


Data-Driven Discovery Science with FAIR Knowledge Graphs



Michel Dumontier, Ph.D.

Distinguished Professor of Data Science

Director, Institute of Data Science



Maastricht University




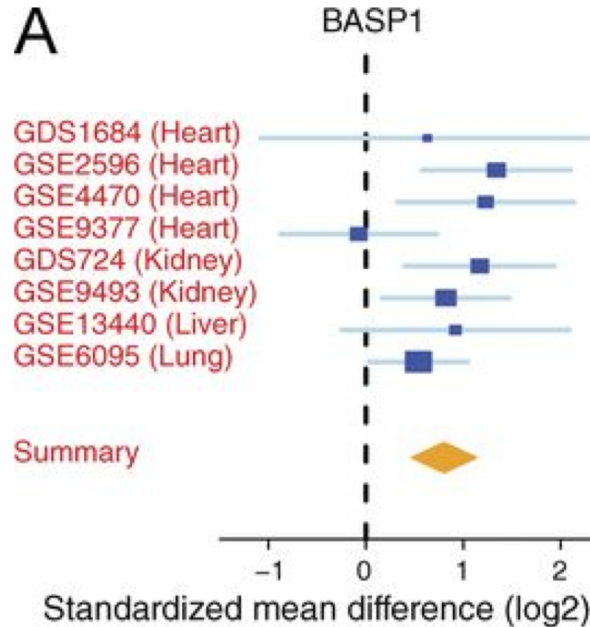

**Vast amounts of (open) data
are enabling and supporting Discovery Science**



A common rejection module (CRM) for acute rejection across multiple organs identifies novel therapeutics for organ transplantation

Khatri et al. JEM. 210 (11): 2205

DOI: 10.1084/jem.20122709



Main Findings:

1. CRM of 11 overexpressed genes **predicted future injury** to a graft
2. Mice treated with **existing drugs** against specific CRM genes **extended graft survival**
3. Retrospective **EHR data analysis** supports **treatment prediction**

Key Observations:


1. **Meta-analysis** offers a **more reliable estimate** of the direction and magnitude of the effect
2. Existing data can be used to **generate and validate new hypotheses**





However, *significant effort* is still needed to find the right dataset(s), make sense of them, and use for a new purpose

Data scientists could be more productive



What data scientists spend the most time doing

- Building training sets: 3%
- Cleaning and organizing data: 60%
- Collecting data sets; 19%
- Mining data for patterns: 9%
- Refining algorithms: 4%
- Other: 5%

Low reproducibility of landmark studies

39% (39/100) in psychology¹

21% (14/67) in pharmacology²

11% (6/53) in cancer³

unsatisfactory in machine learning⁴

¹[doi:10.1038/nature.2015.17433](https://doi.org/10.1038/nature.2015.17433) ²[doi:10.1038/nrd3439-c1](https://doi.org/10.1038/nrd3439-c1) ³[doi:10.1038/483531a](https://doi.org/10.1038/483531a) ⁴<https://openreview.net/pdf?id=By4I2PbQ->

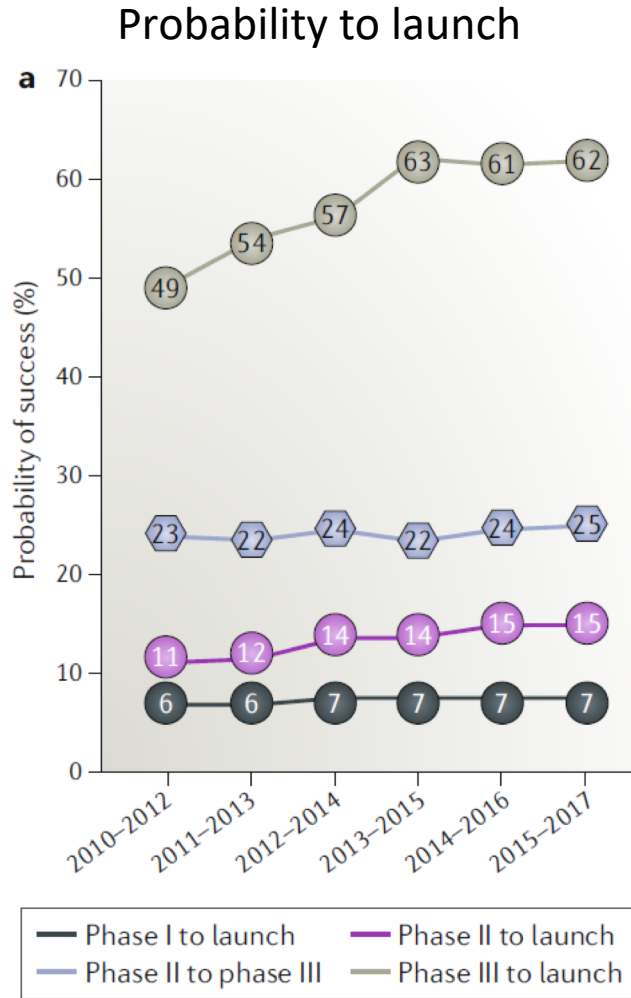
Most published research findings are false.

- John Ioannidis, Stanford University

PLoS Med 2005;2(8): e124.

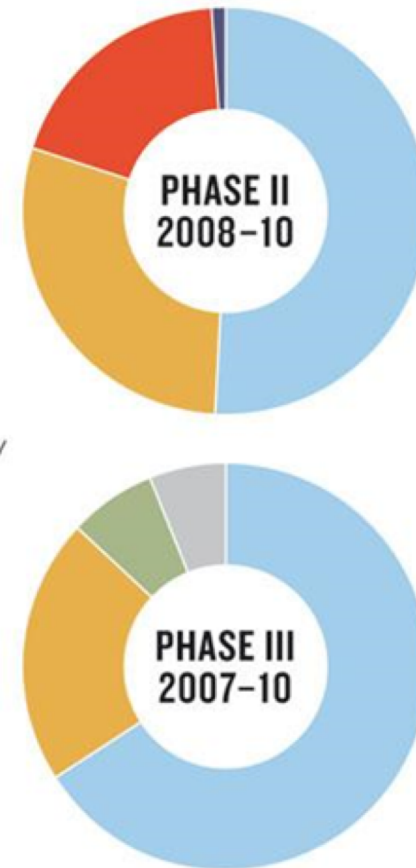
THE CLINICAL-TRIAL CLIFF

Drug companies are removing more compounds from the pipeline at all levels of testing than ever before.



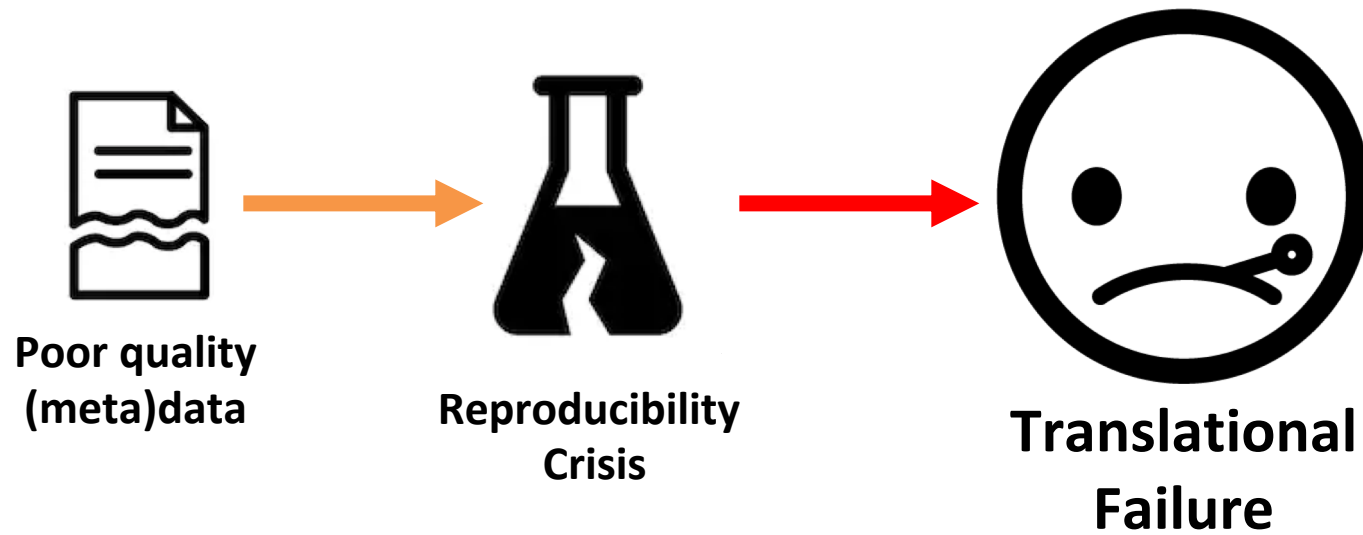
Most of the product failures in phase II and III trials are because researchers are unable to demonstrate efficacy or sufficient safety.

