# The open data challenge: An analysis of 124,000 data availability statements, and an ironic lesson about data management plans

Chris Graf[1], Dave Flanagan[1], Lisa Wylie[1], and Deirdre Silver[1]

[1]Wiley

April 28, 2020

## Abstract

Data availability statements can provide useful information about how researchers actually share research data. We used unsupervised machine learning to analyse 124,000 data availability statements submitted by research authors to 176 Wiley journals between 2013 and 2019. We categorised the data availability statements, and looked at trends over time. We found expected increases in the number of data availability statements submitted over time, and marked increases that correlate with policy changes made by journals. Our open data challenge becomes to use what we have learned to present researchers with relevant and easy options that help them to share and make an impact with new research data.

## Introduction

Researchers, particularly those with grant funding, are increasingly required to share the data they create (Hahnel, 2015) (Popkin, 2019). More than a decade ago the US National Institutes of Health said "data sharing is essential for expedited translation of research results into knowledge, products, and procedures" and began requiring data sharing plans in all grants greater than USD500,000 (*NIH Guide: FINAL NIH STATEMENT ON SHARING RESEARCH DATA*, n.d.). Since 2018, the General Office of the State Council in China has required that "all scientific data derived from the science and technology plans . . . shall be archived in the relevant science data centers" (*China Open Science and Open Data Mandate Released - Enago Academy*, n.d.). State Council governs the funding made available by the Chinese Academy of Sciences and also the Ministry of Science and Technology, under which sits National Natural Science Foundation of China. The EUR100billion Horizon Europe funding programme from the European Commission will challenge funded researchers to deliver "open access to publications, data, and to research data management plans" starting in 2021 (EuropeanUnion, 2018). Private not-for-profit funders, like the Wellcome Trust and the Bill & Melinda Gates Foundation, are among those with the most progressive requirements. Funders argue that sharing data creates more value and impact from every grant they award and enhances trustworthiness and potential reproducibility. This is the open data challenge: Sharing data first, and then realising the value from doing so.

While the arguments for data sharing are compelling and the rhetoric often exciting ("researchers are creating, gathering and using data in hitherto unimagined-volumes" (*Realising the potential – Final report of the Open Research Data Task Force*, n.d.)) reservations have been expressed that reflect researchers' concerns, including the absence of infrastructure and incentives, and the presence of disincentives like the fear of getting scooped and concerns about misinterpretation or misuse of shared data (Bezuidenhout & Chakauya, 2018) (*To Share or not to Share? That is the (Research Data) Question. . .*, n.d.). Whatever your position in those arguments, it is fair to say that the open data challenge is a big challenge, and it is important to

recognise that communities of researchers are ready to meet it in different ways and to different degrees (*Two Competing Visions for Research Data Sharing*, n.d.). For example, life scientists have been reported as the most ready to share the data they create (Jones et al., 2019) (Grant & Hrynaszkiewicz, 2018), and early career researchers may be particularly well-prepared to do so (Sholler et al., 2019) (Campbell et al., 2019). In fact, researchers in most data-oriented disciplines have embraced the challenge to an extent, and even where research data are associated with complex ethical issues like consent and privacy, the obligation to share data is recognised (Taichman et al., 2017). It helps that researchers can look forward to greater impact (Colavizza, 2019) (Vasilevsky, 2017). It also helps that understanding of what 'open' data really means has become quite sophisticated, aided by the FAIR Data Principles (Wilkinson et al., 2016) and promoted with the much-used soundbite that research data should be "as open as possible, as closed as necessary" (*H2020 Programme, Guidelines on FAIR Data Management in Horizon 2020*, n.d.).

Looking at this through a publishing-focused lens, when researchers are ready to share data, publishers and journals can play a useful role in enabling and realising the benefits. They help communicate and explain standards and expectations (Wu et al., 2019) (Hrynaszkiewicz, 2019). They help researchers meet the data sharing requirements including those set by their funders (*Why Share and Cite My Research Data? A Guide to Making Open Research Easier*, n.d.). They increase the discoverability of shared data, perhaps 1000-fold (although there may be more correlation in that number than simply causation) (Vines et al., 2013). They prompt researchers to "plan for the longevity, reusability, and stability of the data" ("Belmont Forum Data Accessibility Statement Policy and Template - Endorsed 18 October 2018", 2018).

Opinions about data sharing among researchers continue to be widely surveyed (Wiley, 2017) (Fecher et al., 2015). Actual data sharing practices have been investigated by looking at data availability statements published in journal articles (Federer et al., 2018) (Stall et al., 2019) (Hardwicke et al., 2018). Data availability statements describe whether and how researchers' newly analysed research data have been made available, and the conditions under which they can be accessed. When research authors have shared data in appropriate research data repositories, their data availability statements can include permanent identifiers to link between the journal article and the data. Studies of data availability statements have been used to assess the impact and increasing effectiveness of data sharing policies at journals (Federer et al., 2018) (Stall et al., 2019) (Hardwicke et al., 2018), and how data sharing practices are changing. Some conclude that practices fall short of the study authors' expectations (Wallach et al., 2018) (Federer et al., 2018) (Rowhani-Farid & Barnett, 2016). This reminds us how important it is for publishers and journals to set reasonable expectations, and to support those expectations with robust policy and process (*How and Why We're Making Research Data More Open*, n.d.). It also speaks to the pace of change in different communities as they become familiar with, interested in, able to, and required to share research data. Publishers and journals need to match that pace, and they can also lead change. Measuring, interpreting, and acting on data sharing trends ensures that publishers and journals continue to serve researchers well.

