We Need Improved Data Curation and Attribution in AI for Scientific Discovery

Mara Graziani[1,2], Antonio Foncubierta[1], Dimitrios Christofidellis[1], Irina Espejo-Morales[1], Malina Molnar[1,2], Marvin Alberts[1,2], Matteo Manica[1], and Jannis Born[1]

[1]IBM Research Europe, Zürich, Switzerland
[2]NCCR Catalysis, Switzerland

Abstract

As the interplay between human-generated and synthetic data evolves, new challenges arise in scientific discovery concerning the integrity of the data and the stability of the models. In this work, we examine the role of synthetic data as opposed to that of real experimental data for scientific research. Our analyses indicate that nearly three-quarters of experimental datasets available on open-access platforms have relatively low adoption rates, opening new opportunities to enhance their discoverability and usability by automated methods. Additionally, we observe an increasing difficulty in distinguishing synthetic from real experimental data. We propose supplementing ongoing efforts in automating synthetic data detection by increasing the focus on watermarking real experimental data, thereby strengthening data traceability and integrity. Our estimates suggest that watermarking even less than half of the real world data generated annually could help sustain model robustness, while promoting a balanced integration of synthetic and human-generated content.

1. Introduction

The improvements seen in generative models during the past decade have transformed the way users interact with digital technologies (Pera & Aiello, 2024), together with the role of new data sources in the training of machine learning models. Synthetic data is generated at training time to enhance the datasets, filling gaps in data availability, reducing the validation error of machine learning models and compensating for the under-representation of rare, unlikely events (Sudalairaj et al., 2024; Ge et al., 2024; Taloni et al., 2023). The generated content is becoming increasingly realistic and semi-synthetic content is already widespread in the form of generated image or video captions (Wang et al.; Mizrahi et al., 2023; Grauman et al., 2022) and as text co-edited with generative AI (Astarita et al., 2024). While this demonstrates the high quality of the existing models, it also raises concerns about the reliability and traceability of data, posing risks to the stability of future model trainings (Shumailov et al., 2024; Dey & Donoho, 2024; Dohmatob et al., 2024; Kazdan et al., 2024) and to the factual accuracy of potential discoveries, especially in domains such as media integrity (Vosoughi et al., 2018; Lazer et al., 2018) and scientific research (Cyranoski, 2017; Sheridan, 2007; Hopf et al., 2019). The hypothesis that the availability of human-generated public text data may become a limiting factor for scaling model training in the coming decade (Villalobos et al., 2024) further strengthens the idea that we will increasingly depend on the existing fraction of human-generated data (Dohmatob et al., 2024; Kazdan et al., 2024; Ashok & May, 2024; Dey & Donoho, 2024). It is, therefore, increasingly important to reinforce the FAIR principles of dataset Findability, Accessibility, Interoperability and Reusability for the sharing of data (Hodson et al., 2018; Marsolek et al., 2023; Chen et al., 2019).

In such context, we analyze the trends of data releases in data sharing platforms such as HuggingFace (HF) datasets and Zenodo, evaluating the trend and the relative incidence of synthetic data in textual form or other modalities[1]. We draw enough evidence that the generation of synthetic data has increased over the past two years, but that it has not been sufficiently well documented. The prevalence of AI-favored terms in scientific pre-prints and publications indicates a

---

[1]Data repositories as accessed on Jan, 2025. `https://huggingface.co/` and `https://zenodo.org`.
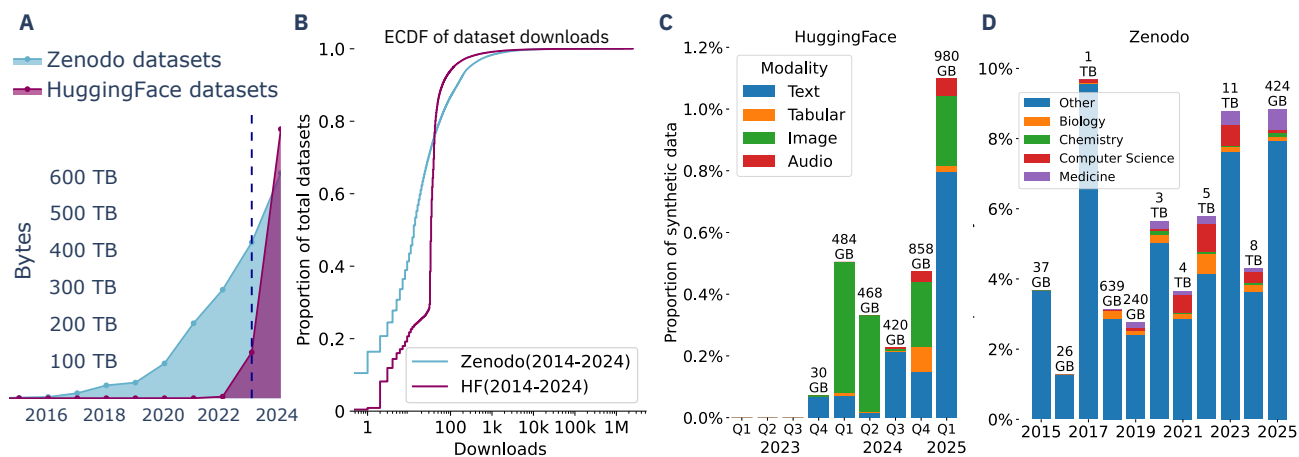
Figure 1. **Analysis of dataset uploads, adoption rates, and the percentage of synthetic content in data-sharing platforms.** A. Trends in dataset uploads, showing the number of bytes uploaded over the past decade (2014-2024) B. Dataset adoption rates in 2024, represented as the cumulative distribution function (ECDF) of dataset downloads. This shows that 87% of Zenodo datasets and 83% of HF dataset have fewer than 100 downloads. C and D. Estimates of the proportion of synthetic data in C. HF datasets and D. Zenodo datasets, reported as the fraction of synthetic bytes compared to the total uploaded bytes.

shifting linguistic landscape, presumably influenced by generative models and AI-assisted writing tools. Distinguishing between human-written and AI-assisted text remains a challenge that requires further research. For other human-generated datasets, we observe a systematic under-utilization of roughly three-quarters of online data repositories, which could be reconsidered as valuable sources of fully human generated content if their quality were improved to meet machine learning-ready standards. Therefore, a shift of focus is needed towards improving the curation and traceability of human and synthetic data, which would consequently enhance the reliability and robustness of the models. Two proactive approaches can be derived from our observations:

(1) Automation of dataset curation: Leveraging agentic workflows for labour-intensive metadata annotation to keep up with the rapid growth of research data.

(2) Person authentication: Ensuring the release and traceability of high-quality human-generated data by watermarking human-generated content.