- Efficacy
- Safety
- Strategic
- Pharmacokinetics/ bioavailability
- Commercial/ financial
- Not disclosed




Nature Reviews | Drug Discovery

Nature Reviews Drug Discovery **18**, 495-496 (2019)
<https://doi.org/10.1038/d41573-019-00074-z>




**It's time to completely rethink how we
perform and document empirical research**


Human Machine collaboration is crucial to our future success




Machines

need to be able to discover and reuse data







**An international, bottom-up paradigm for the discovery and reuse of digital content
*for the machines that people use***

The FAIR Guiding Principles for scientific data management and stewardship

[Mark D. Wilkinson](#), [Michel Dumontier](#), ... [Barend Mons](#) 

[+ Show authors](#)

[Scientific Data](#) **3**, Article number: 160018 (2016) | [Cite this article](#)

402k Accesses | **3123** Citations | **1939** Altmetric | [Metrics](#)

This article is in the 99th percentile (ranked 42nd) of the 273,306 tracked articles of a similar age in all journals and the 1st percentile (ranked 1st) of the 1 tracked articles of a similar age in *Scientific Data*

Box 2 | The FAIR Guiding Principles

To be Findable:

- F1. (meta)data are assigned a globally unique and persistent identifier
- F2. data are described with rich metadata (defined by R1 below)
- F3. metadata clearly and explicitly include the identifier of the data it describes
- F4. (meta)data are registered or indexed in a searchable resource

To be Accessible:

- A1. (meta)data are retrievable by their identifier using a standardized communications protocol
 - A1.1 the protocol is open, free, and universally implementable
 - A1.2 the protocol allows for an authentication and authorization procedure, where necessary
- A2. metadata are accessible, even when the data are no longer available

To be Interoperable:

- I1. (meta)data use a formal, accessible, shared, and broadly applicable language for knowledge representation.
- I2. (meta)data use vocabularies that follow FAIR principles
- I3. (meta)data include qualified references to other (meta)data

To be Reusable:

- R1. meta(data) are richly described with a plurality of accurate and relevant attributes
 - R1.1. (meta)data are released with a clear and accessible data usage license
 - R1.2. (meta)data are associated with detailed provenance
 - R1.3. (meta)data meet domain-relevant community standards



European Commission > Press releases database > Press Release details

European Commission - Statement

G20 Leaders' Communique Hangzhou Summit

Hangzhou, 5 September 2016

1. We, the Leaders of the G20, met in Hangzhou, China on 4-5 September 2016.



**Annex 4:
G7 Expert Group on Open Science**

Turin, Italy, September 28, 2017



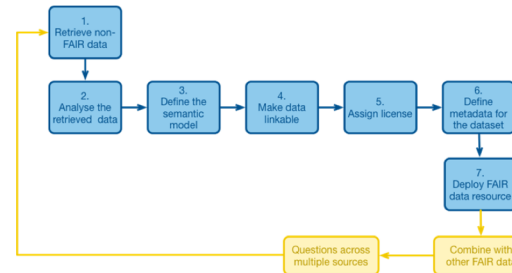
<http://www.nature.com/articles/sdata201618>

FAIR in a nutshell

FAIR aims to create **social, scientific** and **economic impact** by facilitating the discovery and reuse of **digital resources** through a set of requirements:

- **unique identifiers** to reliably retrieve all forms of digital content and knowledge
- **high quality meta(data)** to make more efficient the discovery of relevant digital resources
- **shared vocabularies** to power structured query and statistical analysis
- **community standards** to reduce the effort in wrangling (meta)data
- **detailed provenance** to provide context and support reproducibility
- **repositories** to make content available to others in the long term
- **standardized terms of use** to clarify expectations and intensify innovation

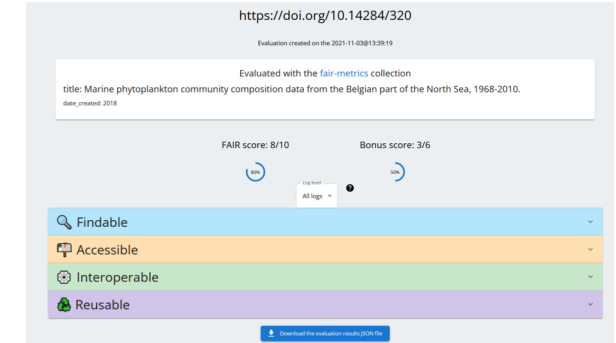
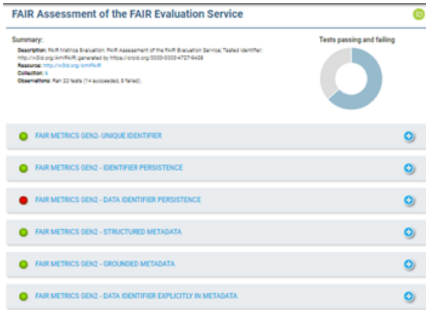
Learn how to create and publish FAIR data



Tools to check the correctness of FAIR Implementations


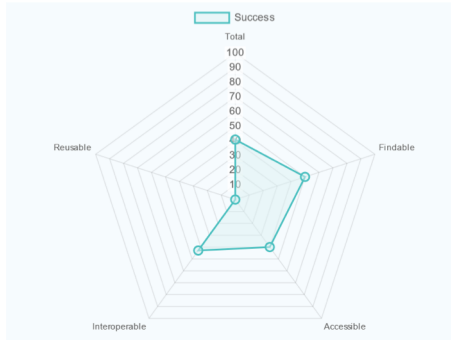
FAIR Evaluator

<https://w3id.org/AmlFAIR>




FAIR Checker

https://fair-checker.france-bioinformatique.fr/base_metrics



F-UJI

<https://www.f-uji.net>

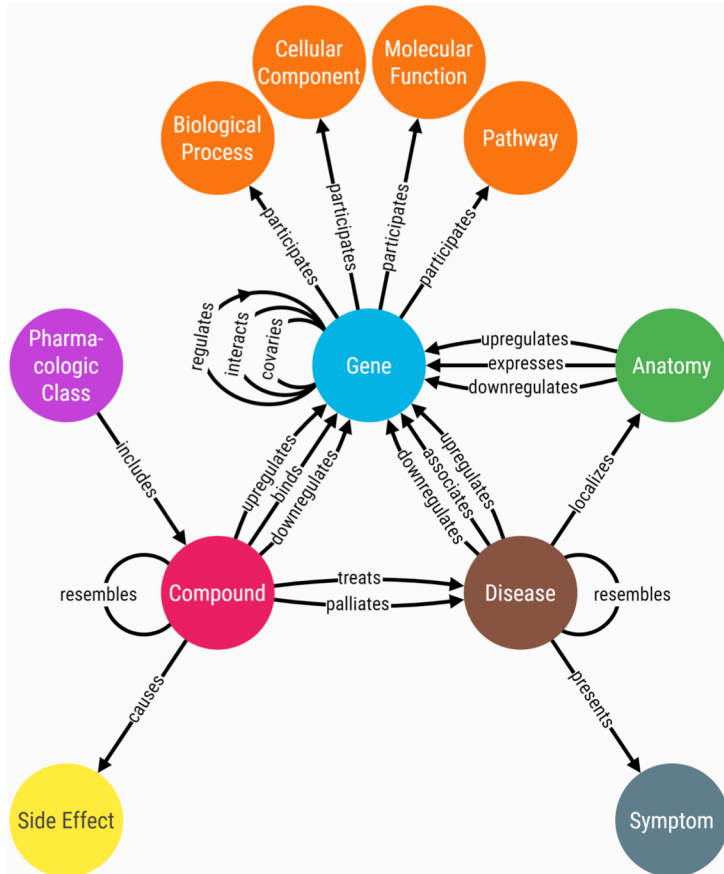


<https://fair-enough.semanticscience.org>

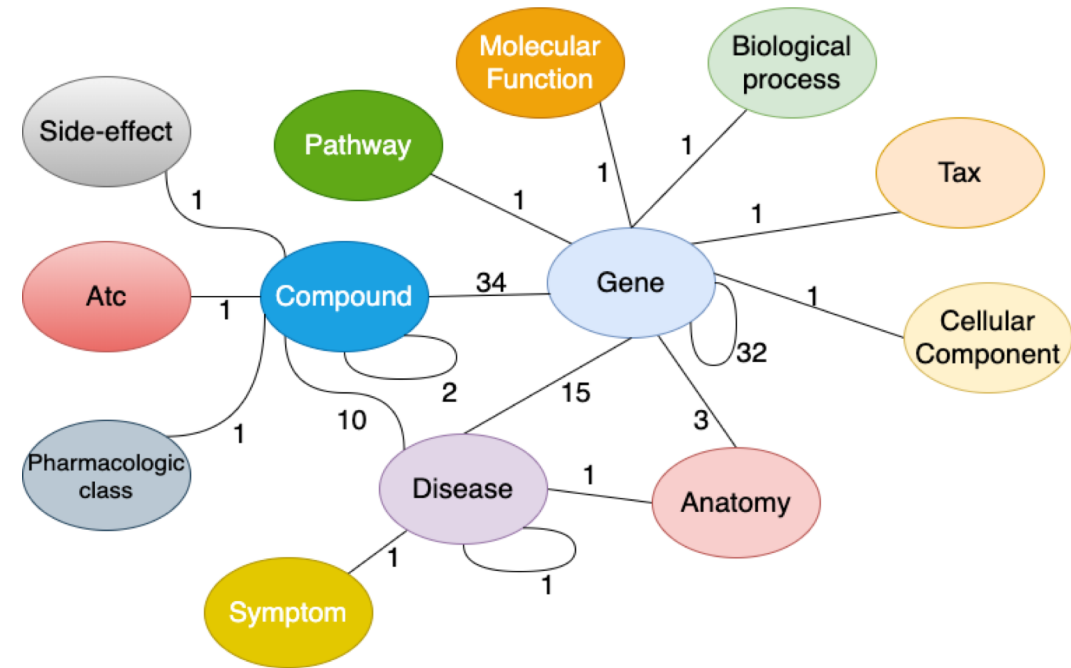
Towards FAIR Knowledge Graphs

A number of biomedical KGs have been developed

See recent review Callahan et al. 2020. Knowledge-Based Biomedical Data Science.

