin St, Houston, Texas 77030, USA. <sup>5</sup>National Cancer Institute, National Institutes of Health, 9609 Medical Center Dr, Rockville, Maryland 20850, USA. \*These authors contributed equally to this work. Correspondence and requests for materials should be addressed to S.-A.S. (email: susanna-assunta.sansone@oerc.ox.ac.uk).

## Introduction

The biomedical and healthCARE Data Discovery Index Ecosystem (bioCADDIE, <https://biocaddie.org>) is an international joint effort of researchers, informaticians, data scientists, IT professionals, governmental agencies, services providers, publishers and the NIH to facilitate the discovery of available biomedical datasets that are spread across different databases, repositories and on the cloud, through the development of a Data Discovery Index (DDI). The NIH BD2K<sup>1</sup> DDI Consortium is a direct response to the key recommendations made in 2012 by the Data and Informatics Working Group (<https://acd.od.nih.gov/diwig.htm>), set by the Advisory Committee to the Director of the NIH, that requested the development of a catalog of biomedical data. bioCADDIE started effectively when strategies and activities needed to implement the DDI were released as a white paper in 2015 (ref. 2).

The DDI prototype, named DataMed (<https://datamed.org>), was first launched in 2016. Currently, it enables browsing and searching over 60 biomedical data sources (as of March 2017), getting recommendations on datasets that best satisfy the users' specific interests, preferences, and needs. Also the ISA-formatted metadata<sup>3</sup> associated to each Data Descriptor article in *Scientific Data* (<http://scientificdata.isa-explorer.org>) are also indexed and available in the current DataMed release.

DataMed focuses on Findability (F) and Accessibility (A) of datasets. Along with Interoperability (I) and Reusability (R), F&A compose the four elements of the widely-endorsed FAIR principles<sup>4</sup>, which put a specific emphasis on enhancing the ability of individuals and machines to automatically discover and (re)use digital objects, throughout their life cycle. DataMed's architecture encompasses: (i) a repository ingestion and indexing pipeline, which maps the disparate metadata from the diverse repositories, into the unified DATS model; (ii) a terminology server to ensure metadata consistency at a semantic level; and (iii) a web application based on the search engine, which uses the discovery index for locating the appropriate datasets from the set of repositories, and also the terminology server to expand user queries.

DATS metadata elements are used for indexing and searches in DataMed, and are presented as two modules: a core and extended set of elements. The DATS model has been specifically named to echo the JATS metadata elements and format (<https://jats.nlm.nih.gov>) required by PubMed (<https://www.ncbi.nlm.nih.gov/pubmed>) to index publications. Like JATS, the core DATS elements are generic and applicable to any type of dataset. As described in the DataMed paper in *Nature Genetics*<sup>5</sup>, the extended DATS includes additional elements. Some of these elements are specific for life, environmental and biomedical science domains and can be further extended as needed. The similarity and complementarity between DataMed and PubMed are not a simple convenience: it reflects the ecosystem of interconnected resources that are necessary to support and enable biomedical research as a digital research enterprise, known as the NIH Commons<sup>6</sup>. As the Data Science culture grows, all digital research outputs (not just publications and datasets, but also software, workflows, training materials, standards, etc.) must be FAIR. The goal is to empower scientists to fully utilize existing outputs to accelerate their research.

Throughout the project, the bioCADDIE team has undertaken a variety of community-driven activities to design and deliver DataMed; a detailed description of our journey and the resulting implementation are the focus of the complementary DataMed article. Here we focus on presenting the DATS model, its design principles, the methodologies used, the serializations and implementations to date.

## Results

The DATS model describes the metadata elements and the structure for datasets, and powers DataMed's ingestion and indexing pipeline, as well as its search functionality. The model also follows the FAIR best practices for data management, and the recently released 'Candidate Recommendations on Data on the Web Best Practices' (<https://www.w3.org/TR/dwbp>) by the Data on the Web Best Practices Working Group (WG) under the World Wide Web Consortium (W3C). Both community practices support that data on the web should be discoverable and understandable by humans and machines; and that such usage should also be discoverable and the efforts of the data publishers recognized.

