

# On Reproducible AI: Towards Reproducible Research, Open Science, and Digital Scholarship in AI Publications

*Odd Erik Gundersen, Yolanda Gil, David W. Aha*

■ *Artificial intelligence, like any science, must rely on reproducible experiments to validate results. Our objective is to give practical and pragmatic recommendations for how to document AI research so that results are reproducible. Our analysis of the literature shows that AI publications currently fall short of providing enough documentation to facilitate reproducibility. Our suggested best practices are based on a framework for reproducibility and recommendations for best practices given by scientific organizations, scholars, and publishers. We have made a reproducibility checklist based on our investigation and described how every item in the checklist can be documented by authors and examined by reviewers. We encourage authors and reviewers to use the suggested best practices and author checklist when considering submissions for AAAI publications and conferences.*

Reproducibility is a cornerstone of the scientific method. The ability and effort required from other researchers to replicate experiments and explore variations depends heavily on the information provided when the original work was published. Reproducibility is challenging for many sciences, for example when the variability of physical samples and reagents can significantly affect the outcome (Begley and Ellis 2012; Lithgow, Driscoll, and Phillips 2017). In computer science, a large portion of the experiments are fully conducted on computers, making the experiments more straightforward to document (Braun and Ong 2014). Most AI and machine learning research also falls under this category of computational experimentation. However, reproducibility in AI is not easily accomplished (Hunold and Träff 2013; Fokkens et al. 2013; Hunold 2015). This may be because AI research has its own unique reproducibility challenges. Ioannidis (2005) suggests that the use of analytical methods which are still a focus of active investigation is one reason it is comparatively difficult to ensure that computational research is reproducible. For

| Factor     | Variable               | Description  |
|------------|------------------------|--|
| Method     | Problem                | Is there an explicit mention of the problem the research seeks to solve?           |
|            | Objective              | Is the research objective explicitly mentioned?                                    |
|            | Research method        | Is there an explicit mention of the research method used (empirical, theoretical)? |
|            | Research questions     | Is there an explicit mention of the research question(s) addressed?                |
|            | Pseudocode             | Is the AI method described using pseudocode?                                       |
| Data       | Training data          | Is the training set shared?  |
|            | Validation data        | Is the validation set shared?  |
|            | Test data              | Is the test set shared?  |
|            | Results                | Are the relevant intermediate and final results output by the AI program shared?   |
| Experiment | Hypothesis             | Is there an explicit mention of the hypotheses being investigated?                 |
|            | Prediction             | Is there an explicit mention of predictions related to the hypotheses?             |
|            | Method source code     | Is the AI system code available open source?                                       |
|            | Hardware               | Is the hardware used for conducting the experiment specified?                      |
|            | Software dependencies  | Are software dependencies specified?   |
|            | Experiment setup       | Are the variable settings shared, such as hyperparameters?                         |
|            | Experiment source code | Is the experiment code available open source?                                      |

Table 1. Description of All Variables and Their Factors.

example, Henderson et al. (2017) show that problems due to nondeterminism in standard benchmark environments and variance intrinsic to AI methods require proper experimental techniques and reporting procedures. Acknowledging these difficulties, computational research should be documented properly so that the experiments and results are clearly described.

The AI research community should strive to facilitate reproducible research, following sound scientific methods and proper documentation in publications. Concomitant with reproducibility is open science, which involves sharing data, software, and other science resources in public repositories using permissive licenses. Open science is increasingly associated with FAIR principles to ensure that science resources have the necessary metadata to make them findable, accessible, interoperable, and reusable (Wilkinson et al. 2016). Modern digital scholarship promotes proper credit to scientists who document and share their research products through citations of datasets, soft-

ware, and innovative contributions to the scientific enterprise.

The focus in this article is on best practices for documentation and dissemination of AI research to facilitate reproducibility, support open science, and embrace digital scholarship. We begin with an analysis of recent AI publications that highlights the limitations of their documentation in support of reproducibility.

## State of the Art: How AI Research Is Currently Documented

Gundersen and Kjensmo (2018) analyzed how well empirical AI research is documented to facilitate reproducibility. Empirical AI research involves evaluating how well computational AI methods solve a problem. An AI method refers to an abstract method for solving such problems. Examples include agent systems that perform collaborative tasks and neural network architectures trained using backpropagation.

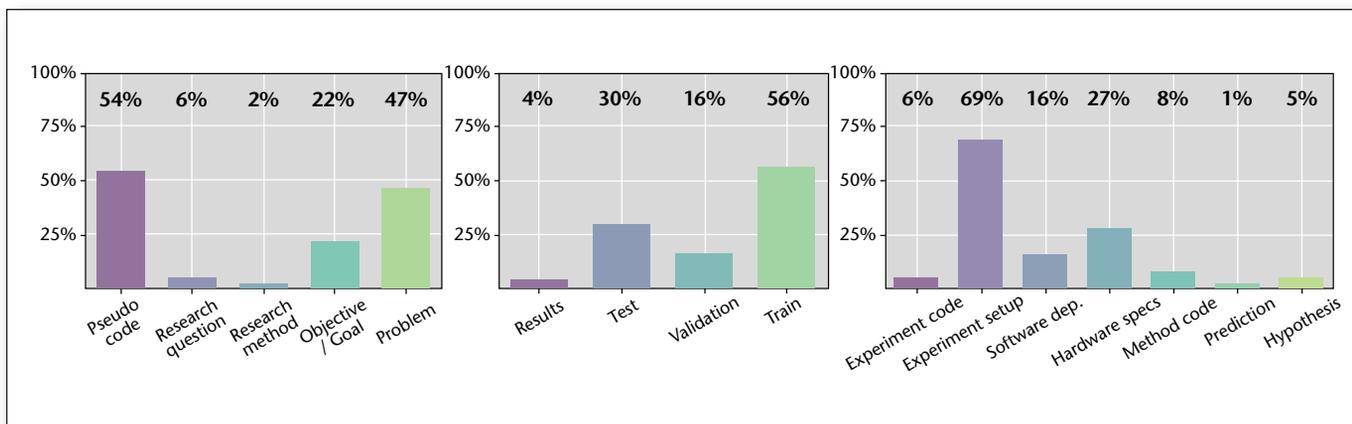


Figure 1. Percentage of Papers Documenting Each Variable for the Three Factors Method (left), data (middle), and experiment (right). Gundersen and Kjensmo (2018).