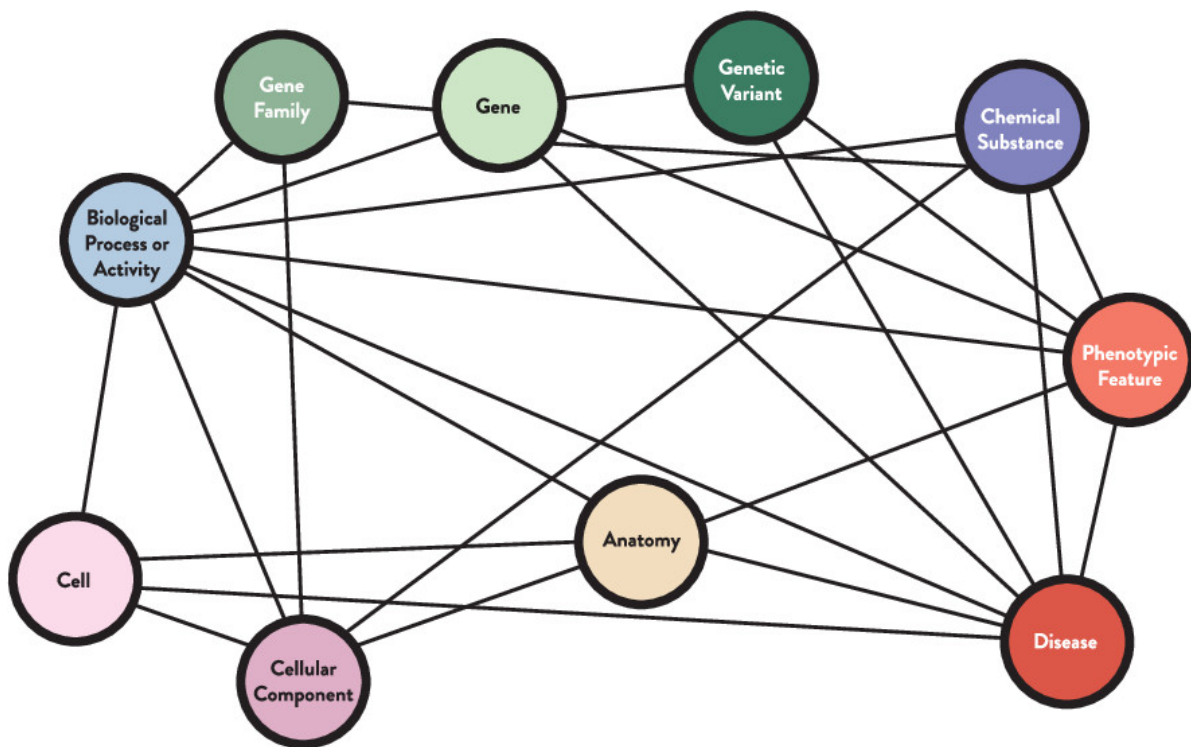


metagraph for het.io



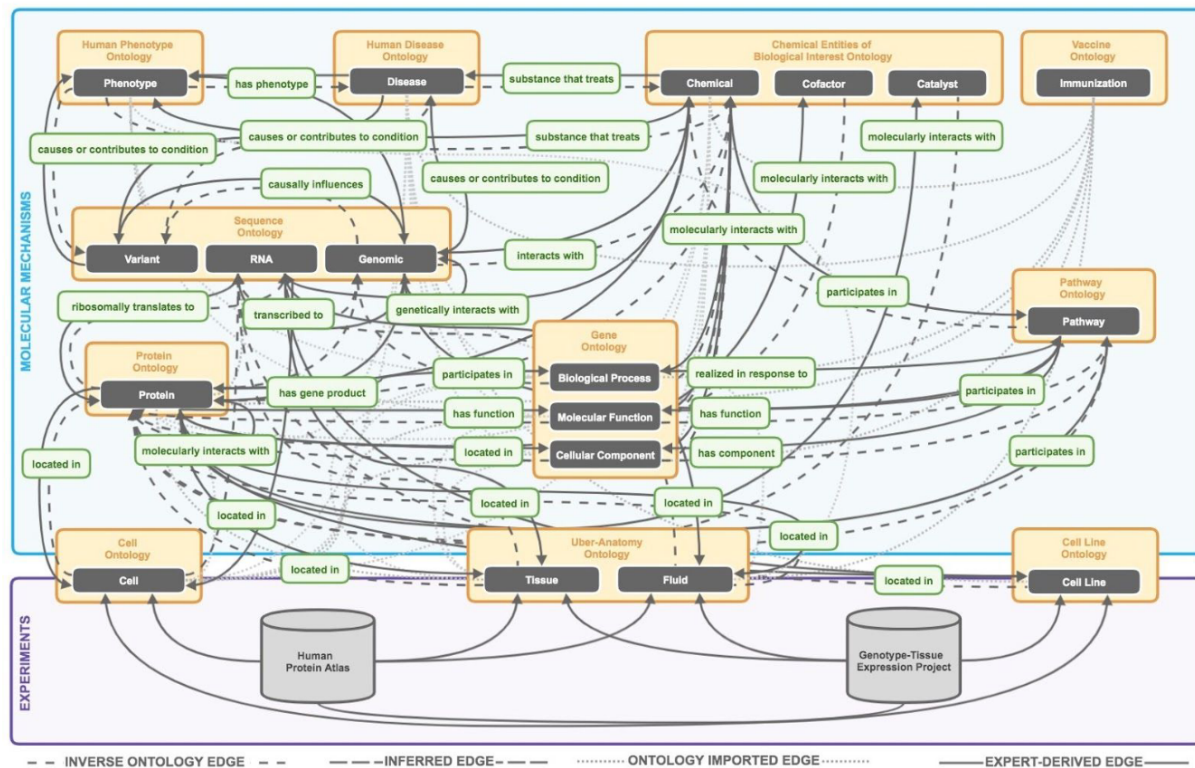
metagraph for [DRKG](https://drkg.org)

Including ontology guided Knowledge Graphs



ROBOKOP

<https://robokopkg.renci.org>



PheKnowLator. Callahan et al.


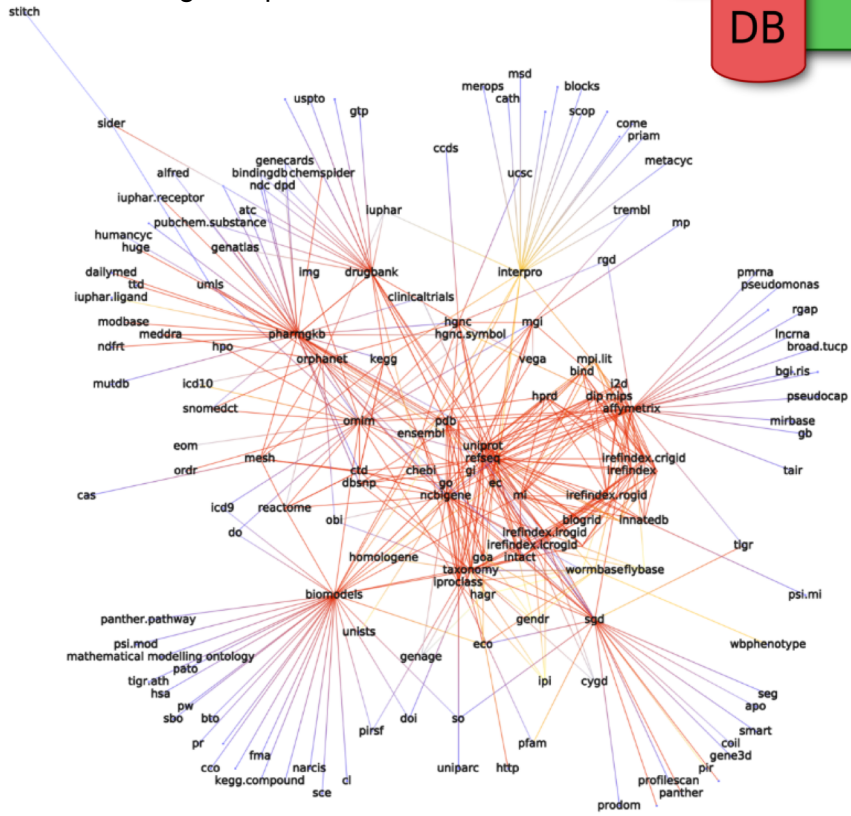
<https://github.com/callahantiff/PheKnowLator>