2. Results

2.1. Datasets grow unused

HuggingFace's vision is centered around democratizing access to artificial intelligence and promoting transparency in the development and use of AI models, and it is currently one of the largest platforms for sharing and collaborating on ML models and datasets. Built and operated by CERN and OpenAIRE, Zenodo is a centralized data storage with the goal of ensuring that everyone can join in Open Science, sharing a powerful vision that aims at simplifying the upload, versioning and tracking of data. The combined efforts of these two platforms have ensured the sharing of over four million records, for slightly more than one petabyte of data (Tamarit, 2024). Figure 1 illustrates a breakdown of the data upload trends of the two platforms over the past decade (2014-2024). The number of bytes released over the past two years has seen an important increase, exponential for HF in the number of bytes available to train models. At the moment of writing this article, we estimate nearly one petabyte (10^{15} bytes) of data available only on HF datasets, which translate to approximately 300 trillion (3×10^{14}) tokens when we consider only images and text data[2]. Despite the massive dataset sizes, we observe that the dataset uptake rate has been small for most of the datasets in the two platforms. In 2024, at least three quarters of all datasets had less than hundred downloads (i.e. 83% of HF datasets and 87% of Zenodo). Datasets with high-quality curation (Koch et al., 2021) and clear usage instructions (Yang et al., 2024) tend to be more frequently adopted by researchers, highlighting the importance of structured data documentation. If we look at the most popular datasets in Zenodo, only 0.18% of all

---

[2]The approximation follows the estimates of byte-to-token (4 bytes per token) and image-to-token (34 tokens per image) conversion in Villalobos et al. (2024)
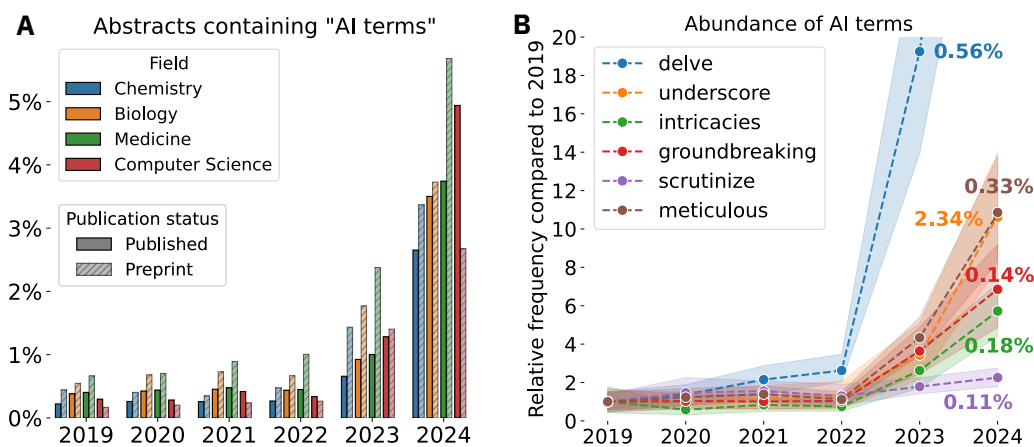
Figure 2. **Frequency of AI-favored terms in scientific paper abstracts.** A. Percentage of abstracts containing at least one of six common AI-favored words: *delve*, *underscore*, *intricacies*, *groundbreaking*, *scrutinize* and *meticulous*. Results are categorized by field and publication status, based on an analysis of the metadata from 1.5M preprints from arXiv, bioRxiv, chemRxiv and medRxiv, as well as 5.5M published papers. B. Breakdown of the relative frequency of each word, compared to that observed in the year 2019.

datasets have been downloaded more than 10K times. Additionally, we observe that the creation-to-publication time is large for Zenodo records (cf. Appendix Figure 5). When retrospective uploads of datasets are excluded from the analysis, one to three years pass on average from the record creation to its actual publication date. While the reasons for these delays are not fully understood, automation presents an opportunity to streamline metadata processing and improve data readiness. These methods could conduct extensive curation, ultimately enhancing data tagging and description practices at the time of publication.

## 2.2. Can we distinguish synthetic from real data?

Figure 1B. and C. show the amount of synthetic data uploaded on HF and Zenodo over the years, measured in bytes to ease comparison across datasets and data types. Although clear tagging of synthetic content was not consistently applied across all dataset entries, leading to some uncertainty in the estimates, our results indicate that the fraction of synthetic content is still negligible in both HF and Zenodo with 0.46% and 1.84% of the entire platforms and the records, respectively. We expect this quantity to steadily increase over time, both as a response to growing demand of LLMs (Villalobos et al., 2024) and as a natural consequence of generative models becoming increasingly more performant and accessible (Sudalairaj et al., 2024; Ge et al., 2024). This hypothesis is in line with the increase of synthetic content uploads observed in both platforms, which for HF datasets, in particular, has already attained half the amount of bytes seen in 2024 during these first days of 2025 (as of January 24th, 2025). These results likely underestimate the actual presence of synthetic data, as the use of dedicated tags is not yet strictly enforced on either platform, and their adoption depends on the culture and data curation standards of each community. The absence of explicit tagging makes it more difficult to distinguish synthetic data, emphasizing the need for improved documentation and attribution standards (cf. Appendix Figure 8).

For text, the presence of AI-generated or co-generated text in scientific papers could have implications for research integrity. Scientific journals provide guidelines on AI-generated text (Thorp, 2023), and studies have already shown some adoption of LLMs as writing aids (Astarita et al., 2024). To that end, we analyze the extent to which LLM-abundant terms[3] found their way into scientific literature. We found an exponential increase in all research fields since the launch of ChatGPT in November 30, 2022, as seen from Figure 8. Notably, the distribution shift was first visible on preprint servers (chemRxiv, bioRxiv, medRxiv and arXiv) but since 2024, the same trend appears in academic journals across all major scientific fields. Among the six terms, *delve* showed by far the strongest trend and *scrutinize* was the only one without a clear increase (for word-specific plots see Appendix Figure 6). Interestingly, while we see a strong increase in LLM-abundant terms in abstracts, we also observed a mild decrease of terms that are frequently associated to human writing such as *show*, *using*, *based* and more (cf. Appendix Figure 7). While it is unlikely for researchers to solely rely on AI-generated text to write

---

[3]We picked the six words with highest frequency-per-token difference between AI and human text (Astarita et al., 2024): *delve*, *underscore*, *intricacies*, *groundbreaking*, *scrutinize*, *meticulous*.

Table 1. **Classifying real versus synthetic data.** Performance of different generative models across multiple scientific modalities measured as the balanced accuracy score for distinguishing real versus synthetic data. Data sizes range from 2K to 2.3M samples, with results reported as mean $\pm$ standard deviation from 5-fold cross-validation. Low accuracy means that it is challenging to distinguish real from synthetic samples.

