



OPEN

DATA DESCRIPTOR

PreprintToPaper dataset: connecting bioRxiv preprints with journal publications

Fidan Badalova ¹✉, Julian Sienkiewicz ^{2,3} & Philipp Mayr ¹

The PreprintToPaper dataset connects bioRxiv preprints with their corresponding journal publications, enabling large-scale analysis of the preprint-to-publication process. It comprises metadata for 145,517 preprints from two periods, 2016–2018 (pre-pandemic) and 2020–2022 (pandemic), retrieved via the bioRxiv and Crossref APIs. We selected the two periods to capture preprint-publication dynamics before and during the COVID-19 pandemic while avoiding transitional years. Each record includes bibliographic information such as titles, abstracts, authors, institutions, submission dates, licenses, and subject categories, alongside enriched publication metadata including journal names, publication dates, author lists, and further information. In addition to the main dataset, a version-history subset provides all available versions of preprints within the two selected periods, enabling analysis of how preprints evolve over time. Preprints are categorized into three groups: Published (formally linked to a journal article), Preprint Only (posted on a preprint server), and Gray Zone (potentially published in a journal but unlinked). To enhance reliability, title and author similarity scores were computed, and a human-annotated subset of 299 records was created to evaluate Gray Zone cases. The dataset supports diverse applications, including studies of scholarly communication, open science policies, bibliometric tool development, and natural language processing research on textual changes between preprints and their corresponding journal articles.

Background & Summary

Preprints have become an essential channel for the rapid dissemination of scientific findings, particularly in the life sciences, where timely access to results can accelerate discovery and collaboration. Their importance was intensified during the COVID-19 pandemic, when preprints served as a primary means of sharing the newest research across domains before formal peer review^{1,2}.

The PreprintToPaper dataset³, which we describe in this paper, links bioRxiv preprints with their subsequent journal publications, allowing for large-scale analysis of the preprint-to-publication process. It includes metadata on more than 145,000 preprints from two distinct periods (2016–2018, the pre-pandemic period, and 2020–2022, the COVID-19 pandemic period), with information on titles, authors, abstracts, institutions, licenses, submission and publication dates, and subject categories. Preprints are categorized into “Published”, “Preprint Only”, and “Gray Zone” cases, based on automated linking procedures combining Digital Object Identifier (DOI) information, title similarity, and author similarity. In addition to the main dataset, two subset files are provided: (i) a human-annotated subset of Gray Zone cases is provided to support evaluation of the automated matching approach (this subset includes records with title similarity (*title_match_score* = 0.75) and was independently checked by two annotators; further details on the annotation procedure are provided in the Methods section - Step 5: Gray Zone Verification) and (ii) a version-history subset that lists all available versions for preprints that have version one and at least one additional version. This subset is intended for detailed, version-level analyses of preprint evolution.

The motivation for creating the PreprintToPaper dataset is to provide a comprehensive resource for studying how preprints evolve into published journal articles or remain as preprints posted only on a preprint server^{4,5}. While some preprints are eventually linked to formal publications, many others stay as preprints without ever appearing in journals. This dataset is the systematic effort to automatically collect and link metadata

¹GESIS - Leibniz Institute for the Social Sciences, Cologne, 50667, Germany. ²Warsaw University of Technology, Faculty of Physics, Warsaw, 00-662, Poland. ³Warsaw University of Technology, Centre for Credible AI, Warsaw, 00-614, Poland. ✉e-mail: fidan.badalova@gesis.org

from bioRxiv⁶ with publication records, making it possible to analyze patterns in the preprint-to-publication process at scale⁷. We selected the 2016–2018 (pre-pandemic) and 2020–2022 (pandemic) periods to capture preprint-publication dynamics before and during the COVID-19 pandemic while avoiding transitional years.

We identify the following potential reuse cases for the PreprintToPaper dataset.

- Policy and evaluation studies: Provides evidence for assessing the role of preprints in accelerating scientific dissemination, informing open science policies, and tracking uptake during events such as the COVID-19 pandemic^{2,8,9}. To enable comparison, we included preprints from 2016–2018, which we designate as the pre-pandemic period.
- Scholarly communication research: Enables analysis of publication delays, changes in titles, abstracts, and authorship between preprints and journal articles, and differences across research fields or time periods^{10,11}.
- Tool and method development: Serves as a reference dataset for testing algorithms in informetrics, metadata linking, and similarity matching (e.g., title and author disambiguation).
- Natural language processing and text analysis: Supports research on linguistic and structural changes between preprints and their corresponding journal articles, such as title reformulations, abstract rewriting, and shifts in author order or contribution statements.

In the following, we summarize related work that has used preprint data from bioRxiv.

Research on bioRxiv has shown that articles posted as preprints tend to receive more citations and online attention than comparable articles not deposited, although the causal mechanisms behind this effect remain uncertain¹². Surveys of authors indicate that the main motivations for posting preprints are to increase visibility and accelerate dissemination, while concerns about peer review and lack of awareness are the main barriers². Large-scale bibliometric analyses further confirm a citation and altmetric advantage for bioRxiv-linked publications, with preprints themselves increasingly cited and shared on social media¹². During the COVID-19 pandemic, preprints on bioRxiv and medRxiv played an unprecedented role, being accessed, cited, and shared at higher rates, and influencing both journalistic reporting and policy discussions¹. An analysis of biomedical publishing during the first months of the COVID-19 pandemic reveals a sharp rise in publication volume, faster acceptance times for COVID-19 papers, and reduced international collaboration, with these shifts occurring largely at the expense of non-COVID-19 research⁸. Other studies of bioRxiv preprints have found that approximately 30% remain as preprints posted only on a preprint server, that many are posted too late to enable substantive feedback, and that nearly half of preprints published in a journal are concentrated among four major publishers, suggesting that the platform is used more for priority setting and visibility than for pre-publication peer review^{4,13}.