## Methods

We used topic modeling, an unsupervised machine learning technique, to identify topics from 124,000 data availability statements submitted by research authors to 176 Wiley journals between 2013 and 2019. The complete workflow is available at GitHub ([https://github.com/DWFlanagan/data-availability-statements](https://github.com/DWFlanagan/data-availability-statements)). The workflow is managed with Snakemake (Koster & Rahmann, 2012).

Wiley's electronic editorial office systems allow for the inclusion of custom questions on a journal-by-journal basis. We first extracted all of the records that contained the term "data" in either the question or the answer, but then limited the selection to questions that mentioned "data availability" or "data accessibility".

We then used SpaCy (*spaCy · Industrial-strength Natural Language Processing in Python*, n.d.) to tokenize the answers, limiting the tokens to nouns, proper nouns, and adjectives. We also added some custom stop words to ignore like "Wiley", "url", "et", and "al".

Then, we used scikit-learn (Pedregosa et al., 2011) to create a term frequency-inverse document frequency (TF-IDF) matrix of the tokenized answers, followed by Latent Dirichlet Allocation (LDA) (Blei et al., 2003). We initially used 20 topics to cluster the documents. We used pyLDAvis (Sievert & Shirley, 2014) to visualize the topics estimated by the model.

Finally, we labelled the topics where possible for further analysis and discussion, using Wiley's Data Sharing Policy Author Templates (*Data Sharing Policy — Wiley*, n.d.) as a starting guide.

## Results and Analysis

Simply counting the number of answers to the custom questions that contain the terms "data availability" or "data accessibility" shows a dramatic uptick in volume starting in early 2019 (Figure 1). This coincides with the rollout of Wiley's Expects Data policy, which added data availability statement requirements to 100+ journals starting in December 2018 (Wu et al., 2019).
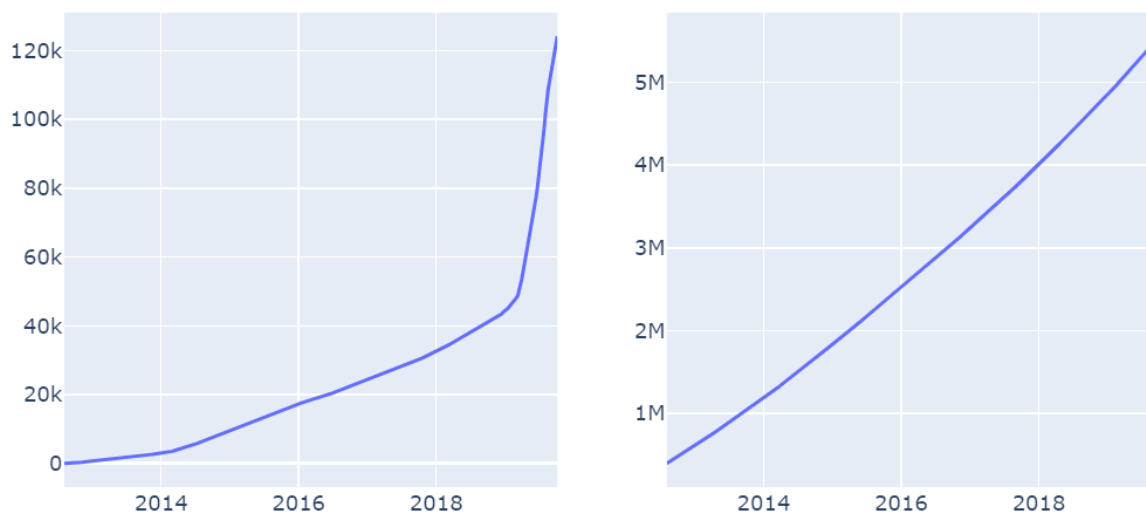


Figure 1: Cumulative data availability statements (left) and approximate cumulative number of submissions (right) from August 2012 to October 2019. The number of data availability statements increased dramatically in the first half of 2019.