BIO2RDF

Linked Data for the Life Sciences



Bio2RDF is an open source project that uses semantic web technologies to make it easier to reuse biomedical data

chemicals/drugs/formulations,
genomes/genes/proteins, domains
Interactions, complexes & pathways
animal models and phenotypes
Disease, genetic markers, treatments
Terminologies & publications



github

Conversion
Scripts



- Resolve URIs
- Lookup services
- Query Federation

- SPARQL endpoint
- Search
- Faceted Browser

- Raw Linked Data
- Virtuoso DB files




- **30+** biomedical data sources
- **10B+** interlinked statements
- EBI, SIB, NCBI, DBCLS, NCBO, and many others produce this content

Alison Callahan, Jose Cruz-Toledo, Peter Ansell, Michel Dumontier:
Bio2RDF Release 2: Improved Coverage, Interoperability and
Provenance of Life Science Linked Data. ESWC 2013: 200-212


The Semantic Web is a portal to the **web of knowledge**

standards for publishing, sharing and querying
facts, expert knowledge and services

scalable approach for the discovery
of *independently constructed,*
collaboratively described,
distributed knowledge
(in principle)



The Linked Open Data Cloud



<https://lod-cloud.net/>

The Linked Open Data Cloud Project (LODCloud)



›micheldumontier::HCLS@KGC:2022-05-02

About: <http://bio2rdf.org/drugbank:DB00586> [Sponge](#) [Permalink](#)

An Entity of Type : http://bio2rdf.org/drugbank_vocabulary:Small-molecule, within Data Space : drugbank.bio2rdf.org associated with source [dataset\(s\)](#)

Type: http://bio2rdf.org/drugbank_vocabulary:Small-molecule





Attributes	Values
rdf:type	http://bio2rdf.org/drugbank_vocabulary:Drug http://bio2rdf.org/drugbank_vocabulary:Resource http://bio2rdf.org/drugbank_vocabulary:Small-molecule
rdfs:label	Diclofenac [drugbank:DB00586]
rdfs:seeAlso	http://www.drugbank.ca/drugs/DB00586 http://www.drugs.com/cdi/diclofenac-drops.html http://www.rxlist.com/cgi/generic/diclofen.htm
owl:sameAs	http://identifiers.org/drugbank/DB00586
dcterms:title	Diclofenac
dcterms:description	A non-steroidal anti-inflammatory agent (NSAID) with antipyretic and analgesic actions. It is primarily available as the sodium salt. [PubChem]
dcterms:identifier	drugbank:DB00586
void:inDataset	http://bio2rdf.org/drugbank_resource:bio2rdf.dataset.drugbank.R3
http://bio2rdf.org...bulary:identifier	DB00586
http://bio2rdf.org...abulary:namespace	drugbank
http://bio2rdf.org...df_vocabulary:uri	http://bio2rdf.org/drugbank:DB00586
http://bio2rdf.org...x-identifiers.org	http://identifiers.org/drugbank/DB00586
http://bio2rdf.org...bulary:absorption	http://bio2rdf.org/drugbank_resource:af3a8b347e732d3c3b48a5428a6160e0
http://bio2rdf.org...affected-organism	http://bio2rdf.org/drugbank_vocabulary:Humans-and-other-mammals

Information Retrieval: Phenotypes of knock-out mouse models for the targets of a selected drug

Endpoint: Output:

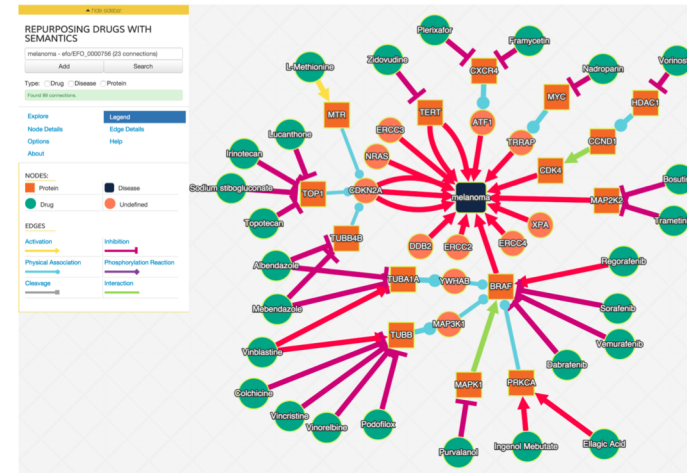
```

1 PREFIX dct: <http://purl.org/dc/terms/>
2 SELECT DISTINCT ?phenotype_label
3 WHERE {
4   SERVICE <http://drugbank.bio2rdf.org/sparql> {
5     ?drug <http://bio2rdf.org/drugbank_vocabulary:target> ?target .
6     FILTER (?drug = <http://bio2rdf.org/drugbank:DB00619>)
7     ?target <http://bio2rdf.org/drugbank_vocabulary:x-hgnc> ?hgnc .
8   }
9   SERVICE <http://hgnc.bio2rdf.org/sparql> {
10    ?hgnc <http://bio2rdf.org/hgnc_vocabulary:x-mgi> ?marker .
11  }
12  SERVICE <http://mgi.bio2rdf.org/sparql> {
13    ?model <http://bio2rdf.org/mgi_vocabulary:marker> ?marker .
14    ?model <http://bio2rdf.org/mgi_vocabulary:allele> ?allele .
15    ?allele <http://bio2rdf.org/mgi_vocabulary:allele-attribute> ?allele_type .
16    ?model <http://bio2rdf.org/mgi_vocabulary:phenotype> ?phenotypes .
17    FILTER (str(?allele_type) = "Null/knockout")
18  }
19  SERVICE <http://bioportal.bio2rdf.org/sparql> {
20    ?phenotypes rdfs:label ?phenotype_label .
21  }
22 }
    
```

phenotype_label
"hemorrhage [mp.0001914]"@en
"intracranial hemorrhage [mp.0001915]"@en
"perinatal lethality [mp.0002081]"@en

Exploration: drug-target-disease networks

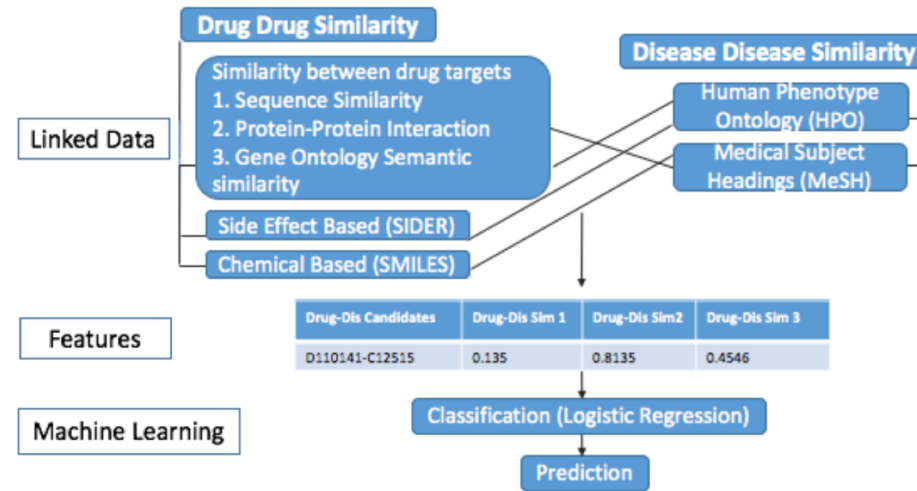


<https://doi.org/10.7717/peerj-cs.106>

Custom Knowledge Portal: EbolaKB

<https://doi.org/10.1093/database/bav049>

Reproducible ML: new uses for existing drugs



<https://doi.org/10.7717/peerj-cs.281>

Predict new drug applications in a documented and reproducible manner

Mol Syst Biol. 2011; 7: 496.

Published online 2011 Jun 7. doi: [10.1038/msb.2011.26](https://doi.org/10.1038/msb.2011.26)

PMCID: PMC3159979

PREDICT: a method for inferring novel drug indications with application to personalized medicine

Assaf Gottlieb,¹ Gideon Y Stein,^{2,3} Eytan Ruppin,^{1,2} and Roded Sharan^{a,1}

