

Tool-based Support for the FAIR Principles for Control Theoretic Results: The “Automatic Control Knowledge Repository”

Carsten Knoll
Institute of Control Theory
Dresden University of Technology
 Dresden, Germany
 carsten.knoll@tu-dresden.de

Robert Heedt
Institute of Control Theory
Dresden University of Technology
 Dresden, Germany
 robert.heedt@tu-dresden.de

Abstract—In 2016 a collection of guiding principles for the management of scientific data was proposed by a consortium of scientists and organizations under the acronym FAIR (Findability, Accessibility, Interoperability, Reusability). As many other disciplines, control theory also is affected by the (mostly unintended) disregard of these principles and to some degree also suffers from a reproducibility crisis. The specific situation for that discipline, however, is more related to software, than to classical numerical data. In particular, since computational methods like simulation, numeric approximation or computer algebra play an important role, the reproducibility of results relies on implementation details, which are typically out of scope for written papers. While some publications do reference the source code of the respective software, this is by far not standard in industry and academia. Additionally, having access to the source code does not imply reproducibility due to dependency issues w.r.t. hardware and software components. This paper proposes a tool based approach consisting of four components to mitigate the problem: a) an open repository with a suitable data structure to publish formal problem specifications and problem solutions (each represented as source code) along with descriptive metadata, b) a web service that automatically checks the solution methods against the problem specifications and auxiliary software for local testing, c) a computational ontology which allows for semantic tagging and sophisticated querying the entities in the repo and d) a peer-oriented process scheme to organize both the contribution process to that repository and formal quality assurance.

Index Terms—computational methods in control theory, test driven development, semantic technology, computational ontology, reproducibility, FAIR principles

I. INTRODUCTION

The scientific advances of the past have been possible because researchers either added their own ideas to earlier ones or modified them to draw different conclusions. This process is mainly facilitated by the publication of results. For the creative individual it is necessary to fully understand published results in their important details like assumptions, reasoning, conclusions – both for improving and for criticizing. Despite omnipresent digital technology, the main medium for the necessary information transport is still (written) natural language. Like in many other disciplines, scientific communication in control theory typically is enriched with an ample stock of

mathematical symbols and other formalisms such as block diagrams, but natural language still provides the necessary context.

However, since more and more computational methods play a crucial role in systems and control research, it becomes increasingly difficult to fully understand the essence of many publications in contrast to merely taking notice of their statements and claims. The difficulties of software in relation to reproducibility have been recognized early on [1] and are still a frequent point of discussion [2–6]. Recently, the journal *ReScience C* [7] was founded specifically to document examples of these issues. From the perspective of the authors, full understanding implies the ability to independently reproduce the results from a ground of common assumptions. For (conventional) mathematical proofs and similar reasoning, natural language and written formal symbols therefore still seem appropriate. On the other hand, for results that practically cannot be reproduced (i.e. understood) without using computation (like simulation, numerical linear algebra, or symbolic calculations with computer algebra systems), it appears naturally that all information that is necessary to reproduce these results should be provided along with the published description in natural language. This includes experimental data [8–10] but also, and arguably more important in our field, source code for simulations and other types of calculations [11–17]. Even with the information pieces seemingly available in published texts and source code, interested readers still often run into road blocks caused by differences in the environment, configuration, or missing documentation. Recent efforts therefore propose employing so called Continuous Integration (CI) [18–20], a concept explained below.

However, reproducibility is only one of several problem areas in which significant research resources are dissipated. In 2016 an international consortium of scientists and organizations published a joint appeal to “improve the infrastructure supporting the reuse of scholarly data” [21]. In this widely recognized publication a set of guiding principles for the management of scientific data was proposed under the acronym FAIR (Findability, Accessibility, Interoperability,

Reusability)¹.

This perspective reflects the fact that before reproduction of scientific results can be considered, those results have to be found, accessed and evaluated w. r. t. relevance. While classical search engines are doubtlessly helpful, they operate mainly on the *syntactical* level (e. g. searching for matching words) and thus are of limited use. An improvement can be expected by the use of *semantic* information as this would allow searching via the *meaning of words* instead of their syntactical representation as a string of characters. However, publishing results with such suitable semantic information requires certain standardization. Additionally, efficient searching would benefit greatly from the existence of a *central* knowledge base.

While the present paper makes extensive use of the term 'knowledge', the exact meaning of that term is an (open) philosophical issue and thus far beyond the scope of this paper. In this contribution we mainly refer to instrumental and to conceptual aspects of knowledge: The *know how* suitable to solve certain specific problems and the capability to express relevant statements about such problems and solutions in a precise and understandable way. This kind of knowledge can beneficially be represented as a combination of source code, a so called (computational) ontology (see below) and some additional metadata.

Assuming that an upcoming publication owes (part of) its content to computational methods, the present contribution proposes the following: a) As a supplement to the usual natural-language-based publications (typically published as PDF file), the related source code along with semantic metadata should be published to a central and open repository with a suitable data structure. To facilitate this, we further suggest b) a web service connected to that repository that automatically checks the solution methods against the problem specifications and a suite of auxiliary software for local testing. To enable users to better find existing solutions, independently of linguistic ambiguity, we propose c) the coupling of said repo to an ontology which provides both a controlled vocabulary usable for tagging, as well as an interface for sophisticated querying based on semantic web standards. Finally, we propose d) a peer-oriented process scheme to organize the contribution process (acceptance/refusal) and formal quality assurance without the need for a central authority.

From the viewpoint of the authors, the proposed approach of an "active knowledge repository" interconnects innovative technical and social aspects. However, to have any noteworthy effect on the above mentioned FAIR-related problems it must gain some significant support and adoption by other researchers and therefore must be open to ideas, criticism and improvement suggestions by the community from an early stage. Consequently, the presented approach should not be

¹See section VI-A for further details. While reusability (of data) and reproducibility (of scientific results) in general are not identical, for the scope of this contribution we consider it as similar goals. This is due to the fact that we aim to apply the FAIR principles to software-based results instead of mere (numerical) data. For software to be reusable it is necessary to be functioning and thus its results must be reproducible.

considered a "final solution" but instead a conceptual draft, which might be the point of departure for further discussion and iterative development.

The remainder of this paper is structured as follows: Section II summarizes relevant established technologies and a concept for how to integrate them. Section III-A proposes the above mentioned dedicated structure of the repository, while Section IV describes how the repository and the proposed web service can be combined to perform a *contribution* as the basic use case of the whole platform. In Section V we briefly comment on what elements of the concept are already implemented (see also [22]) and in Section VI-B answer some critical questions, which are expected to arise for the reader.

