

Overview of the CLEF 2023 SimpleText Track

Automatic Simplification of Scientific Text

Liana Ermakova Éric SanJuan Stéphane Huet
Hosein Azarbondyad Olivier Augereau Jaap Kamps



TURN IT SIMPLE



CLEF 2023, Thessaloniki, Greece, September 18, 2023

SimpleText Track Motivation



- Track focuses on **Scientific Text Simplification**
- In a world of fake news and misinformation
 - Objective, evidence based science is more important than ever before
 - But hard to find and hard to understand for non-experts
 - Text simplification can remove some of these barriers!
- Also useful for:
 - Promoting science literacy
 - Science journalism and communication
 - Education

Zero-shot Text Simplification



- *Scientific Abstract (FKGL 17.0 ~ University graduate school)*

Searching scientific literature and understanding technical scientific documents can be very difficult for users as there are a vast number of scientific publications on almost any topic and the language of science, by its very nature, can be complex. Scientific content providers and publishers should have mechanisms to help users with both searching the content in an effective way and understanding the complex nature of scientific concepts. . . .

- *GPT-2 revisions (FKGL 12.9 ~ High school diploma)*

Searching for scientific literature and ~~understanding~~ technical scientific documents can be very ~~difficult~~ time-consuming for users as there are a vast number of scientific publications on almost any topic and the language of science , by its very nature , can be ~~complex~~ very confusing . Scientific content providers and publishers should have mechanisms to help users ~~with both searching~~ find the content right information in an effective way , and understanding the ~~complex~~ nature of scientific concepts

Goals



- **Use case:** non-expert searching for popular science topics in the scientific literature
- **Overall task:**
 - provide an overview of the scientific literature on this topic
 - create a simplified summary of multiple scientific documents
- We break down this problem in **three IR+NLP subtasks** and provide appropriate reusable **data** and **benchmarks**.
- Track provides a forum to collaborate on **technical & evaluation** challenges of scientific text simplification

CLEF 2023 SimpleText: Three Tasks



Three IR and NLP tasks forming together a pipeline:

- Task 1 (*What is in, or out?*):
 - Selecting passages to include in a simplified summary
 - → Retrieval of scientific articles in response to a popular science query
- Task 2 (*What is unclear?*):
 - Difficult concept identification and explanation
 - → Complex terminology spotting in sentences from scientific abstracts
- Task 3 (*Rewrite this!*):
 - Rewriting scientific text
 - → Rewriting/simplifying sentences from scientific abstracts

SimpleText'23 Submission Stats



Team	Task 1	Task 2.1	Task 2.2	Task 3	Total runs
<i>Elsevier</i>	10				10
<i>Maine (Aiirlab)</i>	10	3	3	2	18
<i>uninib_DoSSIER</i>	2				2
<i>UAMs</i>	10	1		2	13
<i>LIA</i>	7				7
<i>MiCroGerk</i>		4	4	3	11
<i>Croland</i>		2	2		
<i>NLPalma</i>		1	1	1	3
<i>Pandas</i>				6	6
<i>QH</i>				3	3
<i>SINAI</i>		4	2		
<i>irgc</i>				4	4
<i>CYUT</i>				4	4
<i>UOL-SRIS</i>		1			1
<i>Smroltra</i>		10	10	1	21
<i>TeamCAU</i>		3	3	1	7
<i>TheLangVerse</i>		1	1	1	3
<i>ThePunDetectives</i>		2	2	2	6
<i>UBO</i>		7	1	1	9
<i>RT</i>				1	1
Total runs	39	39	29	32	139

Task 1: What is in, or out?



- *Task 1: Selecting passages to include in a simplified summary*
- Given a popular science article targeted to a general audience, this task aims to retrieve passages that can help understand this article from a large corpus of academic abstracts and bibliographic metadata
- Citation Network Dataset: DBLP+Citation, ACM Citation network
 - 4,232,520 abstracts in English
- Topics based on 40 press articles + 114 manually extracted queries
 - 20 articles from *The Guardian*
 - 20 articles from *Tech Xplore*

Task 1: Examples



- Text of news articles as context (the **topic**)
 - 1 Patient data from GP surgeries sold to US companies
 - 2 Baffled by digital marketing? Find your way out of the maze
- Input: **query** based on these articles
 - 1 patient data
 - 2 digital marketing
 - 2 advertising
- Output:
 - Given the corpus of 4M articles (metadata+abstracts)
 - rank a list of abstracts relevant to the topic/query
 - in JSON format (\sim trec_eval + passage)