Below, we provide an overview of related datasets based on bioRxiv preprints and collections of COVID-19-related preprints.

The Rxivist database provides a snapshot of preprints from bioRxiv and medRxiv collected via a custom web crawler, enabling readers to sort, filter, and analyze tens of thousands of preprints based on a PostgreSQL-importable database^{7,14}.

The covid19_preprints dataset compiled weekly updated metadata on COVID-19-related preprints from multiple sources (Crossref, DataCite, arXiv, and RePEc), using keyword-based matching and deduplication procedures to track their distribution over time^{1,15}.

Europe PMC¹⁶ maintains a large corpus of preprints and journal articles, including bioRxiv preprints, and links them to their corresponding publications.

Compared with the three related preprint resources discussed above, the PreprintToPaper dataset offers two main advantages:

- It stores both the initial and final versions of bioRxiv preprints, along with detailed metadata (e.g., title, authors, journal, and publication date) for their corresponding journal articles. In addition, a separate version-history file contains all available preprint versions for cases with more than one version within our study periods.
- It identifies articles that are not formally linked but may have been published in a journal (Gray Zone cases) by using title and author similarity, which tolerate minor title changes and variations in author order, thereby going beyond exact-match linking.

Find a detailed comparison of the four datasets in Table 1.

Two further resources provide data related to preprint-publication links: PreprintMatch¹⁷ and the Crossref relationships dataset¹⁸. In contrast to PreprintToPaper, they focus on DOI-to-DOI linkage and matching infrastructure, rather than delivering a bioRxiv-specific dataset with enriched metadata.

Methods

Dataset generation workflow. Figure 1 illustrates the data generation workflow for the PreprintToPaper dataset. The process consists of six main steps: (1) collecting metadata from the bioRxiv preprint platform; (2) retrieving additional metadata from Crossref based on the published journal article DOIs listed in the bioRxiv records; (3) categorizing preprints into published and preprint-only; (4) identifying unlinked publications in bioRxiv (Gray Zone); (5) verifying Gray Zone preprints using author-based metrics (author match score and author count difference); and (6) generating the final dataset.

Step 1: bioRxiv Metadata Retrieval. The first step in this study was to obtain metadata from the bioRxiv preprint platform using the official bioRxiv API (<https://api.biorxiv.org/details/server/interval/cursor/format>).

Aspect	covid19_preprints	Rxivist	Europe PMC corpus	PreprintToPaper (Ours)
Time span	2017–2021	Nov 2013 - Feb 2023	2013–2023	2016–2018 (pre-pandemic) and 2020–2022 (pandemic)
Coverage	Multiple preprint servers (bioRxiv, arXiv, OSF, ESSOAR, Figshare, Zenodo, etc.)	bioRxiv and medRxiv	Multiple 32 life-science preprint servers	bioRxiv
Scope	COVID-19 related preprints only	All subjects (bioRxiv + medRxiv)	All subjects (32 servers)	All subjects (bioRxiv)
Size	66,130 records	223,541 records	761,123 records	145,517 records
Preprint information	Source, identifier (DOI), posted date, title, abstract	Title, abstract, DOI, category, submission date, authors, institutions	Title, abstract, DOI, submission date, authors, full text of COVID-19 and Europe PMC funder preprints in JATS XML, author affiliations, license, funding, versions, status (Withdrawal/removal)	Title, abstract (initial + last), DOI, category, submission date (initial + last), authors (initial + last), corresponding authors, institutions, license, version number, category, fulltext JATS XML
Publication information	None	Published journal article DOI, journal name	Published journal article DOI, journal name, publication date, citations, references	Published journal article DOI, title, authors, journal name, publication date (from Crossref)
Versioning	Only one version per preprint	Only one version per preprint	All versions (except from bioRxiv, medRxiv, SSRN)	All versions at the relevant period (bioRxiv)
Publication linking	None	Published journal article DOIs partially incomplete (crawling stopped early)	Official bioRxiv links enriched with journal metadata + an investigation of unlinked preprints where an exact match of title and first author is required. Even minor changes in title or author names are ignored; as a result, many preprints with slightly different titles are not linked to their journal papers	Official bioRxiv links enriched with Crossref + additional Gray Zone investigation where both titles and full author lists are compared using similarity scores. Preprint-paper pairs with title similarity between 75% and 100% are identified and added as potential links ($\approx 19,000$)

Table 1. Comparison of the covid19_preprints, Rxivist, Europe PMC, and PreprintToPaper datasets.

The API has the following structure: *server* indicates the source of the data (e.g., bioRxiv or MedRxiv), *interval* indicates the desired date range, *cursor* indicates the pagination of the results, and *format* indicates the output format (e.g., JSON). The following data fields were collected through the API, and the retrieved metadata fields were mapped to the following structure in the dataset (see Table 2).

Data were collected for two periods: **2016–2018** (pre-pandemic period) and **2020–2022** (COVID-19 pandemic period). The collected metadata was then processed. Because most preprints have multiple versions, only records that met the following criteria were retained during the processing stage:

1. If the preprint is in the relevant period and both the first and last versions of the preprint are available, then both versions were kept.
2. If the preprint has only one version and this version is the first, then this version was also retained.

If only non-initial versions were available during the relevant period, the corresponding records were deleted. Additionally, a pivoting step was applied to the entire dataset to prevent DOI duplication. To avoid this, each preprint was stored only once, but non-repeating fields were stored in separate columns. For example, the *title*, *authors*, *corresponding author*, *submission dates*, and *abstract* of both the first and last versions were stored in new columns. For this purpose, the *_1st* and *_last* suffixes were systematically added to the column names (e.g., *biorxiv_title_1st* and *biorxiv_title_last*).