Note that the present contribution is an extended version of the conference paper [23], where the Automatic Control Knowledge Repository (ACKREP) was firstly introduced. The extension mainly consists in the development of an ontology for the domain of control systems engineering (OCSE) and its coupling to the original ACKREP approach, which enables better support for the FAIR principles, especially for the findability aspect². Furthermore, the actual knowledge repository has also been extended significantly, from originally 9 entities to currently about 50.

II. EXISTING SOLUTIONS AND INTEGRATION CONCEPT

This section briefly summarizes already existing technologies and techniques. Then these elements are combined to form an integrated process concept.

A. Collaboration via Distributed Version Control Systems

A key feature for digitally represented information is the ability to perform changes (including adding or removing content) with very little effort. However, keeping track of such changes is a challenge on its own, especially if changes are asynchronously made by several individuals. In software engineering, version control systems (VCS) are used to solve this problem. During the last decade, decentralized version control systems (DVCS) and most prominently the software *git* have arguably become the de facto standard. A set of information that is under version control is called a *repository*. It typically represents a file system (directories and files) along with its evolution history. This history is represented by a collection of incremental snapshots, stored along with metadata like time stamp or author as a collection of *commits*. The order relation inside such a collection is ensured by using a hash tree (also called Merkle tree [24]) as the data structure: For each commit, a unique cryptographic hash is calculated and stored as part of the metadata. Additionally, each commits metadata also contains the hash code(s) of its parent(s) commit(s), i. e. the state of the repository to which the incremental snapshot applies.

This data structure allows for temporary parallelism. Each node (commit) in the tree can have multiple child nodes

²The aspects accessibility, interoperability and mainly reusability (via reproducibility) were already covered by the original approach to some extent.

(branches). Moreover, each node can have multiple parent nodes too, which enables branches to be *merged* together after a period of independent evolution, see Fig. 1. Git provides mechanisms to resolve merge conflicts automatically or manually.

This technology facilitates a distributed development model, which is adopted by many open source projects: Potential contributors *clone* the original repository, i.e. create a local copy of it, make their own contributions in the form of commits, and turn in a *merge request* for the original repository. Typically, some review iterations take place before the external contribution is merged into the main repository.

B. Automated Tests and Continuous Integration Services

An *automated test* is a piece of software that executes other pieces of software and compares their results to (hardcoded) expected results [25]. Roughly, tests can be subdivided into *unit tests*, checking the isolated functioning of a small piece of software, and *black box tests*, checking correct input-output behavior of a larger piece of software.

While creating and maintaining such “unproductive” code takes additional effort, this investment usually is justified even for projects of modest complexity by the increased efficiency of identifying unintended behavior (“bugs”). Additionally, automated tests are an important technology if heterogeneous execution environments (e.g. different version of dependency libraries) should be supported and greatly help to facilitate distributed software development, see section II-A. Due to this surplus value of suitable automated testing, so called *continuous integration services* were developed. Such a service basically monitors the state of a repository and whenever this state changes (e.g. by some uploaded commits) it executes the set of automated tests that are defined in the repository itself and reports the result in a suitable way.

C. Semantic Technologies: OWL-Ontologies and SPARQL-Queries

Currently the most common methods of representing complex knowledge computationally are semantic triples of

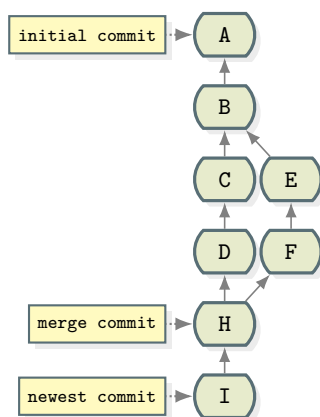


Fig. 1. Schematic visualization of branching and merging in the distributed version control system git.

subject-predicate-object-relations and (computational) ontologies. In this context an *ontology* is a formal (i.e. machine-actionable) specification of a shared conceptualization (semantic coverage) of a knowledge domain [26]. In other words, an ontology specifies which concepts do exist in a domain and how they are related to another [27–29].

Semantic triples are often represented in RDF³ file format while for ontologies the *Web Ontology Language (OWL)* is prevalent, which can be interpreted as an additional layer on top of RDF and enables knowledge representation by means of so called Description Logics [30, 31] which are decidable fragments of first order predicate logics. The added value of such an ontological formalization is that it enables automatic reasoning for a knowledge base.

To be of any use, knowledge bases, such as Wikidata [32] have to provide search access to potential users. The *SPARQL Protocol and RDF Query Language* (recursive acronym) is a standard approach to this regard. This query language explicitly is designed to make use of the triple structure, where the predicate typically is given by an ontological relation and subject and object can either be abstract OWL concepts or an instance of such. The strength of this approach is that via boolean combination of such atomic queries, results can be retrieved which are not explicitly present in the knowledge base. This is especially true if an automatic reasoner has been applied to infer “new” relations between the entities of the knowledge base.

Interestingly, such technologies, which originated in the context of the so called semantic web [33, 34] have had huge impact on life sciences, e.g. cf. [27], but as of yet seem to be almost unknown or rarely used in engineering. For example, [35] is one of the very few attempts to apply semantic technologies to the domain of control theory of which the authors are aware of⁴.

D. Integration Concept: “Automatic Control Knowledge Repository”

The aforementioned technologies and approaches do exist and have been applied for years or even decades. However, to the best of the authors’ knowledge, there does not exist a combination of these components suitable for the needs of the scientific community and in particular the requirements of researchers in the field of automatic control and systems theory. Thus, the purpose of this contribution is to propose such a combination and to expose these ideas to critical comments from the research community.

As mentioned in the introduction, we propose in particular

- 1) to establish a central repository in the sense of Sec. II-A, which collects content-related software (e.g. control theoretic algorithms), organizational software (the automated tests) and metadata in a formal structure,

³Abbreviation for *Ressource Description Framework*.

⁴See also [36] for “Methodnet”, a complementary approach by the authors which has a different focus than the proposal from this paper but is designed to explicitly benefit from interlinking data with ACKREP.

- 2) to establish a supportive web service for that repository, providing CI functionality tailored to its structure and to the needs of systems and control engineering, plus a suite of software tools to simplify contributions to that repository e.g. by allowing for local testing,
- 3) to establish an ontology to formalize control theoretic knowledge in order to harmonize language use and enable SPARQL-searches on the repository,
- 4) to establish a set of rules that distribute the decision making about submitted contributions (merge requests) and other maintenance work among the previous active contributors.

We call this concept *active knowledge repository* to reflect the fact that it is more than a passive store of information. Due to the testing carried out by the web service, it is guaranteed that the contained software had worked under clearly specified conditions on computational infrastructure that is independent from the content contributor (and thus not prone to undiscovered dependency issues). The term “knowledge repository” distinguishes from the term “knowledge base” because it suggests dynamic evolution by contributions from multiple sources and from the mere “repository” because it hints that it does not just contain source code but by the kind and structure of this source code (and metadata), the repository content can be regarded as a representation of domain specific problem solving knowledge (which can be transferred to other problems).