Task 1: Evaluation



Qrels	Topics	#Queries	# 0	# 1	# 2
2023 train	G01-G15	29	728	338	237
2023 test	G16-G20, T01-T05	34	2,260	357	1,218

- Train data for system development:
 - 15 popular news topics, with 29 specific queries.
 - Judgments on top 50 abstracts retrieved by ElasticSearch
- Test data:
 - 10 popular news topics, with 34 specific queries
 - Judgments on top 10 abstracts retrieved by all runs
- Evaluation measures:
 - traditional IR metrics (relevance): NDCG, MAP, ...
 - additional complexity/credibility aspect evaluation
 - can we rank relevant and accessible results first?

Task 1: Results on Test Data



Run	MRR	Precision		NDCG		Bpref	MAP
		10	20	10	20		
ElsevierSimpleText_run8	0.8082	0.5618	0.3515	0.5881	0.4422	0.2371	0.1633
ElsevierSimpleText_run7	0.7136	0.5618	0.4103	0.5704	0.4627	0.2626	0.1915
maine_CrossEncoder1 ^{rel}	0.8106	0.5382	0.4456	0.5675	0.4908	0.3317	0.2810
maine_CrossEncoderFinetuned1 ^{rel}	0.7691	0.5559	0.4441	0.5542	0.4840	0.3433	0.2572
maine_CrossEncoder1 ^{comb}	0.7309	0.5265	0.4500	0.5455	0.4841	0.3337	0.2754
ElsevierSimpleText_run5	0.6600	0.4765	0.3838	0.4826	0.4186	0.2542	0.1828
UAms_CE100 ^{rel}	0.7050	0.4912	0.4044	0.4782	0.4236	0.2616	0.2011
UAms_CE1k_Filter	0.6403	0.4765	0.3559	0.4533	0.3743	0.2727	0.1936
UAms_CE1k ^{rel}	0.6329	0.4735	0.4044	0.4448	0.4049	0.2797	0.2051
Elastic baseline	0.6424	0.4059	0.3456	0.3910	0.3541	0.2501	0.1895
unimib_DoSSIER_2	0.5201	0.2853	0.2515	0.2980	0.2683	0.1898	0.1141
unimib_DoSSIER_4	0.5202	0.2853	0.2441	0.2972	0.2632	0.1873	0.1111
run-LIA.bm25	0.4536	0.1912	0.1338	0.2192	0.1700	0.1384	0.0515
run-LIA.all-MiniLM-L6-v2.query	0.3505	0.2000	0.1662	0.2019	0.1767	0.1956	0.0667
run-LIA.all-MiniLM-L6-v2.query-topic	0.3655	0.1765	0.1485	0.1912	0.1647	0.2043	0.0591

- Neural rankers outcompete lexical systems by a large margin
- In particular precision gains, some also recall
- Some submissions prioritized other aspects than relevance

Task 1: Complexity Aware Ranking?



Run	NDCG		FKGL	
	10	20	Mean	Median
ElsevierSimpleText_run8	0.5881	0.4422	14.1	14.0
UAMS_CE1k_Filter	0.4533	0.3743	12.2	12.4
Elastic baseline	0.3910	0.3541	14.0	13.9
UAMS_CE1k_Combine	0.3706	0.3334	11.2	11.2

- Runs targeting relevant and more accessible abstracts
 - Performing competitive on retrieval effectiveness
 - readability level from "university" to "high school"
 - → avoiding very complex (but relevant) abstracts
 - *What would our user prefer?*

Task 1: Findings



- Scientific passage retrieval test collection constructed in 2022-2023
 - High pooling diversity
 - Reusable with limited pooling bias
- Almost all submissions based on neural rankers
 - Crosscoders and biencoders popular and **very effective**
 - Training on scientific text helps
- Promising results for runs prioritizing credibility/complexity
 - Possible to factor the text complexity into the ranking
 - Guide users to accessible content first, and more complex text later

Task 2: What is unclear?



