

Shoulders of Giants: A Look at the Degree and Utility of Openness in NLP Research

Surangika Ranathunga¹, Nisansa de Silva², Dilith Jayakody², Aloka Fernando²

¹School of Mathematical and Computational Sciences, Massey University, New Zealand
s.ranathunga@massey.ac.nz

²Dept. of Computer Science & Engineering, University of Moratuwa, 10400, Sri Lanka
{NisansaDdS, dilith.18, alokaf}@cse.mrt.ac.lk

Abstract

We analysed a sample of NLP research papers archived in ACL Anthology as an attempt to quantify the degree of openness and the benefit of such an open culture in the NLP community. We observe that papers published in different NLP venues show different patterns related to artefact reuse. We also note that more than 30% of the papers we analysed do not release their artefacts publicly, despite promising to do so. Further, we observe a wide language-wise disparity in publicly available NLP-related artefacts.

1 Introduction

The advancement of the Computer Science research field heavily depends on publicly available code, software, and tools. Its sub-fields Machine Learning and Natural Language Processing (NLP) have the additional requirement of datasets - to train and evaluate computational models. Lack of access to these research artefacts has been identified as a major reason for the difficulty in reproducing works of others (Pineau et al., 2021). The data requirement is particularly challenging in NLP - a dataset available for one language usually cannot be used in the context of another language¹.

Therefore, the NLP community is highly encouraged to make their research artefacts publicly available. However, as far as we are aware, there is no quantifiable evidence on (1) the degree of openness in the NLP community or (2) the benefit of openness to the community. Since “*what we do not measure, we cannot improve*” (Rungta et al., 2022), in this paper, we quantify both these aspects. To this end, we semi-automatically analyse a sample of NLP research papers published in ACL Anthology (AA) and corpora/ Language Models

¹Other than in techniques such as multi-tasking and intermediate-task fine-tuning.

(LMs) released in Hugging Face², and answer the following questions:

1. To what degree has the NLP research community been able to reuse open-source artefacts (data, code, LMs) in their research?
2. How much has the community freely shared the artefacts produced by their research?

To answer the first question, we record the number of papers that reuse the artefacts released by past research. Since there is a language-wise disparity in NLP research (Joshi et al., 2020; Ranathunga and de Silva, 2022), this analysis is conducted while separating low- and high-resource languages.

To answer the second question, we record the papers that indicate they would release the newly produced artefacts. We also record whether they have provided a repository URL. We do further analysis to find out whether these repositories have the artefacts they are supposed to have. Finally, we record the number of datasets and LMs available for different language classes on Hugging Face.

We observe that papers published in different venues show different patterns in artefact reuse. We also observe that a worrying percentage of papers that produced an artefact have not publicly released those artefacts. To a lesser degree, broken repository links and empty resource repositories were also noted. Finally, it is noted that the language-wise disparity in LM/data availability (Joshi et al., 2020; Ranathunga and de Silva, 2022; Khanuja et al., 2023) is still staggering.

2 Data Extraction

We use AA as the research paper repository. While AA is the largest NLP-related paper repository, Ranathunga and de Silva (2022) note that many papers related to low-resource languages also

²<https://huggingface.co/>

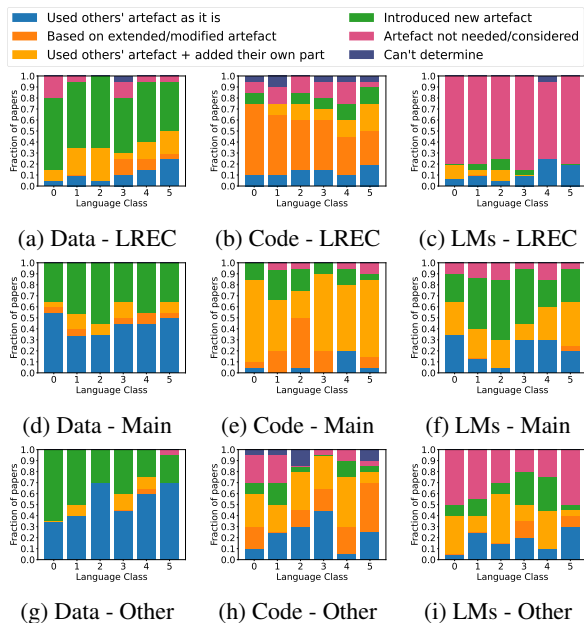


Figure 1: Artefact (Data, Code, and LMs) creation, extension, and reuse across PVs.

get published in other venues such as IEEE conferences or regional journals. However, the popularly used Google Scholar does not have a free API to extract data, and the coverage of Semantic Scholar is rather poor³. Moreover, some conference and journal publications are hidden behind paywalls. While archives such as arXiv are a possible option, they do not contain the meta data for us to carry out a conference/journal-specific analysis. Considering all these factors, we selected AA to extract papers for our analysis. AA has been the common choice for many research related to diversity analysis in NLP research (Rungta et al., 2022; Blasi et al., 2022; Cains, 2019).

