On the management of definitions in fast-paced research – A systematic collection of Uncertainty definitions in Computer Science

Nemi Pelgrom[1,*], Mats Walldén[2] and Karl Andersson[2,3]

[1]Department of Computer Science and Media Technology, Linnaeus University
[2]Skillsta Teknik Design och Kvalitet AB
[3]DHEAR, Skövde University

Abstract

We report a meta-study of computer science literature published in 2024, focusing on how the concept of *Uncertainty* is defined and referenced across the field. While *Uncertainty* is a highly technical term with deep theoretical foundations, its recent rise in prominence—driven in part by the surge of interest in Artificial Intelligence—has not always been matched by a corresponding depth of treatment. To assess how the term is currently used, we conducted a systematic literature review of papers discussing *Uncertainty* across the broader field of Computer Science. For each relevant paper, we analysed whether a definition was provided, whether it was technical or non-technical, and whether it was properly referenced. Our findings confirm two hypotheses: (a) a substantial proportion of papers use the term "uncertainty" without offering a technical definition, and (b) many technical definitions are not properly referenced, even when they are not novel. Specifically, 74% of the papers include non-referenced technical definitions. We also conducted a focused sub-analysis of papers that mention large language models (LLMs) in the same sentence as "uncertainty". In this subset, we observed an even higher proportion of papers lacking definitions altogether, and similarly high rates of non-referenced technical definitions. We present our methodology and findings in detail and discuss their implications, particularly the risks of conceptual ambiguity in a field increasingly reliant on shared but often unstated assumptions.

Keywords

Uncertainty, Research Paradigm, Research Methodologies, Terminology

1. Introduction

Computer Science has been growing faster and wider than anyone (arguably) had predicted [1]. This growth has led to a large number of new results and people interacting with the field, but academic standards have not been taught or followed in the same ways as before the start of this asymmetric expansion [2, 3, 4, 5]. Significantly, publication is done mainly through conferences rather than journals [6, 7]. Papers are still peer-reviewed before publication, but not by teams of well-experienced journal editors, but by the ad hoc reviewers selected by the conference committees, which can have other priorities and incentives than journal editors (admittedly

SAIS2025: Swedish AI Society Workshop 2025, 16-17 June 2025, Halmstad, Sweden.
*Corresponding author.
nemi.pelgrom@lnu.se (N. Pelgrom); mats@skillsta.com (M. Walldén); karl@skillsta.com (K. Andersson)
https://lnu.se/personal/nemi.pelgrom/ (N. Pelgrom)
0009-0004-0150-665X (N. Pelgrom); 0000-0001-9141-9242 (K. Andersson)

overlapping). Furthermore, papers are often restricted to 6-10 pages, which is a length which would be rejected due to insufficiency by other channels. This is further changed by many papers that are published on arXiv before conventional publication, allowing researchers to be influenced by many ideas and results that might not be of a publishable standard but still available to cite and use [2, 8].

The change in publication practice is motivated by immediate access to information, by many seen as a necessity in a fast moving field such as Computer Science. The literature on peer-review however, is stating that the current state of the practice in all of Academia, not only Computer Science, is not working as a way to keep a high standard for published papers [9, 5, 10, 11, 12, 13].

While waiting for a paradigm shift to happen in publication practices, the field of Computer Science has started a paradigm shift of its own, where different research rules and standards are reigning [14, 15, 2]. Without stating hierarchical preferences between these paradigms, it is possible to claim that the Computer Science paradigm adds a haste requirement that was not part of the original peer-review paradigm and which prioritizes direct and current application over general and long-term relevance [2, 3, 4, 5].

On this background, we wanted to explore how this haste is affecting the citation/reference standards of the field. In particular, we wanted to observe whether there continue to be close relationships between this field, its foundations, and the possible connecting fields, or whether this speed makes such connections looser. It is not inherently a problem if the latter is the case; it is a natural development of a field to get a specialized vocabulary and less interaction with the background [16, 17, 18]. Nevertheless, it is problematic if a paradigm shift happens so fast that the relevant parties (connected fields and experts) do not get the opportunity to question the methods and comment what may be connected to overlapping subjects and problem areas [19].