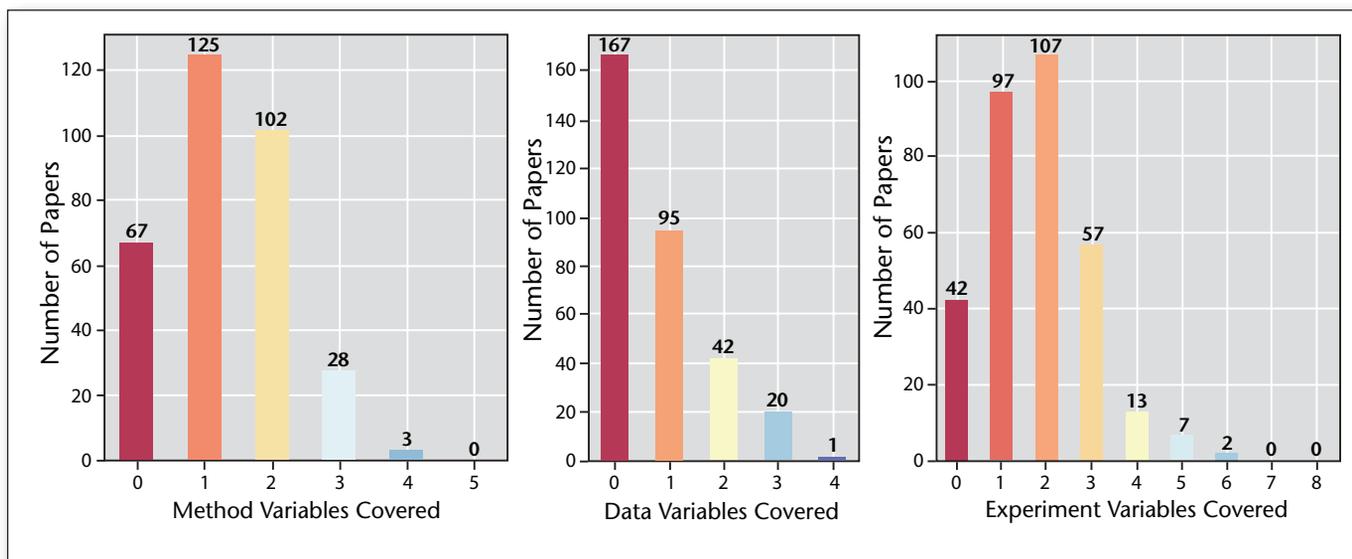


Figure 2. The Number of Variables for the Three Factors as Documented for All the Papers Describing Empirical Research in the Study. Method (left), data (middle), and experiment (right). Gundersen and Kjensmo (2018).

The analysis by Gundersen and Kjensmo (2018) is based on a literature review and a framework for reproducibility. Their framework divides documentation into three factors: (1) method, which specifies the AI method under investigation and the problem to be solved; (2) data, which describes the data used for conducting the empirical research; and (3) experiment, which documents how the experiment was conducted. How well these three factors are documented is indicated by 16 yes/no variables (see table 1) that are directly relevant for facilitating reproducibility.

A publication that documents an empirical

research study can be scored using these variables. Three reproducibility metrics are proposed. The three metrics are: (1) R1D, which calculates the average of all variables for all three factors (method, data, and experiment); (2) R2D, which computes the average of the variables for the method and data factors; and (3) R3D, which calculates the average of all variables for the method factor. These three metrics provide an indication of how well the scored papers document the research for three different degrees of reproducibility (we provide more detail on these degrees of reproducibility later on in the article).

In total, Gundersen and Kjensmo sampled 400

papers from the AAAI 2014, AAAI 2016, IJCAI 2013, and IJCAI 2016 conferences. Among these, 325 papers describe empirical studies, while the remaining 75 papers do not. Figure 1 displays the percentage of the surveyed papers that documented the different variables, while figure 2 summarizes how many of the variables were documented for each factor per paper.

We make a few observations. As seen in figure 1, few of the papers explicitly mention the research method that is used, and only around half explicitly mention which problem is being solved. Only about a third of the papers share the test dataset and only 4 percent share the result produced by the AI program. Only 8 percent of the papers share the source code of the AI method that is being investigated, while only 5 percent explicitly specify the hypothesis and 1 percent specify their prediction. Figure 2 shows that 67 papers do not explicitly document any of the variables for the factor method; only one paper documents and shares training, validation, and test sets as well as the results; and approximately 90 percent of the papers document no more than three of the seven variables of the factor experiment.

As seen in figure 3, the trends are unclear. Statistical analysis showed that only two of the metrics, R1D and R2D, for IJCAI had a statistically significant increase over time. While R2D and R3D for AAAI decrease over time, the decrease is not statistically significant.

The study by Gundersen and Kjensmo (2018) has some limitations. For example, the study required that for the variable problem to be set to yes (true), the paper must explicitly state the problem that is being solved. Another shortcoming is that all the AI methods that are documented in the research papers are not necessarily described better with pseudocode than without, but this fact was not given any consideration. If a paper described an AI method and pseudocode was not provided, the pseudocode variable was set to false for that paper. Finally, some of the variables might be redundant (for example, problem, goal, or research questions). Still, despite these shortcomings, the findings indicate that computational AI research is not documented systematically and with enough information to support reproducibility.

### Degrees of Reproducibility

Gundersen and Kjensmo (2018) distinguish between three degrees of reproducibility, which are defined as follows:

*R1: Experiment Reproducible.* The results of an experiment are experiment reproducible when the execution of the same implementation of an AI method produces the same results when executed on the same data. This is often called *repeatability*.

*R2: Data Reproducible.* The results of an experiment are data reproducible when an experiment is conducted that executes an alternative implementation of the AI method that produces the same results when execut-

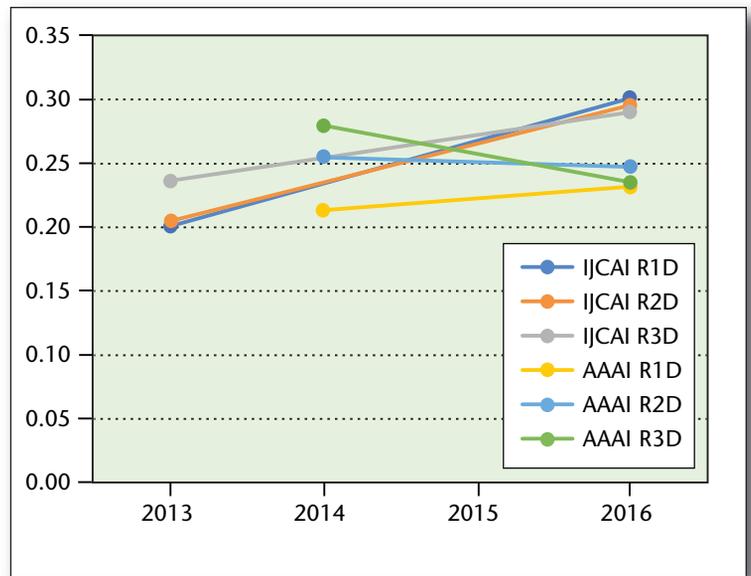


Figure 3. Change Over Time of the Three Reproducibility Metrics for Selected Years of the Two Conferences AAAI and IJCAI.