Visualization of the topics estimated by the LDA (Figure 2) shows that the initial choice of 20 topics is reasonable, with only some significant overlaps (e.g., between topics 5 and 15), but otherwise reasonable spreads. Overlapping topics can be related, as between topics 4 and 19, which both seem to have something to do with genetic sequence databases, but not necessarily, because unrelated topics might use some of the same terms.
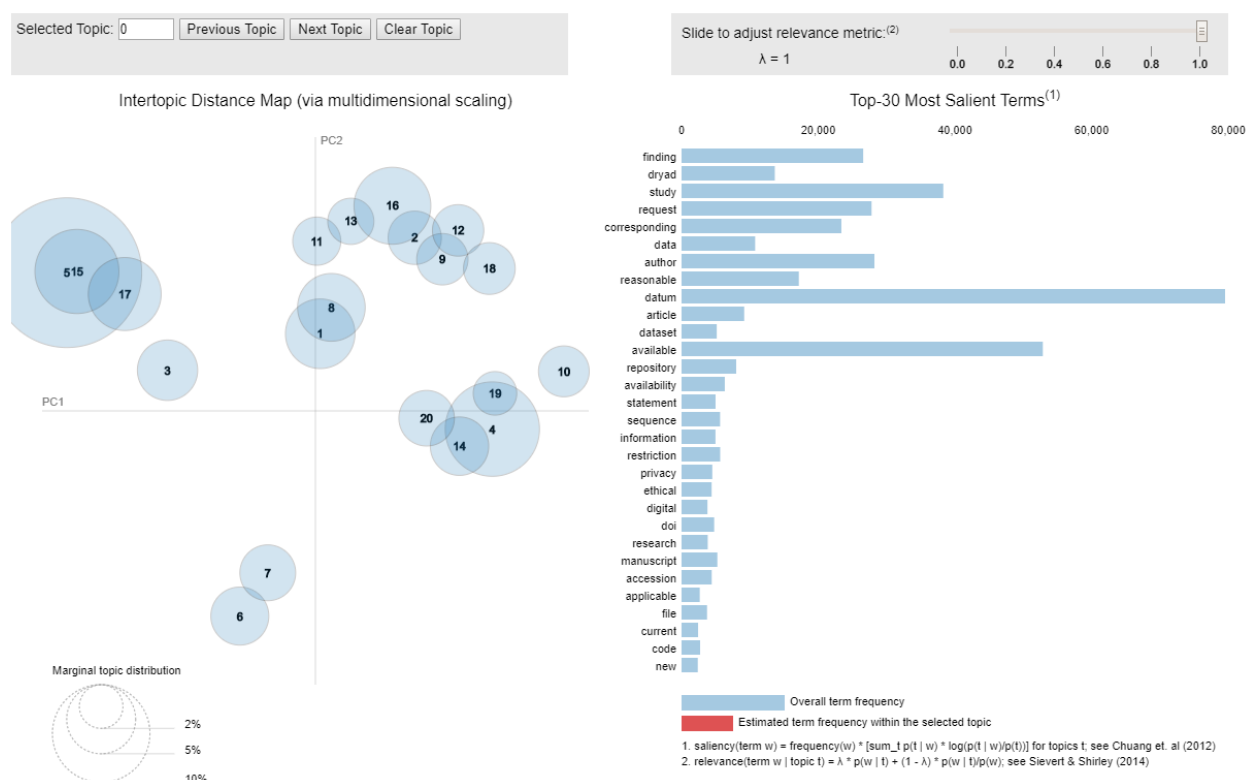
Figure 2: Visualization of the topics estimated by the LDA using pyLDAvis.

After checking that the number of topics seemed reasonable, we labeled the individual topics by hand by reading the top 20 statements for each topic and making our best guess what the cluster was about.

Some of the topics correspond well with one of the standard templates we encourage authors to use on our Data Sharing Policy page (*Data Sharing Policy — Wiley*, n.d.) and were easy to label, such as "#5: Third party restrictions" which matched with "Data subject to third party restrictions".

Other labels were more problematic. "6: Uncharacterizable" was a cluster that included experimental sections and actual data that the authors had copied and pasted into the Data Availability statement, perhaps highlighting the need for better author instructions. "7: Mixed" had many different kinds of statements that the LDA algorithm with the given parameters had combined. Tuning the text preprocessing parameters or the LDA parameters (number of topics, other hyperparameters) might resolve this mixed topic.

Some labels are also repeated. Topics 8 and 9 are both examples of "Available on reasonable request", although the LDA algorithm has resolved them into two separate topics based on the words contained in the statements themselves.

The percentages of each topic are shown in Figure 3.