The work has been driven and executed by the bioCADDIE team with the input and feedback of international researchers, service providers and knowledge representation experts. Community participation was coordinated via the bioCADDIE's Descriptive Metadata WG3 and the Accessibility Metadata WG7; participants are listed in the Acknowledgement section. The Metadata WG is a joint activity with a wider Metadata WG—encompassing other BD2K centers of excellence—and closely connected to the NIH BD2K Center for Expanded Data Annotation and Retrieval<sup>7</sup> and ELIXIR activities in Europe (<https://www.elixir-europe.org>). Currently, the NIH National Library of Medicine (NLM) is exploring DATS and its possible role in ongoing efforts to make a broader range of biomedical data more readily discoverable.

In August 2015, the first version (v1.0) of the DATS specification was released<sup>8</sup>, with the model available as machine readable JSON (<http://json-schema.org>) schemata and with several examples. This initial model and each of the following versions were tested by the DataMed development team, with a variety of data sources, and reviewed by the bioCADDIE's Descriptive Metadata WG3 members and the larger community. As a result of this iterative process, specifications v1.1 (ref. 9), v2.0 (ref. 10) and v2.1 (ref. 11) were released in March, June and September 2016, respectively. The specification v2.0 represents a further evolution of the model, including the additional metadata elements produced by the

Accessibility Metadata WG7, a set of JSON-schemas and JSON examples, and a Schema.org-annotated context file. The DATS v2.1 also includes edits based on feedback from a first bioCADDIE repository workshop and a Schema.org-annotated JSON-LD (<http://json-ld.org>) serialization. The current version v2.2 (ref. 12) includes support for temporal and spatial extension. All specifications, serializations and examples are freely available from the bioCADDIE Github repository (<https://github.com/biocaddie/WG3-MetadataSpecifications>), where enquiries and feedback can be submitted and tracked. Each release is also packaged and released via Zenodo<sup>8-12</sup> to receive a unique identifier for citation purpose and ensure long term preservation.

The following sections provide a description on the current DATS model v2.2.

### Scope and coverage: Focus on discoverability

First and foremost, it is important to clarify that DATS has not been designed to perfectly represent each and every step of an experiment; this level of detail and metadata are left to the specific repositories that DataMed works with to index. The scope of DataMed is discoverability: DATS provides enough metadata descriptors to help researchers find and access datasets they need and that are stored across a variety of repositories.

Developed iteratively, the model is the result of three complementary approaches, detailed in the Methods section; briefly: (i) a review of existing meta-models, (ii) an analysis of the use cases (top-down), and (iii) a mapping of existing metadata schemas (bottom-up), to find convergences and common metadata elements. Such a combined approach has been necessary to define the appropriate boundaries and level of granularity: which queries will be answered in full by the DataMed prototype, which only partially, and which queries are out of scope.