Instances where publication haste leads to uninformed research lead to high risks of reinventing wheels and an unnecessary gap between researchers in the field and the methods available to solve their problems. We created and conducted our study on the idea that it can act as part of identifying if this is what is happening in the field of Computer Science. We wanted to explore the reference practices currently in the field, in relation to the technical details of it. The idea is to identify the attitude towards referencing, and towards making technical details accessible, through identifying a core technical concept and categorizing it into one of the below defined categories.

For our study, we chose the concept *Uncertainty*. It is a popular and highly relevant concept, with many different definitions available in previous research. It is also a performance distinguishing concept that forms the basis for choosing one AI technology over another in early development, putatively leading to immense importance for end results. It is further a concept where there already exists a large pool of useful definitions [20, 21, 22, 23], implying that papers with technical definitions that are not referencing previous relevant research could be considered to be lacking in connection with previous research. We used only papers that had some variant of "uncertain" in the title, to exclude those that only make references to the concept in passing. Whereas it is not compulsory to explicitly define every word that is used in the title of a paper, within that paper, it would be a result of its own if the term was used in many titles without being further discussed at all in the paper connected to those titles.

In linguistics, the phenomenon of isolated groups developing their own specialized language

is often referred to as sociolectal variation or idiolectal innovation [24, 25]. These new linguistic forms, whether jargon, or technical terms, emerge as part of a group's in-group language. The process is driven by linguistic innovation, where speakers modify existing words or introduce entirely new terms to reflect shared experiences or activities. This can involve semantic shift, where a word takes on a new meaning within a specific context, or neologism, the creation of entirely new lexical items. In the context of our study, a similar dynamic occurs when new definitions are introduced to the field, functioning like lexical innovation in a sociolect. These new definitions may not always be referenced back to previous work, much like how slang may not always trace its origins to formal language use, yet both are essential for expressing complex, context-specific ideas. This parallels how the development of new terms within a community enhances communication and adapts to evolving needs, even if those innovations are not directly linked to prior academic discourse [26]. However, most languages are developed solely to facilitate communication, and publishing papers have several objectives in addition to communicating the new results to the relevant scientific community, such as being basis for funding [27, 28, 29]. The parallel gets broken if we find that the newly introduced jargon gets introduced for other reasons than to communicate ideas within the group, since adding new terminology for concepts that are already known in the group is hindering communication rather than facilitating it [30, 31].

We will revisit possible alternative explanations for our results later in the paper, but we want to mention upfront that there are valid reasons for developing new definitions of this concept within the context we explore here. Therefore, not all new definitions, nor those that are not explicitly referenced, should be considered ignorant or unnecessary. However, completely new definitions may be seen as expanding the field's vocabulary without connecting with previous research, even though they are intended to contribute to the field's development. From the perspective of this study, both types of new definitions (restating previously known definitions in a new way, and fully new definitions) can be grouped together under the category "non-referenced definitions."

2. Methodology

To examine how the concept of the *Uncertainty* is invoked in contemporary Computer Science research, we collected a dataset consisting of all papers submitted to arXiv under the Computer Science (cs) classification during the calendar year 2024 that included some version of the word "uncertainty" in their titles. ArXiv was selected over other digital libraries and publication databases (such as ACM Digital Library or IEEE Xplore) because our interest lies in analyzing current representational practices within the field. While arXiv hosts preprints that may not always be peer-reviewed or eventually published in formal venues, the platform has nonetheless become a de facto part of the scholarly citation landscape in Computer Science [32, 33]. Today, it is common to cite arXiv papers in both peer-reviewed and preprint publications without explicit justification for doing so. Although some researchers may be selective, only citing arXiv when the author is well known or the topic is especially fast-moving, these motivations are rarely articulated in the citing papers themselves. The norm has shifted such that citing arXiv preprints is treated comparably to citing peer-reviewed work.

To collect this dataset, we executed the following search query on arXiv:

```
order: -submitted_date;
size: 200;
page_start: 1000;
date_range: from 2024-01-01 to 2025-01-01;
classification: Computer Science (cs);
include_cross_list: True;
terms: AND title=user
```

We retrieved a total of 1,090 papers. Among these, three papers had been withdrawn from arXiv, and their PDF versions were no longer accessible, thus they were excluded from subsequent analysis. Following the primary data collection, we developed a script to automate the extraction of relevant textual content from each paper. The script accessed the HTML version of each arXiv submission and parsed its content to identify and extract all sentences containing the terms "uncertain" or "uncertainty". These sentences were then compiled into an Excel spreadsheet, grouped by the paper the sentences were extracted from (one cell per paper).