4
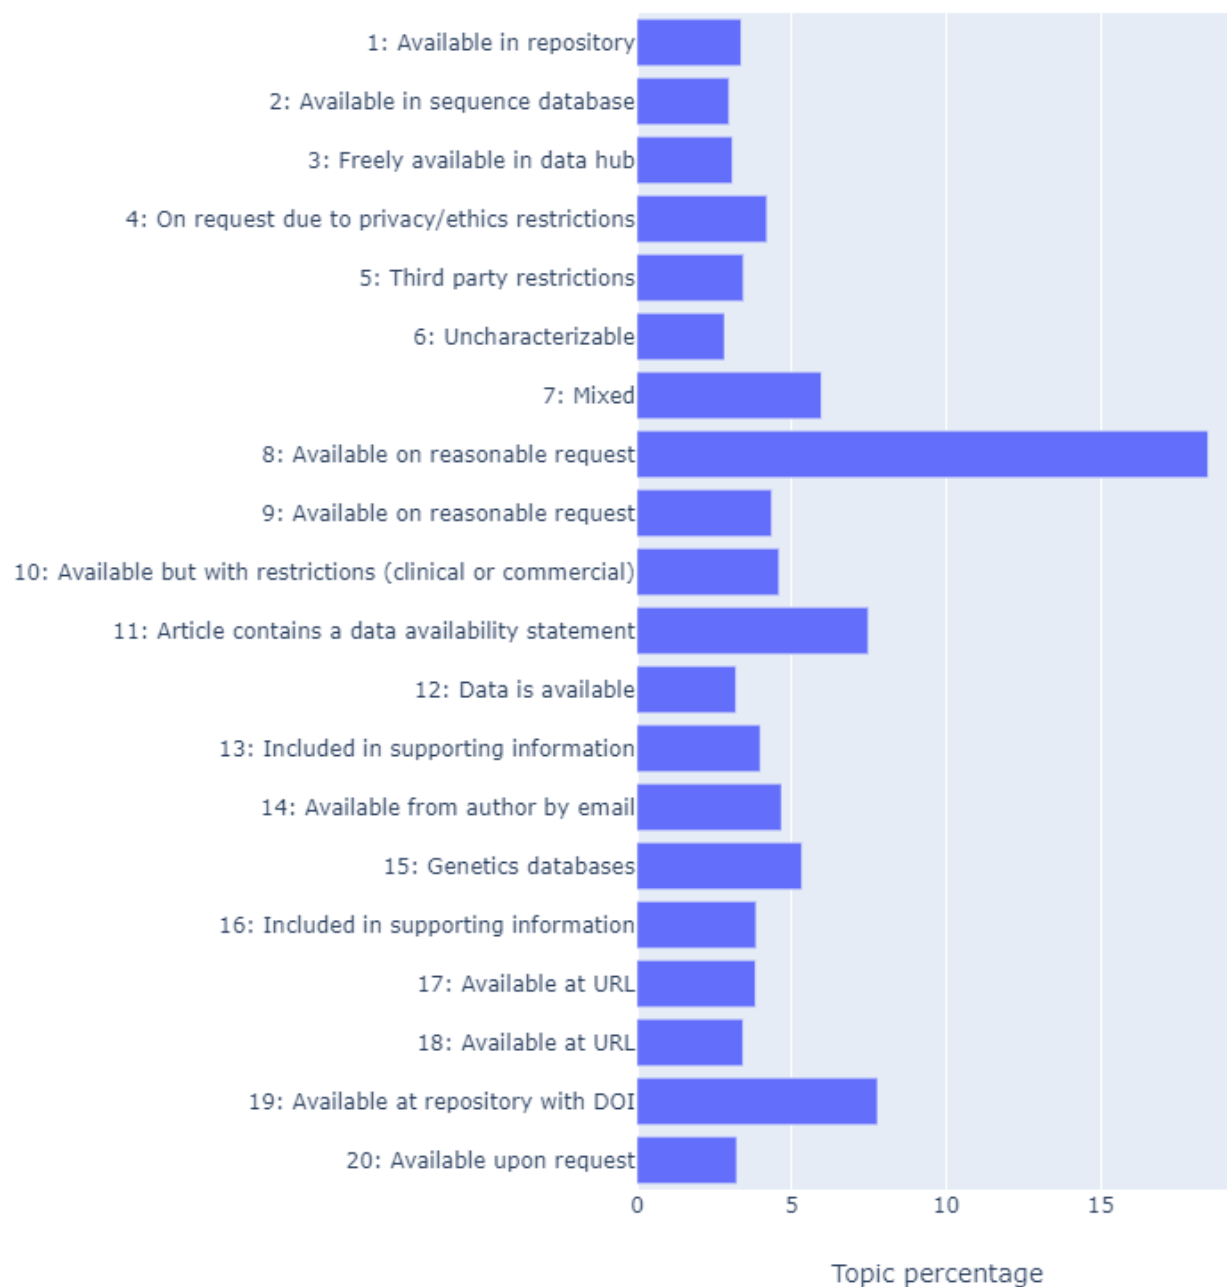
We can visualize which topics correspond to which document, as shown in Figure 3. In Figure 3, each row corresponds to a single document, and each column to a topic. The intensity of the color corresponds to the topic weighting for that document as