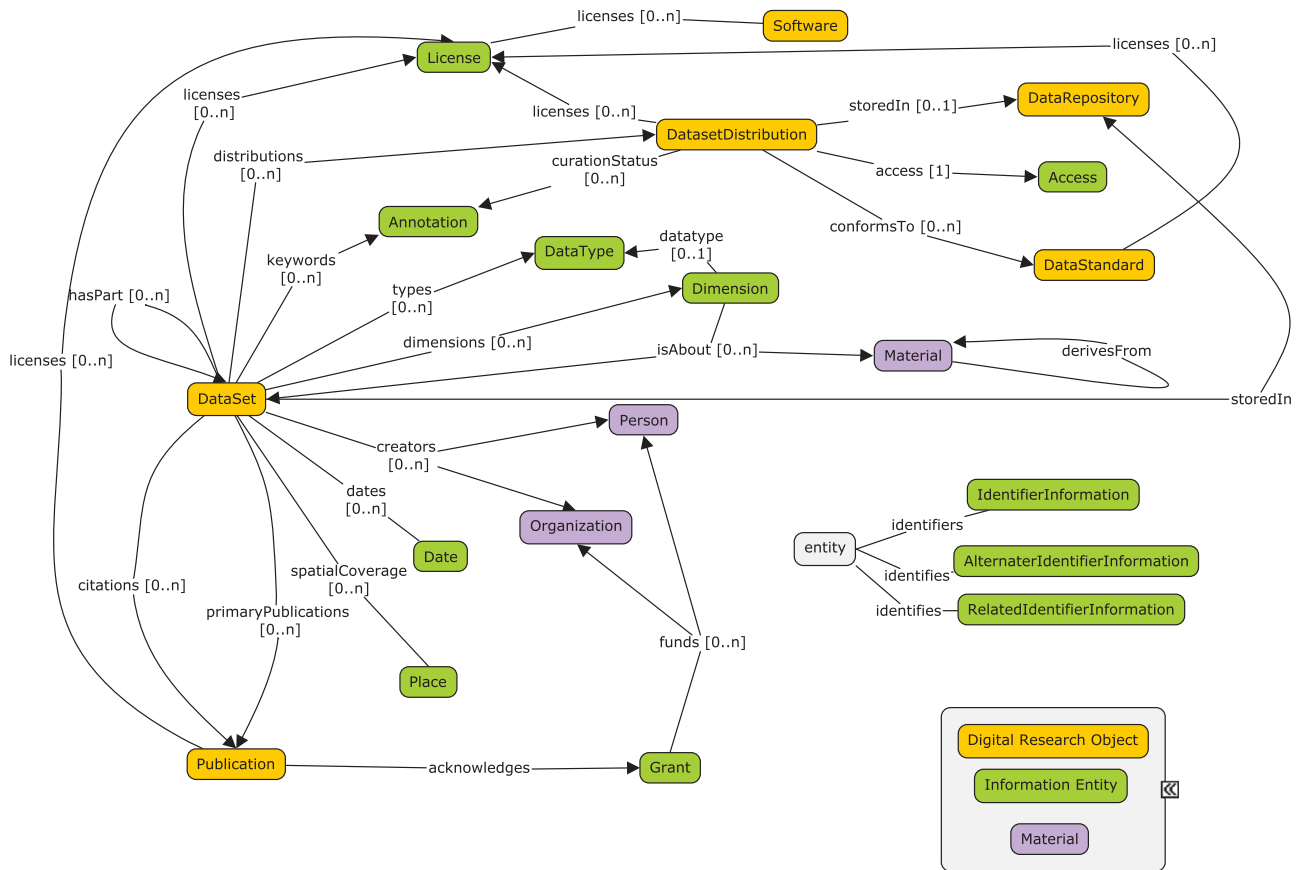
The DATS model is designed around the *Dataset* element, catering for any unit of information stored by repositories. *Dataset* covers both (i) experimental datasets, which do not change after deposition to the repository, and (ii) datasets in reference knowledge bases, describing dynamic concepts, such as ‘genes’, whose definition change over time. To implement the concept that DataMed is built as part of an interlinked ecosystem of resources, the *Dataset* element is linked to other digital objects, which are the focus of other indexes; specifically *Publication* (e.g., PubMed), *Software* (e.g., BD2K Aztec: <https://aztec.bio>), *DataRepository* and *DataStandard* (e.g., BioSharing<sup>13</sup>, <https://biosharing.org>). The latter is especially important in the light of the vast swathes of data that still remain locked in esoteric formats, are described using *ad hoc* or proprietary terminology or lack sufficient contextual information. Knowing if a repository uses open community standards to harmonize the reporting of its different datasets, will provide researchers with some confidence that these datasets are (in principle) more comprehensible and reusable.

Key information about the *Dataset* element is its accessibility, which is represented by the *Access* metadata element that encompasses information on authorization, authentication, and access type. Many types of data in the biomedical domain are restricted to protect confidential information about human subjects. Researchers want to know which data sets are readily available and which ones require prior approval, institutional data use agreements, and security clearances. The *DatasetDistribution* element, linked to the *DataRepository*, *DataStandard* and *Access*, also informs researchers about which data can be accessed directly by machines through an application programming interface (API).

### Modular model: Domain-agnostic core with extended elements

The Google JSON style guide (<https://google.github.io/styleguide/jsoncstyleguide.xml>) has been used to name relevant elements in the DATS model; the requirement levels, indicating what elements are compulsory or desired in the specification, follow the standard RFC 2,119 (<https://tools.ietf.org/html/rfc2119>). The descriptors for each metadata element (or Entity), include: Property (attributes describing the Entity or relating the entity with other entities), Definition (of each Entity and Property), Value(s) (allowed for each Property). In both core and extended DATS, Entities are not mandatory, but applicable only if relevant to the dataset to be described; when an Entity is used, only some of its Properties are defined as mandatory. For example, for the 18 entities of the core there are only a total of 10 mandatory properties. An overview of the core and extended entities, their types and relations are shown in Figs 1 and 2, respectively; Fig. 3 provides a view of the core entities highlighting the few properties defined as mandatory. DATS has been designed to specifically drive discoverability: to find and access datasets via key metadata descriptors. DATS is not meant to perfectly model an experiment, which is the scope of many repository-level models and schemas. The DATS model may appear quite detailed in places as a consequence of (i) the combined approaches used to identify the required metadata elements (see section 2), and (ii) the attempt to aim for the maximum coverage of use cases with minimal number of metadata elements. Nevertheless, it is anticipated that not all use cases can be fulfilled, and that it is difficult to foresee all types of data sources the DataMed prototype should retrieve information from.