When collecting data from AA, we reuse data and code from Ranathunga and de Silva (2022) who in turn had used code and data from Blasi et al. (2022) and Rohatgi (2022) (respectively). However, we had to collect data post 2022 by ourselves.

We use the URLs of papers from the ACL Anthology Bibliography to extract the title and abstract of each paper. We then allocate the papers to different languages, following the language list (of 6419 languages) given by Ranathunga and de Silva (2022). For each language name, we check for matches in both the title and abstract and download the matched papers using their respective URLs (where a URL to the PDF is available). Of these,

³For example, the search query "english+nlp" returns 4312 results on Semantic Scholar as opposed to the 495,000 results returned by Google Scholar.

130 languages are ignored due to the high count of false positives caused by matches with existing words and author names⁴. Next, we convert each paper to its text format.

Then we further group these language-wise papers according to language category. The commonly used language category definition that is based on language resources is Joshi et al. (2020) (see Table 4 in Appendix). This definition can be used to categorise languages into six classes, with class 5 being the highest resourced, and class 0 being the least resourced. Joshi et al. (2020) used this definition to classify about 2000 languages. However, this categorisation was conducted in 2020 and it has considered only ELRA⁵ and LDC⁶ as data repositories. Ranathunga and de Silva (2022) showed that these repositories have very limited coverage for low-resource languages. They reused Joshi et al. (2020)'s language category definition and categorised 6419 languages considering the Hugging Face data repository in addition to ELRA and LDC. In this research, we use this newer language categorisation.

3 Analysis

3.1 The degree of artefact reuse in NLP research

We extract a paper sample of 355 (papers published between 2015-2023) from the dataset downloaded above. To analyse the effect of the publishing venue, these papers are then separated into three categories (henceforth referred to as *PV* categories). These categories are selected based on the suggestion of Ranathunga and de Silva (2022).

- **Main:** Main ACL conferences/journals where NLP researchers publish (Full list in Appendix B).
- **LREC** (Language Resources and Evaluation Conference). It was given a separate category as it is a venue specifically focusing on language resources.
- **Other** - Everything else. Usually, these PVs refer to shared tasks, workshops and regional conferences such as RANLP and ICON.

⁴Examples of languages that were ignored include: *Are, As, Even, One, So, To, Apache, U, Bit, She*.

⁵<http://www.elra.info/en/>

⁶<https://www ldc.upenn.edu/>

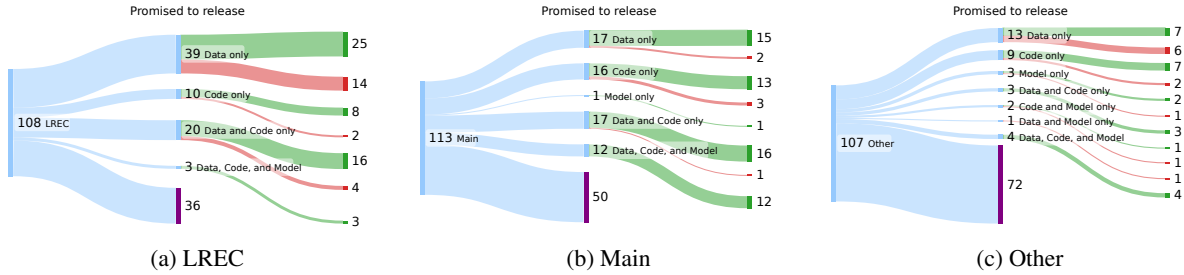


Figure 2: Artefact releasing promise vs artefact link availability across PVs. Green - Artefact Released, Red - Claimed to release the relevant artefact but no link given, Purple - No promise was given to release any artefact.




Figure 3: Analysis on artefact release.

Artefact	Status
Data	Used dataset from some previous research
	Extended an existing dataset
	Used dataset from some previous research but created new data as well
	Introduce new dataset
	Data not needed
	Cannot determine

Table 1: Possible options for use, and reuse of data

For each PV, the resulting paper sample has 20 papers per language class⁷. We manually read each of these papers to find out whether they created/used data, code⁸ and/or LMs⁹. The possible options for data-related mentions in a paper are shown in Table 1. Similar options are considered for code and LMs (see Table 5 in Appendix). Note that the first three entries in Tables 1 and 5 suggest the reuse of artefacts from previous research in some manner.

Out of the 355 papers we analysed, 98.9% has reused some form of artefact from previous research. Further language class-wise analysis on this is shown in Figure 1 (In the Appendix we have a larger version in Figure 5 as well as a chronological breakdown of the data in Figure 6).

⁷Except for language class 1 in *Main* PV, where we could find only 15 papers.

⁸We considered NLP related tools/libraries/code repositories such as NLTK and Huggingface libraries but did not consider generic libraries such as Pandas.