AUC 0.90 across all therapeutic indications

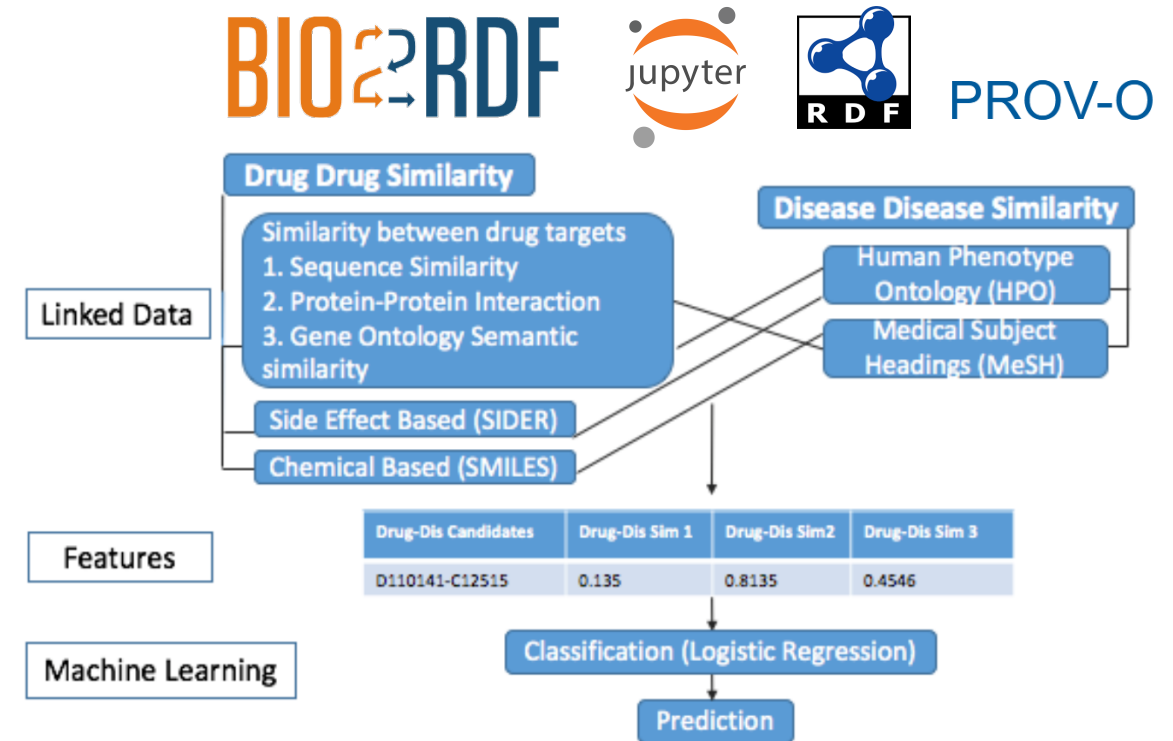
Scripts not available. Feature tables available.
Not reproducible!



**Towards FAIR protocols and workflows:
the OpenPREDICT use case**

Remzi Celebi^{1,*}, Joao Rebelo Moreira^{2,*}, Ahmed A. Hassan³, Sandeep Ayyar⁴,
Lars Ridder⁵, Tobias Kuhn² and Michel Dumontier¹

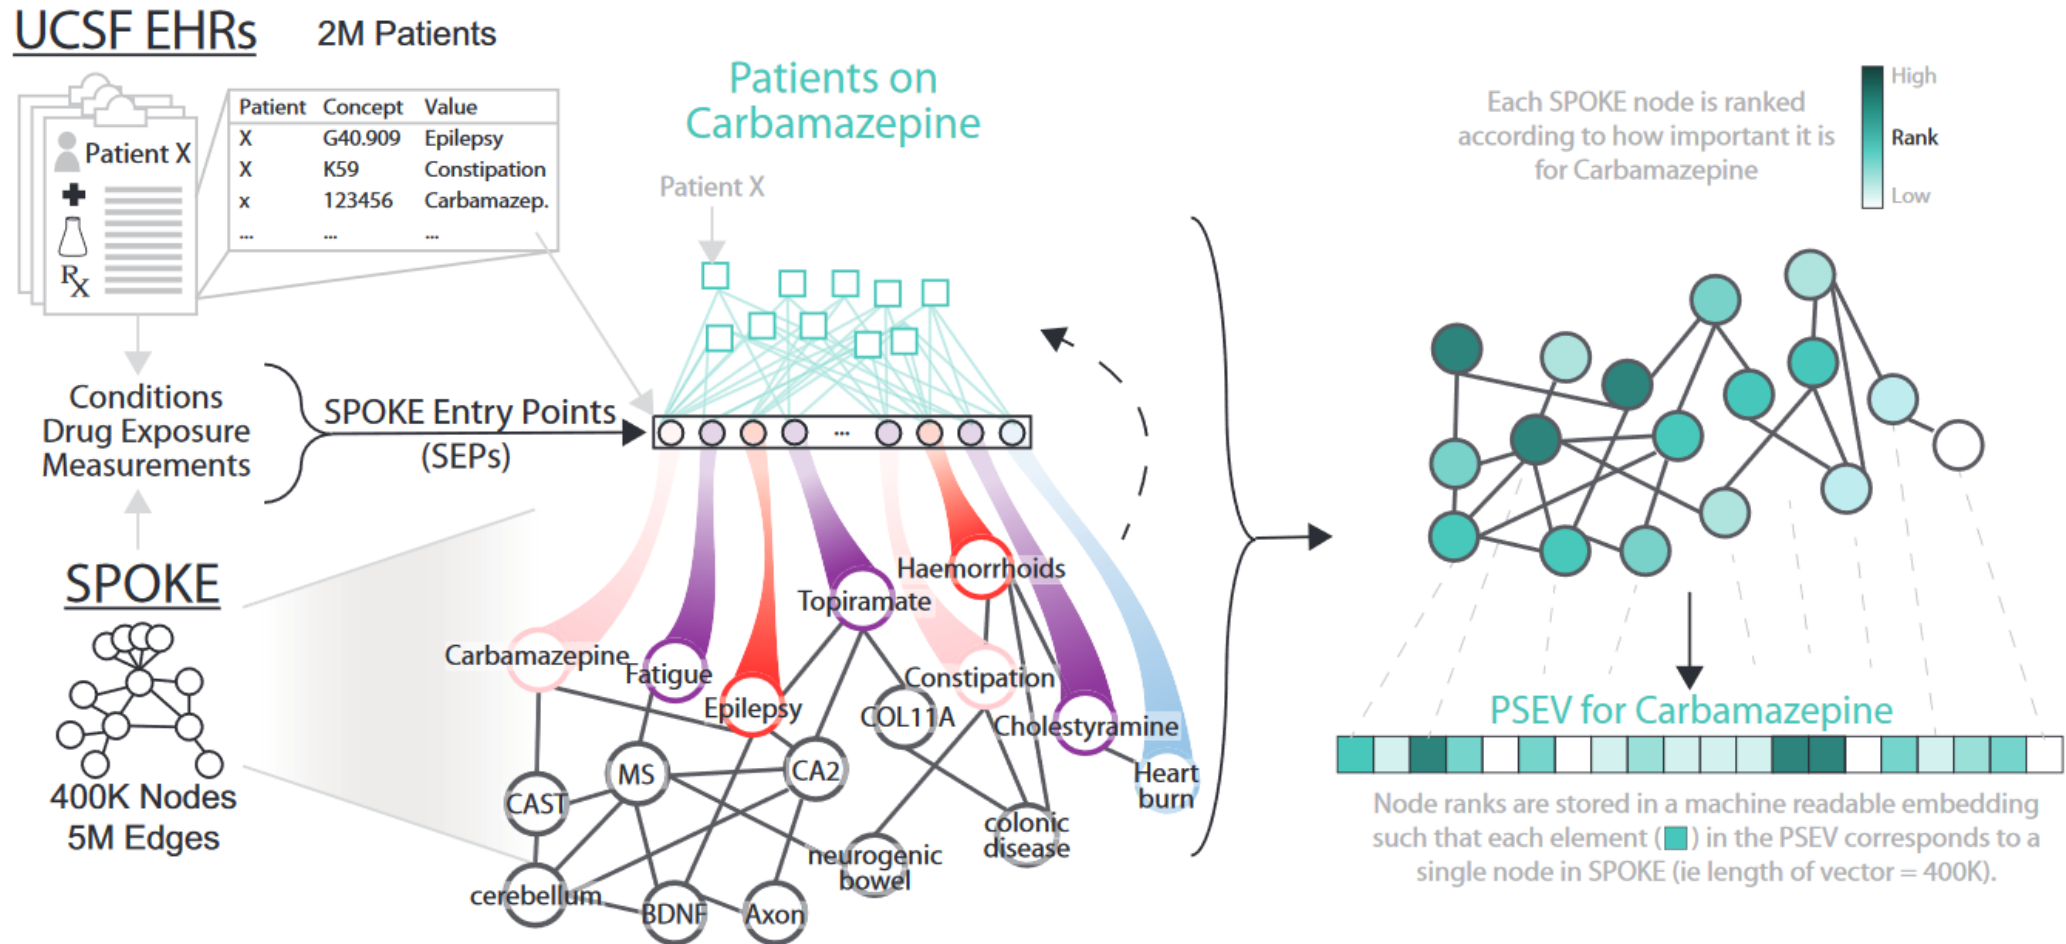
Result: ROCAUC 0.83



Celebi R, Rebelo Moreira J, Hassan AA, Ayyar S, Ridder L, Kuhn T, Dumontier M. 2020. Towards FAIR protocols and workflows: the OpenPREDICT use case. PeerJ Computer Science 6:e281
<https://doi.org/10.7717/peerj-cs.281>