In addition to the main table, a separate version-history subset was created that contains all available versions for preprints with a first version and at least one subsequent version within our study periods. The filtering applied to the main dataset is designed to track changes between the first and last preprint versions, as well as changes that occur upon journal publication, whereas the version-history subset supports more detailed, version-level analyses. However, because the bioRxiv API was observed to return identical abstract text across different versions of the same preprint, abstracts were not included in the version-history subset.

Step 2: Crossref Metadata Retrieval. For preprints that were later published as journal articles, the corresponding journal DOI is stored in the bioRxiv dataset. About two weeks after journal publication, bioRxiv automatically adds a link to the corresponding journal article (<https://www.biorxiv.org/about/FAQ>). Detailed metadata for these articles was retrieved via the official Crossref API (<https://api.crossref.org/works/doi>), by replacing the doi field with the DOI of the corresponding article. This query returns an extensive metadata package for the corresponding journal article. The following fields were retrieved from the Crossref database (see Table 3).

The official publication dates of 500 journal articles were manually compared with the *online_publication_date* and *issue_online_date* obtained from Crossref. In most cases, the *online_publication_date* matched with the actual publication date of the article.

Based on the results of manual checks, we retained the *online_publication_date* when available and used the *issue_online_date* otherwise. Both values were then merged into a single column, *crossref_publication_date*. Additionally, a new column, *crossref_publication_type*, was created to indicate whether the date refers to the online or issue version.

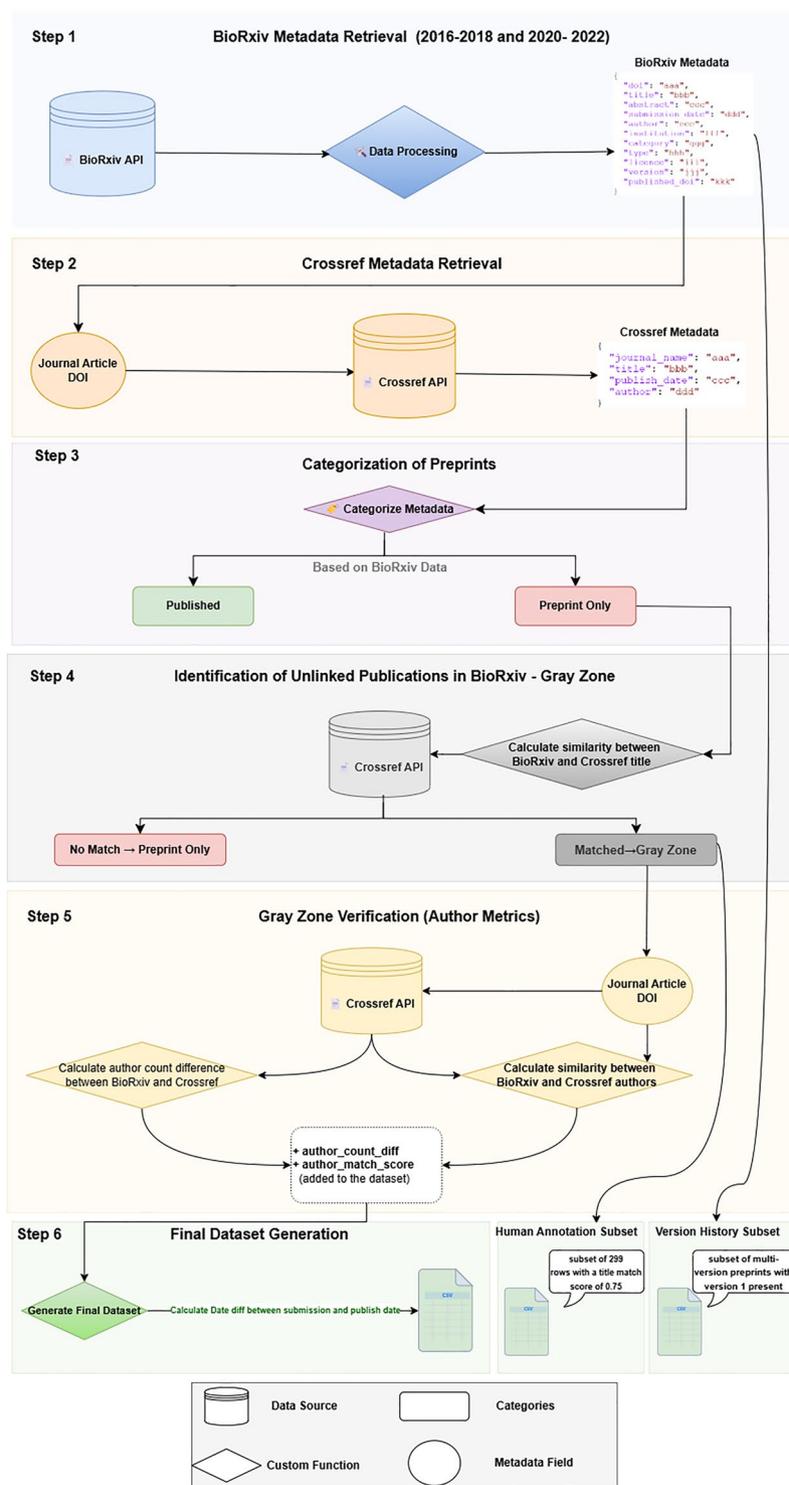


Fig. 1 PreprintToPaper dataset creation workflow. Final enriched dataset and two derived subsets are generated: (i) a human-annotated Gray Zone subset of 299 title-matched cases (title match score = 0.75) and (ii) a version history subset, in which multi-version preprints that retain version 1 and at least one later version within the study periods are included.