calculated by the model. We would expect that a document would correspond strongly to one or two topics. For example, Document 1 corresponds most strongly to topic 11, while Document 2 corresponds most strongly to Topic 15, but also to Topic 16.
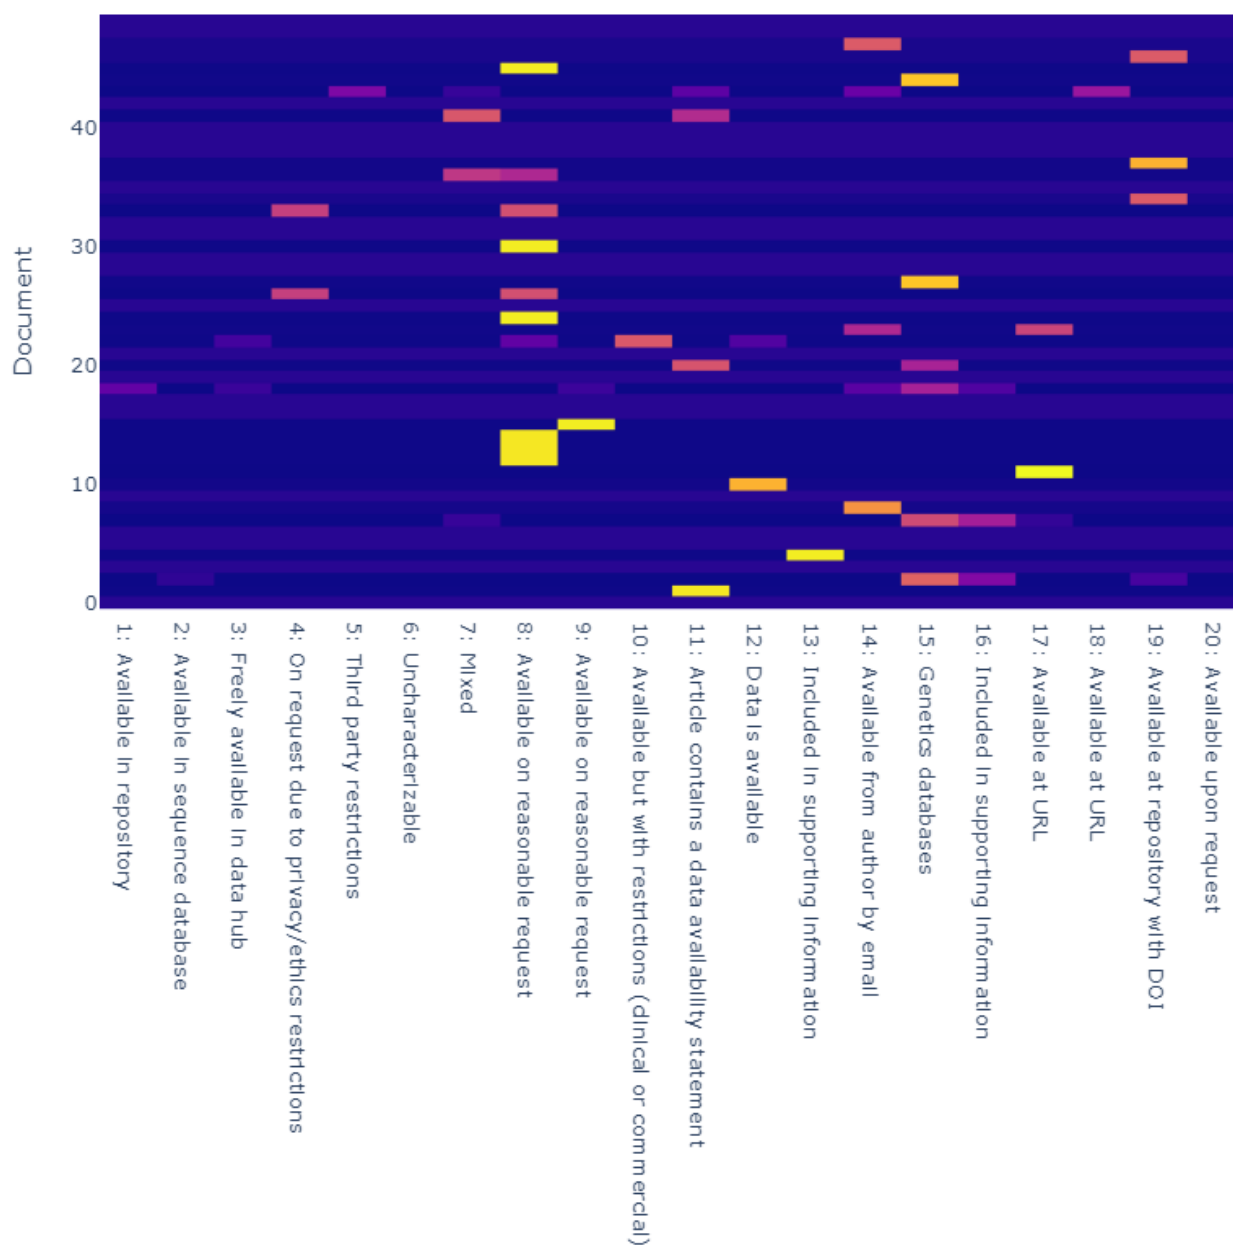
5