III. PROPOSED REPOSITORY STRUCTURE

As described above, the repository should contain instrumental as well as conceptual knowledge. The following subsections detail both aspects.

A. Instrumental Knowledge Represented as Software + Metadata

To reproduce a result, it is necessary to reduce it to a binary yes-no question. Our approach is based on the observation that the typical (computational) result in control theory can be split up into a precise description of the problem (or class of problems), e.g. computing an input trajectory for an equilibrium transition of a state-space system, and a description of how to solve these problems. This corresponds to the concept of a unit test, where the problem description is the test case and the solution description is the software to be tested. As a test case can succeed or fail, a problem solution can either solve a problem (i.e. calculate the expected result from the provided data) or fail to do so. To achieve flexibility and modularity, the repository structure requires a formal split of contributions into independent (but linked) *entities* of different types, see Fig. 2. The most important entities are `ProblemSpecification` and `ProblemSolution`⁵.

Note that in this structure a `ProblemSpecification` is also given by executable code (indicated by the gear wheel).

⁵Short forms: `Problem` and `Solution`. Other short forms are similarly self-explaining.

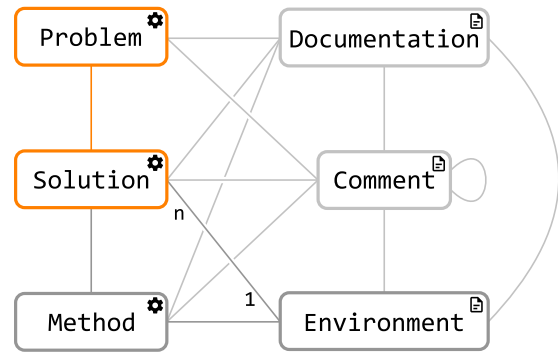


Fig. 2. Relations between the various entities. All cardinalities are m -to- n with one exception: Each `ProblemSolution` is associated with exactly one `EnvironmentSpecification`. Note that entities of type `Comment` can also reference entities of their own type. The gear wheel symbolizes executable code associated with the entity, other (meta)data is marked by the file icon.

The solution entity is expected to represent the concrete application of zero or more `MethodPackages`, which for example contain the code for solving boundary value problems or adaptive grid refinements. Such methods might be complex pieces of software, which are compatible with one or more runtime `EnvironmentSpecifications`.

Finally there are entity types for `Documentation` (with obvious use) and `Comments`. The latter provides a formal possibility for transparent long term communication, such as improvement suggestions or criticism related to other entities.

In the repository, each entity is represented by a directory containing at least one file named `metadata.yml`, which must specify some defined⁶ attributes in the widespread and easy to edit YAML format. The most important of these attributes is the *entity key*, which is required to be a five character string of numbers and capital letters and serves to uniquely identify every entity in the repository. While most attribute fields are generic to all entity types, some are specific to one type such as the field `estimated_runtime` for a `ProblemSolution`. Entities that contain source code obviously need to include the respective files in their entity directory and reference it in the metadata, e.g. `solution_file="solution.py"` for the type `ProblemSolution`.

While the type for each entity is specified in the metadata, it is a convention to have a top level directory structure that maps to the different types. Inside these top level directories, arbitrary subdirectories can exist, but entities are not allowed to be nested.

B. Conceptual Knowledge Represented as Ontology

While the entities from the previous subsection are suitable to describe the precise details of problems and solutions, they

⁶For details see the `README.md` file of [37].

(alone) are not condign to represent their interrelations and how they fit into “the big picture” of control theory⁷. To this end, we propose to include a computational ontology in the repository, which contains all terms relevant to properly characterize the ACKREP entities. This so called *Ontology of Control Systems Engineering (OCSE)* mainly consists of a taxonomy sorting concepts into a hierarchy of distinct generality. E. g., a `Linear_State_Space_System` is a special case of a `State_Space_System` which is a special case of an `ODE_System` etc. In addition to these `is_a`-relations, OWL allows for arbitrary further relation types. Thus it is possible to represent, e. g., the fact that a `Transfer_Function` (special kind of dynamical system model) can be *represented by* a `Rational_Function` (special kind of scalar mathematical mapping).

A major challenge in this regard is the proper formal representation of mathematical knowledge. Obviously, a significant conceptual coverage of mathematics is needed for control theory, but on the other hand a too deep representation of mathematics inside OCSE would require an infeasible effort. While there are several existing approaches of formal mathematical knowledge representation (see [38, 39] and references therein), many of which are focused on automated theorem proving, unfortunately none of them seems suitable to provide the necessary concepts for control theory. As a preliminary pragmatic solution mathematical concepts are therefore covered in OCSE only as needed with a rather flat relation structure.

The coupling between the ontology and the software-related entities is achieved by using the concept names from the OCSE (i. e. names of ontological classes) as tags in the metadata files for each entity. Thus the OCSE provides a *controlled tag vocabulary* and prevents linguistic ambiguity. But more than a mere word list the ontology also encodes semantic relationships between the concepts and allows for SPARQL queries (see Section IV-D).

For the sake of understandability the OCSE is represented as a YAML file in the repository, i. e. in the same format as the defining metadata files of the instrumental entities. This is motivated a) as we assume little to no prior knowledge of OWL and b) to better track and review changes in version control. To be compatible to established standards we also provide a version in RDF/XML which is automatically generated from the OCSE YAML source by means of [40].

Fig. 3 shows an excerpt of the OCSE taxonomy.

IV. GENERAL ARCHITECTURE AND PROCESS SCHEME FOR CONTRIBUTION AND INFORMATION RETRIEVAL

A. Architecture Overview

So far, we only have a passive repository with a suggested structure (in the following named `ackrep_data`). However,

⁷In the original approach [23] this task was referred to an additional entity type `ProblemClass`. However, it turned out that representing the necessary relations solely via the respective `metadata.yml`-files would either be very limited or would require reinventing established technologies from knowledge representation.

the main advantage of the proposed approach should be that a contribution to the repo will only be merged if it is *consistent*, i. e. obeys the formal structure, and all (old and new) `ProblemSolutions` are able to solve the related `Problems`. In other words: all tests must pass. To check this, we propose a suitable interoperation between different software components. This architecture is visualized in Fig. 4.

