

Results of SemTab 2024

Okkie Hassanzadeh^{1,*}, Nora Abdelmageed², Marco Cremaschi³, Vincenzo Cutrona⁴,
Fabio D’Adda³, Vasilis Efthymiou⁵, Benno Kruit⁶, Elita Lobo⁷,
Nandana Mihindukulasooriya¹ and Nhan H. Pham¹

¹IBM Research, USA

²Friedrich Schiller University Jena, Germany

³University of Milan - Bicocca, Italy

⁴University of Applied Sciences and Arts of Southern Switzerland, Switzerland

⁵Harokopio University of Athens & FORTH-ICS, Greece

⁶Vrije Universiteit Amsterdam, The Netherlands

⁷University of Massachusetts Amherst, USA

Abstract

SemTab 2024 marked the sixth iteration of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, held in conjunction with the 23rd International Semantic Web Conference (ISWC). SemTab serves as a platform for the systematic evaluation of state-of-the-art semantic table interpretation systems. This paper provides an overview of the 2024 challenge and highlights the key outcomes.

Keywords

Tabular data, Knowledge Graphs, Matching, SemTab Challenge, Semantic Table Interpretation

1. Introduction

Tabular data is ubiquitous across the Web, enterprise data lakes, data catalogs, and other repositories, serving as a foundational format in data science and analytics. However, a significant gap often exists between those producing tabular data and those consuming it. Data producers focus on storing, maintaining, and ensuring the availability of raw data, frequently sharing it with minimal metadata or metadata in non-standard or textual forms. In contrast, data consumers must locate the data they need, extract relevant subsets, and refine and integrate the raw data to render it suitable for their applications. Achieving this transformation is often impractical without automated solutions. A cornerstone of such automation is the annotation of data elements with entities, classes, and relationships from a knowledge graph (KG). These annotations facilitate knowledge-based data discovery [1, 2, 3, 4], organization [5], integration [6, 7], and augmentation [8]. Automating the task of linking tabular data to KGs, commonly known as Semantic Table Interpretation (STI), has been extensively studied in the literature [9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19].

The Semantic Web Challenge on Tabular Data to Knowledge Graph Matching (SemTab) started in 2019 with the goal of providing an avenue for benchmarking and evaluation of various STI solutions. Over the years, the SemTab participants have proposed a range of solutions incorporating a variety of approaches to automated matching, with their key strengths and weaknesses analyzed using different datasets and rounds of each of the SemTab editions. In this paper, we provide a high-level summary of the 2024 edition of the SemTab challenge, along with the results.

SemTab’24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024, co-located with the 23rd International Semantic Web Conference (ISWC), November 11-15, 2024, Baltimore, USA

*Corresponding author.

✉ hassanzadeh@us.ibm.com (O. Hassanzadeh); nora.abdelmageed@uni-jena.de (N. Abdelmageed); marco.cremaschi@unimib.it (M. Cremaschi); vincenzo.cutrona@supsi.ch (V. Cutrona); fabio.dadda@unimib.it (F. D’Adda); vefthym@hua.gr (V. Efthymiou); b.b.kruit@vu.nl (B. Kruit); elobo@umass.edu (E. Lobo); nandana@ibm.com (N. Mihindukulasooriya); nhp@ibm.com (N. H. Pham)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

2. The Challenge

The SemTab 2024 challenge comprised three tracks. The **Accuracy Track**, which evaluated the accuracy of semantic table interpretation solutions, continued as in previous years. Two new tracks were introduced this year: the **STI vs LLMs Track**, which focused on assessing cell entity annotation solutions leveraging large language models (LLMs), and the **Table Metadata to KG Track**, which addressed the challenge of matching tabular data using only table metadata. Although a call for datasets was issued, no submissions were received, and the datasets track was consequently omitted from this year's challenge.

2.1. Accuracy Track

The Accuracy Track consisted of two rounds, each featuring three datasets. This year, all datasets were aligned with the same target knowledge graph, Wikidata [20]. Similar to last year, participants submitted their solutions via a submission form, and the results were evaluated at the conclusion of each round. This year we also performed additional rounds of evaluation, with the last round right before the conference.