Figure 3: Document/topic matrix for the first 50 documents in the dataset.

We can also visualize which words correspond most strongly to each topic, as shown in Figure 4. As expected, words like "datum" (the tokenized version of "data"), "available", "reasonable", and "request" correspond strongly for Topic 8, "Available on reasonable request". Topic 19, "Available at repository with DOI", is strongly associated with terms like "Dryad" and "repository".
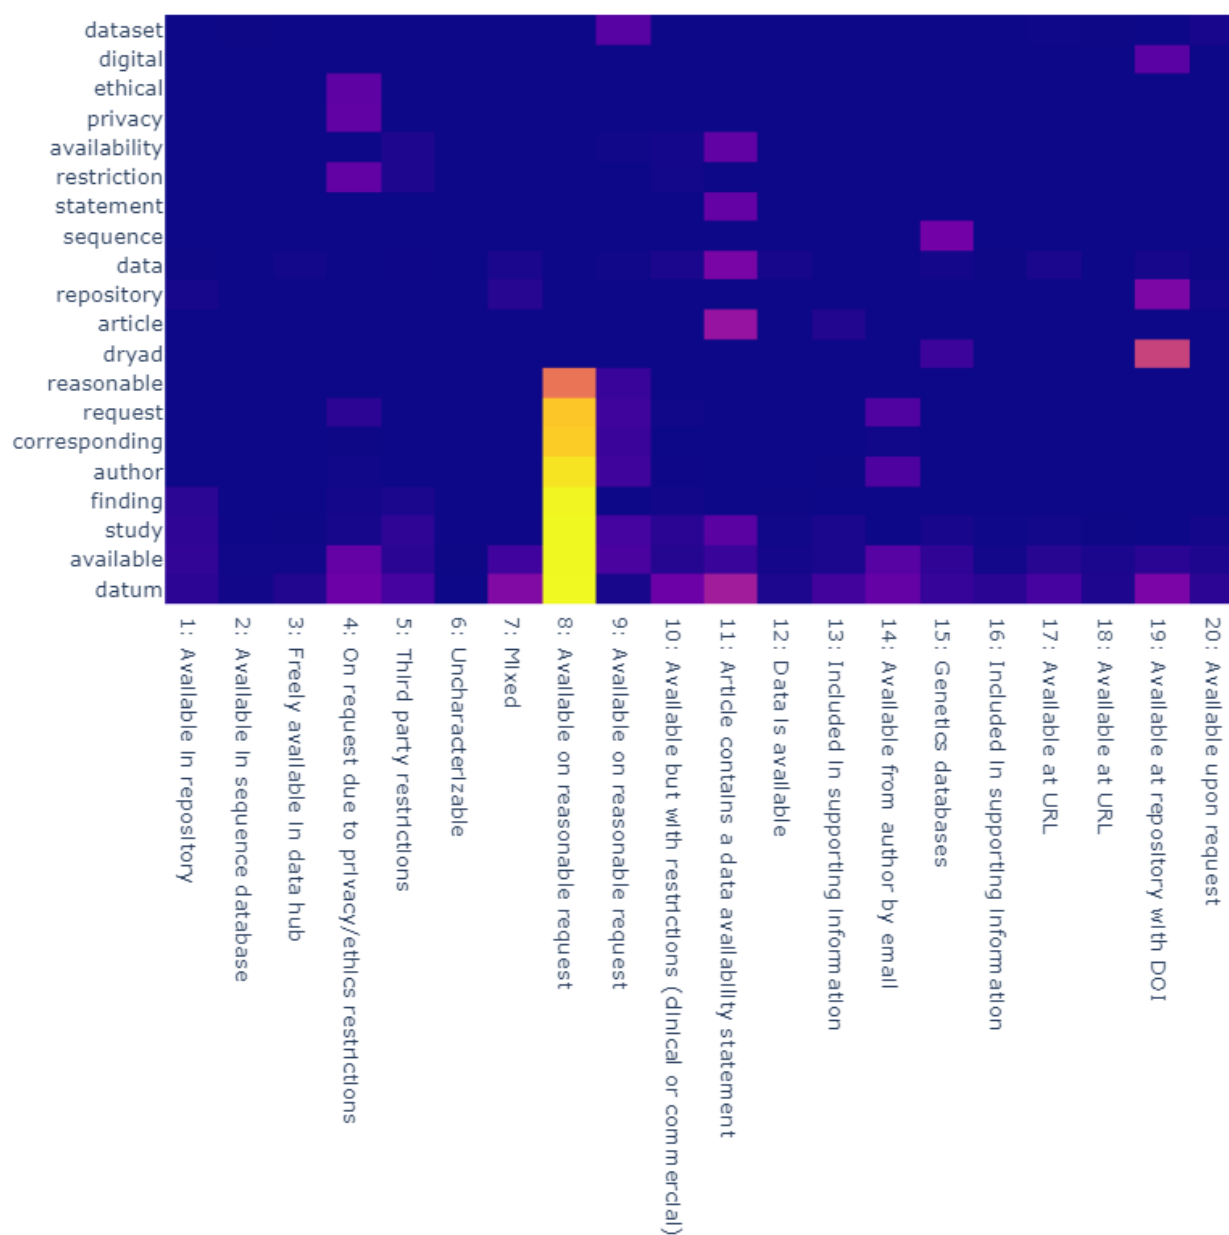
Figure 4: Word/topic matrix for the 20 most popular overall words in the dataset.

Finally, we can visualize the trends in topic growth over time (Figure 5) by assigning each document to the highest-weighted topic. A significant number of statements (approximately 46,000) were not predicted to belong to any topic category, so we filtered these out. This could be investigated in future work – it might be noise spread out over the entire data set that could be reduced by tuning the model. We can see that after December 2018 the number of authors making data available on reasonable request (Topic 8) made a sharp increase.
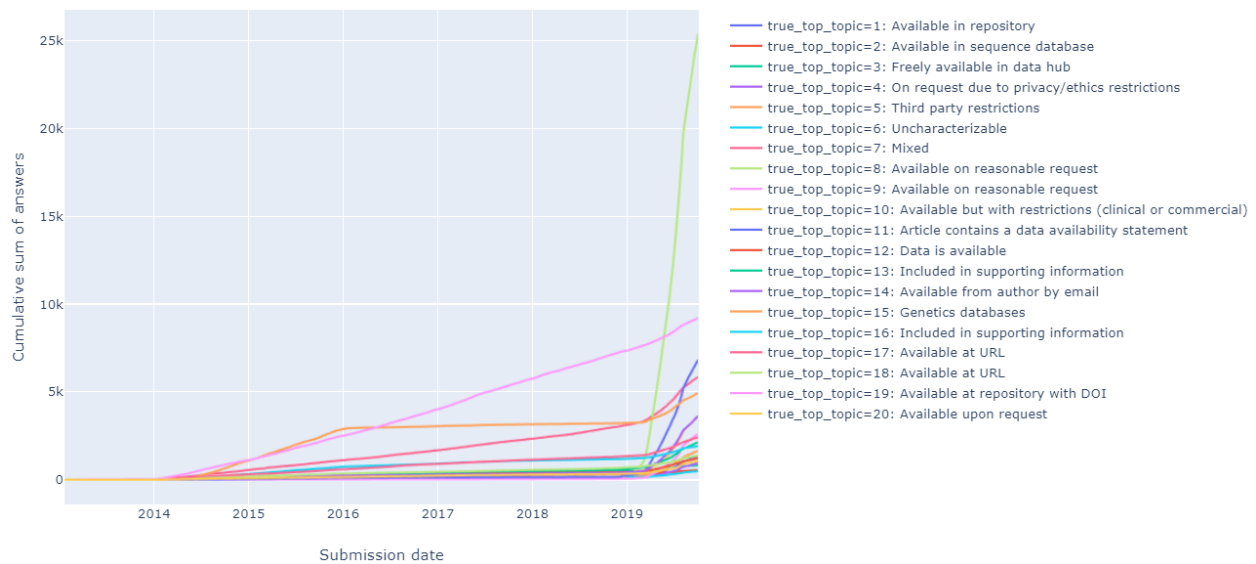
Figure 5: Growth of individual topics over time.