Like the JATS, the DATS core elements are generic and applicable to any type of datasets, and cover basic information, for example: what the dataset is about (using the *Material* entity), how it was produced and which information it holds (*DataType*, *Information*, *Method*, *Platform*, *Instrument*), where it can be found (*DatasetDistribution*, *Access*, *Licence*), when it was released and by whom was it created and funded (*Person*, *Organization* and their roles, *Grant*). A set of these DATS core elements also map to the



**Figure 1.** A schematic overview of the DATS core elements, their types and relations.

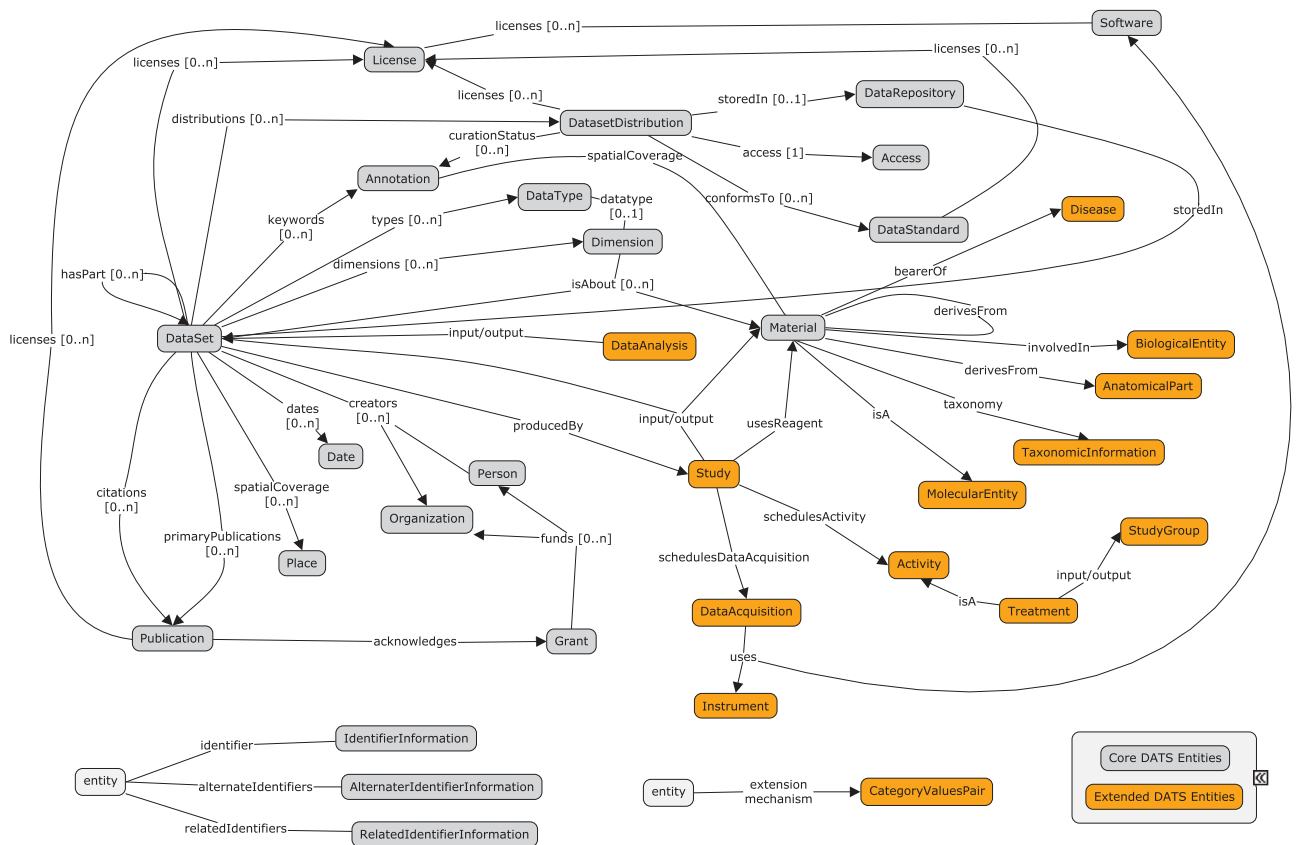
minimal metadata requirements for repositories' landing pages to support data citation; this work was done by the Repositories Early Adopters Expert Group<sup>14</sup>, part of the Data Citation Implementation Pilot (DCIP) project, an initiative of FORCE11 (<https://www.force11.org>) and bioCADDIE. The DATS extended elements have been created to progressively accommodate domain-specific metadata for more specialized data types. The current extended set is tailored to DataMed use cases, and therefore specific for life, environmental and biomedical science domains, but it can be further extended as needed.