Step 3: Categorization of Preprints. In the third step, the dataset was categorized based on whether the preprints had DOIs linked to journal articles on the bioRxiv preprint server.

1. If a journal article DOI was available for the preprint, the preprint was included in the *published* group.
2. If a DOI was not available, the preprint was categorized as *preprint only*.

bioRxiv API field	Dataset column	Description
DOI	<i>biorxiv_doi</i>	Unique digital identifier of the preprint
Title	<i>biorxiv_title</i>	Title of the preprint
Authors	<i>biorxiv_authors</i>	List of all authors
Corresponding author	<i>biorxiv_author_corresponding</i>	Name of the corresponding author
Corresponding author institution	<i>biorxiv_author_corresponding_institution</i>	Institution of the corresponding author
Date	<i>biorxiv_submission_date</i>	Submission date of the preprint
Version	<i>biorxiv_version</i>	Version number of the preprint
Type	<i>biorxiv_type</i>	Document type of the preprint
License	<i>biorxiv_license</i>	License
Category	<i>biorxiv_category</i>	Subject category
JATS XML path	<i>biorxiv_jatsxml</i>	Link to the XML file
Abstract	<i>biorxiv_abstract</i>	Abstract text
Published	<i>biorxiv_published_doi</i>	Publication DOI in a journal

Table 2. Mapping of bioRxiv API fields to the dataset columns.

Crossref API field	Dataset column	Description
Container-title	<i>crossref_journal_name</i>	Name of the journal in which the article was published
Title	<i>crossref_title</i>	Official title of the corresponding journal article
Author	<i>crossref_authors</i>	List of authors
Published-online	<i>crossref_online_publication_date</i>	Date the journal article was first made available online
Published-print	<i>crossref_issue_online_date</i>	Date associated with the volume or issue of the journal in which the article appeared

Table 3. Mapping of Crossref API fields to the dataset columns.

The categorization result was stored in the column *custom_status*.

Step 4: Identification of unlinked publications in bioRxiv - Gray Zone. Although bioRxiv metadata is expected to include DOIs for all preprints that have been published as journal articles, verification revealed that in some cases the DOI of the corresponding journal article is missing, even when the preprint has been published in a journal. For example, the fourth version of a preprint with <https://doi.org/10.1101/170381> was submitted to bioRxiv on August 6, 2020, but the same article was later published in Brain Structure and Function with <https://doi.org/10.1007/s00429-020-02136-0> on September 18, 2020. Nevertheless, the preprint still appears as preprint only in the bioRxiv record.

Previous studies have also documented this issue^{1,5,7}. For example, Abdill and Blekhman (2019) found that 37.5% of 120 bioRxiv preprints were missing publication links. Fraser *et al.* (2021) reported that 7.6% of 12,788 preprints had missing publication links. Cabanac *et al.* (2021) reported that 60.3% of preprints marked as journal-published on medRxiv and three other servers lacked corresponding publication links.

To identify such cases, we compared the titles of works marked as *preprint only* with the titles of the corresponding journal articles in the Crossref database. First, a separate request was sent to the Crossref database for each DOI, and matching titles were identified. Python's **SequenceMatcher** library was then used to assess title similarity. The algorithm determines the longest common subsequence between two texts, accounts for the total length of matching parts, and calculates a similarity coefficient ranging from 0 to 1. A coefficient close to 1 indicates a complete match, whereas a value near 0 indicates substantial differences.

To assess our choice of similarity measure, we compared Python's SequenceMatcher with two FuzzyWuzzy metrics on a random sample of 100 preprint and publication title pairs from our dataset, all of which were confirmed as published in a journal. In 58% of the cases, all three measures produced identical scores, and in the remaining cases the differences were small and did not change the classification of any pair with respect to our 0.75 threshold. We therefore retained SequenceMatcher as our main measure, as it is part of the Python standard library.

Matches with a similarity index of 0.75 or higher were classified into the *Gray Zone* category. This threshold was chosen because, in tests comparing preprints already marked as *journal-published* in bioRxiv with their corresponding journal versions, results in this range proved most reliable in Fig. 2.

The calculated title similarity values were stored in the *title_match_score* column of the dataset. Additionally, for the preprints classified as **Gray Zone**, the corresponding metadata retrieved from the Crossref API including *journal name*, *title*, *published date*, *author list*, and *publication date* was stored in these fields *crossref_journal_name*, *crossref_title*, *crossref_authors*, *crossref_publication_date*, *crossref_publication_type*.

Step 5: Gray Zone Verification (Author Metrics). To further improve the accuracy of preprints in the *Gray Zone* category, additional author-based metrics were applied. At this stage, the author names from the bioRxiv and Crossref metadata were compared. Two main metrics were calculated:

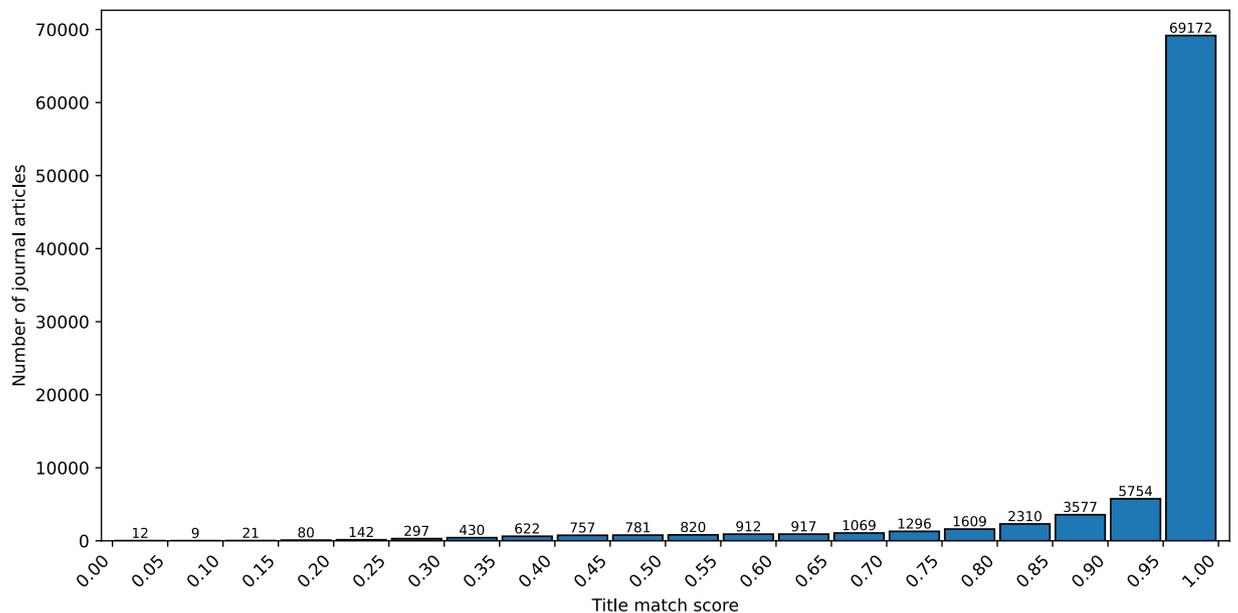


Fig. 2 Distribution of journal papers by title match score.

- **Author Count Difference** – the difference between author counts in the bioRxiv and Crossref records.
- **Author Match Score** – the similarity score between the author lists.

For the second metric, Unicode normalization was first performed to minimize differences in name format, diacritical marks were separated and removed, and all letters were reduced to lowercase. Then, for each record, the author list was first divided into individual authors by the “;” separator, and each author line was tokenized by splitting it into words using the space character. As a result, a nested (i.e., list-within-a-list) structure consisting of authors was constructed for each line. Based on this structure, the SequenceMatcher library was applied, and the author’s match score was calculated. This approach allowed for more accurate results in cases such as variations in the order of first and last names or when only initials were recorded.

Additionally, to validate these metrics, a human-annotated subset with a title match score of 0.75 was created, representing the lower bound of our title similarity threshold. Two annotators manually checked the corresponding bioRxiv and Crossref abstracts, and inter-annotator agreement was high.

Thus, the publication status of preprints in the Gray Zone category was determined more accurately using the author match score and author count difference metrics.

Step 6: Final Dataset Generation. In the final stage, several additional calculations were performed to enrich the dataset with further features. For example, the difference between the submission dates of the first and last versions of each preprint was calculated, as well as the difference between the official journal publication date and the last submission date on bioRxiv for the preprints published in journal. These values were stored in the *custom_biorxivVersion_dateDifference* and *custom_submission&publication_dateDiff* columns.

These operations made it possible to track the time dynamics between preprint versions and their subsequent journal publication. The finalized dataset was then saved in CSV format.

Data Record

The dataset is based on data collected from the bioRxiv and Crossref databases. It is stored in CSV format and is publicly available via Zenodo³ (<https://doi.org/10.5281/zenodo.17992421>). The dataset consists of three main files:

1. The **Main dataset** includes metadata of preprints collected in 2016–2018 and 2020–2022. This file contains primary fields obtained from the bioRxiv API (DOI, title, authors, corresponding institutions, abstract, version, category, license, etc.) and additional fields obtained from the Crossref API (e.g., journal name, title in the journal, publication date, author list). In addition, there are custom fields calculated by us: title match score, author match score, author count difference, the difference between the first and last version submission dates, and the difference between the last preprint submission date and the publication date. Only fields that are not self-explanatory are described in detail below:
 - **custom_status**: Classification indicating the publication status of preprints: *Published* (officially published in a journal and confirmed with a DOI in bioRxiv), *Preprint Only* (posted on a preprint server), or *Gray Zone* (no official link in bioRxiv, but a possible publication identified based on title and author match).
 - **biorxiv_published_doi**: DOI of the corresponding journal articles. This information is collected from two sources: (1) official publication DOIs provided by bioRxiv; (2) DOIs additionally identified in Gray Zone cases.

Period	Gray Zone	Preprint Only	Published	Total Preprints
2016	359	1 019	3 343	4 721
2017	841	2 314	8 191	11 346
2018	1 121	4 115	12 943	18 179
2020	3 746	8 884	26 081	38 711
2021	5 292	8 987	22 539	36 818
2022	7 731	10 494	17 517	35 742
Total	19 090	35 813	90 614	145 517

Table 4. Distribution of preprints in the PreprintToPaper dataset.

- **custom_submission&publication_dateDiff:** The difference in days between the last preprint submission date and the publication date of the corresponding journal article.
- **custom_biorxivVersion_dateDifference:** The difference in days between the submission dates of the first and last bioRxiv versions.
- **biorxiv_version_last:** The serial number of the last version of the preprint available on bioRxiv (e.g., a value of “5” indicates that five versions have been published).
- **biorxiv_jatsxml:** A URL that references a JATS XML file containing the full text and metadata of the preprint, if available.
- **crossref_publication_date:** The publication date retrieved from Crossref. This date can refer either to the online publication or the issue publication.
- **crossref_publication_date_type:** Indicates the type of date given in the *crossref_publication_date* column: online publication or issue publication. If the online publication date is available, it is preferred; otherwise, the issue publication date is used.
- **author_count_diff:** The difference between the number of authors in the preprint and the number of authors in the corresponding journal article. Positive values indicate that the preprint published in journal has more authors, while negative values indicate fewer.
- **title_match_score:** The similarity score between the preprint and the journal article titles. Values range from 0 to 1 and are calculated with the SequenceMatcher algorithm. Only matches within the 0.75–1.0 range were considered potential publication candidates.
- **author_match_score:** Similarity score between the preprint and corresponding journal article author lists. Ranges from 0 to 1 and is primarily used to verify Gray Zone cases, with higher values indicating stronger similarity.