A small subset of the papers (17) did not include an HTML version of the text on arXiv. Because this subset represented a minority of the dataset, we handled them manually. For each such paper, we downloaded the PDF and used in-browser search (Ctrl+F) to locate sentences containing the target keywords. These papers where then annotated in the same way as the other papers. In another 16 of the papers, "uncertainty" was mentioned only ≤ 5 times. These papers may be interpreted as not having *Uncertainty* as main topic and simply used the term descriptively rather than referring to it as a technical concept. They were included in the analysis.

Finally, we manually annotated each paper using a set of predefined categories. The categorization scheme was hierarchical, with clearly delineated distinctions between levels, minimizing ambiguity and reducing the potential for subjective interpretation during annotation.

Each item in the dataset was annotated along two independent dimensions:

1. **Definition Type** – one of the following categories:
   - Definition, technical
   - Definition, non-technical
   - No definition
2. **Reference Status** – whether or not the definition (if present) was accompanied by a citation to previous work.

This two-layered annotation allowed us to not only assess the types of definitions used in the field, but also whether they were grounded in existing literature.

Definition, technical A formal definition of the word "uncertain" or "uncertainty", or of a term in which the word appears. These definitions often include LaTeX notation, mathematical expressions, or clearly formalized criteria that may be applied outside of the paper's immediate context.

Definition, non-technical An informal or interpretive explanation of how the term is used within the specific context of the paper. These are less formal than technical definitions and may be vague or specific only to the particular use case discussed.

The technical category is a subset of this broader non-technical category: all technical definitions are inherently definitions, but not all definitions meet the criteria to be considered technical.

> *"The picture itself is uncertain, that is, it does not contain enough information to infer what it shows."*

This sentence, taken from one of the reviewed papers, illustrates a non-technical definition. It provides an interpretive framing of the term in context, but lacks the structure required for generalizability or formal precision.

No definition The term "uncertain" or "uncertainty" is used without any clarifying explanation, formal or informal.

We initially attempted to use a generative language model to assist in the annotation process. However, it was unable to reliably distinguish between technical and non-technical definitions. As a result, all annotations were performed manually.

After annotation, we used standard spreadsheet functions to compute the counts and proportions in each category, for both definition type and reference status. This enabled an analysis not only of the definitional practices within the field, but also the degree to which definitions are formally grounded in prior research.

Our primary hypothesis was that a significant number of papers in the dataset would discuss *Uncertainty* without citing prior work that defines or theorizes the concept. This hypothesis was confirmed. While the dataset will not be made public—our focus is on identifying a broader paradigm shift rather than assigning responsibility to individual researchers—we are willing to share it upon request.

3. Results

We present two main findings from our analysis of definition usage and referencing practices in Computer Science papers. Firstly, we observed that all combinations of definition types and referencing status were represented, indicating a broad variation in how definitions are introduced across the field. As seen in Table 1, technical definitions dominate, which aligns with expectations given the nature of the subject matter. However, 74% of technical definitions were not properly referenced, a higher proportion than anticipated. This category includes both genuinely novel definitions and those reused without explicit citation. This is suggesting that a considerable portion of the literature either assumes shared understanding or neglects to clarify sources, and no clear way to tell the difference between the two, which raises concerns regarding both clarity and scholarly rigour.

Secondly, we examined a subset of papers that explicitly use the term LLM in the same sentence as "uncertainty", aiming to analyze whether this specific intersection reveals different

Table 1
Distribution of Referenced and Not Referenced Items by Definition Type (Counts and Percentages). Not referenced includes both novel definitions, and definitions that are not properly referenced.

	Technical Def.	Non-tech Def.	No def.
Referenced (count)	184	53	44
Not referenced (count)	537	62	86
Referenced (%)	25.52%	46.08%	33.85%
Not referenced (%)	74.48%	53.92%	66.15%

patterns. As shown in Table 2, this subset contains a disproportionately high number of papers with no definitions at all (22.45%), compared to the overall average (13.46%). It also contains a lower share of technical definitions (66.67%).