⁹By LMs, we refer to LMs starting from Word2Vec, GloVe and FastText, coming to currently used Large LMs

Other PV category is the highest in reusing data as-it-is. This is not surprising, as this category has many papers referring to shared tasks. *Main* category also uses existing data as-it-is to a higher degree, but there is some emphasis on data extension as well. *LREC*, due to its focus on language resources, sees more papers introducing new datasets or extending existing datasets than those that reuse existing data as-it-is.

The *Main* category sees the highest level of code reuse to introduce new implementations - most papers extend code from already existing research. This has to be due to the highly competitive nature of PVs in this category, where reviewers emphasise technical novelty. *Other* PV category is high in reusing code as well, but it has a relatively higher portion of papers using existing code as-it-is.

As mentioned earlier, since most *LREC* papers focus on dataset release, they seem not to have paid attention to the use of state-of-the-art solutions involving LMs. In contrast, papers from *Main* heavily emphasise using LMs, and this PV category seems to be the venue to introduce new LMs.

Overall, the most reused artefact is code, spanning from early APIs/toolkits such as NLTK (Bird et al., 2009) and Kaldi (Povey et al., 2011) to modern-day Hugging Face libraries.

3.2 Percentage of papers that promise to share the newly created artefacts

Next, we focus on papers that create new artefacts (created from scratch or extended existing artefacts) and report the percentage of papers that promise to share the newly created artefacts. If they do promise, then we check whether they have provided the URL of the public repository containing the artefact(s).

This analysis was done in a semi-automated manner on the same 355 paper sample as before, using a keyword-based method to filter papers.

To identify keyword matches, we first replace all non-letter characters of the paper full text with spaces and convert the text to lowercase. To match keywords containing a single term, we split the text by the space character and look for exact matches between the keyword and the words in the resulting array. To match keywords containing multiple terms, we do a direct search over the text (without splitting). We make this distinction between single-word and multi-word keywords due to the false positives caused by matching substrings (for example, "public" would match a text that contains the word "republic"). For each matched keyword, we extract the paragraph in which it was identified and create text files using these paragraphs. These filtered text files assist in identifying the claims of the papers during the manual analysis.

The keywords consist of words that indicate availability. The complete set of keywords is as follows: release, released, public, publicly, github, gitlab, huggingface co, osf io, open source, accessible. Note that the non-letter characters of the keywords are also replaced by spaces to facilitate the matching. Also, note that we do not include keywords such as available and http due to the high number of false positives that they cause. In order to quantify the impact of avoiding these keywords, we look at the false omission rate of a sample of 100 papers. We randomly select 100 papers from the data set and run them through our keyword-based search algorithm. This predicted 69 papers to contain promises of releasing artefacts. We then manually checked the remaining 31 papers in full, to see whether they promised the release of an artefact. Of these 31 papers, one paper has promised and shared the data and code. This results in a false omission rate of approximately 0.03.

We manually read the filtered papers to further verify whether a paper has produced an artefact,

and if so, whether it has promised to release that artefact.

Results are shown in Figure 2. Interestingly, out of the *Main* PV papers that produced some new artefacts, 44% have not mentioned whether that artefact will be released. In the *Other* category, this value is 67%. *LREC* has the lowest percentage at 33%. However, in *LREC*, 36% of the papers that have promised to release data have not given a repository URL.

3.3 Further Analysis into Artefact Availability

In the above analysis, we can only determine whether a paper mentions that research artefacts are publicly released, and if so, a link to a repository is given. However, that analysis does not tell us the type of these repositories, whether they are accessible, or whether they contain the artefact. Therefore, we carry out a second, more detailed analysis.

To get an insight into more recent trends, we consider papers published between 2020-2023. Following the same semi-automated approach discussed above, we extract a list of papers that promised to release at least one of the following artefacts: *data*, *source code*, *LM*, or *tool*. Then the extracted papers are grouped according to the language class. Classes 5, 4, 3 and 2 have a considerable number of papers, so we sampled 75 from each class. Class 1 and 0 only have 71 and 59 papers, respectively, thus all of those papers were included in our analysis. Altogether, this sample contains 430 papers.

The aggregated result is shown in Figure 3. Be reminded that in this analysis, we omitted the papers that do not refer to an artefact type or those that do not promise to release the artefact they produced. A 'No' is marked if a link was not given, a given link is not working, or the repository corresponding to the link does not have the promised artefact (we clicked through and followed all the links mentioned in the papers).

We notice that a considerable portion of papers that promised to release *data* have 'dead-ends' when trying to locate it. This count is higher in low-resource languages. Most tools are hosted on personal or institutional websites, and a portion seems to have fallen out of maintenance in the intervening years. The 'dead-end' problem exists to a lesser degree concerning *code* availability. However, even for *code*, class 0 has a noticeable number of 'dead-ends'. Overall, most of the links to *code* are active and have the artefact, followed by those that promise to release an *LM*.

We also record the common repositories used by NLP researchers and provide a summary in Table 2 (A breakdown of the same data across language classes is available in Figure 7 in the Appendix). According to this, *GitHub* seems to be the most favourite option to release data and code. Some research has considered *Zenodo* and *Hugging Face* for data release¹⁰. In contrast, Hugging Face seems to be the favourite choice for LM releases. Most of the tools have their own unique web link, hence the ‘other’ category is the highest for this type.