# Discussion and Conclusions

We see growth across all the topics we used to categorise data availability statements submitted to 176 journals between 2013 and 2019 (Figure 1).

We see a particularly sharp increase in growth in early 2019 after launching the Expects Data policy which, for journals that adopt it, requires a data availability statement in every article (*How and Why We're Making Research Data More Open*, n.d.). Implementing this requirement for data availability statements correlates with many more submitted data availability statements (as is to be expected of a successful implementation). It also correlates with many more declarations made by researchers that data are available on request (for example, topic 8 in Figure 5, and the related topics 4, 9, 10, 12, 14, and 20). This is progress, and is an improvement over the absence of any statement about data. It seems reasonable to anticipate that as researchers become familiar with, interested in, able to, and required to share research data the high proportion of data availability statements categorised as topic 8 (and related topics listed above) will gradually be replaced by data availability statements that describe shared data (like topics 1, 2, 3, 15, 17, 18, and 19).

For data that have been shared, topic 19 is a good standard to aspire to. It indicates that data are shared in a repository with a permanent digital object identifier (DOI). The number of data availability statements categorised as topic 19 shows steady growth over 6 years. This is reassuring, but topic 19 does not show the sharp increase in 2019 that we might expect to correlate with launch of our Expects Data policy. Several related topics that also describe data that have been shared online (topics 1, 2, 3, 15, 17, and 18) do show the expected sharp increase in early 2019. For Topic 19, consistent growth may be real and based on author behaviour, or it may be an artefact of the analysis that we could investigate in future work.

Data that are available in genetics databases, per topic 15, also show an interesting trend: steeper growth between 2014 and 2016; a distinct flat period between 2016 and 2018; and then steep growth in 2019,

correlating with launch of our Expects Data policy. We wonder why this might be, and perhaps refining the modelling would help.

It is also interesting to note continuing presence and moderate growth in topics 13 and 16, which indicate data that have been shared in journal supporting information.

So what have we learned?

If our goal is simply to enable research authors to describe in their journal articles whether or not they have shared the new data they have created then this can be achieved using a policy that requires data availability statements.

If our goal is to increase data sharing, then launching a policy and studying the data collected from it may also be valuable: it can create insights into how to enable and support better experiences for researchers, more data sharing, and higher-quality data sharing. For example, with data from this study we could compute which kinds of articles without shared data are similar to those with shared data, and to which journals both are submitted. With that information we could design and launch supportive policies and services where they are, in theory at least, most likely to be welcomed by researchers, and therefore where they are most likely to have a positive impact. Perhaps this is what our open data challenge becomes at publishers like Wiley: to present researchers with relevant and easy options that help them to share new research data, and that help them to make an impact with it.

It is ironic to write an article about data availability while not sharing the data set. The data availability statements we analysed were submitted by researchers to Wiley as part of journal articles, some of which we published. We analysed this information to improve our understanding of researchers' practices. We'll use what we've learned to improve our products and services so that we can better serve – and even delight – researchers who choose to publish their research with Wiley. That kind of use is, for articles we publish, covered by the license researchers give us to publish their work or, for articles submitted to but not eventually published by Wiley, by our privacy policy. But we did not ask those researchers when they submitted their articles whether we could share data about their data availability statements, and so somewhat ironically we find ourselves unsure about whether we have the right to share our data set alongside this article. We chose not to because of this uncertainty.

With that in mind, the final and perhaps most important lesson for us from this study is an appreciation for the value of careful study designs and data management plans, created before starting a study. We salute you, open researchers, you're doing amazing things.

## Data availability statement

Research data are not shared.

## Disclosure of conflicts of interest

All authors are employed by Wiley and benefit from the company's success.

