

OCL on life support: can we revitalize the community for a stronger future?

Mauro Dalle Lucca Tosi¹, Jordi Cabot^{1,2}

¹Luxembourg Institute of Science and Technology (LIST), 5 Avenue des Hauts-Fourneaux, 4362, Esch-sur-Alzette, Luxembourg

²University of Luxembourg (UNILU), 2 Av. de l'Universite, 4365, Esch-sur-Alzette, Luxembourg

Abstract

The Object Constraint Language (OCL) is a key element of model-driven engineering (MDE). However, despite its significance, there are indications that research on OCL may be in decline. In this paper, we conduct a metascience study of the OCL research field, examining its size, impact, and outreach to assess whether this perceived decline is supported by the data. Through this analysis, we aim to assess the vitality of OCL research and provoke discussion on how the community can adapt to ensure its continued relevance.

Keywords

OCL, Object Constraint Language, Metascience, Scientometrics, Community

1. Introduction

The Object Constraint Language (OCL) has long been a fundamental component of model-driven engineering (MDE). By extending the expressiveness of graphical notations, OCL enables a range of functionalities, including code generator templates, model transformation instructions, and the specification of use-case constraints, among others [1].

Despite its important role, we have the impression that interest in OCL is declining across various aspects, including research output, community engagement, the development of new tools, etc. This perception aligns with broader trends in MDE, which has also experienced a period of decline but is now experiencing a resurgence with the rise of Low-Code approaches [2].

In this study, we investigate whether this impression holds by analyzing multiple indicators over time to determine when OCL research reached its peak and whether it is now following a downward trend. Specifically, we aim to answer the following question: *What is the current size, impact, and outreach of the OCL research community compared to its former peak?*

To address this question, we conduct a metascience study on the evolution of the OCL research field. Based on our findings, we present key discussion points and highlight critical questions concerning the future of OCL research.

2. Methodology

This section outlines the methodology used for addressing our research question.

2.1. Data Collection & Data Cleaning

Data Collection. We began by searching for academic documents containing the terms “OCL” or “Object Constraint Language” in their titles, abstracts, or keywords. To avoid confusion with other contexts, we excluded documents that referred to “Online Collaborative Learning”, which shares the

Joint Proceedings of the STAF 2025 Workshops: OCL, OOPSLE, LLM4SE, ICMM, AgileMDE, AI4DPS, and TTC. Koblenz, Germany, June 10-13, 2025

✉ mauro.dalle-lucca-tosi@list.lu (M. D. L. Tosi); jordi.cabot@list.lu (J. Cabot)

id 0000-0002-0218-2413 (M. D. L. Tosi); 0000-0003-2418-2489 (J. Cabot)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

OCL acronym but is unrelated to our study. We utilized LENS.ORG¹ to retrieve metadata for relevant documents, resulting in a dataset of 2,898 academic publications.

Data Cleaning. Upon reviewing the retrieved metadata, we identified a number of irrelevant documents, including those that referenced terms such as “Optimum Control Logic” or other unrelated acronyms. We refined our dataset by manually removing these entries. Additionally, we excluded non-English documents, duplicates, retracted articles, and preprints that had been subsequently published. After this process, we retained a final set of 2,376 OCL-related publications.

We also manually corrected the publication venues for all documents to ensure the accuracy of venue names and types.

Data Analysis. To identify trends in OCL publications, we categorized and counted the documents by publication year, venue type, and venue. To gain insights into the OCL research community, we constructed co-authorship networks that illustrate the relationships among researchers. To track the evolution of the community, we measured the number of OCL authors each year, distinguishing between new and recurring authors. Finally, to assess OCL’s impact within the broader research landscape, we calculated both the total and average number of citations received by OCL papers published each year.

3. Results

3.1. Trends in the Number of Publications

Figure 1 shows the number of OCL-related publications per year, categorized by publication type. The results reveal a clear decline in the total number of OCL publications over time. Notably, conference and workshop publications, which played a significant role in sustaining OCL research until 2020, have nearly disappeared. Today, the field is primarily supported by a small number of journal articles, and even smaller number of conference proceedings, and preprints.




Figure 1: Number of OCL-related publications over the years.

3.2. Top Venues

Figure 2 highlights the main venues where OCL-related research is published. Among the top 10 venues, five are conferences: MODELS, UML, ER, ICSE, and ECMFA. Additionally, the OCL workshop and the Software and Systems Modeling (SoSyM), the Electronic Communications of the EASST (ECEASST), and the Journal of Object Technology (JOT) journals are also prominent publication outlets.