- *Task 2: Difficult concept identification and explanation*
- Identify **key concepts** that need to be contextualized with a definition, example, and/or use-case and provide useful and understandable explanations for them.
- Two subtasks:
 - 1 to retrieve up to 5 difficult terms in a given passage from a scientific abstract;
 - 2 to provide an explanation of these difficult terms (e.g. definition, abbreviation deciphering, example, etc.).

Task 2.1: Example



- Input:

```
{ "query_id": "G14.2",  
  "query_text": "end to end encryption",  
  "doc_id": "2884788726",  
  "snt_id": "G14.2_2884788726_2",  
  "source_snt": "However, in information-centric networking (ICN)  
the end-to-end encryption makes the content caching ineffective since  
encrypted content stored in a cache is useless for any consumer  
except those who know the encryption key." }
```

- Output (1):

```
{ "snt_id": "G14.2_2884788726_2",  
  "term": "content caching",  
  "difficulty": 1.0,  
  "run_id": "team1_task_2.1_TFIDF",  
  "manual": 0 }
```

Task 2.1: Results for the official runs



	Total	Test		Train	
		Evaluated	+Limits	Score	+Limits
SINAI_task_2.1_PRM_ZS_TASK2_1_V1	11081	1185	507	94	56
UAMS_Task_2_RareIDF	675090	1145	241	40	21
SINAI_task_2.1_PRM_FS_TASK2_1_V1	10768	1122	405	81	40
Smroltra_task_2.1_keyBERT_FKgrade	11099	1061	341	41	4
Smroltra_task_2.1_keyBERT_F	11099	1061	171	41	7
UOL-SRIS_2.1_KeyBERT	23757	1061	0	42	0
MiCroGerk_task_2.1_TextRank	21516	1002	391	87	61
Smroltra_task_2.1_TextRank_FKgrade	10056	1002	363	87	29
SINAI_task_2.1_PRM_ZS_TASK2_1_V2	10952	965	330	94	53
SINAI_task_2.1_PRM_FS_TASK2_1_V2	8836	915	316	76	41
Smroltra_task_2.1_YAKE_D	11112	905	422	71	21
MiCroGerk_task_2.1_YAKE	23790	905	362	71	51
Smroltra_task_2.1_YAKE_Fscore	11112	905	209	71	32
MiCroGerk_task_2.1_GPT-3	15892	889	459	79	43
UBO_task_2.1_FirstPhrases	14088	831	161	49	19
UBO_task_2.1_PositionRank	13881	825	181	71	29
UBO_task_2.1_SingleRank	14088	748	151	67	19
UBO_task_2.1_TfIdf	14340	740	187	50	13
UBO_task_2.1_TextRank	14088	722	139	67	16
Smroltra_task_2.1_RAKE_AUI	10660	713	288	48	25
Smroltra_task_2.1_RAKE_F	10660	713	170	48	21
UBO_task_2.1_TopicRank	13912	663	144	61	21
UBO_task_2.1_YAKE	14337	576	116	44	11
MiCroGerk_task_2.1_BLOOM	9600	535	218	64	34
Aiirlab_task_2.2_KBIR	4797	429	135	38	11
TeamCAU_task_2.1_ST5	2234	418	201	90	79
Smroltra_task_2.1_SimpleT5	2234	406	239	82	74

Conclusions from Task 2.1



- Used methods:
 - LLMs
 - unsupervised methods
- Results of difficult term detection by LLMs are comparable to RareIDF, TextRank and YAKE!
- Results of the same methods depend heavily on implementation

Task 2.2: Example



- Output (2):

```
{  
  "snt_id": "G14.2_2884788726_2",  
  "term": "content caching",  
  "difficulty": 1.0,  
  "definition": "Content caching is a performance optimization  
mechanism in which data is delivered from the closest servers for  
optimal application performance.",  
  "run_id": "team1_task_2.2_TFIDF_BLOOM",  
  "manual": 0  
}
```