Table 2
Comparison of Overall vs. LLM Definition Distributions

	Technical Def.	Non-tech Def.	No def.
Percentage of total	74.64%	11.90%	13.46%
Percentage of total (LLM)	66.67%	10.88%	22.45%

Further breakdown in Table 3 shows that in this LLM-related subset, the proportion of non-referenced technical definitions remains high at 67.35%, consistent with the general trend. Interestingly, papers with no definitions in this group are more likely to reference external work (51.52%), which may suggest a reliance on assumed external framing rather than self-contained explanations.

Table 3
Distribution of Referenced and Not Referenced Items by Definition Type, on the subset of papers that mention LLMs in the same sentence as "uncertainty"

	Technical Def.	Non-tech Def.	No def.
Referenced (count)	32	6	16
Not referenced (count)	66	10	17
Referenced (%)	32.65%	37.50%	48.48%
Not referenced (%)	67.35%	62.50%	51.52%

These findings highlight an area for potential improvement in scientific writing within this domain. The high frequency of non-referenced definitions, especially technical ones, could lead to ambiguity or misinterpretation, particularly in interdisciplinary contexts or when newer concepts such as LLMs are involved. We suggest that future work should investigate whether this trend persists in other domains, and whether clearer definition practices correlate with better reproducibility or clarity in downstream applications.

4. Discussion

We have reviewed how the concept of *Uncertainty* was used in Computer Science in ArXiv during 2024. The purpose was to shed light on the referencing principles, which in turn conveys information about the velocity of motion in this field.

Approximately 2/3 of the publications where the term ”uncertainty” is defined do not cite previous research. In case the discussed *Uncertainty* concept in a particular publication is novel, this is of course acceptable. The publications reviewed in this study were however, in vast majority, such that concepts known from or similar to the past were discussed. This introduces two critical problems.

The first problem is the risk of reinventing the wheel, putatively due to the velocity of the field, where researchers fail to stay updated while in need of an *Uncertainty* concept and hence design a concept that is functionally similar to what has already been discussed and evaluated previously. This leads to unnecessary work and introduces confusion through redundant terminology and overlapping formalisms.

The second problem is the lack of stringency. Any researcher in Computer Science would, with few exceptions, be able to identify prior work similar to their own definition of *Uncertainty*. When isolating the sub-set of publications mentioning LLMs in the same sentence as ”uncertainty”, the level of rigour was even lower: 22% of these papers lacked any form of definition of *Uncertainty*, compared to 13% in the general dataset. The omission of researching and referencing the background of a core concept being discussed in a scientific publication is troubling. Hence, a call for increased referencing stringency is warranted.

Together, these two problems suggest that the field’s use of *Uncertainty* as a concept lacks maturity. The absence of a few central, shared citations indicates that Computer Science as a whole remains uncertain about *Uncertainty*. Researchers writing about LLMs appear even less grounded. Given the rapid expansion of these areas, this is not entirely surprising, but it is of critical importance that fundamental concepts be clearly defined and consistently reused. Without this, cross-comparison, validation, and cumulative knowledge building become difficult or impossible.

Our results show a wide mix of items across definition and reference categories. This spread suggests the presence of multiple paradigms operating simultaneously in the field. In a more stable field with clear citation norms, we would expect to see less variation in how and whether definitions are referenced.

One potential counter-explanation is that a large number of genuinely novel definitions of *Uncertainty* were introduced in 2024. However, we do not find this plausible. Given the applied nature of most of the work in Computer Science, it is unlikely that hundreds of new and fundamentally distinct definitions would all be necessary or beneficial. Instead, it is more plausible that many of these definitions are restatements of similar ideas in slightly different terms.

This reflects a broader challenge in Computer Science: distinguishing between genuine conceptual contributions and mere reformulations of existing ideas. Terms such as *Pipeline*, *Network*, *Framework*, *Graph*, *Formal model*, and *Logic* are often used to describe structurally similar concepts, but their terminology varies across subfields and use-cases. This results in a landscape where different names are used for essentially the same constructs, depending on

the epistemic traditions of different communities. The root of this problem is that, outside of formal languages with explicit syntactic and semantic rules, it is not possible to definitively determine whether two definitions or constructs are identical. Without a shared syntax and semantics, objects may always be interpreted differently depending on context.