2.1.1. Datasets

Table 1 provides an overview of the datasets used in the Accuracy Track, along with their corresponding statistics. Similar to the previous year—and in contrast to the earlier editions where ground truth was kept hidden—participants were provided with partial ground truth data during the challenge in the form of training and/or validation sets. These labels enabled teams to evaluate their methods locally. All datasets are openly available on Zenodo. Across the two rounds, three groups of datasets were utilized:

- **WikidataTables** <https://doi.org/10.5281/zenodo.14207232>
This dataset comprises tables generated using an enhanced version of our data generator, which produces realistic-looking tables through SPARQL queries [21]. The target knowledge graph (KG) for this dataset is Wikidata, and, as in previous years, the tasks include Cell Entity Annotation (CEA), Column Type Annotation (CTA), and Column Property Annotation (CPA). As detailed in Table 1, the test set for Round 1 consists of 30,000 tables with an average of 2.5 columns and 61.7 rows, while the Round 2 dataset consists of 78,745 tables with an average of 2.5 columns and 11.6 rows. For these collections, the dataset generator was configured to produce a large number of small to medium-sized tables with high ambiguity in entity columns. This ambiguity was introduced by filtering for labels that can refer to multiple entities in Wikidata.
- **tBiodiv** <https://doi.org/10.5281/zenodo.10283015>
tBiodivL <https://doi.org/10.5281/zenodo.10283083>
These datasets, generated using KG2Tables [22] for the biodiversity domain, include two types of tables: 1) "horizontal" relational tables, where each table represents a collection of entities, and 2) "entity" tables, where each table represents a single entity. Ground truth mappings to Wikidata were provided for the CEA, CTA, and CPA tasks, as well as for the Topic Detection (TD) task, which focuses on annotating an entire table to instances/entities or types/classes, and the Row Annotation (RA) task, which involves mapping each row to an entity. As shown in the statistics in Table 1, the relational table datasets are wider and exhibit greater variation in the number of columns.
- **tBiomed** <https://doi.org/10.5281/zenodo.10283103>
tBiomedL <https://doi.org/10.5281/zenodo.10283119>
These datasets, also generated using KG2Tables [22] for the biomedical domain, include both relational and entity tables. They are accompanied by ground truth mappings for the CEA, CTA, CPA, TD, and RA tasks. As indicated in Table 1, the tBiomed datasets contain a larger number of tables but are comparable to the tBiodiv datasets in terms of the average number of rows and columns.

Table 1

Statistics of the Accuracy Track test datasets in each SemTab 2024 round.

	Tables #	Avg. # Cols	Avg. # Rows
WikidataTablesR1	30,000	2.54 ± 0.78	61.72 ± 178.71
WikidataTablesR2	78,745	2.49 ± 0.67	11.56 ± 7.63
tBiodiv-Relational	421	12.53 ± 14.27	19.01 ± 22.09
tBiodiv-Entity	154	2.00 ± 0.00	7.62 ± 9.32
tBiodiv-Large-Relational	1,616	9.34 ± 10.03	16.19 ± 22.15
tBiodiv-Large-Entity	609	2.00 ± 0.00	7.96 ± 6.68
tBiomed-Relational	1,621	8.32 ± 7.69	16.74 ± 26.08
tBiomed-Entity	1,056	2.00 ± 0.00	5.91 ± 4.71
tBiomed-Large-Relational	5,496	6.98 ± 5.92	16.48 ± 24.06
tBiomed-Large-Entity	3,110	2.00 ± 0.00	6.39 ± 5.42

2.1.2. Evaluation measures

As in prior editions, systems were evaluated based on a single annotation for each specified target across all tasks. For CEA, this meant annotating target cells with a single entity from the target KG. In CPA, the task involved assigning a single property to the target column pairs. For CTA, the goal was to annotate target columns with a single type from the target KG, selecting the most specific or fine-grained type in the hierarchy. Similarly, the TD and CQA tasks required a single annotation to be provided as output.

The evaluation metrics for CEA, CPA, and CTA were Precision, Recall, and F1-score, defined as follows in Equation 1:

$$P = \frac{|\text{Correct Annotations}|}{|\text{System Annotations}|}, R = \frac{|\text{Correct Annotations}|}{|\text{Target Annotations}|}, F1 = \frac{2 \times P \times R}{P + R} \quad (1)$$

In this context, target annotations refer to the designated target cells for CEA, target columns for CTA, and target column pairs for CPA. An annotation is considered *correct* if it matches any entry in the ground truth set. Due to redirect links or same-as links in KGs, some target cells may have multiple valid annotations in the ground truth.