Term explanation results



Run	Evaluated	BLEU	ROUGE	Semantic match
UBO_task_2.1_FirstPhrases_Wikipedia	393	29.73	0.41	0.80
Croland_task_2_PKE_Wiki	43	33.68	0.46	0.70
MiCroGerk_task_2.2_GPT-3_Wikipedia	932	26.38	0.41	0.75
Smroltra_task_2.2_Text_Wiki	547	17.59	0.33	0.75
Smroltra_task_2.2_RAKE_Wiki	337	16.95	0.32	0.74
Smroltra_task_2.2_YAKE_Wiki	436	16.94	0.32	0.73
TeamCAU_task_2.1_BLOOM	10	10.46	0.27	0.48
MiCroGerk_task_2.2_GPT-3_BLOOMZ	1,108	9.07	0.40	0.83
Smroltra_task_2.2_keyBERT_Wiki	302	8.60	0.23	0.69
MiCroGerk_task_2.2_GPT-3_GPT-3	1,108	7.73	0.38	0.83
NLPalma_task_2.2_BERT_BLOOMZ	537	7.22	0.39	0.76
Smroltra_task_2.2_Bloomz	23	7.15	0.30	0.69
TeamCAU_task_2.1_AI21	22	6.38	0.31	0.78
TheLangVerse_task_2.2_openai-curie-finetuned	444	5.03	0.25	0.74
Croland_task_2_GPT3	69	4.83	0.27	0.77
SINAI_task_2.1_PRM_FS_TASK2_2_V1	649	4.23	0.21	0.78
MiCroGerk_task_2.2_GPT-3_simpleT5	1,108	4.22	0.28	0.77
TeamCAU_task_2.1_ST5	379	3.33	0.20	0.60
Smroltra_task_2.2_SimpleT5	392	3.09	0.22	0.72
SINAI_task_2.1_PRM_ZS_TASK2_2_V1	649	3.08	0.19	0.69
Smroltra_task_2.2_keyBERT_dict	120	2.07	0.14	0.51
Smroltra_task_2.2_YAKE_WN	48	1.88	0.15	0.44
Aiirlab_task_2.2_KBIR	556	1.62	0.15	0.50
Smroltra_task_2.2_keyBERT_WN	328	1.33	0.14	0.45
Aiirlab_task_2.2_YAKEIDF	179	1.13	0.14	0.41
Aiirlab_task_2.2_YAKE	165	1.10	0.15	0.43
Smroltra_task_2.2_RAKE_WN	70	0.00	0.14	0.46

Abbreviation expansion results



Run	Evaluated	BLEU	ROUGE	Semantic	Exact	Partial
MiCroGerk_task_2.2_GPT-3_BLOOMZ	854	13.87	0.68	0.76	326	185
MiCroGerk_task_2.2_GPT-3_GPT-3	855	11.86	0.64	0.73	294	166
MiCroGerk_task_2.2_GPT-3_Wikipedia	855	4.68	0.43	0.60	205	109
MiCroGerk_task_2.2_GPT-3_Wikipedia	618	5.01	0.56	0.64	198	109
NLPalma_task_2.2_BERT_BLOOMZ	345	6.83	0.39	0.52	50	47
Smroltra_task_2.2_SimpleT5	185	0.00	0.12	0.39	8	7
TeamCAU_task_2.1_ST5	141	1.48	0.14	0.40	6	3
TheLangVerse_task_2.2_openai-curie-finetuned	204	1.60	0.14	0.42	1	2
SINAI_task_2.1_PRM_ZS_TASK2_2_V1	228	1.61	0.13	0.55	1	0
TeamCAU_task_2.1_AI21	10	1.87	0.14	0.38	0	0
SINAI_task_2.1_PRM_FS_TASK2_2_V1	228	1.35	0.10	0.53	0	0
UBO_task_2.1_FirstPhrases_Wikipedia	116	5.09	0.19	0.47	0	0
Aiirlab_task_2.2_KBIR	202	1.17	0.07	0.44	0	0
Smroltra_task_2.2_RAKE_Wiki	27	0.54	0.04	0.14	0	0
Smroltra_task_2.2_Bloomz	4	0	0.22	0.61	0	0
Aiirlab_task_2.2_YAKEIDF	19	0	0.10	0.40	0	0
Smroltra_task_2.2_keyBERT_WN	188	0	0.04	0.27	0	0
Smroltra_task_2.2_keyBERT_Wiki	163	0.21	0.02	0.13	0	0
Smroltra_task_2.2_keyBERT_dict	46	0	0.04	0.34	0	0
Smroltra_task_2.2_RAKE_WN	21	0	0.04	0.24	0	0
Smroltra_task_2.2_YAKE_WN	32	0	0.02	0.21	0	0
Smroltra_task_2.2_YAKE_Wiki	31	0	0.03	0.11	0	0
Smroltra_task_2.2_Text_Wiki	50	0	0.02	0.10	0	0
Aiirlab_task_2.2_YAKE	9	0	0.13	0.36	0	0
TeamCAU_task_2.1_BLOOM	3	0	0	0.14	0	0

Conclusions from Task 2.2



- High semantic similarity between participants' definitions and the ground truth
- Wikipedia-based runs have the highest similarity
- LLMs (BLOOMz, GPT-3) have the best performance for abbreviation expansion

Task 3: Rewrite this!