The basic component is a software that crawls through the repository, checks the consistency of each entity, collects all solutions and related problems, runs the individual test cases and reports the results. In the current prototype implementation this software is called `ackrep_core`. It is available with a command line interface, which allows local testing on personal IT infrastructure. However, the relevant test run, which decides whether a contribution is acceptable, must take place on independent infrastructure, i. e. on a server. To comfortably interact with `ackrep_core` on that server, there is a web interface (`ackrep_web`). To simplify the processing and interconnection of the entities, the metadata from the repository is loaded into a database. When a solution check is triggered via the web service, `ackrep_core` creates an executable script from a template. This script is then executed by a component named `code executor`, which is configured according to the relevant `EnvironmentSpecification`. The script passes the `ProblemSpecification` object to the `ProblemSolution` function. Then the script passes the result into the `evaluate_solution` function, which finally compares the numerical/symbolical values of the solution with the expected ones and decides whether the test was passed or failed. The result is displayed to the user via `ackrep_web` along with optional graphical results that were created by the solution, cf. Fig. 6.

Obviously, there must be a copy of the `ackrep_data` on the server. However, the “canonical place to live” for this repo should rather be a public repository host such as `github.com`, which we chose for the prototype. Alternatives would be e. g. `gitlab.com`, `bitbucket.com` or `codeberg.org` (non-profit). Such a service provides features for browsing the repository history and creating a personal *clone*, which can evolve into a *fork*, i. e. a version of the repository that shares the history but after a branching point differs from the official repo state, e. g. due to additional commits. For such a fork (which gets its own URL from the repo host), a potential contributor can then file a merge request via `ackrep_web`, which triggers the consistency and solution checking described above.

B. Basic Contribution Scheme – Constitutional Rules

The merge request described above only has to be triggered by a potential contributor and the formal quality assurance in the form of consistency and solution checking takes place automatically. However, these formal conditions are only necessary but not sufficient. The final decision whether a contribution should be merged to the official (i. e. canonical) `ackrep_data` should be made by humans based on relevance of the contribution and its scientific and educational value.

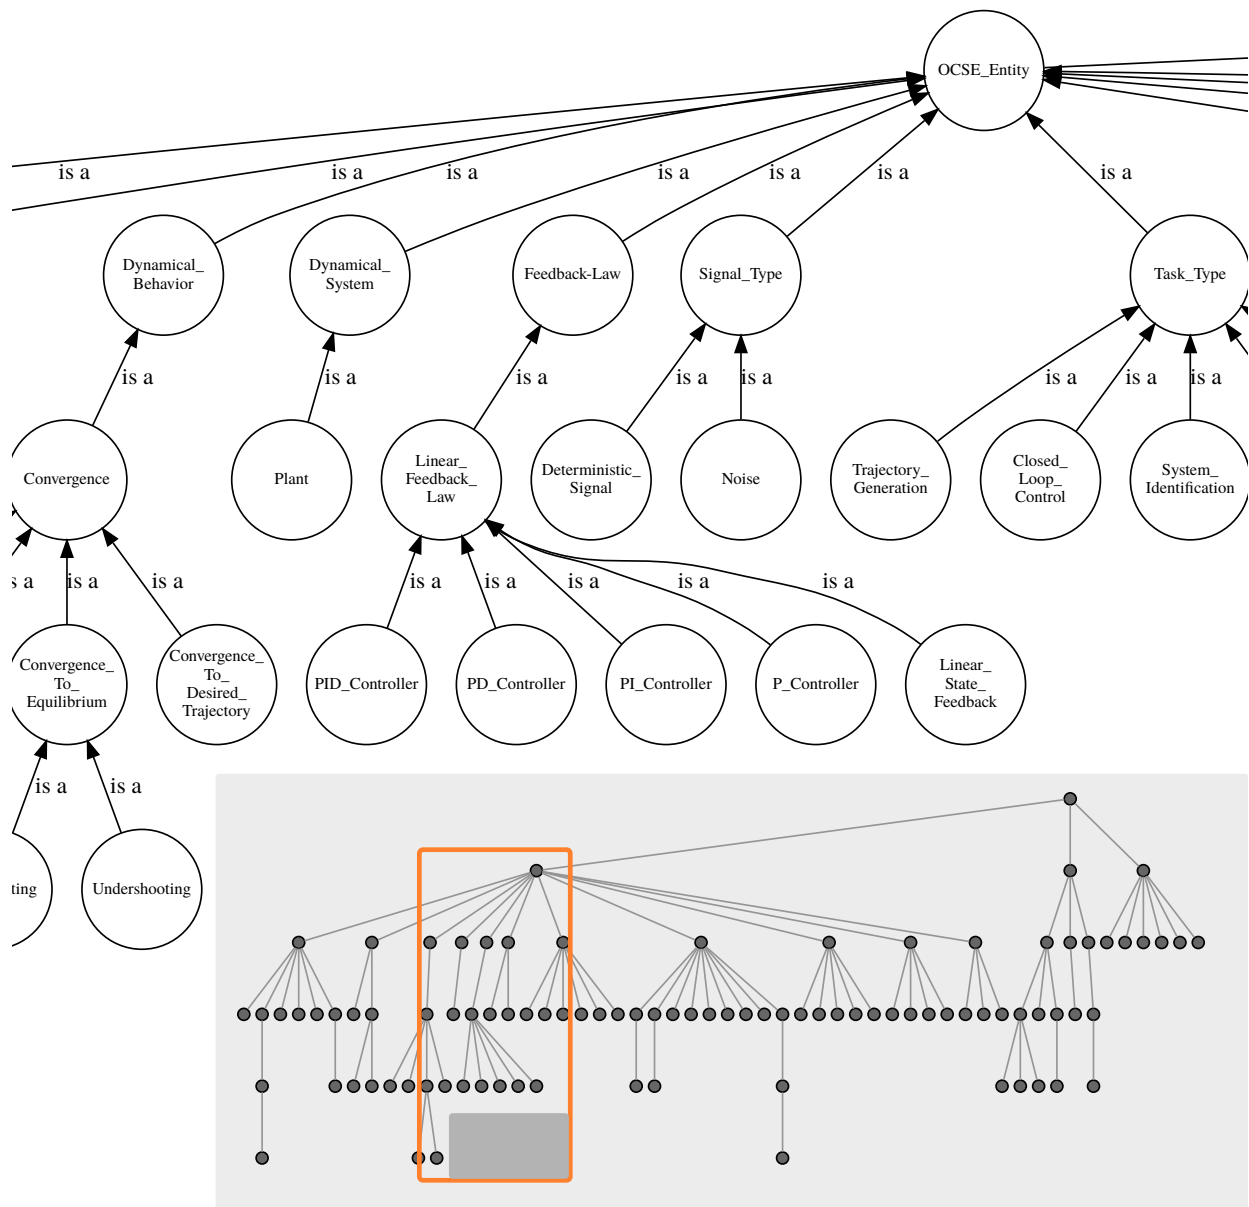


Fig. 3. Section of the taxonomy (subclass hierarchy) of the Ontology of Control Systems Engineering (OCSE). Note that the depicted *is_a*-relation is only one of several relations between the concepts. The inset in the lower part shows an unlabeled overview of the whole taxonomy graph with the visible section highlighted.