MODALITY	GENERATIVE MODEL	DATA	DATA SIZE	BAL. ACCURACY $\pm$ STD
SCIENTIFIC TEXT	TEXT+CHEM T5 CHRISTOFIDELLIS ET AL. (2023)	CHEBI-20 EDWARDS ET AL. (2021)	2.3M	$65.9_{\pm 0.01}$
MOLECULES	GCPN YOU ET AL. (2018)	ZINC-250K STERLING & IRWIN (2015)	250K	$78.8_{\pm 0.003}$
OMICS – SCRNA-SEQ	CMONGE HARSANYI ET AL. (2024)	SCIPLEX SRIVATSAN ET AL. (2020)	5K	$83.8_{\pm 0.02}$
INFRARED SPECTRA	4L MLP CF. APPENDIX 4	NIST-IR NIST (2010)	2K	$85.5_{\pm 5.7}$
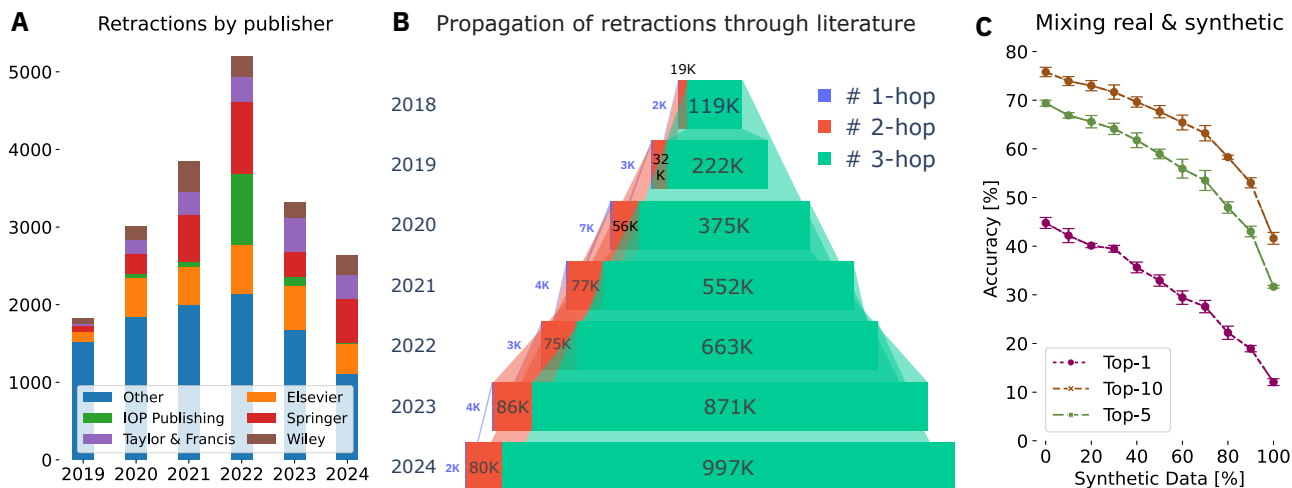
Figure 3. **Impact of inaccurate data on scientific literature and model training.** A. Number of paper retractions categorized by publisher, with retractions from Hindawi excluded from the analysis. B. Propagation effect. 1-hop: papers citing at least one retracted paper; 2-hops: papers citing at least one paper that cites a retracted paper; and 3-hops: papers that continue this chain, citing a paper that cites 2-hop paper. C. Observed performance degradation as increasing fractions of synthetic data are introduced during model training to predict molecular structure from IR spectra. Accuracy is measured by comparing the ground truth to the Top-1, -5 and -10 ranked predicted molecules. The quality of the synthetic data introduces a bias, which impacts the model's performance.

scientific papers, this result suggests that the use of AI for assisted copy editing is an option often taken by researchers. This creates semi-synthetic text, which blurs the border between real and synthetic text. Determining exactly the quantity of synthetic data in scientific research data remains a complex challenge hat requires further research on its own.

We extend the analysis of identifying synthetic against human-generated content to four abundant scientific data modalities (molecules, scientific text, single-cell transcriptomics and spectroscopy) and measure the difficulty of differentiating real and generated samples of that modality as the balanced classification accuracy of classifiers of real against synthetic data. The results in Table 1 consistently show that classifiers with high accuracy are challenging to build. The reported performance likely represents an upper bound, as it assumes that sufficient samples are available to train a classifier and that the labels for both synthetic and real samples are fully known. Additionally, our result confirms the apparent hypothesis that training the generative model on more samples makes it harder to train accurate classifiers. To identify AI-generated text, numerous attempts were made but failed (Kirchner et al., 2023) and we expect a similar trend on scientific data as generative models continue to improve.

2.3. What happens if synthetic is mistaken for real

A useful analogy can be drawn between the challenge of attributing synthetic data and the persistence of outdated or retracted research, highlighting the importance of clear data provenance. In fact, timely retraction and clear labeling of problematic research are essential to maintaining scientific integrity, as previous studies have shown that misinformation can persist even after retraction. Similarly, synthetic data could propagate if no clear labeling framework is put in place. Figure 3A. shows

Figure 4. **Estimated log-likelihood response to increasing fractions of watermarked data for human-generated content.** Computed following Kazdan et al. estimates on the relationship between real content and model performance for models with 512 tokens of context length.

the number of retracted papers from the years between 2019 and 2024[4]. Figure 3B. illustrates that citations to such papers are likely to spread in other literature works over time. This result gives insights on the actual difficulty of retroactively retrieving misleading or incorrect information once it has been published. Very much the same could happen with synthetic data if no distinction is made promptly. We further extend this analysis by considering a simulated scenario where we are unable to label varying fractions of data as synthetic. As shown in Figure 3 C., downstream task performance deteriorates, for example, when an increasing fraction of AI-generated synthetic data (generated as described in subsection 4.4) is added to the training data of a model that is performing molecular structure elucidation from infrared spectra (Priessner et al., 2024; Alberts et al., 2024; Devata et al., 2024).

2.4. What happens if we start watermarking human-generated content

Watermarking content can be a viable solution to track data provenance during training and regulate the synthetic-to-real data ratio. Along this line, the community has discussed the watermarking of synthetic content (Dathathri et al., 2024; Kirchenbauer et al., 2023; Zhao et al., 2024b), acknowledging that it may be challenging to implement robustly (Xing et al., 2025). An alternative approach, inspired by watermarking for currency, government documents or postage stamps, is that of enhancing the attribution and verification of human-generated content through standardized provenance tracking. This solution may ensure traceability and authenticity for better data integrity. Without having to go as far as the High-bandwidth Digital Content Protection (aHDCP[5]) mechanism to revoke access to non human generated content, proper attribution and trust would be supported by provenance-tracking mechanisms that attribute human generated data to the respective authors. Assuming that such a mechanism existed, Figure 4 shows the test log-likelihood evolution in response to the watermarking of various fractions of real data over two years. The obtained values are based on Kazdan et al. estimates using data sourced from HuggingFace datasets (details in Section 6). While this estimate is not highly accurate (the reported R^2 value of the predicted relationship is 0.59), it provides a useful indication of how much progress could be achieved with minimal effort, compared to the daunting task of annotating all synthetic content. Only a small fraction of watermarks would be sufficient, as opposed to content-based watermarking of all digital contents.