Repository	Code	Data	LMs	Tools	Total
GitHub	153	188	17	12	370
Hugging Face	0	6	11	2	19
Zenodo	1	10	1	0	12
Google Drive	0	5	3	1	9
Bitbucket	4	0	0	1	5
GitLab	3	2	0	0	5
Codeberg	1	1	0	0	2
Dropbox	0	1	0	0	1
Mendeley	0	1	0	0	1
Other	5	58	6	44	113
Total	167	272	38	60	537

Table 2: Repository usage across all classes

3.4 Analysis Based on NLP Tasks


Next, we carry out an analysis based on NLP tasks, to understand whether artefact release has any relationship to the type of NLP task¹¹. This analysis was conducted using the paper sample used in Section 3.3. Table 6 in the Appendix shows the raw counts. *Translation* is the NLP task¹² that has the highest number of artefact releases (this artefact is usually parallel data), followed by *morphological analyzer* and *Automatic Speech Recognition (ASR)*. In particular, having morphological analysis as the prevalent NLP domain seems to be common for extremely low-resource languages. This is not surprising - these languages have never had such linguistic resources, and such research is essential in understanding their linguistic properties. The high amount of ASR-related artefacts could be due to the existence of languages that do not have a writing system¹³.

¹⁰This result tallies with the survey results published by [Ranathunga and de Silva \(2022\)](#) to a good extent.

¹¹Initial categorisation of tasks come from Hugging Face task list and a survey paper on NLP research ([de Silva, 2019](#))

¹²As shown in Table 6, *Corpora* has the highest raw counts but is not an *NLP Task* per se.

¹³[Eberhard et al. \(2024\)](#) notes that around 41% of the languages they list may be unwritten.



(a) HF Dataset Counts

(b) HF LM Counts

Figure 4: Number of resources for the language classes on Hugging Face (HF).

3.5 Dataset and LM Availability

Our final analysis is based on the datasets and LM counts reported in Hugging Face¹⁴, which is the fastest-growing repository for NLP-related artefacts. Figure 4 shows¹⁵ the language class-wise distribution of data and LMs. Further, Table 3 shows relevant numerical values, which demonstrates the language class-wise disparity.

Artefact type	Median of Language Class					
	0	1	2	3	4	5
Data set counts	0.0	12.0	53.0	147.5	246.0	657.0
LM counts	0.0	3.0	171.5	443.5	881.0	2601.0

Table 3: Hugging Face Resource Counts

The disparity between different language classes is evident from the medians, despite some outliers. Most notably, out of the 6135 languages in class 0, most have no data or LMs, therefore the handful of languages that have some data/LM have become outliers. The correlation between the class-wise LM and data availability is evident - a Pearson correlation value of 0.9972 is reported between the data and LM counts on languages listed in Hugging Face.

4 Conclusion

We hope our findings would help the NLP community to better appreciate the benefit of openness and to commit to releasing the artefacts they produce. We further hope these statistics will be useful to ACL in making informed decisions. It would be interesting to run this same experiment 5 or 10 years down the line, to see if there are any changes in releasing and reusing artefacts. In hopes to assist in such efforts, our code is publicly released¹⁶.

¹⁴<https://huggingface.co/languages>

¹⁵A larger version is available as Figure 8 in the Appendix.

¹⁶<https://bit.ly/ACL2024ShouldersOfGiants>

5 Limitations

We considered only a fraction of the papers published in AA. Our keyword-based paper filtering mechanism might have missed some papers that have made their artefacts available. If a paper does not mention the language name in its abstract, our algorithm does not pick it up. Thus we highly encourage the community to adhere to ‘Bender Rule’ (Bender, 2019). If a research published their artefact without mentioning that in their paper, or if the link to the artefact was included in a different version of the paper (e.g. ArXiv), such are missed. We might have missed some information on artefacts while manually reading hundreds of research papers, which might have impacted the statistics we present. When checking if a repository link is live, we clicked on that link only once. There could have been instances where the link was momentarily down. In certain instances, we noticed that a URL is not working due to a change in the web repository directory structure. However, we did not try to manually figure out the correct link. We consider an artefact to be available in a repository if we note the availability of files (e.g. python files in a code base) inside the repository. We cannot guarantee the repository has all the artefacts the paper promised (e.g. all the promised data files or whether the given code is working).

6 Ethics Statement

We only used the AA paper repository, which is freely available for research. Our implementation is based on publicly available code. We do not release the paper-wise information we recorded, nor do we re-publish the papers we downloaded from AA.