# Acknowledgements

# References

Global funders who require data archiving as a condition of grants. (2015). *Figshare.* https://doi.org/10.6084/m9.figshare.1281141.v1

Data sharing and how it can benefit your scientific career. (2019). *Nature, 569*(7756), 445–447. https://doi.org/10.1038/d41586-019-01506-x

https://grants.nih.gov/grants/guide/notice-files/NOT-OD-03-032.html. https://grants.nih.gov/grants/guide/notice-files/NOT-OD-03-032.html

https://www.enago.com/academy/china-open-science-open-data-manadate-released/. https://www.enago.com/academy/china-open-science-open-data-manadate-released/

*EU budget for the future: Horizon Europe. EU funding for research and innovation 2021-2027.* (2018). https://doi.org/doi:10.2777/101500

https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/775006/Realising-the-potential-ORDTF-July-2018.pdf. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/775006/Realising-the-potential-ORDTF-July-2018.pdf

Hidden concerns of sharing research data by low/middle-income country scientists. (2018). *Global Bioethics, 29*(1), 39–54. https://doi.org/10.1080/11287462.2018.1441780

https://scholarlykitchen.sspnet.org/2014/11/11/to-share-or-not-to-share-that-is-the-research-data-question/. https://scholarlykitchen.sspnet.org/2014/11/11/to-share-or-not-to-share-that-is-the-research-data-question/

https://scholarlykitchen.sspnet.org/2019/10/14/competing-visions-research-data/. https://scholarlykitchen.sspnet.org/2019/10/14/competing-visions-research-data/

Implementing publisher policies that inform, support and encourage authors to share data: two case studies. (2019). *Insights the UKSG Journal, 32.* https://doi.org/10.1629/uksg.463

*The impact on authors and editors of introducing Data Availability Statements at Nature journals.* (2018). https://doi.org/10.1101/264929

Enforcing public data archiving policies in academic publishing: A study of ecology journals. (2019). *Big Data & Society, 6*(1), 205395171983625. https://doi.org/10.1177/2053951719836258

Early Career Researchers Embrace Data Sharing. (2019). *Trends in Ecology & Evolution, 34*(2), 95–98. https://doi.org/10.1016/j.tree.2018.11.010

Data Sharing Statements for Clinical Trials. (2017). *JAMA, 317*(24), 2491. https://doi.org/10.1001/jama.2017.6514

The citation advantage of linking publications to research data. (2019). *ArXiv*. https://arxiv.org/abs/1907.02565v2

Reproducible and reusable research: are journal data sharing policies meeting the mark?. (2017). *PeerJ*. https://peerj.com/articles/3208

The FAIR Guiding Principles for scientific data management and stewardship. (2016). *Scientific Data*, *3*(1). https://doi.org/10.1038/sdata.2016.18

https://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/hi/oa_pilot/h2020-hi-oa-data-mgt_en.pdf. https://ec.europa.eu/research/participants/data/ref/h2020/grants_manual/hi/oa_pilot/h2020-hi-oa-data-mgt_en.pdf

Paving the Way to Open Data. (2019). *Data Intelligence*, *1*(4), 60–72. https://doi.org/10.1162/dint_a_00021

Developing a research data policy framework for all journals and publishers. (2019). *Figshare*. https://doi.org/10.6084/m9.figshare.8223365.v1

https://www.wiley.com/network/latest-content/why-share-and-cite-my-research-data-a-guide-to-making-open-research-easier. https://www.wiley.com/network/latest-content/why-share-and-cite-my-research-data-a-guide-to-making-open-research-easier

Mandated data archiving greatly improves access to research data. (2013). *The FASEB Journal*, *27*(4), 1304–1308. https://doi.org/10.1096/fj.12-218164

(2018). *Zenodo*.

Wiley Open Science Researcher Survey 2016. (2017). In *Figshare*.

Phillips, R. S. (Ed.). (2015). What Drives Academic Data Sharing?. *PLOS ONE*, *10*(2), e0118053. https://doi.org/10.1371/journal.pone.0118053

Wicherts, J. M. (Ed.). (2018). Data sharing in PLOS ONE: An analysis of Data Availability Statements. *PLOS ONE*, *13*(5), e0194768. https://doi.org/10.1371/journal.pone.0194768

Make scientific data FAIR. (2019). *Nature*, *570*(7759), 27–29. https://doi.org/10.1038/d41586-019-01720-7

Data availability reusability, and analytic reproducibility: evaluating the impact of a mandatory open data policy at the journal Cognition. (2018). *Royal Society Open Science*, *5*(8), 180448. https://doi.org/10.1098/rsos.180448

Dirnagl, U. (Ed.). (2018). Reproducible research practices transparency, and open access data in the biomedical literature, 2015–2017. *PLOS Biology*, *16*(11), e2006930. https://doi.org/10.1371/journal.pbio.2006930

Has open data arrived at the British Medical Journal (BMJ)? An observational study. (2016). *BMJ Open*, *6*(10), e011784. https://doi.org/10.1136/bmjopen-2016-011784

https://www.wiley.com/network/researchers/licensing-and-open-access/how-and-why-we-re-making-research-data-more-open. https://www.wiley.com/network/researchers/licensing-and-open-access/how-and-why-we-re-making-research-data-more-open

Snakemake–a scalable bioinformatics workflow engine. (2012). *Bioinformatics*, *28*(19), 2520–2522. https://doi.org/10.1093/bioinformatics/bts480

https://spacy.io. https://spacy.io

Scikit-learn: Machine Learning in Python. (2011). *Journal of Machine Learning Research*, *12*, 2825–2830.

Latent dirichlet allocation. (2003). *Journal of Machine Learning Research*, *3*(Jan), 993–1022.

LDAvis: A method for visualizing and interpreting topics. (2014). *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*, 63–70.

https://authorservices.wiley.com/author-resources/Journal-Authors/open-access/data-sharing-citation/data-sharing-policy.html. https://authorservices.wiley.com/author-resources/Journal-Authors/open-access/data-sharing-citation/data-sharing-policy.html