Gundersen and Kjensmo (2018).

|    | Method | Data | Experiment |
|----|--------|------|------------|
| R1 |        |      |            |
| R2 |        |      |            |
| R3 |        |      |            |

Figure 4. The Three degrees of Reproducibility Are Defined by Which Documentation Is Used to Reproduce the Results.

The three degrees of reproducibility each require a different set of factors to be documented.

ed on the same data. This is often called *replicability*.

*R3: Method Reproducible.* The results of an experiment are method reproducible when the execution of an alternative implementation of the AI method produces consistent results when executed on different data. This is often called *reproducibility*.

Empirical research that is R1 (experiment reproducible) must document the AI method, the data used to conduct the experiment, and the experiment itself including the source code for the AI method and the experiment setup, while R2 (data reproducible) research must only document the AI method and the data. R3 (method reproducible) research must only document the AI method. Figure 4 illustrates the different factors that must be documented for the three reproducibility degrees.

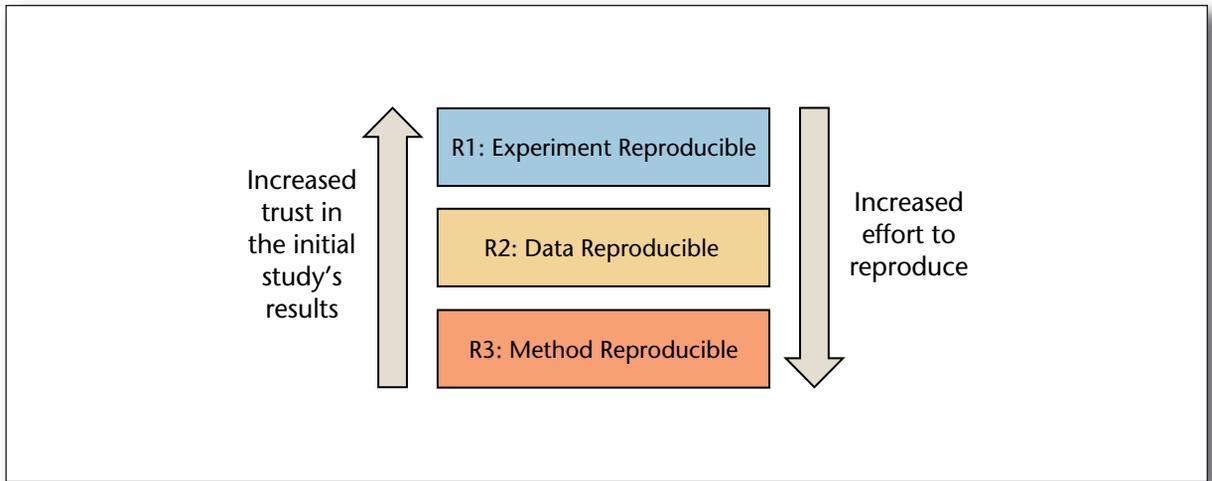


Figure 5. Effects of Documentation as Seen from the Perspective of Independent Researchers.

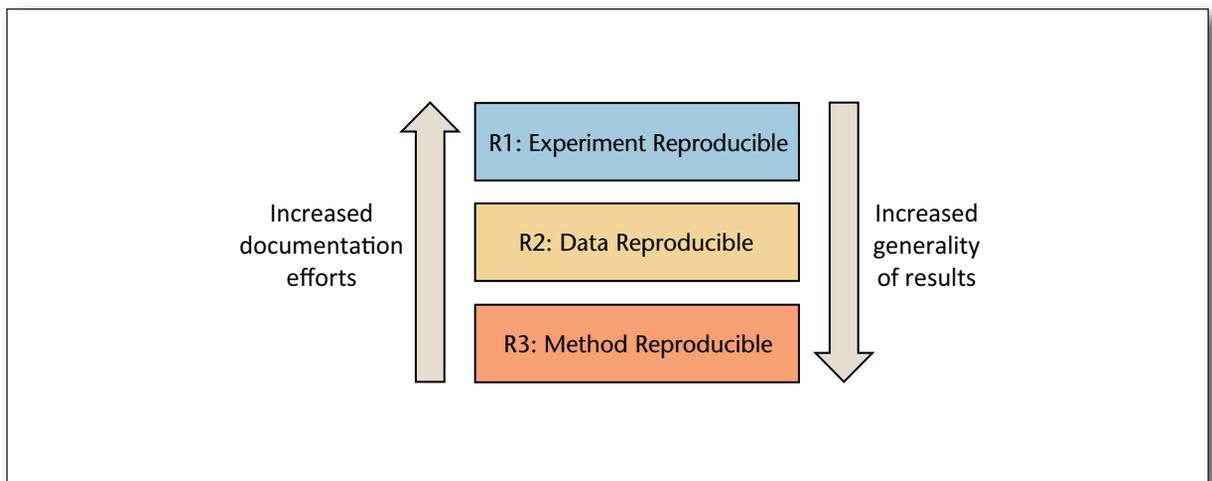


Figure 6. Effects of Documentation as Seen from the Perspective of the Original Researchers.

If an independent team reproduces research and gets results that are consistent with the original results, the generality of the results depends on the level of documentation that was provided to the independent team. If the original research was R1 (experiment reproducible), the independent team has confirmed that the specific implementation of the AI method provided by the original research team achieved the reported results on the specific data that also was provided by the original research team. Hence, the generality of the results is limited to that specific implementation and that specific data. However, if the independent team reproduces the results of some research that was R3 (method reproducible) and gets consistent results, the results are more general, as they apply to a reimplemention and other data. This factor leads to different incentives for the researchers who conducted the initial

empirical study and the independent researchers who attempt to reproduce the results.

Independent researchers trust an empirical study's results increasingly with the amount of documentation that is shared with them, while the effort to reproduce the results increases when the amount of documentation is reduced. This situation is illustrated in figure 5. Hence, independent researchers would prefer R1 (experiment reproducible) research.

On the other hand, the effort to document the research increases for the original researchers with the amount of documentation that needs to be shared, while the generality of the method is increased if independent researchers reproduce the results given less documentation. Hence, the original researchers may prefer to document their research to be R3 (method reproducible) (figure 6).