---

[4]Hindawi faced significant challenges due to the quality of its peer review process which has been associated to being run by paper mills and having published thousands of retractions alone (Joelving & Retraction, 2024).
[5]https://www.digital-cp.com/

3. Discussion

While the need for immediate action towards improving data curation and attribution is agreed upon (Zhao et al., 2024b; Hodson et al., 2018; Kirchenbauer et al., 2023; Grinbaum & Adomaitis, 2022), traditional data curation methods seem no longer practical due to the constantly evolving requirements of the field and the rapidly increasing data generation rates. In contrast, there is evidence that manual or semi-automated processes can be used to analyze single files (Jaimovitch-López et al., 2023; Ansari & Moosavi, 2024; Cheng et al., 2023). Sharing raw, unstructured data as-is appears to be a natural and prolific approach to data-sharing that allows researchers to focus on their core tasks, and much of the labor-intensive curation process can be automated. Agentic workflows could be developed, for instance, to capture important metadata before the data is uploaded on the platforms. Moreover, modality-specific agents can be trained to write code for data analysis, while an overseeing agent can reason across the various data reports and file types to compile a final report (Graziani et al., 2025). Following this method, entire data collections can be organized at scale, with structured datasheets being generated for each file and each collection as byproducts of the process. Integrating retrieval-augmented generation (Lewis et al., 2020) to index the reports would enable deeper reasoning on the data, facilitating the retrieval of entire datasets or individual datapoints of interest. This solution would provide researchers with a comprehensive overview of *what exists*, *where it is stored*, *what can be done* with it and *where to find* each file, information that is currently lacking.

At this point, one may argue that relying on automated methods for curating data could come at an increased risk of inducing errors and overlooking important details in the data. While it is true that automation removes traceability of errors and makes it harder to attribute inconsistencies, we see numerous counter-arguments to the risk of introducing mistakes in the curation process. To begin with, as dataset uptake rates suggest, errors in the data curation are less dramatic than no curation at all. Manual data curation is time consuming, costly and relies on highly trained scientists to meticulously gather all the relevant metadata. As a result, this practice is rarely performed and the majority of datasets are left to interpretation (Graziani et al., 2025). Moreover, the curation required by the field changes over time as the techniques evolve, with the flagging of synthetic data being a practical example of the rise of a new required field. It is impractical to retroactively update the curation of all existing datasets manually. As a result, even datasets that were extensively curated turn out to be insufficiently covered once the requirements of the field have changed (Zhao et al., 2024a). The benefits of an automated approach, in contrast would be practical and quantifiable. As datasets continue to grow in size, (with the average repository size increasing from 2.5 GB pre-2024 to 3 GB in 2024), decreasing the curation costs is a priority to ensure the scaling of the data curation coverage. On this note, agentic workflows would dramatically cut the curation costs, bringing them closer to the inference costs of small transformer models (Graziani et al., 2025). Additionally, outputs generated by these workflows can be versioned based on the models used, allowing for the release of open-source versions capable of reproducing the data analysis locally and on the platform at the time of upload. The versions could be compared across the years based on the dataset uptake by researchers, which could become a human-in-the-loop benchmark for data curation. The ability to run the process on local premises provides an additional solution for the standardized and versionable curation of private datasets. This ensures better documentation and curation even for datasets that are not publicly available, safeguarding privacy while maintaining consistency and quality in the tracking of data resources (Sheridan et al., 2021).

Watermarking real experimental data could further contribute to the detection efforts of synthetic data, ensuring the traceability and reliability of the sources. One common view in the community is the watermarking of synthetic content (Dathathri et al., 2024; Kirchenbauer et al., 2023; Zhao et al., 2024b), which can be mainly distinguished in model-based and content-based approaches. The first category, for instance, groups methods that are based on irrecognizably altering the LLM output by intervening the beam search or postprocessing the output, e.g., by enforcing patterns of character repetitions (Grinbaum & Adomaitis, 2022), for which the technical feasibility and benefits were demonstrated in a seminal paper (Kirchenbauer et al., 2023). The community also explored *content-based watermarking systems* which, in essence, build classifiers that detect AI-generated content post-hoc. Numerous academic (e.g., Ghostbusters (Verma et al., 2023)) and commercial (e.g., GPTZero) attempts have been made and some argued that reliable detectors can be built in theory (Chakraborty et al., 2024). For both approaches, however, implementational barriers remain challenging to address. Model-embedded watermarks, beyond the lack of adoption, need prior knowledge about the LLM (e.g., the tokenization) which is typically unknown for proprietary LLMs. Similarly, content-based watermarks are challenging to buld in practice because LLMs and humans may produce the same text verbatim and the more data become available the harder it becomes to maintain robustness against adversarial modifications (Hu et al., 2023). Even in the ideal case, where such systems were widely used and accepted by LLM stakeholders, incentives would grow to build a LLM that escapes such systems.

On the other end of the spectrum, we proposed a complementary approach that suggests to focus on the improved curation and attribution of *human*-generated data. Indeed, this solution, as many others, presents challenges. For the attribution part,

human actors could intentionally deceive AI-generated content as human-generated, thus rendering the watermark useless. Possible mitigation strategies could be considered for scientific data (HuggingFace, Zenodo etc), for example by requiring uploaders to link their account to their unique researcher profile (e.g., ORCID iD), or even by implementing an arXiv-alike system where new users have to be endorsed by a senior user before they can upload data. These measures could result in a system in which human generated content is watermarked by trustworthy peers or an independent authority that certifies the generation process. Nothwithstanding these challenges, increasing the efforts towards improving the traceability and documentation of human-generated data can sustain a balanced integration of human and synthetic content in the training of the models.

Methods

4. Classifier Experiments

4.1. Molecules

As molecular generative model we utilized the GCPN, a Graph Convolutional Policy Network (You et al., 2018), as available via GT4SD (Manica et al., 2023). The GCPN was trained on ZINC-250k (Irwin et al., 2012) and found to mimic well ZINC 250k in terms of distributions of physicochemical properties of generated molecules (Tadesse et al., 2023). We generated 20k molecules from the GCPN and randomly selected 20k real molecules from the ZINC-250k training data, accounting only for similar distributions in SMILES sequence length, to avoid potential shortcuts for the classifer. We then trained a DistillBERT (Sanh, 2019) classifier to distinguish real vs. generated SMILES strings in a 5-fold cross validation without any hyperparameter tuning.