For CTA evaluation, a modified version of Precision and Recall was applied, given the detailed type hierarchy in Wikidata [23]. This adaptation accounts for partially correct annotations, such as those that are ancestors or descendants of the ground truth (GT) classes. The correctness score c_{score} for a CTA annotation α is based on its distance from the GT classes within the hierarchy and is defined as follows:

$$c_{score}(\alpha) = \begin{cases} 0.8^{d(\alpha)}, & \text{if } \alpha \text{ is in GT, or an ancestor of the GT, with } d(\alpha) \leq 5 \\ 0.7^{d(\alpha)}, & \text{if } \alpha \text{ is a descendant of the GT, with } d(\alpha) \leq 3 \\ 0, & \text{otherwise;} \end{cases} \quad (2)$$

Here, $d(\alpha)$ denotes the shortest distance from α to one of the GT classes. CTA ground truth columns can include multiple valid classes. For example, if α is a GT class ($d(\alpha) = 0$), the correctness score is $c_{score}(\alpha) = 1$. If α is a grandchild of a GT class ($d(\alpha) = 2$), the correctness score is $c_{score}(\alpha) = 0.49$. Types from higher levels of the KG type hierarchy, such as Q35120 [entity] in Wikidata, were excluded from the evaluation.

Using the correctness score c_{score} , the approximated Precision (AP), Recall (AR), and F1-score (AF1) for CTA were calculated as follows:

$$AP = \frac{\sum c_{score}(\alpha)}{|\text{System Annotations}|}, AR = \frac{\sum c_{score}(\alpha)}{|\text{Target Annotations}|}, AF1 = \frac{2 \times AP \times AR}{AP + AR} \quad (3)$$

2.2. STI vs LLMs Track

This track investigates the exclusive use of LLMs for performing the CEA task on Wikidata. Participants are tasked with either fine-tuning an LLM or employing prompting techniques on a dataset enriched

with semantic annotations. The task presents several challenges, including integrating factual knowledge from a knowledge graph (KG) into an LLM, devising strategies for handling Wikidata QIDs, enhancing the training dataset to improve disambiguation accuracy, mitigating hallucination issues, and designing effective prompts for fine-tuning or annotation purposes. The primary objective is to leverage the capabilities of LLMs to generate high-quality annotations for the CEA task, advancing their applicability in semantic enrichment. Participants are required to submit their annotations for evaluation on the test set, demonstrating the practicality and effectiveness of their approaches.

The provided tabular datasets consist of columns with entity mentions, which must be annotated with the corresponding Wikidata entities. These annotations should include the entity’s URI, though the prefix <http://www.wikidata.org/entity/> is optional. The evaluation metrics—Precision, Recall, and F1—are consistent with those used for CEA in the Accuracy Track.

2.2.1. Datasets

- **SuperSemtab 24** <https://doi.org/10.5281/zenodo.11031987>
This dataset was created by combining various tables from past SemTab Challenge datasets. It was then split into training and validation sets. The dataset features general-purpose tables as well as intentionally misspelled entities, designed to assess the model’s robustness. The dataset consists of 16,180 training tables and 4,044 test tables.
- **MammoTab 24 (SemTab)** <https://doi.org/10.5281/zenodo.11519643>
MammoTab dataset [24] includes 1 million tables extracted from over 20 million Wikipedia pages and enriched with annotations from Wikidata. It addresses a significant gap in the state-of-the-art by providing a valuable resource for testing and training Semantic Table Interpretation approaches. Designed to tackle critical challenges, MammoTab focuses on issues such as disambiguation, homonymy, and NIL mentions, making it an essential tool for advancing research in this domain. The MammoTab 24 (SemTab) dataset is a subset of the MammoTab dataset composed of 2,500 tables (2,000 for training and 500 for testing).

2.3. Table Metadata to KG Track

This track challenges participants to match limited table metadata, such as table names and column headers, to knowledge graphs without access to the actual table data or content. The task is inherently difficult due to the limited context available for annotation systems to perform semantic linking. LLMs offer a promising solution to address this challenge, providing flexibility in their application. The datasets for this track are adapted from our previous work on matching table metadata with business glossaries using large language models [25].