@micheldumontier::HCLS@KGC:2022-05-02

KG-based data fusion for predictive analytics



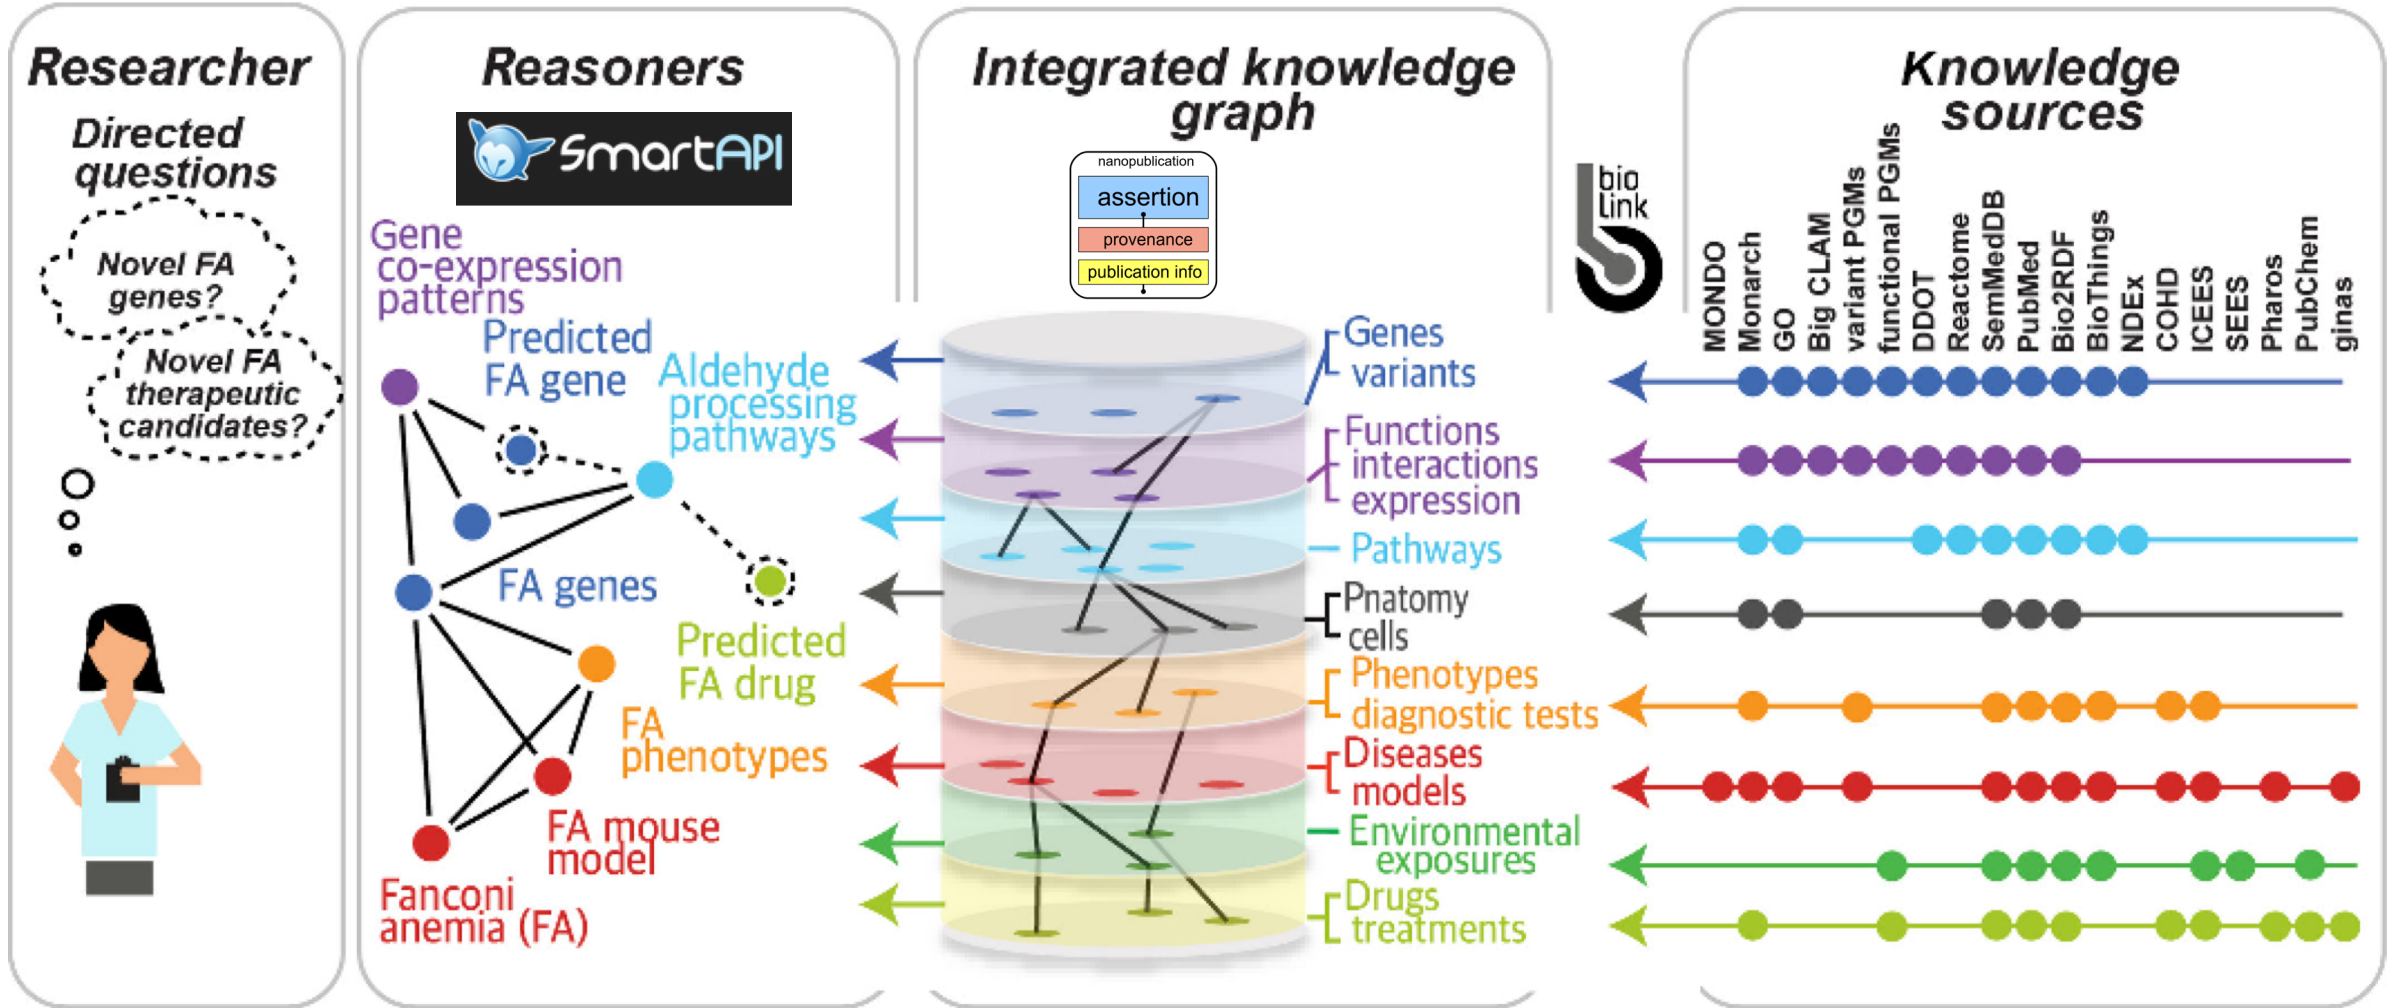
Journal of the American Medical Informatics Association, 29(3), 2022, 424–434