The aim to establish peer-based decision making in potentially controversial situations while minimizing both conflict and bureaucratic overhead poses some social challenges, which are tackled in the form of a set of rules called *constitution*. The basic idea is that any decision, both regarding contributions and the constitution itself, should be made consensually by the *group of active contributors* (see below). For initial simplicity, there is no other body of decision-making, but as this group grows, it can and should (consensually) decide to establish a more elaborated structure. To facilitate this idea, the proposed initial constitution describes a *scalable rejection-minimizing consensus mechanism*, see *Constitution.md* in [37].

From the viewpoint of the authors, rejection minimization is important for the following reason: In contrast to rules implemented in program code, which are applied to inanimate and indifferent pieces of information (“data”), the subject of the constitutional rules are human individuals with their own interests and preferences. Thus, these rules are not followed due to natural laws (as is the case for data processed by software on a digital device) but due to free decision. If a majority vote displeases group members, it is likely they leave and will not contribute anymore.

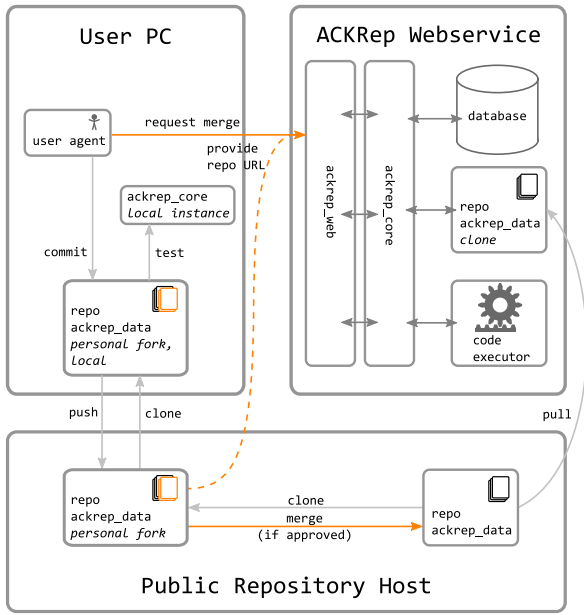


Fig. 4. Schematic visualization of the proposed component architecture and their interaction during the use case of a contribution (see Sec. IV-C). The orange elements are associated directly with a merge request.

C. Basic Use Case 1: Making a Contribution

For the sake of understandability, the use case of contributing is described as a concrete example. We assume a) that there is a set of distinct individuals $\{C_1, \dots, C_n\}$ who have already contributed to the repository (“active contributors”) and b) a *prospective contributor* \mathcal{P} who has developed a method for trajectory planning of nonlinear state-space systems and wants to contribute this method to the `ackrep_data` repository. To accomplish this, the following steps are required:

- 1) \mathcal{P} creates a personal fork of `ackrep_data` and locally installs the supporting software (`ackrep_core`).
- 2) \mathcal{P} splits up their code into entities according to the required data structure⁸, see Sec. III-A. This includes the creation of appropriate `metadata.yml` files containing suitable tags from the OCSE.
- 3) With the help of the `ackrep` command line tool, \mathcal{P} tests locally whether all consistency and solution checks pass and, if necessary, fixes all reported formal issues.
- 4) If no issues are reported, \mathcal{P} pushes the changes to their personal fork.
- 5) \mathcal{P} files a *merge request*⁹ via the `ackrep.org` web service.
- 6) The web service verifies that all tests pass.

⁸The main code will be a `MethodPackage` while `ProblemSpecification` and `ProblemSolution` can easily be created based on existing examples

⁹Often also called *pull request*, although *merge request* better reflects the essence of the concept: the request to merge new content into the existing repository.

- 7) $C_1 \dots C_n$ get notified about the merge request and express their opinion about accepting or rejecting the request and optionally suggest improvements.
- 8) After optional review iterations, the (final) merge request is accepted or refused based on the constitutional rules. If accepted, it is part of the official version of `ackrep_data` and if $\mathcal{P} \notin \{C_1, \dots, C_n\}$ then \mathcal{P} becomes C_{n+1} and $n := n + 1$.

It is worth mentioning that the decision makers $C_1 \dots C_n$, although they could, are not expected to perform a careful review of the code, which would clearly be too much effort in some cases. However, they are expected to check that the code is relevant to the scope of `ackrep_data` and does not contain obvious errors. If a problem with a contribution later arises, this could be documented e. g. via the `Comment` entity.

D. Basic Use Case 2: Information Retrieval

The actual goal of contributing knowledge to ACKREP is to provide access to this knowledge. To this end the ACKREP web service exposes a SPARQL-interface by which queries in that language can be directed to the ontology – or more precisely to the ontology engine which is currently implemented using *owlready* [41]. Because there is a one-to-one correspondence between ACKREP entities as described in Section III-A and *individuals* in the ontology (i. e. instances of `ACKREP_Entity` which is a subclass of `owl:Thing`) this interface offers a convenient way to apply semantic search to the ACKREP entities, cf. Fig. 5.

The main difference to ordinary text-based search is the following:

- **Controlled vocabulary:** Because only search terms which are present in the ontology (OCSE) apart from built-in SPARQL constructs can be used, the problem of linguistic ambiguity is prevented. E.g. the concepts of “trajectory planning” and “open loop control” both refer to essentially the same concept but this fact is completely

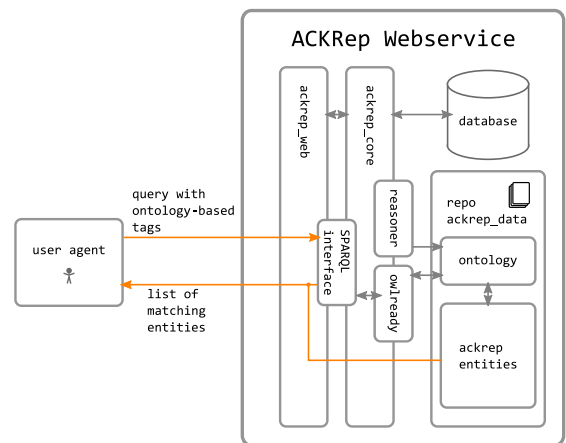


Fig. 5. Schematic visualization of the information retrieval via SPARQL queries and the involved components of the ackrep web service.

invisible on the syntactical level (on which text based search operates).

- Subsumption: Due to the taxonomic *is_a*-relation inherent in the OCSE the knowledge about generalization is represented. If a query asks for *State_Space_System* the result should contain entities with tagged with *Linear_State_Space_System* as the latter is a subclass of the former, although the more general tag is not explicitly present.
- Fine grain control of the results: in contrast to ordinary search dialogues, a SPARQL interface offers many more possibilities to formulate complex queries, e.g. by usage of logical operators such as UNION or NOT IN.