Table 4 presents the distribution of preprints in the dataset by period and category (Gray Zone, Preprint Only, Published) along with the total numbers.

2. The **Version History Subset** is a separate, version-level file that contains all available versions for preprints that have version one and at least one additional version within our study periods (2016–2018 and 2020–2022). Preprints with only a single version or without the first version present are not included in this file. Each row corresponds to a bioRxiv version and includes the DOI, version number, submission date, title, authors, category, license, and other core bioRxiv metadata. This subset allows detailed analyses of how preprints evolve across multiple versions.
3. The **Human-Annotated Subset**, containing 299 records, includes only cases with a *title match score* of 0.75. Preprints from both periods (2016–2018 and 2020–2022) were selected, and two annotators verified potential matches by manually checking the corresponding abstracts from bioRxiv and Crossref. To support evaluation of the Gray Zone category, this file is provided as an additional resource. The dataset includes the following variables: *year*, *biorxiv_doi*, *suspected_published_doi*, *author_match_score*, *annotator1*, and *annotator2*.
 - **suspected_published_doi:** Suspected DOI of the journal-published article. Includes DOIs identified for Gray Zone preprints where a potential corresponding journal article was found but could not be fully confirmed.
 - **annotator.1/annotator2:** Labels provided independently by two annotators to validate Gray Zone matches. Values: *True* (confirmed as correct match), *False* (not a valid match), or *NA* (uncertain).

Technical Validation

To assess the validity of the proposed approach, all the records characterized by the *title_match_score* = 0.75 (in total 299 records) have been selected and annotated by two referees; each of them had an option to mark a record as TRUE (the identified journal publication was based on the preprint), FALSE (there was no connection between the preprint and the publication) or NA (the annotator was unable to make a call). The Cohen’s kappa coefficient¹⁹ for non-NA data among annotators is equal to $\kappa = 0.86$, indicating strong agreement among the annotators²⁰.

The annotations have been then merged by taking the conjunction of two annotations, i.e., only if both annotators deemed the identified publication to be true, it was considered as such; if either of the annotators regarded the publication as NA, it was assigned NA. The resulting distribution gives 69% of positive recognition and 28% of falsely assigned publications, with the remaining 3% that cannot be resolved. These outcomes

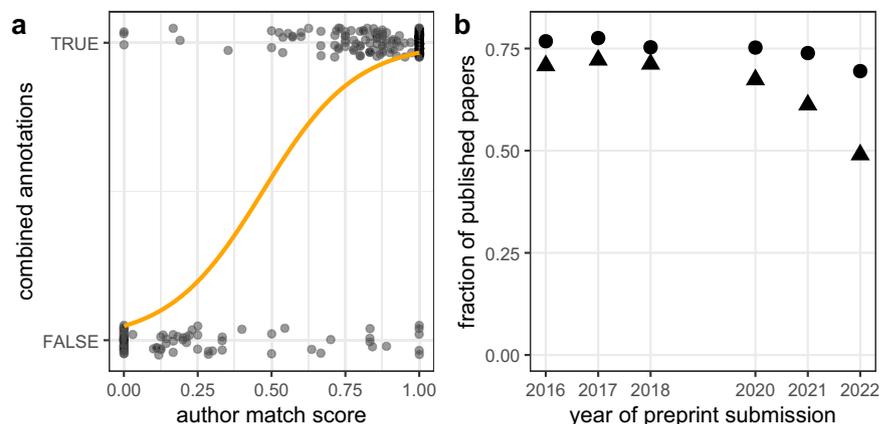


Fig. 3 (a) Combined annotations versus author match score for records characterized by the title match score equal to 0.75 (see text for rationale). Symbols are data points, and the curve comes from the fit to the logistic regression model (1). Data points are artificially jittered around TRUE and FALSE values to avoid overlap. (b) The fraction of preprints that were eventually published as journal articles versus the year of preprint submission. Triangles represent data without the gray zone records, circles take into account gray zone records, with the author match score larger than 0.47 (suggested by logistic regression analysis, see text in Technical Validation) used to distinguish between preprints and eventually published journal papers.

suggest that the title matching score alone might be an insufficient indicator. Figure 3a presents the combined annotation versus author match score, presenting evidence that preprint – publication pairs with a low level of author match score are more likely to be among false positives than ones characterized by an author match score close to 1. To provide a more qualitative description of this fact and motivated by the observed relation between the combined annotation and author match score seen in Figure 3a, we used the logistic regression model²¹ given by

$$\log\left(\frac{\pi}{1-\pi}\right) = \beta_0 + \beta_1 x \quad (1)$$

where π is interpreted as the probability of the dependent variable (annotation) equaling a success, x reflects author match score and β_0, β_1 are regression coefficients. The model was fitted on the set of non-NA data, resulting in $\beta_0 = -2.99 \pm 0.46$ and $\beta_1 = 6.36 \pm 0.67$ (see the curve in Figure 3a). The location parameter $\mu = -\beta_0/\beta_1$ – the midpoint of the curve, where π crosses 1/2 is $\mu = 0.47 \pm 0.12$ where the standard deviation of μ has been calculated using propagation of uncertainty²², i.e., $\sigma_\mu^2 = |\beta_0/\beta_1| \sqrt{\sigma_{\beta_0}^2/\beta_0^2 + \sigma_{\beta_1}^2/\beta_1^2}$. To assess the validity of the logistic regression approach, we performed 10-fold cross-validation²¹, obtaining the accuracy = 0.92 and Cohen’s $\kappa = 0.81$.