¹<https://www.lens.org/>





Figure 2: Number of OCL-related publications per venue.


3.3. Co-authorship analysis

Figure 3 presents the co-authorship network of OCL-related publications. As shown in Figure 3a, senior authors—those with a higher number of publications in the field—tend to occupy central positions in the network, where they are more connected to other researchers.

Figure 3b focuses on authors with at least four OCL-related publications, revealing a highly interconnected research community. This indicates strong collaboration among recurrent OCL authors, a characteristic that is not necessarily present in all research areas, such as Low-Code [2].



(a) All authors.



(b) Authors with at least 4 OCL-related publications.

Figure 3: OCL co-authorship graph. Nodes represent OCL authors, and edges indicate co-authorship in OCL-related publications. Node size and color reflect the number of OCL-related publications, while edge thickness represents the number of co-authored papers.

3.4. Active OCL Authors' over the years

Figure 4 illustrates the number of active authors in the OCL field over time. An author is considered to have left the field if they go three consecutive years without publishing an OCL-related paper. This three-year threshold is based on a log distribution of publication intervals, where a chi-square fit of over 0.999 suggests that the probability of an author publishing again after three years is less than 2.3%.

Figures 4a and 4b reveal a significant decline in the number of active OCL authors. Moreover, Figure 4b shows a slightly sharper decline in recurrent authors compared to new authors entering the field. This suggests that while some newcomers are still publishing OCL-related research, fewer of them remain active long enough to become established contributors.

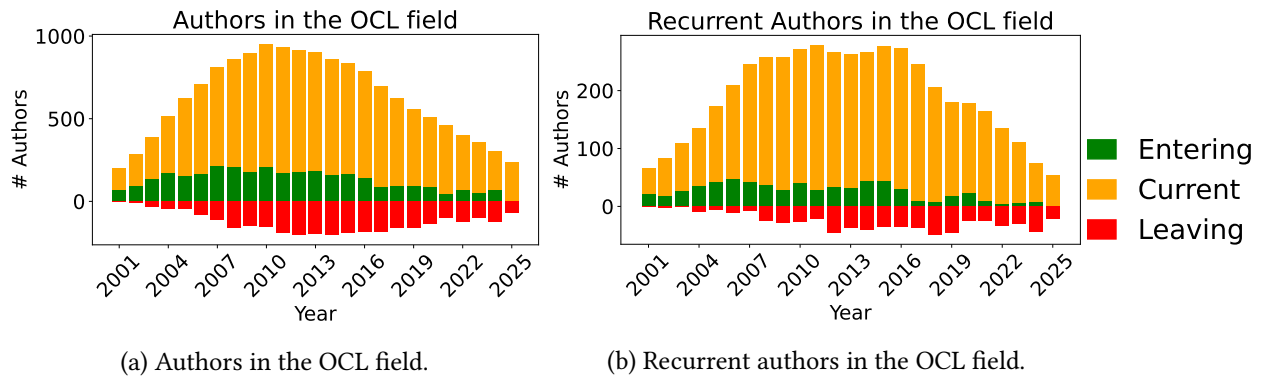


Figure 4: Number of authors that are entering, leaving, and staying in the OCL field over the years.




Figure 5: Total and average number of citations for OCL-related papers over time.

3.5. Evolution of the Number of OCL papers citations

Figure 5 presents the total and average number of citations received by OCL-related papers per publication year. The data clearly show that the most influential OCL papers—those with the highest citation counts—were published between 1998 and 2007. Since then, the impact of OCL research, as measured by citations, has steadily declined.

While it is natural for newer papers to have fewer citations due to their shorter time in circulation, the observed decline in citation rates is not solely attributable to recency. This trend began from 2007, and most scientific publications reach their peak citation rate within five years [3].

4. The Good, The Bad, and The Ugly

4.1. The Bad News

Our analysis reveals a steep decline in the OCL research field across all examined metrics. The number of academic publications has significantly decreased, the community has fewer active authors, and citation rates have dropped. These trends indicate a clear decline in the field, which peaked around 2006 and has been in decline ever since.

With more researchers leaving the OCL field than entering it each year, the community is shrinking. Notably, only 54 recurrent authors remain active in the field. This limited number of experienced contributors likely impacts the declining influence of recent OCL publications. One notable characteristic of the OCL community is its high level of internal collaboration among its most experienced researchers. While it is common for a small subset of authors to contribute disproportionately to a research field, it is less common for them to collaborate so extensively with one another. This strong internal focus may help sustain the field in the short term, but it also raises concerns about its long-term growth, as fewer new researchers are being integrated into the community.