V. RESULTS: CURRENT IMPLEMENTATION STATUS

The aim of the authors is to develop a practically working concept, therefore a prototype of the proposed software `ackrep_core` and `ackrep_web` has been implemented and is publicly available as Free Software (GPLv3). However, not all of the described features are already implemented. For example, currently there is only one default `Environment`, which only supports Python code. Furthermore, the `code executor` component in Fig. 4 is not yet realized as a separate container as it should be for performance and security reasons but instead as a plain process run by the web service. The same holds true for the ontology engine (cf. Fig. 5): This ideally also should be represented by a separate process. Nevertheless, the current state of the code and the running instance at <https://testing.ackrep.org> should be sufficient to illustrate the proposed concept to a certain degree and to generate some feedback.

Since the first launch of the prototype together with the original conference paper [23] the number of entities has increased from originally 9 to currently 47. Among them are basic control tasks (from undergraduate level courses) such as linear state feedback, more sophisticated (and less known¹⁰) control approaches (such as the design of a reduced Luenberger observer or trajectory tracking based on differential flatness) as well as some recent research results of the authors for trajectory planning of closed kinematic loops [42] and numerical estimation of the region of attraction of equilibrium points (see Fig. 6).

VI. DISCUSSION

A. Evaluation of compliance with FAIR principles

The initial proposal of the FAIR principles [21] specifies each of the four goals by several criteria. In this section we review these criteria (quoted text) and how our proposed approach relates to them (unquoted text).

a) "To be Findable":

- "F1. (meta)data are assigned a globally unique and persistent identifier"
 - Implemented via entity key and IRI of ontology.

¹⁰at least to an application focused audience

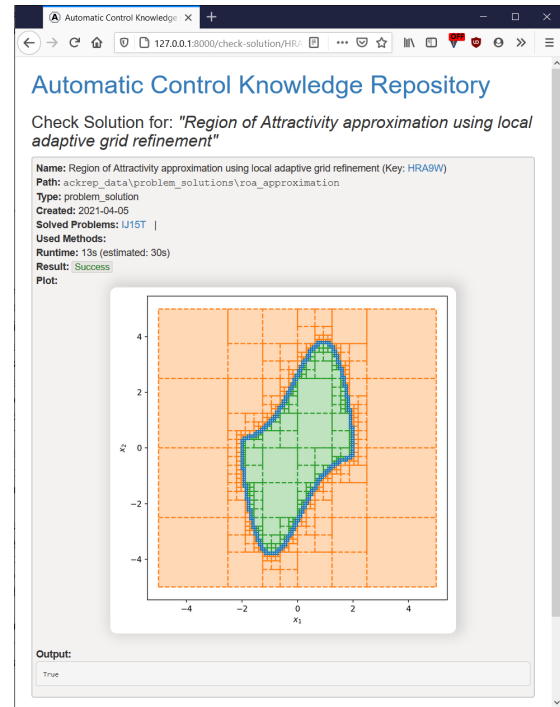


Fig. 6. Screenshot of `ackrep_web`. Currently displayed is the result page of the `check solution` command for a region of attraction approximation of the Van der Pol oscillator. The relevant metadata is shown along with the calculation result.

- "F2. data are described with rich metadata (defined by R1 below)"
 - Providing a `metadata.yml` is required for every entity to be recognized as such.
- "F3. metadata clearly and explicitly include the identifier of the data it describes"
 - Implicitly fulfilled due to the structure of the entities.
- "F4. (meta)data are registered or indexed in a searchable resource"
 - Implemented via the SPARQL interface.
- b) "To be Accessible":
- "A1. (meta)data are retrievable by their identifier using a standardized communications protocol"
 - Partially implemented: The whole repository including data (code) and metadata is available via git which could be interpreted as standardized communication protocol. However, a direct access to individual entities without cloning the whole repo would be more compliant with this criterion.
- "A1.1 the protocol is open, free, and universally implementable"
 - Implemented. git is Free and Open Source Software (FOSS) and so is its protocol.
- "A1.2 the protocol allows for an authentication and authorization procedure, where necessary"
 - Implicitly ensured by the features of git.

- “A2. metadata are accessible, even when the data are no longer available”
 - Not applicable due to the tight coupling of metadata and data. Both types of data have the same risk of becoming unavailable. However, see also Paragraph c in Section VI-B.
- c) “To be Interoperable”:
- “I1. (meta)data use a formal, accessible, shared, and broadly applicable language for knowledge representation.”
 - Implemented due to the usage of Python¹¹ (for data) and YAML (for entity metadata and ontology). Note, that the usage of YAML to represent an owl ontology is not standard, but is likely to be more accessible for the intended audience (assuming no familiarity with OWL or associated tools) due to being human readable plain text but much less redundant than the established XML representations.
- “I2. (meta)data use vocabularies that follow FAIR principles”
 - The controlled vocabulary used in the context of ACKREP is given by the Ontology of Control Systems Engineering (OCSE) – which itself is part of the repo and thus is subject to the general orientation towards these principles.
- “I3. (meta)data include qualified references to other (meta)data”
 - Not yet implemented. Unfortunately, the authors are not aware of any collection of (meta)data relevant to the scope of ACKREP. This is likely due to the rare usage of semantic technologies in the field of control systems engineering. However, both the metadata of the instrumental entities as well as the OCSE allow to add such references in the future.
- d) “To be Reusable”:
- “R1. meta(data) are richly described with a plurality of accurate and relevant attributes”
 - Implemented via tagging with OCSE concepts.
- “R1.1. (meta)data are released with a clear and accessible data usage” license
 - Implemented via the requirement of a Free Software license. The General Public License (GPLv3) is default for all contributions, but contributions with other licenses such as BSD are also admissible.
- “R1.2. (meta)data are associated with detailed provenance”
 - Implemented via attributes `creator` and `editor_list` in the metadata of instrumental entities. Furthermore, documenting the provenance of contributions is an essential feature of a version control system and thus built into ACKREP by construction.

¹¹or other programming languages in the future

- “R1.3. (meta)data meet domain-relevant community standards”
 - Not applicable, as the authors are not aware of existing domain-relevant community standards for sharing (meta)data. However, this contribution aims to conduce to the establishment of such standards.

Note that in the domain of control systems engineering – and probably other scientific domains which also strongly rely on computational results – the term “reusability” implies “reproducibility” (of such results) . Therefore, we want to point out that the automatic testing incorporated in the ACKREP approach facilitates reproducibility and thus reusability even more than required by the official FAIR principles.