Combine this conflict of incentives for the origi-

| Recommendations | <i>Data mentioned in a publication should:</i>  |
|-----------------|---|
| 1.              | Be available in a shared community repository, so anyone can access it  |
| 2.              | Include basic metadata, so others can search and understand its contents  |
| 3.              | Have a license, so anyone can understand the conditions for reuse of the data   |
| 4.              | Have an associated digital object identifier (DOI) or persistent URL (PURL) so that the data is available permanently   |
| 5.              | Be cited properly in the prose and listed accurately among the references, so readers can identify the datasets unequivocally and data creators can receive credit for their work |

Table 2. Author Checklist Part I.

Recommendations for data in publications.

nal and independent researchers with the pressure to publish, and it is easy to see how this situation can lead to research being documented less vigorously. However, by following the recommendations given here, authors can increase the trustworthiness and reproducibility of research results with relatively little effort. Still, changes cannot be expected solely from individual researchers alone. The research community, funding sponsors, employers of researchers, and publishers should each, in their respective roles, incentivize and reward reproducible research.

## Best Practices and Recommendations

The recommendations we introduce are based on best practices put forward by scientific organizations such as the Research Data Alliance;<sup>1</sup> the Federation of Earth Science Information Partners;<sup>2</sup> DataCite;<sup>3</sup> the National Research Council (2012); the Task Group on Data Citation Standards and Practices (2013); the Data Citation Synthesis Group (2014); and scholars such as Ball and Duke,<sup>4</sup> Wilkinson et al. (2016), Stodden et al. (2016), Gil et al. (2016), Nosek et al. (2015), Starr et al. (2015), Downs et al. (2015), Mooney and Newton (2012), Goodman et al. (2014), Garijo et al. (2013), and Altman and King (2007), as well as earth and space science publishers<sup>5</sup> Hanson et al. (2015).

Strong momentum is building in support of FAIR practices, that is, to make data findable, accessible, interoperable, and reusable (Wilkinson et al. 2016). Our recommendations support FAIR principles and extend them to promote reproducible research, open science, and digital scholarship.

Implementing these recommendations requires extra space in publications. We suggest including this additional content in appendices that technical reviewers will not be required to assess but can quickly check. For electronic publications, there should not be any space limitations imposed for such appendices.

When these recommendations cannot be met, a brief explanation should be included about the reasons. Possible reasons may be restricted access (for example, proprietary or sensitive data), ownership by close collaborators who do not wish to disclose certain details, inadequate resources (for example, to house large datasets), or an unreasonable burden on authors.

We begin with recommendations for data and source code as the basic ingredients of a computational experiment. Then we describe recommendations to document AI methods and the experiments themselves. If all recommendations for AI methods (table 4) are implemented, then the publication should in theory be R3 (method reproducible), while if all recommendations for data (table 2) are also implemented, then the research should be R2 (data reproducible). Finally, all four sets of recommendations (tables 2–5) must be implemented for the research to be fully R1 (experiment reproducible).

We will refer to the complete set of 20 recommendations as an author checklist, we provide examples to demonstrate that they are synergistic, and we argue that they can be easily implemented.

## Recommendations for Data

Table 2 summarizes our recommendations for documenting data, which concern (1) repository use, (2) metadata, (3) licenses, (4) persistent unique identifiers, and (5) citations. These recommendations can be easily implemented if researchers use community data repositories that support recommended best practices.

### Data Repositories

Data repositories exist for many domains, and as such they are available to the AI community. Examples of these general repositories are Zenodo,<sup>6</sup> figshare,<sup>7</sup> and Dataverse.<sup>8</sup> These repositories will automatically assign a DOI to any uploaded data and will also accept software, figures, movies, and slide

| Recommendations | Source code used for implementing an AI method and executing an experiment should:  |
|-----------------|---|
| 6.              | Be available in a shared community repository, so anyone can access it  |
| 7.              | Include basic metadata, so others can search and understand its contents  |
| 8.              | Include a license, so anyone can understand the conditions for use and extension of the software  |
| 9.              | Have an associated digital object identifier (DOI) or persistent URL (PURL) for the version used in the associated publication so that the source code is permanently available |
| 10.             | Be cited and referenced properly in the publication so that readers can identify the version unequivocally and its creators can receive credit for their work                   |

Table 3. Author Checklist Part II.

Recommendations for source code implementing AI methods and experiments in publications.

presentations. They will also inquire about choosing a license and about specifying a descriptive name and authors for a submitted dataset. AAAI could, as a service, provide a list of recommended data repositories. This list could be modeled on a service provided by COPDESS, which is a large coalition for publishing data in the earth and space sciences.<sup>5</sup> Universities also offer general repositories, whether developed in-house or as installations of general infrastructure such as Dataverse. University repositories are typically maintained by library departments, and always offer DOIs, licenses, and citations.

We encourage maintainers of data repositories that serve the AI community to adopt mechanisms for assigning DOIs or persistent URLs (PURLs) to datasets that they provide. The management of PURLs or DOIs can be complex. We suggest consulting with organizations such as FORCE11 and the Research Data Alliance, which have working groups with extensive and detailed recommendations on this topic.

#### Metadata

Basic metadata includes a descriptive title, the dataset's authors, and creation date. Additional metadata is always valuable to others in terms of understanding and reusing the dataset.

#### Licenses for Data

Recommended licenses for data are Creative Commons licenses,<sup>9</sup> preferably CC-BY (unlimited reuse as long as there is attribution) or CC0 (unlimited reuse without conditions).

#### Permanent Unique Identifiers for Data

Many authors make data available by providing a URL to their personal or lab pages. These references may not last long due to changes in sites and in author affiliations (Klein et al. 2014). Instead, we encourage authors to use persistent unique identifiers so that their data is always available. DOIs are managed by data repositories and given to individual datasets or to collections (DeRisi et al. 2003).

Most data repositories provide DOIs, and for this they forge an agreement with a DOI authority. Another option that anyone can use is PURLs. PURLs can be assigned by anyone to any web resource using a trusted service such as the W3C's w3id.<sup>10</sup> Data repositories also have the option of using PURLs.

#### Data Citation

A data citation can be directly provided by a data repository, or it can be constructed by hand. A citation for a dataset consists of a descriptive name (or title) for the dataset, its creators, the name of the repository where it can be accessed, and the permanent URL. For example, a citation for a dataset in Gil et al. (2017) is:

Adusumilli, Ravali. (2016). Sample datasets used in (Gil et al. 2017) for AAAI 2017 (Data set). Zenodo. <http://doi.org/10.5281/zenodo.180716>.

Note that by simply uploading the dataset to the Zenodo repository, we obtained the DOI and the citation. Specifying the authors, the name, and the license takes negligible effort. The author checklist for data required little time to implement.