The difficulty in drawing clear boundaries between new contributions and redundant reinventions is probably not unique to *Uncertainty*. We suspect it is a widespread issue in Computer Science, and one that grows more pressing as the velocity of research increases. The current state of the field, where a core concept such as *Uncertainty* can be used in diverse and often unreferenced ways, suggests that this increased research velocity is negatively impacting both the clarity and cumulative progress of scholarly discourse. Understanding how fast-paced research can share findings in a manner not to reinvent wheels and stay with coherent definitions of core concepts is an entire research topic in itself.

A field where fundamental concepts are used inconsistently and without reference cannot build reliably upon itself. The results presented in this paper serve as a call for a more rigorous and reflective approach to how foundational concepts such as *Uncertainty* are defined, cited, and integrated into ongoing research.

5. Conclusion

We found our main hypothesis to be confirmed; in arXiv there is a significant amount of newly published papers in the field of Computer Science that are using *Uncertainty* as a core concept, while not referencing sources for their definitions of *Uncertainty*. While we cannot make claims on the exact reason for *why* this is the case, we make the claim that this is a representative image of the current state of communicability within the field. The amount of new technical descriptions of concepts is out-running any researchers capability to overview them.

6. Citations and Bibliographies

References

[1] X. Tang, X. Li, Y. Ding, M. Song, Y. Bu, The pace of artificial intelligence innovations: Speed, talent, and trial-and-error, Journal of Informetrics 14 (2020) 101094.

[2] H. Nguyen, S. Eger, Is there really a citation age bias in nlp?, 2024. URL: https://arxiv.org/abs/2401.03545. arXiv:2401.03545.

[3] G. Smith, Distrust: Big data, data-torturing, and the assault on science, Oxford University Press, 2023.

[4] C. L. A. Navarro, J. A. Damen, T. Takada, S. W. Nijman, P. Dhiman, J. Ma, G. S. Collins, R. Bajpai, R. D. Riley, K. G. Moons, et al., Systematic review finds "spin" practices and poor reporting standards in studies on machine learning-based prediction models, Journal of clinical epidemiology 158 (2023) 99–110.

[5] M. Shepperd, L. Yousefi, An analysis of retracted papers in computer science, Plos one 18 (2023) e0285383.