4.2. Scientific Text

As scientific text generation model, we utilized Text+Chem T5 (Christofidellis et al., 2023), a multitask chemical language model which has shown compelling results at generating PubChem alike textual descriptions of molecules. The model was, among other tasks, trained on the CheBI-20 (Edwards et al., 2021) dataset, containing molecules as SMILES and corresponding textual descriptions describing their functional role. We utilized the `base-augm` checkpoint from HuggingFace to predict text descriptions for the 3k molecules from the test set. We then trained a Random Forest classifier in a 5-fold cross validation on BERT text embeddings (Devlin et al., 2019) of the generated and the ground truth textual descriptions in a 5-fold cross validation without any hyperparameter tuning.

4.3. Transcriptomics (single-cell RNA-seq)

As generative model for transcriptomcis (gene expression) we utilized the CMonge (Harsanyi et al., 2024; Driessen et al., 2024) model that exhibited state-of-the-art performance at modeling SciPlex (Srivatsan et al., 2020), a single-cell RNA-seq dataset of drug perturbation effects. This model was trained on the entirety of SciPlex (760k cells) with 5k of the cells from the drug that we picked for comparison (belinostat). We generated 1500 synthetic cell profiles, characterized by their 50 most highly variant genes, and then trained a Random Forest classifier in a 5-fold cross validation to distinguish them from 1500 real cells randomly sampled from the 5k real cells treated with belinostat.

4.4. IR

To generate synthetic IR spectra, we employed a four-layer multilayer perceptron (MLP) that reconstructs experimental IR spectra from simulated data. Our training data combined simulated IR spectra from Alberts et al. (2025) with experimental spectra from the NIST database (NIST, 2010). Both simulated and real IR spectra are represented as 1,625-dimensional vectors. ReLU activation is used in the hidden layers of the model and sigmoid activation in the output layer. The model was optimised using the ADAM optimizer (learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$) with spectral information divergence (SID) as the loss function (McGill et al., 2021). We generated 600 synthetic IR spectra and trained a Random Forest classifier to distinguish them from 600 real spectra using 5-fold cross validation. We performed this evaluation using default hyperparameters without optimization.

We pretrain a transformer model to predict the chemical structure from simulated IR spectra. While Alberts et al. (2024) exclusively used experimental data for finetuning, we explored the impact of finetuning the model on varying combinations of real and synthetic spectra. We varied the ratio of real to synthetic spectra in 10% increments, ranging from purely real
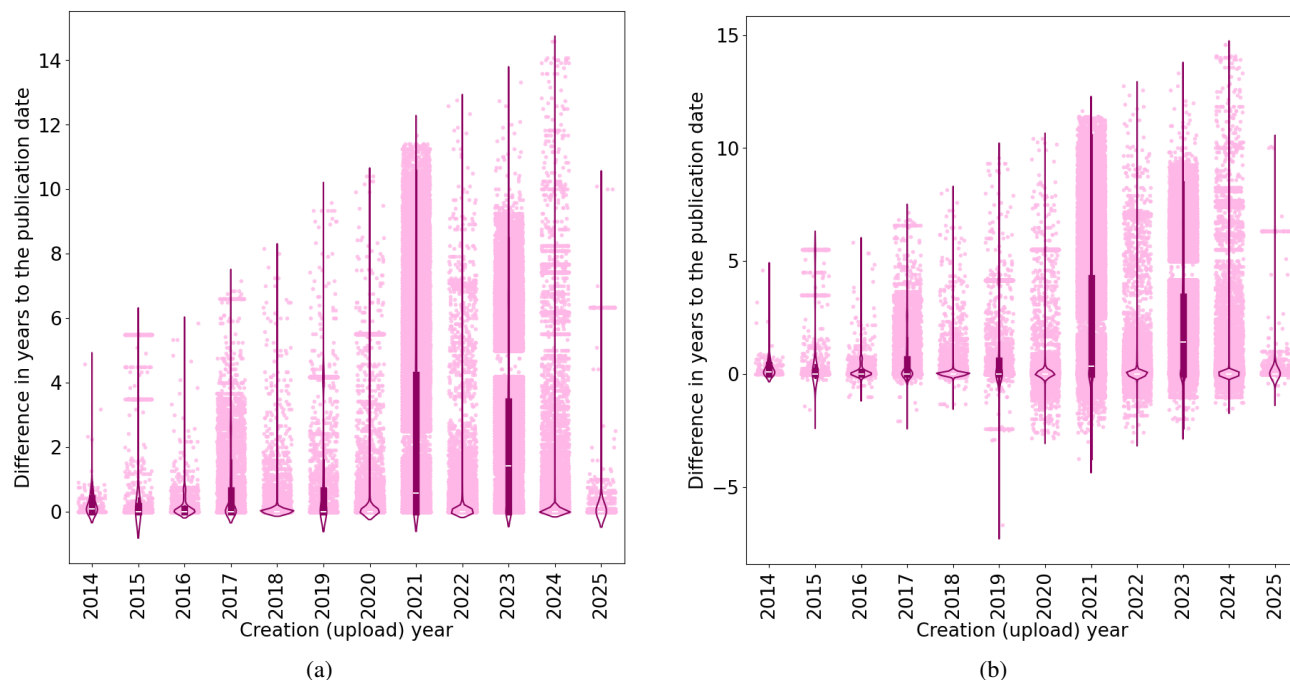
Figure 5. **Record creation to publication time.** Time in years passing from the creation date to the publication date of a record in Zenodo: (a) Absolute time (b) Overall time that takes into account retrospective uploads.

(100% real, 0% synthetic) to purely synthetic (0% real, 100% synthetic) training data. For each ratio, we employed 3-fold cross validation. The results are shown in Figure 3.

5. Zenodo Records

5.1. Data Retrieval

The Zenodo REST API (https://developers.zenodo.org) was used to retrieve the records from the platform. We have used one query to download all available datasets from 2015 to 2025 by specifying the record type and the date. A total of 362715 datasets were downloaded and 824 were deleted due to missing dates or anomalous ones where the year surpassed 2025. For the domain specific datasets, the search string included the record type and the domain (chemistry, biology, computer science, medicine). When comparing the domain specific datasets with all the rest, as in Figure **??**, the domain specific ones were deleted based on the ID from the dataset with all the records to avoid duplicates.

```
1  search_query_all = "type:dataset AND created:[{2015-01-01} TO {2025-01-09}]"
```

5.2. Creation-to-Publication Time

If we take the absolute time passed in days, the average value changes to 385 days between the original publish date of the dataset and the creation (upload) date on Zenodo. Figure 5 shows the high delay in years between the creation and the publication date, sometimes years passing until a dataset is uploaded on Zenodo. On the left side we have the absolute values, while on the right side we have the the actual values. A negative year means the dataset was not yet published, and it would be in the future.