B. Critical Questions

It seems probable to the authors that the proposed approach raises questions and maybe skepticism. In this section we try to anticipate and answer some of them.

a) *What is the added value to related publications?:* As stated in the introduction, earlier publications ceased either with a description (and lamentation) of reproduction hurdles or with an appeal to publish source code and to connect it with CI services. While the latter doubtlessly fosters reproducibility, this appeal is based on mere comprehension of the underlying problem. Given typical resource constraints and incentive structures in the research system, even strong intention for reproducibility might be outweighed by the necessary *additional* effort to transform the code into a “presentable” state and to deploy CI. The proposed approach of the *active knowledge repository*, however, might change this balance in favor of code publication due to the following expected effects:

- 1) Basing own work on existing examples reduces necessary effort and helps adjust the own ambition to coding style and quality.
- 2) Linking to the own code with an `https://ackrep.org/<key>-URL` proves that published code works on independent infrastructure ($\hat{=}$ formal quality assurance), which might provide some extra incentive.
- 3) Making the own results easily findable on a platform dedicated to the own research field increases peer-related visibility much more than just publishing code among millions of unrelated projects on a public repository host (which of course does not conflict with the proposed approach). Increased visibility is expected to be a strong extra incentive.

b) *How does an ACKRep-publication relate to “classical publications”?:* Publishing source code and metadata is supposed to be a *supplement to classical publications* such as scientific reports, conference or journal papers. However, both publication types would benefit from linking each other: the paper/pdf publication becomes much easier to reproduce with linked (working) source code and the published source code becomes easier to comprehend with the background

information from a peer reviewed scientific paper. Indirectly, the proposed approach might also increase the content quality of final publications because the respective peer reviewers will be the first beneficiaries of low-barrier reproducibility and thus are able to assess the submission more precisely and, if necessary, point out weaknesses and possible improvements.

c) *Is the proposed central platform not a single point of failure or otherwise problematic?*: Experience shows that a limited (and often short) lifetime is attached to almost all human creations. This holds even more for digital creations. As the proposed concept does not create hard dependencies it can be considered as expandable as useful¹²: Every contributor is free (and encouraged to) publish the own contributions elsewhere (cf. public repository host in Fig. 4). With the features of distributed version control this redundancy is easily manageable and by default every contributor has at least one local and one public copy of the *whole data* under control. Furthermore, the constitutional rules (see Sec. IV-B) of the platform should prevent that centralization implies concentration of influence and the availability of all the data and all the code under a Free Software License gives anybody the low-effort possibility to clone/fork the whole project – which in turn acts as a feedback in favor of finding compromises.

The most sensitive part would be the URLs hard-coded in publications, which in the worst case might become unreachable. But on the other hand, this could happen with any other web service as well, and those are typically not as simple to replace, do not grant open access to all their content and do not offer easy identification of contributions with unique entity keys.

d) *Does the proposed solution have scalability issues?*: In this paper we suggest to establish *one* repository and a central web service (currently running on a single virtual machine). This results in two obvious limitations when many contributions are made: a) The repository might become quite big, which complicates further contributions and other interactions and b) the computational resources of the current web server will not suffice if confronted with several requests in parallel. However, many other platforms and projects show that such problems are solvable. This contribution only aims to propose the basic concept to raise feedback. If scalability becomes an problem, then there necessarily is a group of several contributors who can then also provide the resources to solve that problem.

e) *How are “older results” supposed to be included?*: The incentive structures discussed in question a) mainly apply to recent results. However, it would also be desirable that established control methods such as, e. g., Kalman filtering, loop shaping or PID design are available as tested reference implementations. This could be achieved by including ACK-REP in educational projects: students could learn from existing implementations (as supplement to reading the respective theory) and then create own contributions or improve existing

ones as merge requests. The potential global visibility of the own work as well as constructive feedback during the merge request is expected to be a significant motivation. In fact, most of the entities currently present in `ackrep_data` were created during or are at least related to student projects.

f) *How can the correctness of the content be ensured in the long term?*: The proposed solution offers two layers of validation: First, an automatic formal verification, i. e. a test whether submitted code runs without errors and produces the asserted results. This is required before inclusion into the repository. Additionally, content-related flaws such as implementation bugs or conceptual errors can only be detected by domain experts – as is the case for classical publications. However, the `Comment` entity (cf. Fig. 2), which by definition is associated with specific other entities, offers an accountable and traceable communication channel that can be used to precisely identify and document issues. In contrast, for classical publications there is no predefined mechanism to inform the audience about errors or caveats.

VII. CONCLUSION

In this contribution, we proposed a concept to increase compliance with the FAIR principles for results in system and control theory based on computational tools. To this end, we first summarized the motivation for the principles and especially the reproducibility problem, which affects our research discipline as many others. After that we gave a brief overview on relevant existing technologies and how to combine them. The core contributions of this paper consist in the description of the repository structure, the overall architecture of the proposed tool, the description of the ontology (OCSE) and the process scheme for the basic use cases (along with the constitutional rule set). To facilitate the further experimental validation of the proposed concept, the prototype implementation was briefly outlined. Finally, we gave a comparison of our suggested approach to the original FAIR principles and some discussion of expected concerns.

The authors want to close with the normative claim that reproducibility should be a key feature of scientific results and thus the FAIR principles deserve more attention. Published results should be as easy as possible to find, to access and to reproduce by interested individuals or groups, e. g. peer researchers, graduate students or business and industry. The proposed approach of the active knowledge repository and much more its current implementation status is not expected to be the optimal solution or even close to that. But nevertheless it might be a relevant step in an ongoing optimization process.

REFERENCES

- [1] J. F. Claerbout and M. Karrenbach. “Electronic Documents Give Reproducible Research a New Meaning”. In: *SEG Technical Program Expanded Abstracts 1992*. Society of Exploration Geophysicists, 1992, pp. 601–604.

¹²Even in the worst case scenario of a sudden irreversible shutdown after a period of successful operation this would only destroy (some of) the benefits it added before but not cause additional problems.