#### Recommendations for Source Code

We refer to source code as the human-readable computer instructions written in plain text and software as computer programs that are executable by a computer. Typically, source code is compiled to software for a computer to run it. Our recommendations for source code are summarized in table 3.

#### Source Code Repositories

Source code repositories can be used by any scientists to share code, and as such they are available to the AI community. These code repositories include general repositories such as GitHub and BitBucket, and language-specific repositories such as CRAN for R code or File Exchange in MATLAB Central. General data repositories such as those mentioned above accept source code as an entry, and as with any dataset they

| Recommendations | <i>AI methods used in a publication should be:</i>  |
|-----------------|---|
| 11.             | Presented in the context of a problem description that clearly identifies what problem they are intended to solve |
| 12.             | Outlined conceptually so that anyone can understand their foundational concepts                                   |
| 13              | Described in pseudocode so that others can understand the details of how they work                                |

Table 4. Author Checklist Part III.

Recommendations for AI methods in publications.

always offer DOIs, licenses, and citations.

For a specific publication, the version of the source code that is being used should be clearly specified, and the source code repository should support the identification and future access of specific versions.

#### Source Code Metadata

Basic metadata includes a descriptive title, the source code's authors, and the creation date. Additional metadata is always valuable to others in terms of understanding and reusing the source code.

#### Licenses for Source Code

Recommended licenses for source code are the standard licenses from the Open Source Initiative. Licenses such as Apache v2 or MIT are recommended because they provide unlimited reuse (as long as there is attribution). Other more restrictive licenses are available to limit commercial use or impose licensing conditions on extensions of the original source code.

#### Permanent Unique Identifiers for Source Code

A separate DOI should be assigned to meaningful versions of the source code, such as a version used for a publication. GitHub offers an option to obtain a DOI for a source code version, which is done by storing that version permanently in the Zenodo data repository. Any source code can be uploaded manually to community data repositories such as Zenodo, figshare, and Dataverse. PURLs can be assigned by anyone to any source code version that has a URL on the web, using a trusted service such as w3id.org.

#### Source Code Citation

Source code citation can be directly provided by a source code repository, or it can be constructed by hand. A citation for a source code version consists of a descriptive name (or title) for the source code, its creators, the name of the repository where it can be accessed, the version, and the permanent URL. For example, a citation for GitHub code in (Gil et al. 2017) is:

Ratnakar, Varun. "DISK software" (v1.0.0). Zenodo. 2016. <http://doi.org/10.5281/zenodo.168079>.

By uploading the source code into the GitHub code repository, we obtained a persistent identifier for the version used in the publication as well as the citation. Specifying the authors, the name, and the license takes negligible effort. Implementing the author checklist for source code required little time.

### Recommendations for AI Methods

Our recommendations for AI methods are listed in table 4.

#### Problem Description

The problem that a conceptual AI method solves should be explicitly described in the publication. In De Weerd et al. (2013) the following example can be found: "To address this problem, we propose a novel navigation system ..." The authors explicitly describe the problem that they address. Another good example of this practice can be found in He et al. (2016). Here the authors state the problem explicitly: "In this paper, we address the degradation problem by introducing a deep residual learning framework." The degradation problem is also properly described in their publication.

#### Conceptual Method

A high-level, textual description of the AI method should be provided to readers to allow them to gain an understanding of it. This description should include a broad overview of how the AI method works and specify input variables and the resulting output. In general, the AI research community excels at providing this information in publications.

#### Pseudocode

Pseudocode for the AI method should also be provided. In cases where detailed pseudocode cannot be provided due to the complexity of the proposed algorithm or system, a more abstract pseudocode summary can be provided that illustrates the AI method's flow.

| Recommendations | Descriptions of experiments in a publication should:  |
|-----------------|---|
| 14.             | Explicitly present the hypotheses to be assessed, before other details concerning the empirical study are presented   |
| 15.             | Present the predicted outcome of the experiment, based on beliefs about the AI method and its application   |
| 16.             | Include the experiment design (parameters and the conditions to be tested) and its motivation, such as why a specific number of tests or data points are used based on the desired statistical significance of results and the availability of data |
| 17.             | Identify and describe the measure and metrics   |
| 18.             | Provide the evaluation protocol   |
| 19.             | Share the results   |
| 20.             | Describe the results and the analysis   |
| 21.             | Be described as a workflow that summarizes how the experiment is executed and configured  |
| 22.             | Include documentation on workflow executions or execution traces that provide parameter settings and initial, intermediate, and final data  |
| 23.             | Specify the hardware used to run the experiments  |
| 24.             | Be cited and published separately when complex, so that others can unequivocally refer to the individual portions of the method that they reuse or extend   |

Table 5. Author Checklist Part IV.

Recommendations for experiments described in publications.

Both a high-level description and pseudocode help independent researchers to decide whether their own implementation of the method is correct. If these are not presented carefully, then the empirical study cannot always be easily reproduced.

### Recommendations for Experiments

Authors should, to the degree possible, detail how their experiments are designed, and indicate the rationale for their design. Our recommendations for documenting experiments are summarized in table 5.

#### Hypotheses and Predictions

Hypotheses and predictions should be stated explicitly before descriptions of the other components of an empirical study to ensure that the results analysis is meaningful (Baker 2016).

#### Experiment Design

A textual description and justification of the experiment's design should be provided, to include a description of each test condition. This description should also explain, for example, why a specific number of tests or data points are used, based on the desired statistical significance of the results and the availability of data.

#### Measure and Metrics

Identify/define the measures and metrics to be used for the results analysis.

#### Evaluation Protocol

A justification should be provided for the chosen protocol when documenting an experiment. To avoid hypothesis myopia, this experiment should not be designed to collect only evidence that is guaranteed to support the stated hypotheses. Instead, to encourage an insightful study, it should include conditions that could lead to the rejection of these hypotheses. Why are the datasets used appropriate for the experiment? Why is the chosen empirical design appropriate for assessing the hypothesis, and why are the metrics and measures appropriate for assessing the results?

#### Results

In order for an independent research team to be able to fully evaluate their reproduction, they would need to compare with the actual results. Hence, the results (the actual output) of the experiment should be shared.

#### Result Description and Analysis

The results should be presented, along with an in-depth analysis of the results based on the specified

measure and metrics. The documentation should provide an explicit indication of whether the analyses support the hypotheses.

#### Workflow

This workflow should describe, in a machine-readable way, how software and data are used to implement the evaluation protocol. A workflow step is an invocation of the software. Each step has input data and parameters as well as output data. Input data and the output of any step can be used as input to subsequent steps. The simplest workflow languages capture methods that are directed acyclic graphs, while other languages can represent iterations and conditionals. A publication that simply mentions what software was used usually leaves out critical information about how the software was configured or invoked.