References

- Emily Bender. 2019. [The #Benderrule: On naming the languages we study and why it matters](#). *The Gradient*, 14.
- Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit*. " O’Reilly Media, Inc."
- Damian Blasi, Antonios Anastasopoulos, and Graham Neubig. 2022. [Systematic inequalities in language technology performance across the world’s languages](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5486–5505, Dublin, Ireland. Association for Computational Linguistics.
- Andrew Cains. 2019. The geographic diversity of NLP conferences. *MAREK REI*.
- Nisansa de Silva. 2019. Survey on publicly available sinhala natural language processing tools and research. *arXiv preprint arXiv:1906.02358*.
- David M. Eberhard, Gary F. Simons, and Charles D. Fennig. 2024. *Ethnologue: How many languages in the world are unwritten?* Dallas, Texas: SIL International.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Simran Khanuja, Sebastian Ruder, and Partha Talukdar. 2023. [Evaluating the diversity, equity, and inclusion of NLP technology: A case study for Indian languages](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1763–1777, Dubrovnik, Croatia. Association for Computational Linguistics.
- Joelle Pineau, Philippe Vincent-Lamarre, Koustuv Sinha, Vincent Larivière, Alina Beygelzimer, Florence d’Alché Buc, Emily Fox, and Hugo Larochelle. 2021. Improving Reproducibility in Machine Learning Research (A Report from the NeurIPS 2019 Reproducibility Program). *Journal of Machine Learning Research*, 22(164):1–20.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al. 2011. The Kaldi Speech Recognition Toolkit. In *IEEE 2011 workshop on automatic speech recognition and understanding*, CONF. IEEE Signal Processing Society.
- Surangika Ranathunga and Nisansa de Silva. 2022. [Some languages are more equal than others: Probing deeper into the linguistic disparity in the NLP world](#). In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 823–848, Online only. Association for Computational Linguistics.
- Shaurya Rohatgi. 2022. [ACL Anthology Corpus with Full Text](#). GitHub.
- Mukund Rungta, Janvijay Singh, Saif M. Mohammad, and Diyi Yang. 2022. [Geographic citation gaps in NLP research](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1371–1383, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A Language Category Definition

Class	Description	Language	
		Count	Examples
0	Have exceptionally limited resources, and have rarely been considered in language technologies.	2191	Slovene Sinhala
1	Have some unlabelled data; however, collecting labelled data is challenging.	222	Nepali Telugu
2	A small set of labelled datasets has been collected, and language support communities are there to support the language.	19	Zulu Irish
3	Has a strong web presence, and a cultural community that backs it. Have highly benefited from unsupervised pre-training.	28	Afrikaans Urdu
4	Have a large amount of unlabelled data, and lesser, but still a significant amount of labelled data have dedicated NLP communities researching these languages.	18	Russian Ukrainian
5	Have a dominant online presence. There have been massive investments in the development of resources and technologies.	7	English Japanese

Table 4: Language Category definition by Joshi et al. (2020)

B Main Conference and Journal List

(1) Annual Meeting of the Association for Computational Linguistics, (2) North American Chapter of the Association for Computational Linguistics, (3) European Chapter of the Association for Computational Linguistics, (4) Empirical Methods in Natural Language Processing, (5) International Conference on Computational Linguistics, (6) Conference on Computational Natural Language Learning (7) International Workshop on Semantic Evaluation, (8) Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, and (9) Conference on Computational Natural Language Learning.

In addition, the following journals are considered: (1) Transactions of the Association for Computational Linguistics and (2) Computational Linguistics.

C Artefact Annotation Scheme

All the annotators involved in this study are coauthors of the paper. In Table 5 we show the annotation scheme we used.

Artefact	Status
Data	Used dataset from some previous research
	Extended an existing dataset
	Used dataset from some previous research but created new data as well
	Introduce new dataset
	Data not needed
Code	Cannot determine
	Used an implementation from some previous research
	Extended an existing implementation (e.g. toolkit, library)
	Used an implementation from some previous research but implemented part of the solution from scratch
	Provided their implementation
LM	Code not needed
	Cannot determine
	Used an existing LM
	Extended an existing LM
	Used an existing LM but trained their LM(s) as well
	Trained their own LM
	LM not needed
	Cannot determine

Table 5: Possible options for Artefacts

D Code and Data Reuse

Code and data from [Ranathunga and de Silva \(2022\)](#) and [Rohatgi \(2022\)](#) are released under **CC BY-NC 4.0** licence. The authors obtained permission from [Blasi et al. \(2022\)](#) to use the code on their public repository¹⁷.

E NLP Task Breakdown Across Language Classes

We show the NLP task breakdown across the five language classes in Table 6.

F Code and Data Intended Use

All the code use was consistent with their intended use as specified on the relevant research publications ([Ranathunga and de Silva, 2022](#); [Blasi et al., 2022](#)) and the readme files on the repositories ([Rohatgi, 2022](#)).