<https://doi.org/10.1093/jamia/ocab270>

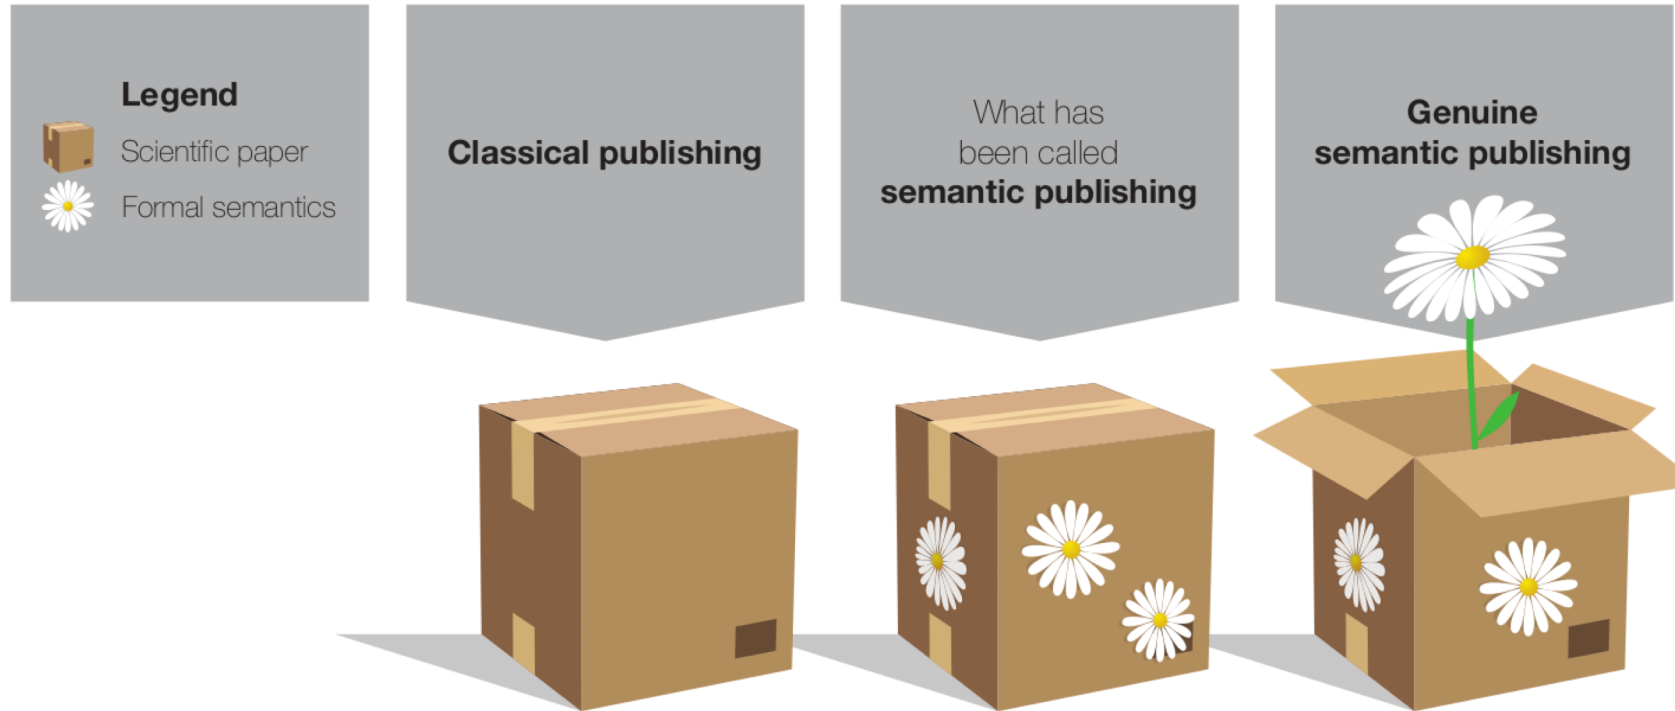
@micheldumontier::HCLS@KGC:2022-05-02



Biomedical Data Translator



Rethinking Publishing Data-Driven Research



Genuine Semantic Publishing

by Tobias Kuhn and Michel Dumontier

Content:

- as PDF
- as HTML/Dokieli
- as HTML/RASH
- as RDF/Turtle
- as RDF/TriG

Data Science. 2017 1(1-2):139-154. DOI: 10.3233/DS-170010
<http://www.tkuhn.org/pub/sempub/>

A formalization of one of the main claims of “The FAIR Guiding Principles for scientific data management and stewardship” by Wilkinson et al. 2016¹

Michel Dumontier

Maastricht University, The Netherlands

E-mails: michel.dumontier@maastrichtuniversity.nl, michel.dumontier@gmail.com; ORCID:

<https://orcid.org/0000-0003-4727-9435>

53

```
← → ↻ https://np.petapico.org/RA22JAQihYeiJkNIjvwnxLpMjuG74yPcRXpPyVX8DV6FA
@prefix ns1: <https://w3id.org/np/o/ntemplate/> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix this: <http://purl.org/np/RA22JAQihYeiJkNIjvwnxLpMjuG74yPcRXpPyVX8DV6FA#> .
@prefix sub: <http://purl.org/np/RA22JAQihYeiJkNIjvwnxLpMjuG74yPcRXpPyVX8DV6FA#> .
@prefix np: <http://www.nanopub.org/nschema#> .
@prefix npx: <http://purl.org/nanopub/x/> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix orcid: <https://orcid.org/> .
@prefix prov: <http://www.w3.org/ns/prov#> .

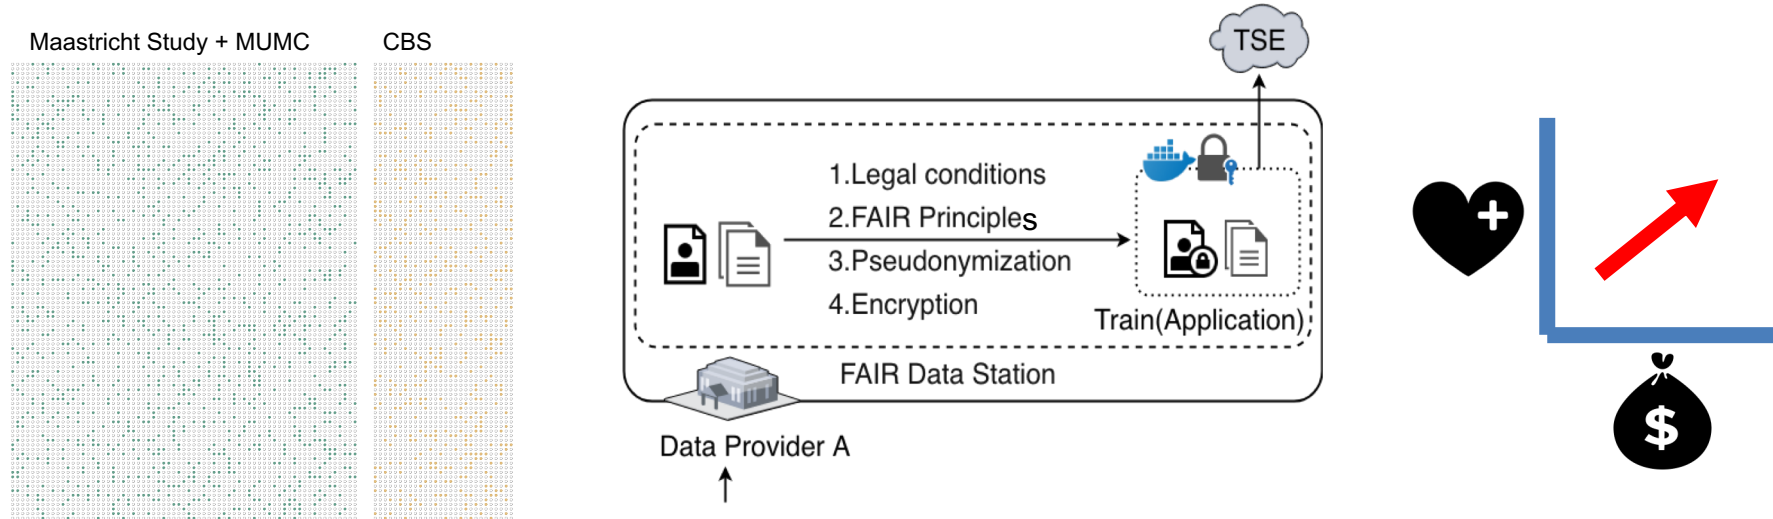
sub:Head {
  this: np:hasAssertion sub:assertion ;
  np:hasProvenance sub:provenance ;
  np:hasPublicationInfo sub:pubinfo ;
  a: np:Nanopublication .
}

sub:assertion {
  sub:spi a <https://w3id.org/linkflows/superpattern/terms/SuperPatternInstance> ;
  rdfs:label "Adherence of a dataset to the FAIR Guiding Principles enables its automated discovery." ;
  <https://w3id.org/linkflows/superpattern/terms/hasContextClass> <http://www.wikidata.org/entity/Q1172284> ;
  <https://w3id.org/linkflows/superpattern/terms/hasObjectClass> <http://purl.org/np/RAFQovt9yQD7n22td29_Uhpb7CsfT3k64pK7dh63xd-50#automatedDiscovery> ;
  <https://w3id.org/linkflows/superpattern/terms/hasQualifier> <https://w3id.org/linkflows/superpattern/terms/canGenerallyQualify> ;
  <https://w3id.org/linkflows/superpattern/terms/hasRelation> <https://w3id.org/linkflows/superpattern/terms/enables> ;
  <https://w3id.org/linkflows/superpattern/terms/hasSubjectClass> <http://purl.org/np/RAodU4AmRjFzyjwJK31u001yRJJUPBjkjKwDlMHvack#adherenceToTheFAIRGuidingPrinciples> .
}

sub:provenance {
  sub:activity a <https://w3id.org/linkflows/superpattern/terms/FormalizationActivity> ;
  prov:used sub:quote , <https://doi.org/10.1038/sdata.2016.18> ;
  prov:wasAssociatedWith orcid:0000-0003-4727-9435 .
  sub:assertion prov:wasGeneratedBy sub:activity .
  sub:quote prov:value "the FAIR Principles put specific emphasis on enhancing the ability of machines to automatically find and use the data" ;
  prov:wasQuotedFrom <https://doi.org/10.1038/sdata.2016.18> .
}

sub:pubinfo {
  sub:siq npx:hasAlgorithm "RSA" ;
  npx:hasPublicKey "MIGIMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCTQs+mANCSHw1W/YP1o468UdGNHsPvADpjfaW8um/v2L4AodIANqinfoU65VnbPT5D0AD1y0uFNne3VEMe9Y+I2HFaz6IKj+LdYmJk6VUf5WJoImRHIX6B2QwUc22cBTBfYxvqvp3UmmHrCchLIzjDSyutExK3tOTRoMDjGowIDAQAB" ;
  npx:hasSignature
  "hHeN9qAHbRXqs1k6ztdFWPTPTPYrI1c1GL+nH6YX7A88Qq170dJFXyBFGcv70pO1EmEvsAlNs2Xn7oe1CmpsFBT11vwVPLf8SswzXrpnDU2p9na1r6YMLyrNJ3wLq1pXwAOH82n1sVr1GMtL7v0VGw8cCmhdvzAS1o"
}
```