Scripts or electronic notebooks can be an effective way to document workflows, although the organization of source code is more modular in a workflow structure.

#### Executions

A general workflow can be run many times with different datasets or parameter settings and generate different results. Execution traces of executed workflows provide a complete provenance trail of how each result was generated.

#### Hardware Specification

The hardware that is used should be specified if it is important for the experiment. This documentation may include specification of the processor type, the number of cores and processors, and RAM and hard disk requirements. Also, the provider of the cloud solution that is used, if any, should be specified. The machine architecture and operating system may need to be specified, so that any discrepancies in results can be properly diagnosed. Library dependencies should also be described. Virtualization technologies, such as Docker and Kubernetes, facilitate these specifications through artifacts called containers. Containers can be provided as appropriate to share the experiment hardware setup.

#### Workflow Citation

Citing a publication does not make explicit whether the citation is to its AI method, source code, data, empirical design, workflows, execution traces, results, or a general body of work or contributions. If it is important for others to be explicit about what aspects of the work are being reused, then separate citations should be given to each, as appropriate. Although workflow repositories are not as common as data and software repositories, many general data repositories accept any research product and can be used for this purpose.

For example, a citation for a bundle containing workflows and execution details for Gil et al. (2017) is:

Adusumilli, Ravali, Ratnakar, Varun, Garijo, Daniel,

Gil, Yolanda, and Mallick, Parag. (2016). Additional materials used in the paper “Towards Continuous Scientific Data Analysis and Hypothesis Evolution” on the Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17) (Data set). Zenodo. <http://doi.org/10.5281/zenodo.190374>.

By organizing the workflows and executions described into this publication and bundling them to upload to a general data repository, these authors obtained a persistent identifier as well as a citation. The author checklist for experiments was implemented quickly.

## Benefits to Authors

Recognizing that our recommendations will require effort from authors, we want to highlight the following 10 benefits: (1) Practice open science and reproducible research. This approach ensures the kinds of checks and balances that lead to better science. (2) Receive credit for all your research products (that is, through citations for software, datasets, and other products). (3) Increase the number of citations to your publications. Studies have shown that well-documented articles receive more citations (Piwowar et al. 2007). (4) Improve your chances of being funded (that is, by writing coherent and well-motivated empirical study and data management plans). (5) Extend your curriculum vitae. Include data and software sections with citations. Maintaining datasets and writing code are important contributions to the field of AI. (6) Improve the management of your research assets (for example, so your new students, and others, can more easily locate materials generated by your earlier students). (7) Allow for the reproduction of your work (for example, so you and others can leverage it in new studies, even if it was conducted many years ago). (8) Address new sponsor and journal requirements. They are steadfastly driving research towards increased reproducibility and open science. (9) Attract transformative students. They strive for a rigorous research methodology. (10) Demonstrate leadership. Step into the future.

By explicitly citing datasets and source code, and by providing workflows that are machine readable, we create the structure needed for the development of AI systems that can analyze and reason about our literature (Gil 2017). These AI systems would have access to a vast amount of structured scientific knowledge with comprehensive details about experimental design and results. This change could revolutionize how we approach the scientific research process.

## Discussion

It is reasonable to expect a limited release of data and source code until the creator has completed the research for which the data was collected, or for

which the source code was written, or until their draft is published. Many journals impose this, such as *Science* and *Nature*. See Joly et al. (2012) for a review of data retention policies.

The creation and documentation of additional information we recommend should be done by researchers who publish their studies. Documenting and sharing code and data in such a way that this information can be easily used and cited by others gives researchers credit for a larger portion of their research effort. For academic researchers, we advocate that tenure committees give weight to the publication of data and source code when evaluating candidates for tenure. Thus, the publication velocity should not be reduced, but include research products other than publications.

The recommendations we suggest should be a part of daily research practices. According to Irakli Loladze, despite increasing the work load by 30 percent, “Reproducibility is like brushing your teeth. It is good for you, but it takes time and effort. Once you learn it, it becomes a habit” (Baker 2016).

Another recommendation for improving the readability and comparability of research papers is to require structured abstracts, which are commonly used in medical journals. Structured abstracts can be used to efficiently communicate a research objective, the motivation for and process by which an empirical study was conducted, and what results were achieved. Structured abstracts also require researchers to structure their own thoughts about their research. We suggest a five-part structured abstract containing (1) the research motivation, (2) the research objective, (3) the method used to conduct any empirical studies, (4) the results of the research, and (5) the conclusion. This structure enforces a coherent research narrative, which is not always the case for unstructured abstracts. The abstract for this article is an example of the proposed structure, while Gundersen and Kjensmo (2018) provides an abstract for empirical research that follows these recommendations and includes an explicit description of the hypothesis and an interpretation of the results.

## Call to Arms

As a community, we should ensure that the research we conduct is properly documented. To make AI research reproducible and more trustworthy, we proposed best practices that should be adopted by editors and program chairs and incorporated into the review forms of AAAI publication venues.

Publishers should provide extra space to document and cite data, source code, and empirical study designs. AAAI leadership should encourage AI researchers to increase the reproducibility of their published work. This support could include providing structured templates to organize appendices and

making available extra space in publications to accommodate the needed documentation.

For AI research to become open and more reproducible, the research community and publishers have to establish suitable practices. Authors need to adopt these practices, disseminate them to colleagues and students, and help develop mechanisms and technology to make it easier for others to adopt them.

Our objective with this article is to highlight the benefits of reproducible science and to propose initial, modest changes that can increase the reproducibility of AI research results. There are many additional actions that could and should be taken, and we look forward to further dialogue with the AI research community on how to increase the reproducibility and scientific value of AI publications.

## Acknowledgements

This research was funded in part by the National Science Foundation under grant ICER-1440323. This work has in part been carried out at the Telenor-NTNU AI Lab, Norwegian University of Science and Technology, Trondheim, Norway. The recommendations proposed are based on the Geoscience Paper of the Future and the Scientific Paper of the Future best practices developed under that award. Thanks to Sigbjørn Kjensmo for all the effort put into surveying the state of the art of reproducibility of AI.

## Notes

1. [rd-alliance.org/outcomes](http://rd-alliance.org/outcomes).
2. [wiki.esipfed.org/index.php/Interagency\\_Data\\_Stewardship/Citations/provider\\_guidelines](http://wiki.esipfed.org/index.php/Interagency_Data_Stewardship/Citations/provider_guidelines).
3. [www.datacite.org](http://www.datacite.org).
4. [www.dcc.ac.uk/resources/how-guides/cite-datasets#sthash.MJQjNn3i.dpuf](http://www.dcc.ac.uk/resources/how-guides/cite-datasets#sthash.MJQjNn3i.dpuf).
5. [www.copdess.org/statement-of-commitment](http://www.copdess.org/statement-of-commitment).
6. [zenodo.org](http://zenodo.org).
7. [figshare.com](http://figshare.com).
8. [dataverse.org](http://dataverse.org).
9. [creativecommons.org](http://creativecommons.org).
10. [www.w3id.org](http://www.w3id.org).