G Artefact Creation, Extension, and Reuse

In Figure 5 we have the larger version of the Figure 1 for improved readability. Further, given that the information in Figure 5 is presented after aggregating across time but separated into language classes, we also include a set of cumulative percentage graphs in Figure 6 where we show the same data aggregated across the language classes but spread out over the publication years to better

¹⁷<https://github.com/neubig/globalutility>

NLP Task	Language Class						Total
	0	1	2	3	4	5	
Corpora	19	22	11	11	11	29	103
Translation	10	12	8	6	10	6	52
Morphological Analyzer	11	8	2	3	1	0	25
Automatic Speech Recognition (ASR)	5	1	10	4	3	0	23
Language Model	1	2	10	1	2	5	21
Parsers	4	5	3	1	3	4	20
Data Sets	6	1	5	3	4	0	19
Dictionary/Lexicon	6	4	1	1	3	3	18
Named-Entity Recognition (NER)	1	0	3	5	7	2	18
Text Classification	1	2	1	2	1	9	16
Part of Speech (PoS)	1	6	3	2	2	1	15
Cross-Lingual Applications	2	1	3	6	2	0	14
Text Generation	0	0	0	0	6	4	10
Hate Speech Detection	0	0	2	6	1	0	9
Misinformation Detection	0	0	0	4	3	1	8
Wordnets/Ontology/Taxonomy	3	1	0	2	0	1	7
Discourse Analysis	0	2	1	1	2	1	7
Question and Answer (QnA)	0	1	2	1	3	0	7
NLP Tools	1	4	1	0	0	0	6
Semantic (Other)	0	0	0	0	1	5	6
Tokenizer	0	0	1	2	0	2	5
Semantic Similarity	0	0	0	3	1	1	5
Multiple Tasks	0	0	1	3	0	0	4
Spelling and Grammar	0	1	1	0	1	1	4
Summarizing	0	0	0	3	1	0	4
Phonological Analyzer	0	1	0	2	1	0	4
Sentiment Analyzer	0	0	1	1	2	0	4
Text-to-Speech	0	2	1	0	0	0	3
Transliteration	0	0	2	0	1	0	3
Lexical Inference	0	0	0	0	3	0	3
Coreference Resolution	0	0	0	0	3	0	3
Information Extraction	0	0	1	1	0	0	2
Bilingual Lexicon Induction (BLI)	0	1	0	0	1	0	2
Optical Character Recognition (OCR)	0	1	0	0	0	1	2
Language Identification (LangID)	0	0	1	0	0	0	1
Intent Detection	1	0	0	0	0	0	1
News/Social Media Recommendation	0	0	0	0	1	0	1
Text Classification	0	0	1	0	0	0	1
Stemming	0	0	0	0	1	0	1
Total	72	78	76	74	81	76	457