[6] G. Vrettas, M. Sanderson, Conferences versus journals in computer science, J. Assoc. Inf. Sci. Technol. 66 (2015) 2674–2684. URL: https://doi.org/10.1002/asi.23349. doi:10.1002/asi.23349.
[7] M. Franceschet, The role of conference publications in cs, Communications of the ACM 53 (2010) 129–132.
[8] D. van Ravenzwaaij, M. Bakker, R. Heesen, F. Romero, N. van Dongen, S. Crüwell, S. Field, L. Held, M. Munafò, M.-M. Pittelkow, et al., Perspectives on scientific error, Royal Society Open Science 10 (2023) 230448.
[9] M. Thelwall, J. A. Hołyst, Can journal reviewers dependably assess rigour, significance, and originality in theoretical papers? evidence from physics, Research Evaluation 32 (2023) 526–542.
[10] S. Y. Hwang, D. K. Yon, S. W. Lee, M. S. Kim, J. Y. Kim, L. Smith, A. Koyanagi, M. Solmi, A. F. Carvalho, E. Kim, et al., Causes for retraction in the biomedical literature: a systematic review of studies of retraction notices, Journal of Korean medical science 38 (2023).
[11] L. Waltman, W. Kaltenbrunner, S. Pinfield, H. B. Woods, How to improve scientific peer review: Four schools of thought, Learned Publishing 36 (2023) 334–347.
[12] D. M. Herron, Is expert peer review obsolete? a model suggests that post-publication reader review may exceed the accuracy of traditional peer review, Surgical Endoscopy 26 (2012) 2275–2280. doi:10.1007/s00464-012-2171-1, epub 2012 Feb 21.
[13] J. S. Trueblood, D. B. Allison, S. M. Field, A. Fishbach, S. D. M. Gaillard, G. Gigerenzer, W. R. Holmes, S. Lewandowsky, D. Matzke, M. C. Murphy, S. Musslick, V. Popov, A. L. Roskies, J. ter Schure, A. R. Teodorescu, The misalignment of incentives in academic publishing and implications for journal reform, Proceedings of the National Academy of Sciences of the United States of America 122 (2025) e2401231121. doi:10.1073/pnas.2401231121.
[14] T. S. Kuhn, et al., Second thoughts on paradigms, The structure of scientific theories 2 (1974) 459–482.
[15] J. P. Sturmberg, Changing the paradigm of research, Journal of Evaluation in Clinical Practice 29 (2023) 726–729.
[16] R. K. Merton, The sociology of science: Theoretical and empirical investigations, University of Chicago press, 1973.
[17] R. Marius, Genre analysis: English in academic and research settings, 1991.
[18] J. Harmon, K. Wood, The vocabulary-comprehension relationship across the disciplines: Implications for instruction, Educ. Sci. 8 (2018) 101. doi:10.3390/educsci8030101.
[19] M. T. Soto-Sanfiel, C.-W. Chong, J. I. Latorre, Hype in science communication: Exploring scientists' attitudes and practices in quantum physics, arXiv preprint arXiv:2311.07160 (2023).
[20] F. Wang, Y. Liu, K. Liu, Y. Wang, S. Medya, P. S. Yu, Uncertainty in graph neural networks: A survey, 2025. URL: https://arxiv.org/abs/2403.07185. arXiv:2403.07185.
[21] O. Shorinwa, Z. Mei, J. Lidard, A. Z. Ren, A. Majumdar, A survey on uncertainty quantification of large language models: Taxonomy, open research challenges, and future directions, 2024. URL: https://arxiv.org/abs/2412.05563. arXiv:2412.05563.
[22] C. Yin, R. Liu, D. Zhang, P. Zhang, Identifying sepsis subphenotypes via time-aware multimodal auto-encoder, in: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery amp; Data Mining, KDD '20, ACM, 2020, p. 862–872. URL:

http://dx.doi.org/10.1145/3394486.3403129. doi:10.1145/3394486.3403129.
[23] X. Huang, S. Li, M. Yu, M. Sesia, H. Hassani, I. Lee, O. Bastani, E. Dobriban, Uncertainty in language models: Assessment through rank-calibration, 2024. URL: https://arxiv.org/abs/2404.03163. arXiv:2404.03163.
[24] S.-I. Harada, Ga-no conversion and idiolectal variations in japanese, Gengo Kenkyu (Journal of the Linguistic Society of Japan) 1971 (1971) 25–38.
[25] M. Lewandowski, Sociolects and registers–a contrastive analysis of two kinds of linguistic variation, Investigationes Linguisticae 20 (2010) 60–79.
[26] G. Yule, Language and social variation, Cambridge University Press, 2005, p. 205–215.
[27] N. S. Young, J. P. Ioannidis, O. Al-Ubaydli, Why current publication practices may distort science. the market for exchange of scientific information: the winner's curse, artificial scarcity, and uncertainty in biomedical publication, PLoS Medicine 5 (2008).
[28] N. C. Herndon, Research fraud and the publish or perish world of academia, 2016.
[29] J. E. Bekelman, Y. Li, C. P. Gross, Scope and impact of financial conflicts of interest in biomedical research: a systematic review, Jama 289 (2003) 454–465.
[30] D. Y. Manin, Zipf's law and avoidance of excessive synonymy, Cognitive Science 32 (2008) 1075–1098.
[31] B. Szymanek, Remarks on tautology in word-formation, in: L. Bauer, L. Körtvélyessy, P. Štekauer (Eds.), Semantics of Complex Words, volume 3 of *Studies in Morphology*, Springer, 2015.
[32] J. Lin, Y. Yu, Y. Zhou, Z. Zhou, X. Shi, How many preprints have actually been printed and why: a case study of computer science preprints on arxiv, Scientometrics 124 (2020) 555–574.
[33] M. B. Hoy, Rise of the rxivs: How preprint servers are changing the publishing process, Medical Reference Services Quarterly 39 (2020) 84–89.