## References

- Altman, M., and King, G. 2007. A Proposed Standard for the Scholarly Citation of Quantitative Data. *D-Lib Magazine* 13 (3/4).
- Baker, M. 2016. Is There a Reproducibility Crisis? *Nature* 533(7604): 452–54.
- Begley, C. G., and Ellis, L. M. 2012. Drug Development: Raise Standards for Preclinical Cancer Research. *Nature* 483(7391): 531–33. doi.org/10.1038/483531a.
- Bouquet, P.; Serafini, L.; Zanobini, S.; and Benerecetti, M. 2003. An Algorithm for Semantic Coordination. Paper presented at the Second International Semantic Integration Workshop. Sanibel Island, FL, October 20–23.
- Braun, M. L., and Ong, C. S. 2014. Open Science in

- Machine Learning. In *Implementing Reproducible Research*, edited by V. Stodden, F. Leish, and R. D. Peng, 343. Boca Raton, FL: CRC Press.
- Data Citation Synthesis Group. 2014. Joint Declaration of Data Citation Principles, edited by M. Martone. San Diego, CA: FORCE11. doi.org/10.25490/a97f-egyk.
- DeRisi, S.; Kennison, R.; and Twyman, N. 2003. The What and Whys of DOIs. *PLoS Biology* 1(2): 133–34, e57. doi.org/10.1371/journal.pbio.0000057.
- De Weerd, M. M.; Gerding, E. H.; Stein, S.; Robu, V.; and Jennings, N. R. 2013. Intention-Aware Routing to Minimise Delays at Electric Vehicle Charging Stations. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, 83–89. Palo Alto, CA: AAAI Press.
- Downs, R. R.; Duerr, R.; Hills, D. J.; and Ramapriyan, H. K. 2015. Data Stewardship in the Earth Sciences. *D-Lib Magazine* 21 (7/8). doi.org/10.1045/july2015-downs.
- Fokkens, A.; Erp M. V.; Postma, M.; Pedersen, M.; Vossen, P.; and Freire, N. 2013. Offspring from Reproduction Problems: What Replication Failure Teaches Us. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, 1691–701. Stroudsburg, PA: Association for Computational Linguistics.
- Garijo, D.; Kinnings, S.; Xie, L.; Xie, L.; Zhang, Y.; Bourne, P. E.; and Gil, Y. 2013. Quantifying Reproducibility in Computational Biology: The Case of the Tuberculosis Drugome. *PLoS ONE* 8(11): e80278. doi.org/10.1371/journal.pone.0080278.
- Gil, Y. 2017. Thoughtful Artificial Intelligence: Forging A New Partnership for Data Science and Scientific Discovery. *Data Science* 1(1–2): 119–29. doi.org/10.3233/DS-170011.
- Gil, Y.; David, C. H.; Demir, I.; Essawy, B. T.; Fulweiler, R. W.; Goodall, J. L.; Karlstrom, L.; Lee, H.; Mills, H. J.; Oh, J.; Pierce, S. A.; Pope, A.; Tzeng, M. W.; Villamizar, S. R.; and Yu, X. 2016. Towards the Geoscience Paper of the Future: Best Practices for Documenting and Sharing Research from Data to Software to Provenance. *Earth and Space Science* 3(10): 388–415. doi.org/10.1002/2015EA000136.
- Gil, Y.; Garijo, D.; Ratnakar, V.; Mayani, M. R.; Adusumilli, R.; Boyce, H.; Srivastava, A.; and Mallick, P. 2017. Towards Continuous Scientific Data Analysis and Hypothesis Evolution. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 4406–14. Palo Alto, CA: AAAI Press.
- Goodman, A.; Pepe, A.; Blocker, A. W.; Borgman, C. L.; Cranmer, K.; Crosas, M.; Stefano, R. D.; Gil, Y.; Groth, P.; Hedstrom, M.; Hogg, D. W.; Kashyap, V.; Mahabal, A.; Siemiginowska, A.; and Slavkovic, A. 2014. Ten Simple Rules for the Care and Feeding of Scientific Data. *PLoS Computational Biology* 10(4): e1003542. doi.org/10.1371/journal.pcbi.1003542.
- Gundersen, O. E., and Kjensmo, S. 2018. State of the Art: Reproducibility in Artificial Intelligence. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 1644–51. Palo Alto, CA: AAAI Press.
- Hanson, B.; Lehnert, K.; and Cutcher-Gershenfeld, J. 2015. Committing to Publishing Data in the Earth and Space Sciences. *Eos: Earth and Space Science News* 96. doi.org/doi:10.1029/2015EO022207.
- He, K.; Zhang, X., Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*. Piscataway, NJ: Institute of Electrical and Electronics Engineers.
- Henderson, P.; Islam, R.; Bachman, P.; Pineau, J.; Precup, D.; and Meger, D. 2018. Deep Reinforcement Learning that Matters. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, 3207–14. Palo Alto, CA: AAAI Press.
- Hunold, S. 2015. A Survey on Reproducibility in Parallel Computing. arXiv preprint arXiv:1511.04217 [cs.DC]. Ithaca, NY: Cornell University Library.
- Hunold, S., and Träff, J. S. 2013. On the State and Importance of Reproducible Experimental Research in Parallel Computing. arXiv preprint arXiv:1308.3648 [cs.DC]. Ithaca, NY: Cornell University Library.
- Ioannidis, J. P. 2005. Why Most Published Research Findings Are False. *PLoS Medicine* 2(8): e124. doi.org/10.1371/journal.pmed.0020124.
- Joly, Y.; Dove, E. S.; Kennedy, K. L.; Bobrow, M.; Ouellette, B. F. F.; Dyke, S. O. M.; Kato, K.; and Knoppers, B. M. 2012. Open Science and Community Norms: Data Retention and Publication Moratoria Policies in Genomics Projects. *Medical Law International* 12(2): 92–120. doi.org/10.1177/20968533212458431.
- Klein, M.; de Sompel, H. V.; Sanderson, R.; Shankar, H.; Balakireva, L.; Zhou, K.; and Tobin, R. 2014. Scholarly Context Not Found: One in Five Articles Suffers from Reference Rot. *PLoS One* 9(12): e115253. doi.org/10.1371/journal.pone.0115253.
- Lithgow, G. J.; Driscoll, M.; and Phillips, P. 2017. A Long Journey to Reproducible Results. *Nature* 548(7668): 387–88.
- Mooney, H., and Newton, M. P. 2012. The Anatomy of a Data Citation: Discovery, Reuse, and Credit. *Journal of Librarianship and Scholarly Communication* 1(1): eP1035. doi.org/10.7710/2162-3309.1035.
- National Research Council. 2012. *For Attribution: Developing Data Attribution and Citation Practices and Standards: Summary of an International Workshop*. Washington, DC: The National Academies Press. doi.org/10.17226/13564.
- Nosek, B. A.; Alter, G.; Banks, G. C.; Borsboom, D.; Bowman, S. D.; Breckler, S. J.; Buck, S.; Chambers, C. D.; Chin, G.; Christensen, G.; Contestabile, M.; Dafoe, A.; Eich, E.; Freese, J.; Glennerster, R.; Goroff, D.; Green, D. P.; Hesse, B.; Humphreys, M.; Ishiyama, J.; Karlan, D.; Kraut, A.; Lupia, A.; Mabry, P.; Madon, T.; Malhotra, N.; Mayo-Wilson, E.; McNutt, M.; Miguel, E.; Levy Paluck E.; Simonsohn, U.; Soderberg, C.; Spellman, B. A.; Turitto, J.; VandenBos, G.; Vazire, S.; Wagenmakers, E. J.; Wilson, R.; and Yarkoni, T. 2015. Promoting an Open Research Culture. *Science* 348(6242): 1422–25. doi.org/10.1126/science.aab2374.
- Piwowar, H. A.; Day, R. S.; and Fridsma, D. B. 2007. Sharing Detailed Research Data Is Associated with Increased Citation Rate. *PLoS ONE* 2(3): e308. doi.org/10.1371/journal.pone.0000308.
- Starr, J.; Castro, E.; Crosas, M.; Dumontier, M.; Downs, R. R.; Duerr, R.; Haak, L. L.; Haendel, M.; Herman, I.; Hodson, S.; Hourclé, J.; Kratz, J. E.; Lin, J.; Nielsen, L. H.; Nurnberger, A.; Proell, S.; Rauber, A.; Sacchi, S.; Smith, A.; Taylor, M.; and Clark, T. 2015. Achieving Human and Machine Accessibility of Cited Data in Scholarly Publications. *PeerJ Computer Science* 1(1): e1. doi.org/10.7717/peerj-cs.1.
- Stodden, V.; McNutt, M.; Bailey, D. H.; Deelman, E.; Gil, Y.; Hanson, B.; Heroux, M. A.; Ioannidis, J. P. A.; and Tauber, M. 2016. Enhancing Reproducibility for Computational Methods. *Science* 354(6317): 1240–41. doi.org/10.1126/science.aah6168.
- Task Group on Data Citation Standards and Practices. 2013.