- [2] D. L. Donoho, A. Maleki, I. U. Rahman, M. Shahram, and V. Stodden. "Reproducible Research in Computational Harmonic Analysis". In: *Computing in Science Engineering* Vol. 11 (2009), pp. 8–18.
- [3] R. R. Downs, W. C. Lenhardt, E. Robinson, E. Davis, and N. Weber. "Community Recommendations for Sustainable Scientific Software". In: *Journal of Open Research Software* Vol. 3 (2015), e11.
- [4] "Devil in the Details (Editorial)". In: *Nature* Vol. 470 (2011), pp. 305–306.
- [5] Z. Merali. "Computational Science: ...Error". In: *Nature* Vol. 467 (2010), pp. 775–777.
- [6] A. Morin et al. "Shining Light into Black Boxes". In: *Science* Vol. 336 (2012), pp. 159–160.
- [7] N. P. Rougier et al. "Sustainable Computational Science: The ReScience Initiative". In: *PeerJ Computer Science* Vol. 3 (2017), e142.
- [8] J. C. Molloy. "The Open Knowledge Foundation: Open Data Means Better Science". In: *PLOS Biology* Vol. 9 (2011), e1001195.
- [9] O. J. Reichman, M. B. Jones, and M. P. Schildhauer. "Challenges and Opportunities of Open Data in Ecology". In: *Science* Vol. 331 (2011), pp. 703–705.
- [10] V. Gewin. "Data Sharing: An Open Mind on Open Data". In: *Nature* Vol. 529 (2016), pp. 117–119.
- [11] G. Boulton, M. Rawlins, P. Vallance, and M. Walport. "Science as a Public Enterprise: The Case for Open Data". In: *The Lancet* Vol. 377 (2011), pp. 1633–1635.
- [12] B. Hanson, A. Sugden, and B. Alberts. "Making Data Maximally Available". In: *Science* Vol. 331 (2011), pp. 649–649.
- [13] V. Stodden. "Trust Your Science?: Open You Data and Code". In: *AMSTAT news: the membership magazine of the American Statistical Association* (2011), pp. 21–22.
- [14] D. C. Ince, L. Hatton, and J. Graham-Cumming. "The Case for Open Computer Programs". In: *Nature* Vol. 482 (2012), pp. 485–488.
- [15] D. Bailey, J. Borwein, and V. Stodden. "Set the Default to "Open"". In: *Notices of the Ams, Accepted March* (2013), p. 2013.
- [16] S. M. Easterbrook. "Open Code for Open Science?" In: *Nature Geoscience* Vol. 7 (11 2014), pp. 779–781.
- [17] L. A. Barba. "Praxis of reproducible computational science". In: *Computing in Science & Engineering* Vol. 21.1 (2019), pp. 73–78.
- [18] F. Chirigati, R. Rampin, D. Shasha, and J. Freire. "ReproZip: Computational Reproducibility With Ease". In: *Proceedings of the 2016 International Conference on Management of Data. SIGMOD '16. San Francisco, California, USA: Association for Computing Machinery, 2016, pp. 2085–2088.*
- [19] T. Crick, B. Hall, and S. Ishtiaq. "Reproducibility in Research: Systems, Infrastructure, Culture". In: *Journal of Open Research Software* Vol. 5 (1 2017), p. 32.
- [20] M. Krafczyk, A. Shi, A. Bhaskar, D. Marinov, and V. Stodden. "Scientific Tests and Continuous Integration Strategies to Enhance Reproducibility in the Scientific Software Context". In: *Proceedings of the 2nd International Workshop on Practical Reproducible Evaluation of Computer Systems. P-RECS '19. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 23–28.*
- [21] M. D. Wilkinson et al. "The FAIR Guiding Principles for scientific data management and stewardship". In: *Scientific data* Vol. 3.1 (2016), pp. 1–9.
- [22] *ACKRep Testing Instance*. URL: <http://testing.ackrep.org/>.
- [23] C. Knoll and R. Heedt. "Automatic Control Knowledge Repository — A Computational Approach for Simpler and More Robust Reproducibility of Results in Control Theory". In: *Proc. of the 24th International Conference on System Theory, Control and Computing. IEEE. Sinaia, 2020.*
- [24] R. C. Merkle. "Method of Providing Digital Signatures". U.S. pat. Leland Stanford Junior University, 1982.
- [25] G. Fraser and J. M. Rojas. "Software Testing". In: *Handbook of Software Engineering*. Ed. by S. Cha, R. N. Taylor, and K. Kang. Cham: Springer International Publishing, 2019.
- [26] N. Guarino, D. Oberle, and S. Staab. "What is an ontology?" In: *Handbook on ontologies*. Springer Berlin, 2009, pp. 1–17.
- [27] C. Dessimoz and N. Škunca, eds. *The Gene Ontology Handbook*. Springer New York, 2017.
- [28] M. K. Bergman. *Knowledge Representation Practionary*. Springer, 2018.
- [29] M. Keet. *An introduction to ontology engineering, v1.5*. College Publications, 2020.
- [30] F. Baader, D. Calvanese, D. McGuinness, P. Patel-Schneider, D. Nardi, et al. *The description logic handbook: Theory, implementation and applications*. Cambridge university press, 2003.
- [31] M. Krötzsch, M. Marx, A. Ozaki, and V. Thost. "Attributed description logics: Ontologies for knowledge graphs". In: *International Semantic Web Conference*. Springer. 2017, pp. 418–435.
- [32] D. Vrandečić and M. Krötzsch. "Wikidata: a free collaborative knowledgebase". In: *Communications of the ACM* Vol. 57.10 (2014), pp. 78–85.
- [33] T. Berners-Lee, J. Hendler, and O. Lassila. "The semantic web". In: *Scientific american* Vol. 284.5 (2001), pp. 34–43.
- [34] A. Patel and S. Jain. "Present and future of semantic web technologies: a research statement". In: *International Journal of Computers and Applications* (2019), pp. 1–10.
- [35] C. Benavides, I. García, H. Alaiz, and L. Quesada. "An ontology-based approach to knowledge representation for Computer-Aided Control System Design". In: *Data & Knowledge Engineering* Vol. 118 (2018), pp. 107–125.

- [36] R. Heedt, C. Knoll, and K. Röbenack. “Formal Semantic Representation of Methods in Automatic Control”. In: *Tagungsband VDI Mechatroniktagung*. (in German). 2021.
- [37] C. Knoll and R. Heedt. *ACKRep Data Repository on GitHub*. URL: https://github.com/cknoll/ackrep_data.
- [38] C. Lange. “Ontologies and languages for representing mathematical knowledge on the semantic web”. In: *Semantic Web* Vol. 4.2 (2013), pp. 119–158.
- [39] M. Kohlhase, T. Koprucki, D. Müller, and K. Tabelow. “Mathematical Models as Research Data via Flexiformal Theory Graphs”. In: *Intelligent Computer Mathematics*. Ed. by H. Geuvers, M. England, O. Hasan, F. Rabe, and O. Teschke. Cham: Springer International Publishing, 2017, pp. 224–238.
- [40] C. Knoll. *yamlyowl – a Python based YAML-to-OWL-converter*. URL: <https://github.com/cknoll/yamlyowl>.
- [41] J.-B. Lamy. “Owlready: Ontology-oriented programming in Python with automatic classification and high level constructs for biomedical ontologies”. In: *Artificial intelligence in medicine* Vol. 80 (2017), pp. 11–28.
- [42] C. Knoll, X. Jia, and R. Heedt. “Trajectory Planning for Closed Kinematic Chains Applied to Cooperative Motions in Health Care”. In: *PAMM* Vol. 20.1 (2021).