5.3. Extended synthetic data identification

We extended our search to include domain specific words for each domain (biology, chemistry, computer science, medicine) and analyzed two cases. Figure 8 (a) shows the number of identified synthetic datasets when looking at the description and title of the record, while Figure 8 (b) shows the number of identified synthetic datasets when inspecting only on the

8

keywords provided by the users. The difference in the number of retrieved records obtained by the two approaches suggests that users tend to be less specific when assigning keywords to the data uploads, making it hard to track synthetic datasets.

To produce all of the figures, the list of base keywords used to produce the plots in **??** has been extended with domain-specific jargon that could be used to identify synthetic or simulated data. The keywords extensions are reported below.

> **Base keywords**: synthetic data, artificial data, generated data, synthetic dataset, artificial dataset, generated dataset, ai-generated, ai generated, simulated data, genai, generative, llm, synthetic

> **Chemistry**: simulated, in silico, modeled, generated, virtual experiments, quantum calculations, molecular dynamics, monte carlo, force field, ab initio, crystal prediction, reaction networks, kinetic modeling, therodynamic predictions, machine learning, data augmentation, virtual screening, latent space, surrogate data, synthetic benchmarks, no verification, artificial noise, hypothetical molecules, theoretical spectra, predicted values, DFT, density functional theory, simulat, meta dynamics, molecular simulation

> **Biology**: simulated, in silico, modeled, generated, virtual experiments, synthetic genomes, synthetic sequences, synthetic organisms, data augmentation, machine learning, computational models, hypothetical data, artificial noise, synthetic cells, synthetic benchmarks, no verification, surrogate data, gene predictions, protein folding, simulated pathways, biological models, virtual screening, training data, predicted structures, synthetic traits, simulated populations

> **Medicine**: simulated, in silico, modeled, generated, virtual trials, synthetic cohorts, synthetic patients, machine learning, data augmentation, computational models, prediction models, theoretical data, artificial noise, surrogate data, hypothetical patients, virtual screeing synthetic benchmarks, training data, no verification, simulated outcomes, hypothetical outcomes, synthetic biomarkers, imputed data, synthetic images, predicted outcomes, simulated diagnoses.

> **Computer science**: simulated, generated, synthetic data, data augmentation, machine learning, training data, test data, artificial noise, no verification, synthetic benchmarks, virtual environments, synthetic images, synthetic text, synthetic speech, synthetic audio, synthetic video, computational models, predicted data, imputed data synthetic datasets, deepfake, synthetic graphs, synthetic logs, surrogate data, latent space, augmented data.

6. Kazdan et al. log-likeklihood estimates

Following the lead of Kazdan et al., we estimate model log-likelihood using the number of real and synthetic data points used for training a model with fixed context length to 512 tokens:

$$\ell_k = \log \frac{n_{\text{real}}}{n_{\text{real}} + n_{\text{synthetic}}}$$

To extrapolate the number of real and synthetic data points, we source data from HuggingFace datasets (from late 2022 to end of 2024). We focused on text and image datasets where we leveraged tags ("ai-generated" and "synthetic") to distinguish between real and synthetic datasets. To compute the number data points at fixed context length, we calculated the total number of tokens per month using the relations from Villalobos et al.: 4B equals to 1 token (text) and 1 image equals to 34 tokens. Using a third-degree polynomial interpolation, we extrapolated the trend both for the number of real data points as well as for the total number of data points (real and synthetic combined). Finally, to simulate the different trends building Figure 4, we considered a set of scenarios where the number of real data points added each month is scaled by a fraction of the total estimate: from 0%, i.e., no watermarking or no real data added, to 100%, i.e., watermarking/confirming all the projected real data. Interestingly, this analysis shows how by starting watermarking real data, can effectively alleviate potential model collapse, despite the increasing quantity of synthetic data.

7. Conclusions, implications and limitations

It is increasingly apparent that fully tracking the creation of synthetic content may be an unrealistic goal and that the scientific field will need to find effective ways to deal with this. Even if robust detectors arise, the categorization of semi-synthetic content, which involves human oversight and AI assistance, remains an area that requires further refinement to ensure transparency in data generation. The ambiguity surrounding such cases further complicates the efforts to define a clear strategy that focuses on regulating synthetic content.

Our approach seeks to harness the benefits of synthetic data while proactively addressing its challenges. By assuming that soon most data may be synthetic or semi-synthetic, we can devise strategies to prioritize data attribution and curation,

thus helping in maintaining robust model performance. We can focus on maximizing the re-use of already existing human-generated data and on watermarking the essential portion of future human-generated data to prevent detrimental consequences. Our estimates show that prioritizing on just 40% of high-quality human-generated data can help maintain robust model performance and prevent over-reliance on synthetic content.

There are limitations in our observations. The estimates that we draw are surrounded with uncertainty. We had to infer which datasets are synthetic but there is no exact tag for the majority of datasets that we analyzed, and estimates can vary based on the search option. Moreover, none of our projections corrects for data duplicates, although these have been estimated to account for around 30% of online data (Villalobos et al., 2024).

Impact Statement

No ethical review is required for this work. There are no damaging societal consequences associated with the paper. This work is key for data regulation discussions and for the machine learning argument on watermarks for synthetic data and the risk of model collapse. We raise awareness on the scientific perspective in these matters and advocate for alternative views.

Acknowledgements

References

Alberts, M., Laino, T., and Vaucher, A. C. Leveraging infrared spectroscopy for automated structure elucidation. *Communications Chemistry*, 7(1):1–11, 2024.

Alberts, M., Schilter, O., Zipoli, F., Hartrampf, N., and Laino, T. Unraveling molecular structure: A multimodal spectroscopic dataset for chemistry. *Advances in Neural Information Processing Systems*, 37:125780–125808, 2025.

Ansari, M. and Moosavi, S. M. Agent-based learning of materials datasets from the scientific literature. *Digital Discovery*, 2024.

Ashok, D. and May, J. A little human data goes a long way. *arXiv preprint arXiv:2410.13098*, 2024.

Astarita, S., Kruk, S., Reerink, J., and Gómez, P. Delving into the Utilisation of ChatGPT in Scientific Publications in Astronomy. In *Proceedings of SPAICE2024: The First Joint European Space Agency / IAA Conference on AI in and for Space*, pp. 241–246, October 2024.

Born, J. and Manica, M. Trends in deep learning for property-driven drug design. *Current Medicinal Chemistry*, 28(38): 7862–7886, 2021.

Chakraborty, S., Bedi, A., Zhu, S., An, B., Manocha, D., and Huang, F. Position: On the possibilities of AI-generated text detection. In *Forty-first International Conference on Machine Learning*, 2024.

Chen, X., Dallmeier-Tiessen, S., Dasler, R., Feger, S., Fokianos, P., Gonzalez, J. B., Hirvonsalo, H., Kousidis, D., Lavasa, A., Mele, S., et al. Open is not enough. *Nature Physics*, 15(2):113–119, 2019.