- *Task 3: Rewriting scientific text*
- This task aims to provide a simplified version of sentences extracted from scientific abstracts.
- Data
 - Corpus of 2,234 (S), 4,797 (M), and 152,072 (L) sentences from Task 1
 - Train data: 648 manually simplified sentences
 - Test data: 245 manually simplified sentences
- Evaluation
 - Large-scale automatic evaluation measures (SARI, ROUGE, compression, readability)
 - Small-scale detailed human evaluation of other aspects, including information distortion

Task 3: Examples



- Input format:

```
{"snt_id": "G11.1_2892036907_2",  
  "source_snt": "With the ever increasing number of unmanned aerial  
vehicles getting involved in activities in the civilian and commercial  
domain, there is an increased need for autonomy in these systems  
too.",  
  "doc_id": "2892036907",  
  "query_id": "G11.1",  
  "query_text": "drones"}
```
- Output format:

```
{"run_id": "BTU_task_3_run1",  
  "manual": 1,  
  "snt_id": "G11.1_2892036907_2",  
  "simplified_snt": "Drones are increasingly used in the civilian and  
commercial domain and need to be autonomous."}
```

Task 3: Results on Test data



run_id	count	FKGL	SARI	BLEU	Compression ratio	Sentence splits	Levenshtein similarity	Exact copies	Additions proportion	Deletions proportion	Lexical complexity score
Identity_baseline	245	13.64	15.09	26.22	1.00	1.00	1.00	1.00	0.00	0.00	8.64
Reference	245	12.03	100.00	100.00	0.95	1.10	0.66	0.00	0.33	0.40	8.64
CYUT_run1	245	9.63	47.98	14.81	0.87	1.14	0.56	0.0	0.47	0.55	8.35
AiirLab_davinci	243	11.17	47.10	18.68	0.75	1.00	0.68	0.0	0.20	0.45	8.59
irgc_ChatGPT_2stepTurbo	245	12.31	46.98	16.86	0.94	1.04	0.63	0.0	0.37	0.46	8.46
MiCroGerk_GPT-3	245	10.74	46.90	16.98	0.72	1.01	0.67	0.0	0.19	0.47	8.67
ThePunDetectives_GPT-3	245	7.52	41.56	6.10	0.46	0.97	0.50	0.0	0.16	0.68	8.46
Pandas_submission_ensemble	245	10.51	40.25	17.40	0.77	1.09	0.73	0.0	0.15	0.40	8.52
NLPalma_BLOOMZ	245	9.61	35.66	5.76	0.68	1.00	0.51	0.0	0.35	0.66	8.26
UAms_Large_KIS150_Clip	245	11.12	33.47	16.59	1.01	1.23	0.82	0.0	0.24	0.23	8.48
UBO_SimpleT5	245	12.33	30.89	21.08	0.88	1.05	0.89	0.0	0.10	0.22	8.51
TheLangVerse_openai-curie-finetuned	245	12.21	30.78	18.92	0.86	1.00	0.86	0.0	0.11	0.24	8.49
QH_run3	245	12.74	27.56	20.24	0.90	1.01	0.91	0.0	0.09	0.19	8.50
TeamCAU_ST5	245	12.77	27.19	21.06	0.90	1.00	0.91	0.0	0.10	0.20	8.52
Smoltra_SimpleT5	245	12.88	26.25	21.43	0.90	1.00	0.91	0.0	0.09	0.19	8.54

Task 3: Manual Evaluation



- Comparison of manually simplified and source sentences

Metric (Avg)	Source snt	Simplified snt
FKGL	15.16	12.12
# Abbreviations	0.24	0.13
# Difficult terms	0.41	0.28

- Difficulty levels and information distortion of simplified sentences

	1	2	3	4	5	6	7
syntax complexity	259	51	9				
lexical complexity	93	119	62	26	19		
information loss severity	34	25	29	15	16	4	9
information loss severity %	8	6	7	4	4	1	2

- Generally low lexical and grammatical complexity (low FKGL)
- Significant information distortion: 15% missing essential details, 10% oversimplification, 6% insertion false/unsupported information, ...