Table 6: NLP Tasks Conducted

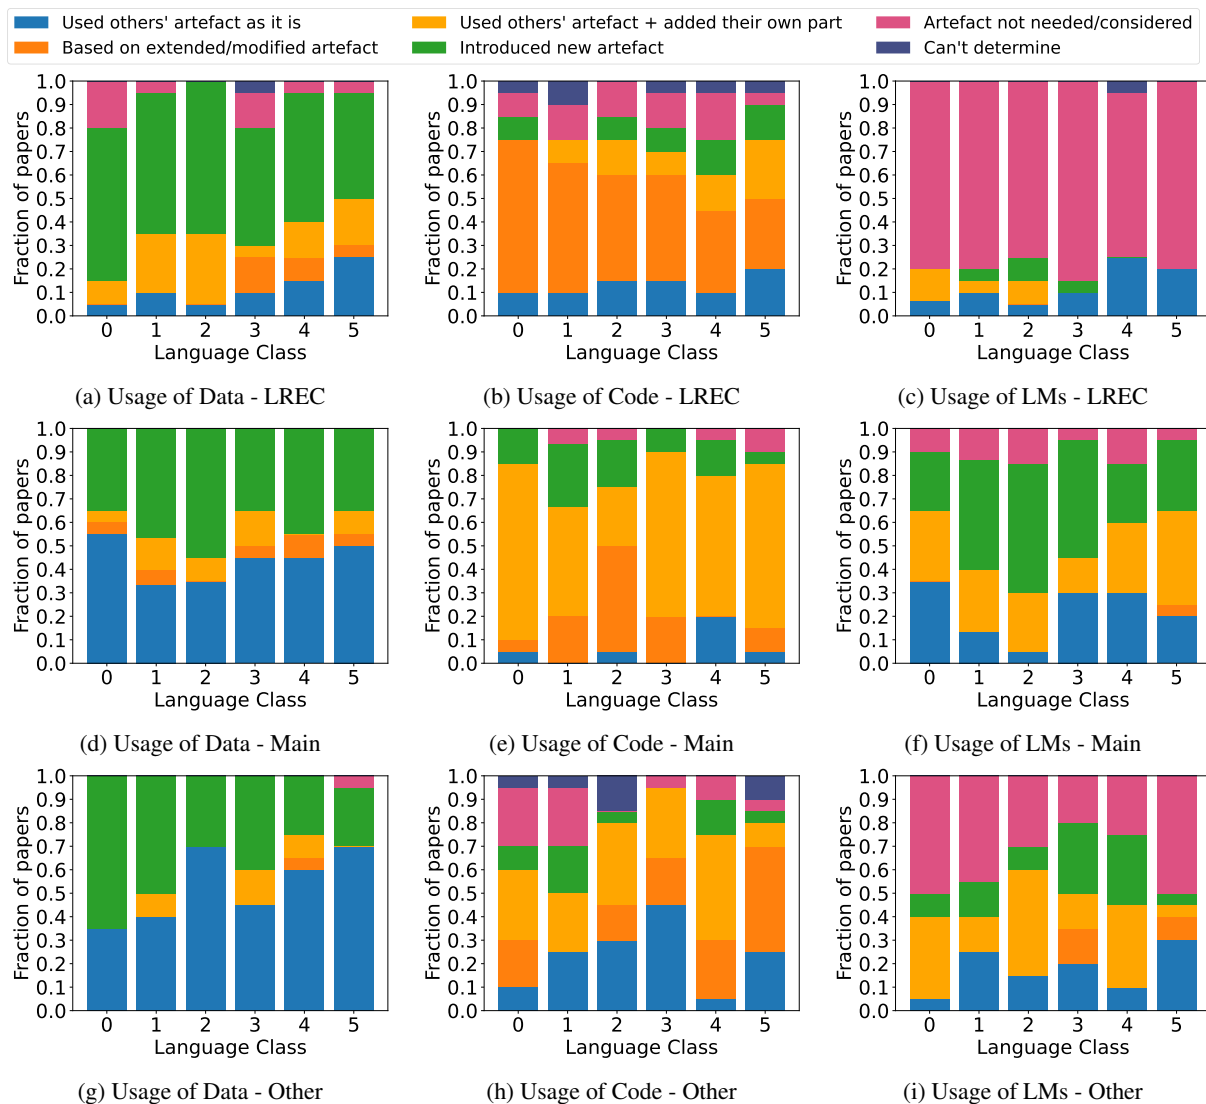


Figure 5: Artefact (Data, Code, LM) creation, extension, and reuse across ACL venues - Aggregated analysis

show the changing trends in resource availability and reuse. Unsurprisingly, as per Figures 6b, 6e, and 6h, we can see that *code* is being re-used the most across all venues. LREC (Figure 6a) stands out among the *data* graphs (Figures 6d and 6g) for consistently being a source of new data sets rather than a venue where existing data is reused. We see that LMs, had a reasonable presence in the main venues (Figure 6f) even before our analysis period while in the *other* venues (Figure 6i), the trend starts just at the beginning of our considered time period. LREC on the other hand, seems to be late to be considered for LMs as it is only in 2018, that we see them becoming noticeable in Figure 6c.

H Artefact Hosting

Table 2 shows a summary of where NLP researchers have published their data, based on the

information mentioned in the research papers. According to this, *GitHub* seems to be the most favourite option to release data and code. Some research has considered *Zenodo* and *Hugging Face* for data release¹⁸. In contrast, Hugging Face seems to be the favourite choice for LM releases. Most of the tools have their own unique web link, hence the ‘other’ category is the highest for this type.

In Figure 7 we show a more detailed view of the artefacts being hosted online; previously discussed in Table 2 as a summary. Here it is possible to note the variations between the language classes. For example, the interesting observation of Figure 7c is that it can be noted that while researchers in all other listed language classes use github to host their trained LMs, the researchers of Class 4

¹⁸This result tallies with the survey results published by Ranathunga and de Silva (2022) to a good extent.