## AAAI Gifts Program

It is the generosity and loyalty of our members that enable us to continue to provide the best possible service to the AI community and promote and further the science of artificial intelligence by sustaining the many and varied programs that AAAI provides. AAAI invites all members and other interested parties to consider a gift to help support the dozens of programs that AAAI currently sponsors. For more information about the Gift Program, please see write to us at [donate18@aaai.org](mailto:donate18@aaai.org).

## Support AAAI Open Access

AAAI also thanks you for your ongoing support of the open access initiative. We count on you to help us deliver the latest information about artificial intelligence to the scientific community. To enable us to continue this effort, we invite you to consider an additional gift to AAAI. For information on how you can contribute to the open access initiative, please see [www.aaai.org](http://www.aaai.org) and click on "Gifts."

*AAAI is a 501c3 charitable organization.  
Your contribution  
may be tax deductible.*

Out of Cite, Out of Mind: Current State of Practice, Policy, and Technology for the Citation of Data. *Data Science Journal* 12:CIDCR1-7 doi.org/10.2481/dsj.OSOM13-043.

Wilkinson, M. D.; Dumontier, M.; Aalbersberg, I. J.; Appleton, G.; Axton, M.; Baak, A.; Blomberg, N.; Boiten, J.-W.; da Silva Santos, L. B.; Bourne, P. E.; Bouwman, J.; Brookes, A. J.; Clark, T.; Crosas, M.; Dillo, I.; Dumon, O.; Edmunds, S.; Evelo, C. T.; Finkers, R.; Gonzalez-Beltran, A.; Gray, A. J. G.; Groth, P.; Goble, C.; Grethe, J. S.; Heringa, J.; Hoen, P. A. C.; Hooft, R.; Kuhn, T.; Kok, R.; Kok, J.; Lusher, S. J.; Martone, M. E.; Mons, A.; Packer, A. L.; Persson, B.; Rocca-Serra, P.; Roos, M.; van Schaik, R.; Sansone, S.-A.; Schultes, E.; Sengstag, T.; Slater, T.; Strawn, G.; Swertz, M. A.; Thompson, M.; van der Lei, J.; van Mulligen, E.; Velterop, J.; Waagmeester, A.; Wittenburg, P.; Wolstencroft, K.; Zhao, J.; and Mons, B. 2016. The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Scientific Data* 3: article 160018. doi.org/10.1038/sdata.2016.18.

**Odd Erik Gundersen** (PhD, Norwegian University of Science and Technology) is an adjunct associate professor at the Norwegian University of Science and Technology (NTNU) in Trondheim, Norway, where he teaches courses and supervises master students in AI. Gundersen has applied AI in the industry, mostly for startups, since 2006. Currently, he investigates how AI can be applied in the renewable energy sector and for driver training.

**Yolanda Gil** (PhD, Carnegie Mellon University) is director of knowledge technologies and an associate division director at the Information Sciences Institute of the University of Southern California, and a research professor in computer science and in spatial sciences. Her research is on intelligent interfaces for knowledge capture and discovery, semantic workflows and metadata capture, social knowledge collection, computer-mediated collaboration, and automated discovery. Gil has served on the Advisory Committee of the Computer Science and Engineering Directorate of the National Science Foundation. She initiated and chaired the W3C Provenance Group that led to a community standard in this area. Gil is also a Fellow of AAAI and was elected as its 24th president in 2016.

**David W. Aha** (PhD, University of California, Irvine) leads a section within NRL's Navy Center for Applied Research in AI, in Washington, DC. In addition to encouraging reproducible research, his interests include mixed-initiative intelligent agents, deliberative autonomy, explainable AI, case-based reasoning, and machine learning. He has co-organized many events on these topics, launched the UCI Repository for ML Databases, served as an AAAI councilor, cocreated AAAI's AI Video Competition, and currently leads the evaluation team for DARPA's Explainable AI program.