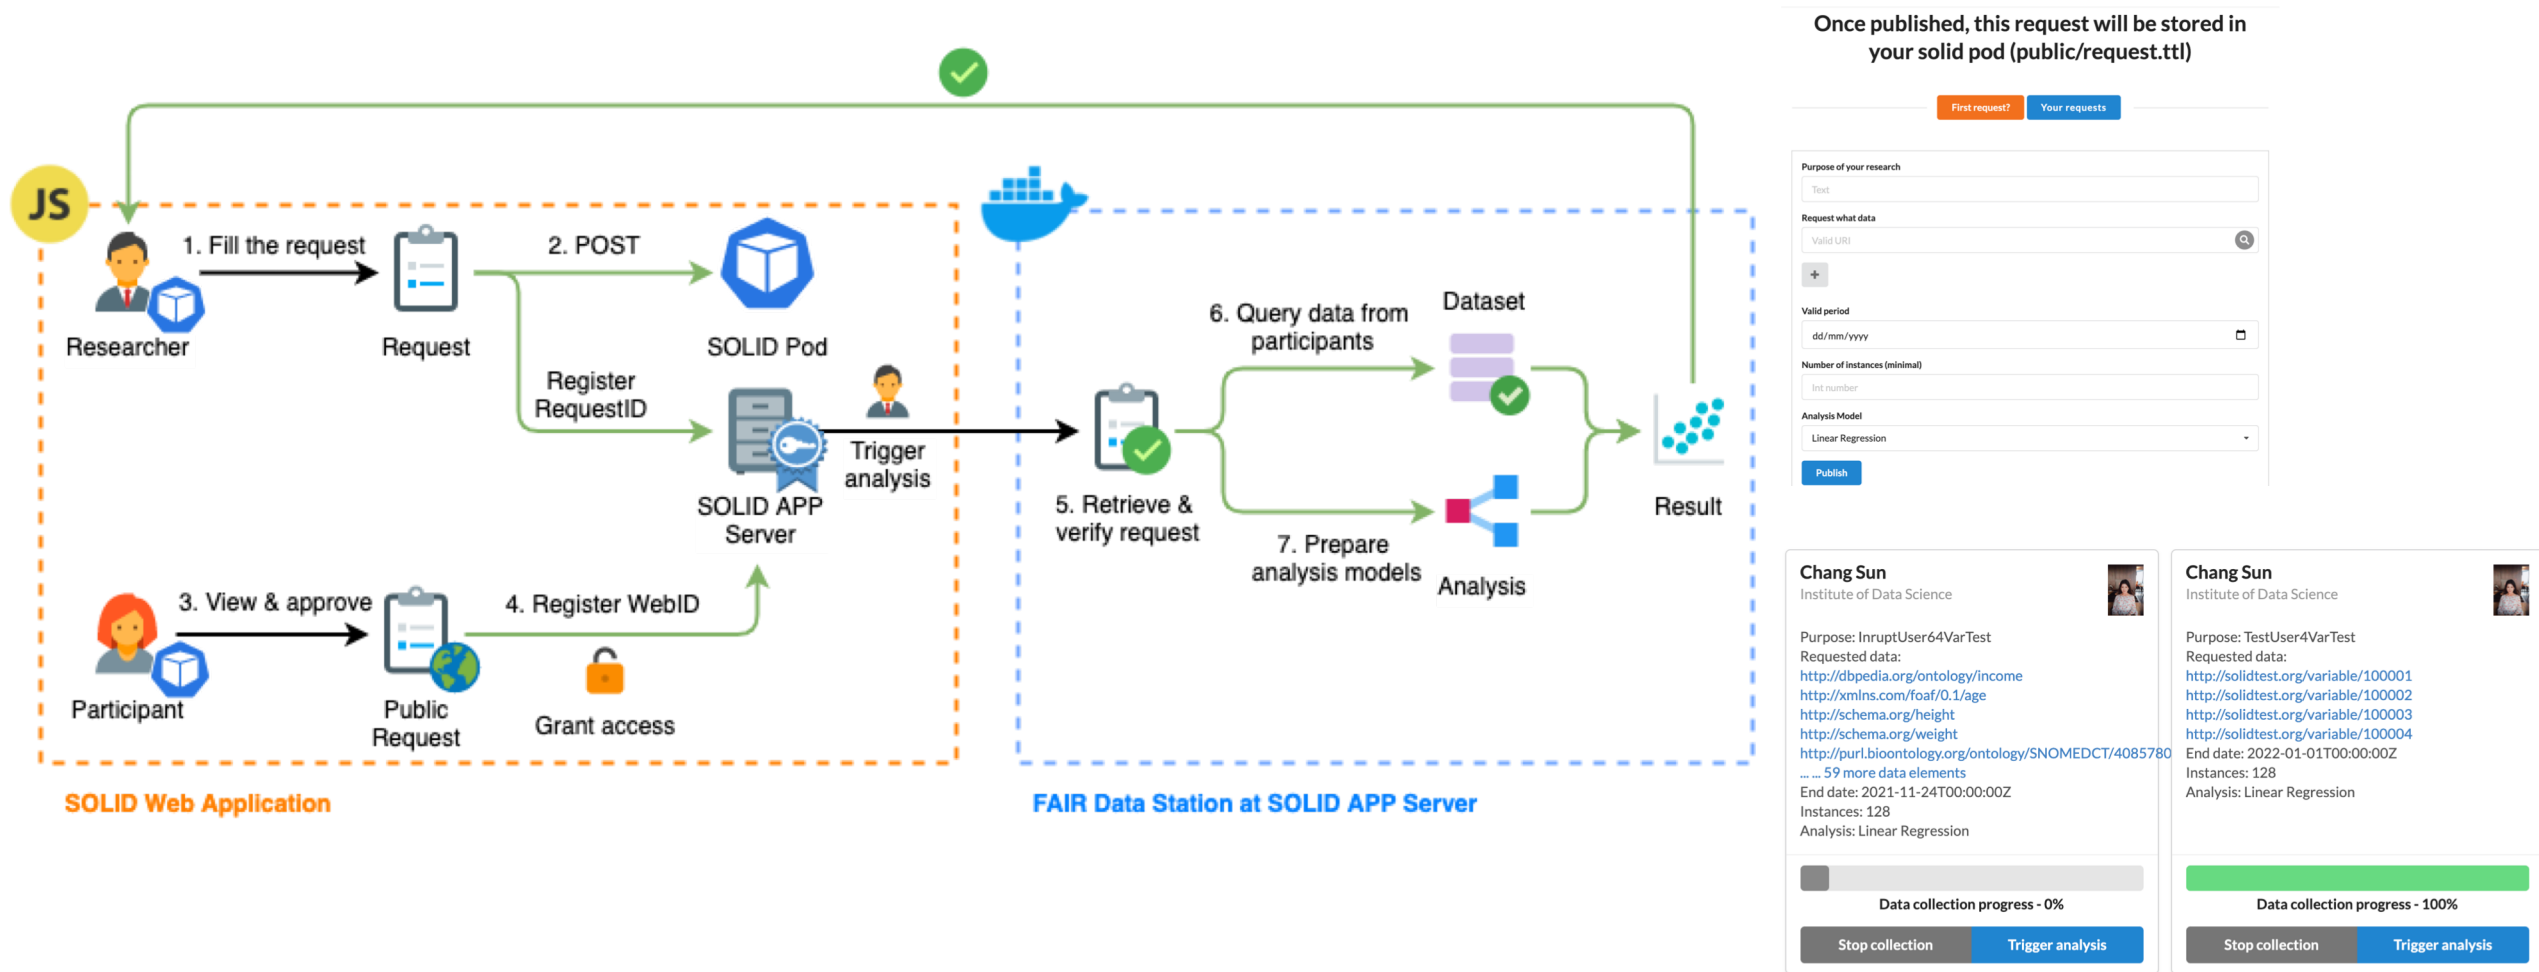

Mine distributed, access restricted FAIR datasets in a privacy preserving manner



Learn **high confidence** determinants of health in a **privacy preserving** manner over **vertically partitioned data**. The data are made available through **FAIR data stations** that provide access to *allowable* subsets of data to *authorized* users of *approved* algorithms.

TIDAL: Citizen-centric data platform using SOLID

using personal data for research




Summary

FAIR requires certain features for machine discovery and reuse, and offers new research possibilities that may reduce effort and improve transparency and confidence in published research

FAIR KGs aim to provide semantically annotated, standardized, and AI-ready data to *facilitate* downstream reuse

We need to foster the construction of FAIR knowledge graphs from the source, and in a manner that supports purpose-driven interrogation in a privacy-preserving manner, where needed.

AI technologies, coupled with semantics, will enable **researchers to exploit an emerging Internet of FAIR data and services in a (semi)automated manner**, and hence to accelerate discovery in biomedicine and in other disciplines, and to help realize the unforeseen value of existing data.



The mission of the **Institute of Data Science at Maastricht University** is to foster a collaborative environment for multi-disciplinary data science research, interdisciplinary training, and data-driven innovation.

We tackle key **scientific, technical, social, legal, ethical issues** that advance our understanding across a variety of disciplines and strengthen our communities in the face of these developments.

michel.dumontier@maastrichtuniversity.nl

Website: <http://maastrichtuniversity.nl/ids>