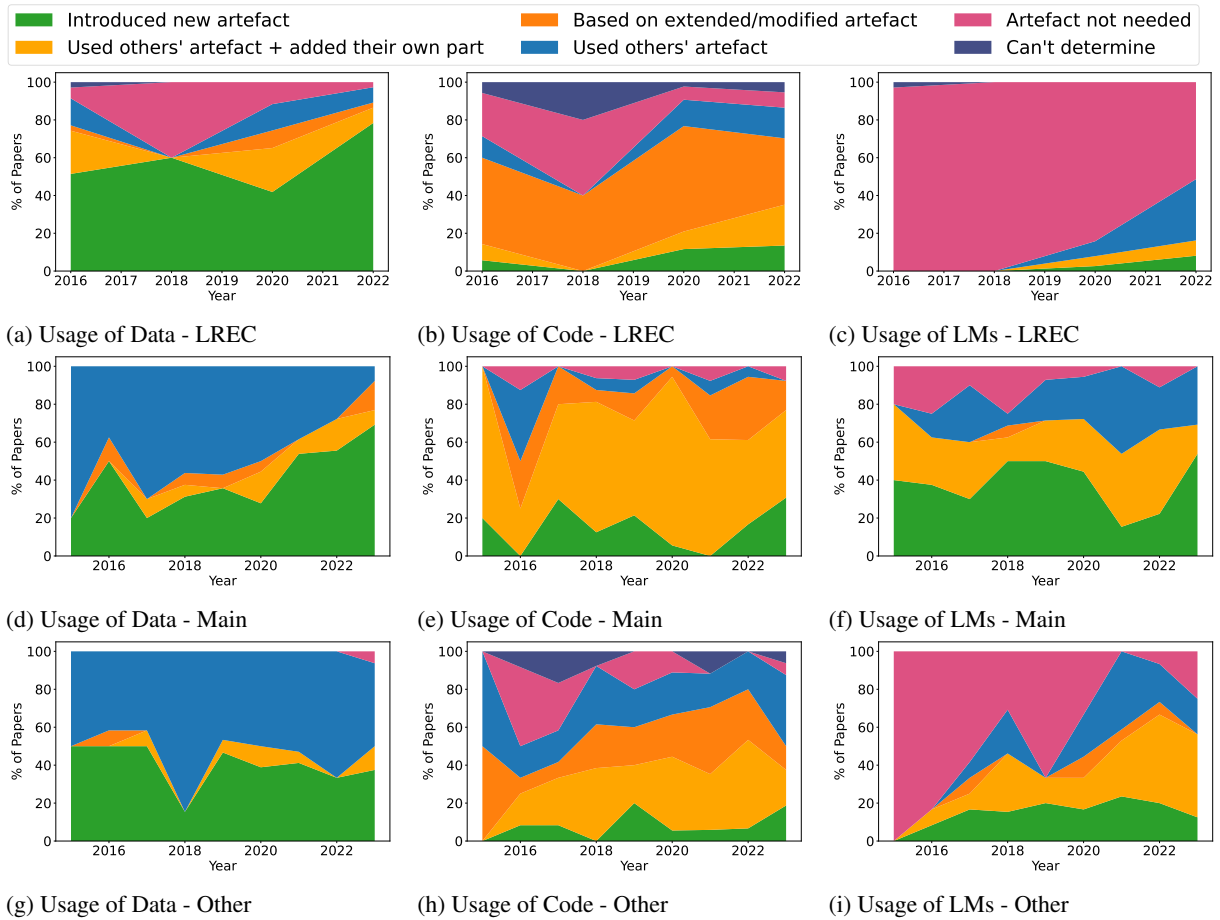


Figure 6: Cumulative percentage graphs - Artefact (Data, Code, LM) creation, extension, and reuse across ACL venues. - Chronological analysis.

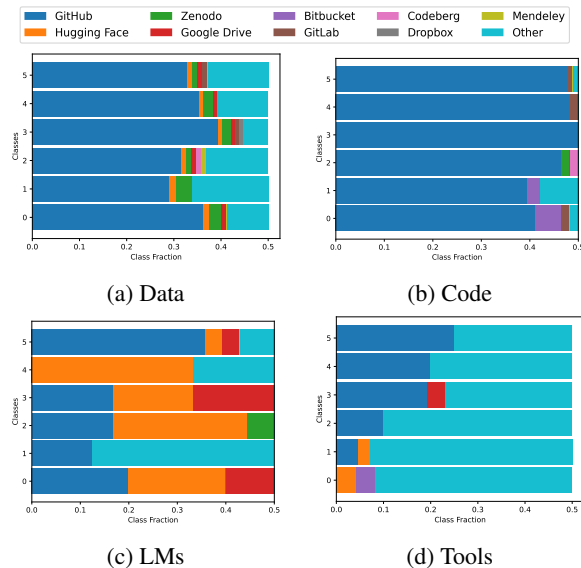


Figure 7: Artefact (Data, Code, LM, Tools) hosting locations.

languages opt for Hugging Face. Conversely, from Figure 7d, it can be noted that in Class 0 languages,

tools are generally not hosted on github. A curious observation in Figure 7c is that for some reason, Class 1 languages do not select Hugging Face as a clear contender to host their language models, something that all other language classes seem to do. The overwhelming prevalence of the *other* option in Figure 7d can be explained by the fact that most tools tend to be hosted on dedicated websites. Even when the actual site is hosted on a service such as github, they are masked with shorter and more market-friendly custom URLs.

I Hugging Face Resources

In Figure 8 we show the resources available on Hugging Face for the 5 language classes. This is a larger version of the Figure 4 for improved readability. Note especially how the entire interquartile range of class 0 is at zero due to the dearth of resources existing for the languages in that class. Thus a language in class 0 with *any* amount of resources gets registered as an outlier. On the opposite end of the spectrum, note class 5 with only




Figure 8: Number of Hugging Face (HF) resources for the language classes.

7 languages in the set even after the reclassification by [Ranathunga and de Silva \(2022\)](#). Despite that, English still manages to be an outlier with its exceptional resource availability.

From Figure 8 and Table 3, it can be observed a considerable jump between the median values when comparing adjacent classes. This may be taken as both: 1) an indication of the visible difference in the resource availability of the language classes, 2) A reaffirmation of the soundness of the class borders proposed by [Ranathunga and de Silva \(2022\)](#) as the distinct medians can be taken as a quality of classes which are internally cohesive and mutually separate.