Let us now see the ramifications of the author match score threshold μ on the possible use of the dataset. Figure 3b presents the fraction of preprints that were eventually published as journal articles f versus the year of preprint submission. The first series, marked by triangles, disregards the gray zone papers, taking into account only records marked as “published”, which leads to a spectacular decrease in f by over 30% between years 2016 and 2022. However, if we take into account that all gray zone records with author match score exceeding $\mu = 0.47$ (obtained as threshold from the logistic regression analysis) can be regarded as journal-published ones, the decrease does not exceed 11%.

Usage Notes

PreprintToPaper dataset can be used in a variety of studies: one of the principal ideas is connected to the issue of information overload²³ in academia – a notion that is widely understood²⁴ (e.g., *paper tsunami*²⁵) but for which no measures have been defined yet. The very construction of our dataset (i.e., a set of preprints submitted before / during the outbreak of the COVID-19 pandemic), along with the information about the period between the preprint submission, can help find the patterns of preprint-to-paper transformations under the stress exerted by the pandemic²⁵. Similarly, the characteristics of “preprints posted on preprint server” – treated in a manner analogous to uncited papers^{26,27} – can be examined using the provided dataset. These include stylometric factors: title and abstract length, text complexity (e.g., Herdan’s measure²⁸), readability scores (e.g., Gunning FOG index²⁹) as well as text sentiment, number of authors, and other that have been shown to play role in the previous studies linking textual factors and paper popularity^{30–32}.

Data availability

The dataset supporting this study has been stored in Zenodo and is openly accessible at <https://doi.org/10.5281/zenodo.17992421>. Further details about the dataset contents are provided in the **Data Record** section of this manuscript.

Code availability

All custom scripts used in this study are publicly available in the GitHub repository: [GitHub \(https://github.com/badalovafidan/preprint_to_paper_dataset\)](https://github.com/badalovafidan/preprint_to_paper_dataset).

The principal code was developed in Python 3.9+. The main external libraries used are **pandas**, **requests**, and **numpy**. Other libraries (**unicodedata**, **re**, **difflib**, **datetime**, **time**) are included in the Python standard library. All packages required for use are listed in the **requirements.txt** file.

The code consists of scripts for different stages: downloading metadata from the bioRxiv API (*download_biorxiv.py*), selecting the first and last version and creating the pivot structure (*filter_1st_and_last_version.py*, *pivot_biorxiv_data.py*), creating additional version-history subset (*create_version_history.py*), getting the corresponding journal article metadata from the Crossref API (*download_crossref.py*), categorizing preprints (*categorize_preprints.py*), detecting “Gray Zone” cases (*find_missing_links_crossref.py*), verifying by author match (*verify_missing_links_by_author_match.py*), standardizing dates and calculating differences (*standardize_crossref_dates.py*, *merge_publication_dates.py*, *calculate_date_difference.py*).

The code for carrying out technical validation and for creating plots shown in Fig. 3 is developed in R 4.5.1³³ using the following external libraries (R packages): **tidyverse**³⁴, **cowplot**³⁵, **psych**³⁶, and **caret**³⁷. The code consists of a single script *technical_validation.R*.

There are no restrictions on obtaining and using the code.

Received: 13 October 2025; Accepted: 9 February 2026;