Cheng, L., Li, X., and Bing, L. Is gpt-4 a good data analyst? In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 9496–9514, 2023.

Christofidellis, D., Giannone, G., Born, J., Winther, O., Laino, T., and Manica, M. Unifying molecular and textual representations via multi-task language modelling. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 6140–6157. PMLR, 2023.

Cyranoski, D. The secret war against counterfeit science. *Nature*, 545(7653):148–150, 2017.

Dathathri, S., See, A., Ghaisas, S., Huang, P.-S., McAdam, R., Welbl, J., Bachani, V., Kaskasoli, A., Stanforth, R., Matejovicova, T., et al. Scalable watermarking for identifying large language model outputs. *Nature*, 634(8035):818–823, 2024.

Devata, S., Sridharan, B., Mehta, S., Pathak, Y., Laghuvarapu, S., Varma, G., and Deva Priyakumar, U. DeepSPInN – deep reinforcement learning for molecular structure prediction from infrared and 13 C NMR spectra. *Digital Discovery*, 3(4): 818–829, 2024.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, 2019.

Dey, A. and Donoho, D. Universality of the $\pi^{2/6}$ pathway in avoiding model collapse. *arXiv preprint arXiv:2410.22812*, 2024.

Dohmatob, E., Feng, Y., Subramonian, A., and Kempe, J. Strong model collapse. *arXiv preprint arXiv:2410.04840*, 2024.

Driessen, A., Born, J., Rueda, R. C., Reddy, S. T., and Rapsomaniki, M. Modeling car response at the single-cell level using conditional ot. In *NeurIPS 2024 Workshop on AI for New Drug Modalities*, 2024. Spotlight talk.

Edwards, C., Zhai, C., and Ji, H. Text2mol: Cross-modal molecule retrieval with natural language queries. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 595–607, 2021.

Ge, T., Chan, X., Wang, X., Yu, D., Mi, H., and Yu, D. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*, 2024.

Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18995–19012, 2022.

Graziani, M., Molnar, M., Morales, I. E., Cadow-Gossweiler, J., and Laino, T. Making sense of data in the wild: Data analysis automation at scale. *arXiv preprint arXiv:2502.15718*, 2025.

Grinbaum, A. and Adomaitis, L. The ethical need for watermarks in machine-generated language. *AI and Ethics*, 2022.

Harsanyi, B., Rapsomaniki, M., and Born, J. Learning drug perturbations via conditional map estimators. In *ICLR 2024 Workshop on Machine Learning for Genomics Explorations*, 2024.

Hodson, S., Jones, S., Collins, S., Genova, F., Harrower, N., Laaksonen, L., Mietchen, D., Petrauskaité, R., and Wittenburg, P. Turning fair into reality. *Final report and action plan from the EC expert group on FAIR data. DOI*, 10:1524, 2018.

Hopf, H., Krief, A., Mehta, G., and Matlin, S. A. Fake science and the knowledge crisis: ignorance can be fatal. *Royal Society open science*, 6(5):190161, 2019.

Hu, X., Chen, P.-Y., and Ho, T.-Y. Radar: Robust ai-text detection via adversarial learning. *Advances in Neural Information Processing Systems*, 36:15077–15095, 2023.

Irwin, J. J., Sterling, T., Mysinger, M. M., Bolstad, E. S., and Coleman, R. G. Zinc: a free tool to discover chemistry for biology. *Journal of chemical information and modeling*, 52(7):1757–1768, 2012.

Jaimovitch-López, G., Ferri, C., Hernández-Orallo, J., Martínez-Plumed, F., and Ramírez-Quintana, M. J. Can language models automate data wrangling? *Machine Learning*, 112(6):2053–2082, 2023.

Joelving, F. and Retraction, W. Paper trail. *Science*, 383(6680):252–255, 2024.

Kazdan, J., Schaeffer, R., Dey, A., Gerstgrasser, M., Rafailov, R., Donoho, D. L., and Koyejo, S. Collapse or thrive? perils and promises of synthetic data in a self-generating world. *arXiv preprint arXiv:2410.16713*, 2024.

Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers, I., and Goldstein, T. A watermark for large language models. In *International Conference on Machine Learning*, pp. 17061–17084. PMLR, 2023.

Kirchner, J. H., Ahmad, L., Aaronson, S., and Leike, J. New ai classifier for indicating ai-written text. *OpenAI. openai. com/blog/new-ai-classifier-for-indicating-ai-written-text*, 2023.

Koch, B., Denton, E., Hanna, A., and Foster, J. G. Reduced, reused and recycled: The life of a dataset in machine learning research. *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, 1, 2021.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., et al. The science of fake news. *Science*, 359(6380):1094–1096, 2018.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020.

Manica, M., Born, J., Cadow, J., Christofidellis, D., Dave, A., Clarke, D., Teukam, Y. G. N., Giannone, G., Hoffman, S. C., Buchan, M., et al. Accelerating material design with the generative toolkit for scientific discovery. *npj Computational Materials*, 9(1):69, 2023.

Marsolek, W., Wright, S. J., Luong, H., Braxton, S. M., Carlson, J., and Lafferty-Hess, S. Understanding the value of curation: A survey of researcher perspectives of data curation services from six us institutions. *PLoS One*, 2023. doi: 10.1371/journal.pone.0293534.

McGill, C., Forsuelo, M., Guan, Y., and Green, W. H. Predicting Infrared Spectra with Message Passing Neural Networks. *Journal of Chemical Information and Modeling*, 61(6):2594–2609, 2021.

Mizrahi, D., Bachmann, R., Kar, O., Yeo, T., Gao, M., Dehghan, A., and Zamir, A. 4m: Massively multimodal masked modeling. *Advances in Neural Information Processing Systems*, 36:58363–58408, 2023.

NIST. National Institute of Standards and Technology / EPA Gas-Phase Infrared Database, 2010. URL nist.gov/srd/nist-standard-reference-database-35. Accessed: 2025-01-10.

Pera, A. and Aiello, L. M. Shifting climates: Climate change communication from youtube to tiktok. In *Proceedings of the 16th ACM Web Science Conference*, pp. 376–381, 2024.

Priessner, M., Lewis, R., Janet, J. P., Lemurell, I., Johansson, M., Goodman, J., and Tomberg, A. Enhancing Molecular Structure Elucidation: MultiModalTransformer for both simulated and experimental spectra, 2024.

Sanh, V. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.

Sheridan, C. Bad medicine. *Nature biotechnology*, 25(7):707–710, 2007.

Sheridan, H., Dellureficio, A. J., Ratajeski, M. A., Mannheimer, S., and Wheeler, T. R. Data curation through catalogs: A repository-independent model for data discovery. *Journal of eScience Librarianship*, 2021. URL https://doi.org/10.7191/jeslib.2021.1203.