The declining interest in OCL research is reflected in its decreasing impact in terms of citations. This reduced impact could be both a cause and a consequence of the decreasing number of active authors. However, while the decline in author participation has been gradual, the total and average number of citations for OCL papers have dropped more sharply, indicating a waning influence of the field within the broader research community.

4.2. The Ugly Context

OCL decline does not occur in isolation but aligns with broader trends in software engineering and modeling. The modeling field, until recently, had reached an all-time low but managed to revitalize itself with the rise of Low-Code approaches [2]. To be seen whether this revitalization of modeling could also positively impact the OCL community.

Another critical factor hindering OCL's vitality is the lack of investment in tools that could facilitate its research and adoption in industrial projects [4]. This issue has been recognized for a long time, yet little progress has been made. In fact, the situation may be worse than before, with B-OCL [5] being one of the few, if not the only, recent OCL-focused projects. It could be that the new generation of AI coding assistants, starting to be used in all types of Software Engineering tasks [6], lowers the barrier for the creation of new OCL tools.

4.3. The Good Hopes

Despite these challenges, OCL remains alive, supported by a highly collaborative community with opportunities for revitalization. The longevity of the field is noteworthy—unlike other fields such as aspect-oriented programming (AOP), which rose quickly but faded soon after, OCL has endured, demonstrating its resilience. This is a major achievement and a great quality for a sustainable future.

One potential path forward is bridging the gap between OCL and conceptually similar communities, such as the Semantic Web community, which has SHACL [7], a constraint language with similarities to OCL. Strengthening such connections could help OCL regain traction and expand its user base.

Moreover, with renewed interest in data models due to AI, challenges like data annotation and mining now need revisiting in the machine learning context. This presents an opportunity for the modeling community, including OCL researchers, to contribute their expertise in structuring, transforming, reusing, and deploying ML artifacts—helping to shape the future of AI through better modeling practices. In order to do so, OCL researchers should go out of their comfort zone (in terms of traditional venues) and target conferences aimed at exploring the intersection between Software Engineering and AI.

5. Conclusion

In this paper, we studied and quantified the decline in interest in the OCL field over the past decades. Our analysis highlights the significant decrease in OCL-related publications, active researchers, and overall impact in terms of citations within the scientific community. These findings emphasize the need for reflection on both the reasons behind this decline and the future of the OCL community.

Given this reality, we raise important questions—not only about why this decline is occurring but also, considering OCL's resilience, about how the community should respond to reverse this trend. Now is a critical moment for the OCL community to decide on its path forward, whether by adapting, forging stronger connections with related research communities, leveraging the current focus on AI research, or pursuing other strategic actions.

Acknowledgments

This project is supported by the Luxembourg National Research Fund (FNR) PEARL program, grant agreement 16544475.

Declaration on Generative AI

During the preparation of this work, the author(s) used GPT-4o and GPT-3.5 in order to: Grammar and spelling check. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] J. Cabot, M. Gogolla, Object Constraint Language (OCL): A Definitive Guide, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 58–90. doi:10.1007/978-3-642-30982-3_3.
- [2] M. D. L. Tosi, J. L. C. Izquierdo, J. Cabot, A metascience study of the impact of low-code techniques in modeling publications, arXiv preprint arXiv:2408.05975 (2024).
- [3] Z. Hu, J. Cui, A. Lin, Identifying potentially excellent publications using a citation-based machine learning approach, Information Processing & Management 60 (2023).
- [4] J. Cabot, E. Teniente, Constraint support in mda tools: A survey, in: A. Rensink, J. Warmer (Eds.), Model Driven Architecture – Foundations and Applications, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 256–267.
- [5] F. U. Haq, J. Cabot, B-ocl: An object constraint language interpreter in python, 2025. URL: <https://arxiv.org/abs/2503.00944>. arXiv:2503.00944.
- [6] X. Hou, Y. Zhao, Y. Liu, Z. Yang, K. Wang, L. Li, X. Luo, D. Lo, J. Grundy, H. Wang, Large language models for software engineering: A systematic literature review, ACM Trans. Softw. Eng. Methodol. 33 (2024). URL: <https://doi.org/10.1145/3695988>. doi:10.1145/3695988.
- [7] P. Pareti, G. Konstantinidis, A review of shacl: from data validation to schema reasoning for rdf graphs, Reasoning Web International Summer School (2021) 115–144.