Issues in Generative LLMs



- Fraction of sentences with hallucination varies from 0 to 55%
- Existing evaluation measures **insensitive to hallucination!**

Run	# Input Sentences	Spurious Content	
		Number	Fraction
AiirLab_run1	245	58	0.24
CYUT_run1	757	12	0.02
MiCroGerk_GPT-3	245	1	0.00
NLPalma_BLOOMZ	245	135	0.55
Pandas_alpaca-lora-both-alpaca-simplifier-tripple_10	245	3	0.01
QH_run3	245	1	0.00
Smroltra_SimpleT5	245	0	0.00
TeamCAU_task_2.2_ST5	245	0	0.00
TheLangVerse_openai-curie-finetuned	245	1	0.00
ThePunDetectives_GPT-3	245	0	0.00
UAms_Large_KIS150	757	213	0.28
UBO_SimpleT5	245	0	0.00
irgc_pegasusTuner007plus_plus	245	57	0.23

Task 3: Main Findings



- Every participant uses LLMs
- Larger models tend to perform better (in particular on test)
 - Very high scores (in particular SARI 0.50)
 - Very good zero shot performance, even on **scientific text**
- Finetuning/training can easily lead to overfitting on train data
 - Unrealistic high SARI and BLEU scores (0.8-0.9)
- Output quality looks very good, useful in practice
 - + Text complexity similar or lower than human simplification
 - + Lexical/grammatical issues very minor
 - – Information loss/distortion issues remain
 - – Complex scientific terminology issues remain
 - – Evaluation measures need to factor in hallucination

SimpleText Program (Room 4)



- Mon 18 Sep 2023 (TODAY)
 - 16:10 - 17:10 Invited talk by Federica Vezzani (Padua) *Easy peasy term squeeze": A terminological perspective on text simplification*
 - 17:10-17:40 CLEF 2023 SimpleText Task 1-3 Overviews
- Tue 19 Sep 2023
 - 09:30 - 11:00 Participant presentations
- Tue 19 Sep 2023
 - 14:00 - 15:00 Participant presentations
 - 15:00 - 15:30 Round table and SimpleText 2024 discussion
Any ideas or volunteers are welcome!



TURN IT SIMPLE



Thanks!

Fully funded 3-years PhD available!

Website : <https://simpletext-project.com>¹

E-mail : contact@simpletext-project.com

Twitter : <https://twitter.com/SimpletextW>

Google group : <https://groups.google.com/g/simpletext>

¹This research was funded, in whole or in part, by the French National Research Agency (ANR) under the project ANR-22-CE23-0019-01.

References

- Liana Ermakova, Eric SanJuan, Stéphane Huet, Hosein Azarbonyad, Olivier Augereau, and Jaap Kamps. “Overview of the CLEF 2023 SimpleText Lab: Automatic Simplification of Scientific Texts”. In: *CLEF'23: Proceedings of the Fourteenth International Conference of the CLEF Association*. Lecture Notes in Computer Science. Springer, 2023.
- Eric SanJuan, Stéphane Huet, Jaap Kamps, and Liana Ermakova. “Overview of the CLEF 2023 SimpleText Task 1: Passage Selection for a Simplified Summary”. In: *Working Notes of CLEF 2023: Conference and Labs of the Evaluation Forum*. CEUR Workshop Proceedings. CEUR-WS.org, 2023.
- Liana Ermakova, Hosein Azarbonyad, Sarah Bertin, and Olivier Augereau. “Overview of the CLEF 2023 SimpleText Task 2: Difficult Concept Identification and Explanation”. In: *Working Notes of CLEF 2023: Conference and Labs of the Evaluation Forum*. CEUR Workshop Proceedings. CEUR-WS.org, 2023.
- Liana Ermakova, Sarah Bertin, Helen McCombie, and Jaap Kamps. “Overview of the CLEF 2023 SimpleText Task 3: Scientific Text simplification”. In: *Working Notes of CLEF 2023: Conference and Labs of the Evaluation Forum*. CEUR Workshop Proceedings. CEUR-WS.org, 2023.