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., and Gal, Y. Ai models collapse when trained on recursively generated data. *Nature*, 631(8022):755–759, 2024.

Srivatsan, S. R., McFaline-Figueroa, J. L., Ramani, V., Saunders, L., Cao, J., Packer, J., Pliner, H. A., Jackson, D. L., Daza, R. M., Christiansen, L., et al. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367(6473): 45–51, 2020.

Sterling, T. and Irwin, J. J. Zinc 15–ligand discovery for everyone. *Journal of chemical information and modeling*, 55(11): 2324–2337, 2015.

Sudalairaj, S., Bhandwaldar, A., Pareja, A., Xu, K., Cox, D. D., and Srivastava, A. Lab: Large-scale alignment for chatbots. *arXiv preprint arXiv:2403.01081*, 2024.

Tadesse, G. A., Born, J., Cintas, C., Ogallo, W., Zubarev, D., Manica, M., and Weldemariam, K. Domain-agnostic and multi-level evaluation of generative models. *arXiv preprint arXiv:2301.08750*, 2023. Workshop at International Conference on Knowledge Discovery and Data Mining.

Taloni, A., Scorcia, V., and Giannaccare, G. Large language model advanced data analysis abuse to create a fake data set in medical research. *JAMA ophthalmology*, 141(12):1174–1175, 2023.

Tamarit, P. Digital Repositories and preservation at Harvard and CERN [Slides]. `https://indico.cern.ch/event/1420862/contributions/6002171/attachments/2877240/5039086/zenodo-2024-06-13.pdf`, 2024. Accessed: 2024-11-30.

Thorp, H. H. Chatgpt is fun, but not an author, 2023.

Verma, V., Fleisig, E., Tomlin, N., and Klein, D. Ghostbuster: Detecting text ghostwritten by large language models. *arXiv preprint arXiv:2305.15047*, 2023.

Villalobos, P., Ho, A., Sevilla, J., Besiroglu, T., Heim, L., and Hobbhahn, M. Position: Will we run out of data? limits of llm scaling based on human-generated data. In *Forty-first International Conference on Machine Learning*, 2024.

Vosoughi, S., Roy, D., and Aral, S. The spread of true and false news online. *science*, 359(6380):1146–1151, 2018.

Wang, Y., He, Y., Li, Y., Li, K., Yu, J., Ma, X., Li, X., Chen, G., Chen, X., Wang, Y., et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. In *The Twelfth International Conference on Learning Representations*.

Xing, X., Shi, F., Huang, J., Wu, Y., Nan, Y., Zhang, S., Fang, Y., Roberts, M., Schönlieb, C.-B., Del Ser, J., et al. On the caveats of ai autophagy. *Nature Machine Intelligence*, pp. 1–9, 2025.

Yang, X., Liang, W., and Zou, J. Navigating dataset documentations in ai: A large-scale analysis of dataset cards on huggingface. In *The Twelfth International Conference on Learning Representations*, 2024.

You, J., Liu, B., Ying, Z., Pande, V., and Leskovec, J. Graph convolutional policy network for goal-directed molecular graph generation. *Advances in neural information processing systems*, 31, 2018.

Zhao, D., Andrews, J. T., Papakyriakopoulos, O., and Xiang, A. Position: measure dataset diversity, don't just claim it. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24, 2024a.

Zhao, X., Gunn, S., Christ, M., Fairoze, J., Fabrega, A., Carlini, N., Garg, S., Hong, S., Nasr, M., Tramer, F., et al. Sok: Watermarking for ai-generated content. *arXiv preprint arXiv:2411.18479*, 2024b.

Additional results

We provide additional results that complement our main findings in the paper.

A. Paper keyword analysis

Figure 6 shows the occurrences of the LLM-related keywords detailed by the individual words. "Scrutinize" is the only word that shows a constant uptake over the years from 2019 to 2024. Figure 7 illustrates the relative change with respect to the previous year of human abundant terms. All the analyses were done with *paperscraper* (Born & Manica, 2021).
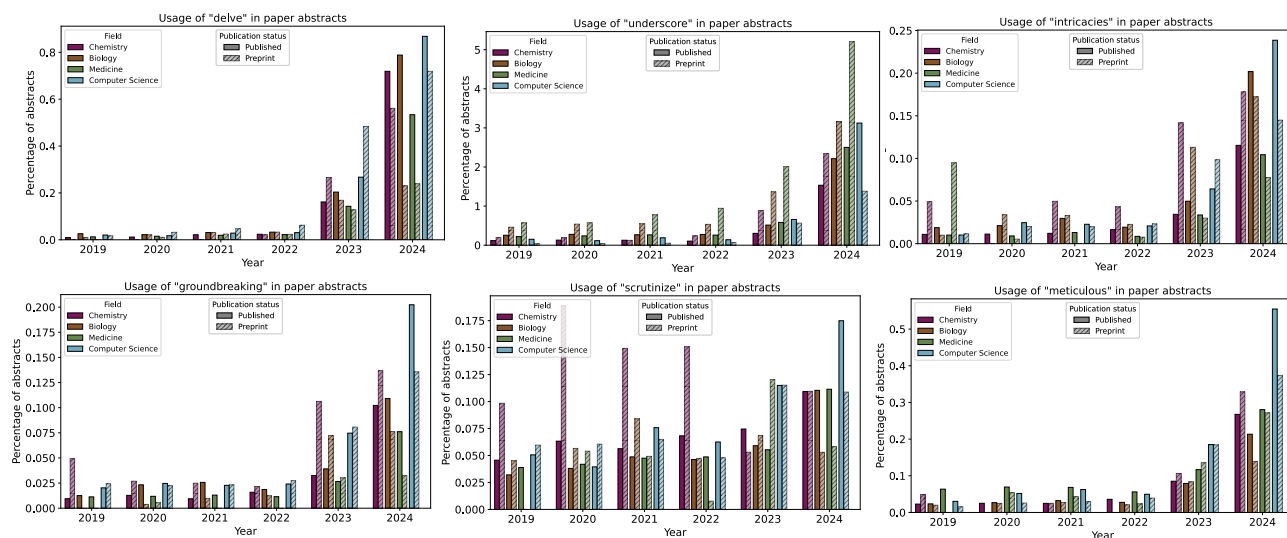

Figure 6. **Occurrence of LLM-related keywords in scientific paper abstracts.** Words taken from Astarita et al. (2024).


Figure 7. **Human-abundant terms in scientific papers.** Relative change (year by year) of human-abundant terms in all scientific paper abstracts. Words taken from the analysis by Astarita et al. (2024).

### A.1. Variability of synthetic estimates based on keyword search

Figure 8. Number of Zenodo records responding to our domain-specific keyword search in the (a) description and title and (b) in the keywords provided by the users