Published online: 24 February 2026

References

- Fraser, N. *et al.* The evolving role of preprints in the dissemination of COVID-19 research and their impact on the science communication landscape. *PLOS Biology* **19**, e3000959, <https://doi.org/10.1371/journal.pbio.3000959> (2021).
- Fraser, N., Mayr, P. & Peters, I. Motivations, concerns and selection biases when posting preprints: a survey of bioRxiv authors. *PLoS ONE* **17**, <https://doi.org/10.1371/journal.pone.0274441> (2022).
- Badalova, F., Mayr, P. & Sienkiewicz, J. PreprintToPaper dataset: Connecting bioRxiv preprints with journal publications, <https://doi.org/10.5281/zenodo.17992421> (2025).
- Anderson, K. R. bioRxiv: Trends and analysis of five years of preprints. *Learned Publishing* **33**, 104–109, <https://doi.org/10.1002/leap.1265> (2020).
- Cabanac, G., Oikonomidi, T. & Boutron, I. Day-to-day discovery of preprint–publication links. *Scientometrics* **126**, 5285–5304, <https://doi.org/10.1007/s11192-021-03900-7> (2021).
- Sever, R. *et al.* bioRxiv: the preprint server for biology, <https://doi.org/10.1101/833400> (2019).
- Abdill, R. J. & Blekhman, R. Tracking the popularity and outcomes of all bioRxiv preprints. *eLife* **8**, e45133, <https://doi.org/10.7554/eLife.45133> (2019).
- Aviv-Reuven, S. & Rosenfeld, A. Publication patterns’ changes due to the COVID-19 pandemic: a longitudinal and short-term scientometric analysis. *Scientometrics* **126**, 6761–6784, <https://doi.org/10.1007/s11192-021-04059-x> (2021).
- Strcic, J. *et al.* Open data and data sharing in articles about COVID-19 published in preprint servers medRxiv and bioRxiv. *Scientometrics* **127**, 2791–2802, <https://doi.org/10.1007/s11192-022-04346-1> (2022).
- Grech, V. Publish or perish, information overload, and journal impact factors - A conflicting tripod of forces. *Saudi Journal of Anaesthesia* **16**, 204, https://doi.org/10.4103/sja.sja_632_21 (2022).
- Abdill, R. J., Adamowicz, E. M. & Blekhman, R. International authorship and collaboration across bioRxiv preprints. *eLife* **9**, e58496, <https://doi.org/10.7554/eLife.58496> (2020).
- Fraser, N., Momeni, F., Mayr, P. & Peters, I. The relationship between bioRxiv preprints, citations and altmetrics. *Quantitative Science Studies* **1**, 618–638, https://doi.org/10.1162/qss_a_00043 (2020).
- Abdill, R. J. & Blekhman, R. Meta-research: Tracking the popularity and outcomes of all bioRxiv preprints. *eLife* **8**, e45133, <https://doi.org/10.7554/eLife.45133> (2019).
- Abdill, R. J. & Blekhman, R. Complete Rvivist dataset of scraped biology preprint data, <https://doi.org/10.5281/ZENODO.2529922> (2023).
- Fraser, N. & Kramer, B. covid19_preprints, <https://doi.org/10.6084/M9.FIGSHARE.12033672.V58> (2021).
- Levchenko, M., Parkin, M., McEntyre, J. & Harrison, M. Enabling preprint discovery, evaluation, and analysis with europe PMC. *PLOS ONE* **19**, e0303005, <https://doi.org/10.1371/journal.pone.0303005> (2024).
- Eckmann, P. & Bandrowski, A. Preprintmatch: A tool for preprint to publication detection shows global inequities in scientific publication. *PLOS ONE* **18**, e0281659, <https://doi.org/10.1371/journal.pone.0281659> (2023).
- Tkaczyk, D. Crossref relationships between preprints and journal articles, <https://doi.org/10.5281/zenodo.15124417> (2025).
- Cohen, J. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* **20**, 37–46, <https://doi.org/10.1177/001316446002000104> (1960).
- McHugh, M. L. Interrater reliability: the kappa statistic. *Biochemia Medica* **22**, 276–282, <https://doi.org/10.11613/BM.2012.031> (2012).
- Hastie, T., Tibshirani, R. & Friedman, J. *The Elements of Statistical Learning*. Springer Series in Statistics (Springer New York Inc., New York, NY, USA, 2001).
- Ku, H. H. Notes on the use of propagation of error formulas, <https://doi.org/10.6028/NBS.RPT.9011> (1966).
- Holyst, J. A. *et al.* Protect our environment from information overload. *Nature Human Behaviour* **8**, 402–403, <https://doi.org/10.1038/s41562-024-01833-8> (2024).
- Landhuis, E. Scientific literature: Information overload. *Nature* **535**, 457–458, <https://doi.org/10.1038/nj7612-457a> (2016).
- Brainard, J. New tools aim to tame pandemic paper tsunamis. *Science* **368**, 924–925, <https://doi.org/10.1126/science.368.6494.924> (2020).
- Kozłowski, D., Andersen, J. P. & Larivière, V. The decrease in uncited articles and its effect on the concentration of citations. *Journal of the Association for Information Science and Technology* **75**, 188–197, <https://doi.org/10.1002/asi.24852> (2024).
- Noorden, R. V. The science that’s never been cited. *Nature* **552**, 162–164, <https://doi.org/10.1038/d41586-017-08404-0> (2017).
- Herdan, G. *Language as choice and chance* (Springer Berlin Heidelberg, Berlin, Germany, 1960).
- Gunning, R. *The technique of clear writing* (McGraw-Hill, New York, NY, 1952).
- Didegah, F. & Thelwall, M. Which factors help authors produce the highest impact research? collaboration, journal and document properties. *J. Informetr.* **7**, 861–873, <https://doi.org/10.1016/j.joi.2013.08.006> (2013).
- Onodera, N. & Yoshikane, F. Factors affecting citation rates of research articles. *J. Am. Soc. Inform. Sci.* **66**, 739–764, <https://doi.org/10.1002/asi.23209> (2015).

32. Sienkiewicz, J. & Altmann, E. G. Impact of lexical and sentiment factors on the popularity of scientific papers. *R. Soc. Open Sci.* **3**, 160140, <https://doi.org/10.1098/rsos.160140> (2016).
33. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria (2025).
34. Wickham, H. *et al.* Welcome to the tidyverse. *Journal of Open Source Software* **4**, 1686, <https://doi.org/10.21105/joss.01686> (2019).
35. Wilke, C. O. *cowplot: Streamlined Plot Theme and Plot Annotations for 'ggplot2'* R package version 1.2.0., <https://doi.org/10.32614/CRAN.package.cowplot> (2025).
36. William Revelle. *psych: Procedures for Psychological, Psychometric, and Personality Research*. Northwestern University, Evanston, Illinois R package version 2.5.6., <https://doi.org/10.32614/CRAN.package.psych> (2025).
37. Kuhn, M. Building predictive models in r using the caret package. *Journal of Statistical Software* **28**, 1–26, <https://doi.org/10.18637/jss.v028.i05> (2008).

Acknowledgements

The authors of this work were funded by the European Union under the Horizon Europe grant OMINO - Overcoming Multilevel INformation Overload (grant number 101086321, <http://ominoproject.eu>). Views and opinions expressed are those of the authors alone and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the European Research Executive Agency can be held responsible for them. F.B. also acknowledges support from Deutsche Forschungsgemeinschaft (DFG) under grant number MA 3964/15-3, the SocioHub project. J.S. also acknowledges support from the Polish Ministry of Education and Science under the programme entitled International Co-Financed Projects and from a special grant for leaders of European projects coordinated by the Warsaw University of Technology.

Author contributions

P.M. conceived the idea, F.B. wrote the code, gathered the data, F.B. and J.S. annotated, validated, and analyzed the data. All authors drafted the first version of the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to F.B.

Reprints and permissions information is available at www.nature.com/reprints.

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



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026