**Serializations: Schema.org and bioschemas**

Currently the core and the extended DATS entities are available as machine readable JSON schemata, and as a JSON-LD serialization annotated with Schema.org (<http://schema.org>); serializations in other formats can also be done, if needed. Schema.org is a collaborative, community activity with a mission to create, maintain, and promote schemas for structured data on the Internet, on web pages, in email messages, and beyond. Sponsored by Google, Microsoft, Yahoo and Yandex, Schema.org is already used by over 10 million sites to markup their web pages and email messages. Therefore, anchoring the metadata model to such a potentially powerful vocabulary was one of the early design decisions, outlined in the bioCADDIE white paper.

As Schema.org is a generic vocabulary, its use in the life and biomedical areas requires some extensions. Gaps in coverage have been identified during the annotation process: notifications of these gaps have been submitted to the Schema.org Github tracker. A file mapping DATS to Schema.org elements is available as part of the DATS model v2.2 release; this may be subject to change as current Schema.org elements evolve. In addition, members of the bioCADDIE team are working to coordinate extensions to Schema.org for datasets by: communicating with the WC3 Healthcore and Schema Vocabulary Community Group (<https://www.w3.org/community/schemed>), focused towards clinical studies; and participating and leading activities under the Bioschemas umbrella (<http://bioschemas.org>), which includes major data repositories, BD2K and ELIXIR resources and is set to cover other digital objects beyond data.

Ultimately, the bioCADDIE team will deliver a Schema.org-annotated DATS model or profile, by packaging relevant Schema.org (and/or Bioschemas extended) vocabularies that are relevant to the DATS elements. By delivering this Schema.org DATS profile, web-based DataMed will benefit from: an increased visibility by major search engines and tools, increased accessibility via common query



**Figure 2.** A schematic overview of the DATS core and extended elements, their types and relations.

interfaces, and possibly, an improvement in search ranking. Conversely, data repositories indexed by bioCADDIE will also get the benefit from being more visible to search engines beyond DataMed.

### Adoption

Developed iteratively with the input and feedback of a community of international researchers, service providers and knowledge representation experts, DATS is currently implemented by the bioCADDIE team in DataMed; this is the intended and primary scope for this metadata model. In the first instance, the bioCADDIE team has mapped the target repositories' model to DATS to inform the implementation of data harvesting converters. This has also enabled the team to test the DATS model in real life, refine it, deliver documentation and create examples that have already been used to enable prospective users to export the relevant metadata for indexing in DataMed. To further the uptake, the bioCADDIE team has gone to great lengths and released mappings of DATS to a number of existing generic and widely used schemas, including Schema.org, and domain specific repository schemas; this is available as part of the DATS model v2.1 release. As the work progresses under Bioschemas, a Schema.org DATS profile will also be released, via the Github and Zenodo platforms.

Especially notable are our successful collaborations with other data aggregators and service providers that are working to implement DATS. These include—but are not limited to: the Inter-university Consortium for Political and Social Research (ICPSR), the world's largest archive of digital social science data and one of the NIH-supported repositories ([https://www.nlm.nih.gov/NIHbmic/nih\\_data\\_sharing\\_repositories.html](https://www.nlm.nih.gov/NIHbmic/nih_data_sharing_repositories.html)); the NIH BD2K OmicsDI<sup>15</sup>, a data discovery index for proteomics, genomics and metabolomics datasets; the NIH Federal Interagency Traumatic Brain Injury Research (FITBIR, <https://fitbir.nih.gov>), an informatics system to share data across the entire TBI research field; ImmPort (<http://www.immport.org>), an informatics system supporting the NIH mission to share data, focused on the immunology data; and DataCite (<https://www.datacite.org>), a global non-profit organisation that supports the creation and allocation of Digital Object Identifiers (DOIs) for research data and accompanying metadata. Another example of use is provided by the NIH BD2K CEDAR metadata authoring tool that works to provide a DATS-compliant template to help researchers to describe and expose their datasets, which are not yet in public repositories, to indexing in DataMed.

Figure 4 provides an overview of the development process, including key deliverables and community engagement. To continue these collaborations and foster new ones, the bioCADDIE team has also set up a WG on Standards-driven Curation Best Practices to: (i) assist repositories in creating DATS metadata









13. McQuilton, P. *et al.* BioSharing: curated and crowd-sourced metadata standards, databases and data policies in the life sciences. *Database* **17**, 2016: baw075 (2016).
14. Fenner, M. *et al.* A Data Citation Roadmap for Scholarly Data Repositories. Preprint at <https://doi.org/10.1101/097196> (2016).
15. Perez-Riverol, Y. *et al.* Discovering and linking public omics data sets using the Omics Discovery Index. *Nat. Biotechnol.* **35**, 406–409 (2017).

### Acknowledgements

This project is funded by grant U24AI117966 from NIAID, NIH as part of the BD2K program. The DATS work in BioSchemas by S.-A.S., A.G.-B. and P.R.-S. is supported by the ELIXIR EXCELERATE (H2020-INFRADEV-1-2015-1 676559) and ELIXIR-UK Node (UK BBSRC BB/L005069/1). The co-authors are the key investigators and chairs/co-chairs of the DATS development and implementation activities in DataMed. We thank all contributors to the bioCADDIE Consortium and list here the members of the NIH BD2K and of the bioCADDIE Steering Committee: Alison Yao, Dawei Lin, Dianne Babski, George Komatsoulis, Heidi Sofia, Jennie Larkin, Ron Margolis; members of the bioCADDIE's Descriptive Metadata WG3: Allen Dearry, Carole Goble, Helen Berman, John Westbrook, Kevin Read, Marc Twagirumukiza, Marcelline Harris, Mary Vardigan, Matthew Brush, Melissa Haendel, Michael Braxenthaler, Michael Huerta, Morris Swertz, Rai Winslow, Ram Gouripeddi, Shraddha Thakkar, Weida Tong; and members of the bioCADDIE's Accessibility Metadata WG7: Alex Kanous, Anne-Marie Tassé, Damon Davis, Frank Manion, Jessica Scott, Kendall Roark, Mark Phillips, Reagan Moore.

### Author Contributions

S.-A.S. led the Descriptive Metadata WG3, contributed to the model and its specification and documentation. A.G.-B. and P.R.-S. co-led on the model development, specification and documentation; A.G.-B. developed the DATS validation code and P.R.-S. focused on competency questions. G.A. led the Accessibility Metadata WG7 and chaired the use cases workshop, and with J.L. contributed to the model and its specification and documentation. J.G. and H.X. led the implementation of the model and their feedback, along with those from H.K., R.L., Y.L., B.O., X.C., I.F. and A.E.G., contributed to its refinement and releases. S.-A.S., A.G.-B., P.R.-S., J.G., and I.F. worked with the Schema.org and BioSchemas collaborators. L.O.-M. led the bioCADDIE consortium and ensured that DATS remained central to the other activities of the NIH Commons ecosystem. She wrote the discussion portion of this manuscript and provided critical edits to other portions of the text. All authors gave iterative feedback on the model and the community engagement and as well as the manuscript. S.-A.S., A.G.-B. and P.R.-S. drafted the first version of the manuscript, with contribution from L.O.-M., G.A. and J.G. All co-authors have contributed to its final version.

### Additional Information

**Competing interests:** The authors declare no competing financial interests. S.-A.S. is *Scientific Data's* Honorary Academic Editor and consultant.

**How to cite this article:** Sansone, S.-A. *et al.* DATS, the data tag suite to enable discoverability of datasets. *Sci. Data* **4**:170059 doi: 10.1038/sdata.2017.59 (2017).

**Publisher's note:** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0>

© The Author(s) 2017