2.3.1. Datasets

Link: <https://doi.org/10.5281/zenodo.14207376>

- **Round 1:** This dataset consists of metadata from selected web tables that need to be mapped to the DBpedia ontology. The target ontology (also referred to as the glossary) contains 2,881 terms from the DBpedia ontology. The test dataset includes metadata (table and column labels) for 141 table columns. A small test set with metadata for 9 table columns, along with an evaluation script, was provided.
- **Round 2:** This dataset consists of metadata from selected open data tables that need to be mapped to a custom glossary containing 1,192 terms, semi-automatically derived from the available metadata. The provided table metadata includes metadata (table and column labels) for 1,192 table columns.

We use “Hit@1” and “Hit@5” as evaluation metrics, representing the percentage of table columns correctly matched to the ground truth glossary item within the top 1 and top 5 predictions in the system outputs, respectively.

Table 2
STI vs LLMs Track Summary Results.

Benchmark	Task	TSOTSA		City System		Kepler ASI	
		F1	Pr	F1	Pr	F1	Pr
SuperSemtab 24 Round1	CEA	0.905	0.905	0.858	0.866	0.764	0.907
MammoTab24	CEA	-	-	0.647	0.648	0.182	0.336

Table 3
Table Metadata to KG Track Summary Results.

Benchmark		Adwan		MetaLinker	
		Hit@1	Hit@5	Hit@1	Hit@5
Metadata2KG Round1	Top Hit@1	0.75	0.92	0.55	0.70
	Top Hit@5	0.75	0.92	0.55	0.70
Metadata2KG Round2	Top Hit@1	0.83	0.98	0.49	0.52
	Top Hit@5	0.83	0.98	0.37	0.68

3. Results

Tables 2, 3, and 4 summarize the results for each of the three tracks. Overall, seven participants submitted solutions to at least one dataset across any round:

- TSOTSA [26] explores building an STI solution using a GPT-3-based model through both few-shot and zero-shot prompting techniques and participated in the Accuracy Track as well as the STI vs LLMs Track.
- DREIFLUSS [27] employs a minimalist approach that carefully utilizes resources such as Wikidata APIs for the annotation process.
- Kepler-aSI [28] leverages SPARQL queries, embeddings, custom index structures, and a NoSQL database to address the CEA, CTA, and CPA tasks.
- MetaLinker [29] investigates the use of various LLMs and sentence embeddings for the Metadata to KG Track.
- Adwan [30] combines Retrieval-Augmented Generation (RAG), Chain-of-Thought (CoT) prompting, Self-Consistency (SC), and Reciprocal Rank Fusion (RRF) to develop an LLM-based solution for the Metadata to KG Track.
- GRAMS+ (ISI KG) [31] constructs a prediction model for CPA and CTA tasks using distant supervision.
- CitySTI [32] participated in the STI vs LLMs Track, utilizing a two-stage approach where LLMs were used for data cleaning and matching, executed entirely through prompting techniques.

In the Accuracy Track, the TSOTSA system participated in the largest number of datasets, while other systems focused on only one or two datasets. TSOTSA demonstrated promising performance on several tasks, including TD, RA, CEA, and CTA, in the tBiodiv-Relational and tBiomed-Relational datasets, as well as the CEA task in the tBiodiv-Entity and tBiomed-Entity datasets. However, it struggled with certain tasks, even on the simpler WikidataTables datasets, suggesting potential scalability challenges in its LLM-based solution. In contrast, ISI-KG delivered exceptional results on the WikidataTables datasets, showcasing the effectiveness of building a prediction model using distant supervision. The DREIFLUSS and Kepler-aSI systems also achieved notable results on the larger tBiodiv-Large-Relational and tBiomed-Large-Relational datasets.

In the STI vs LLMs Track, TSOTSA achieved the best performance on the SuperSemtab 24 Round 1 dataset, while the CitySTI system showed promising results across both datasets.

Finally, the two solutions in the Metadata to KG Track offered valuable insights into how various LLM models and prompting techniques can address the challenge of matching table metadata to knowledge graphs or business glossaries without access to table contents. The Adwan solution achieved outstanding Hit@5 scores of 0.92 in Round 1 and 0.98 in Round 2.

Table 4
Accuracy Track Summary Results.

Benchmark	Task	TSOTSA		Kepler-aSI		DREIFLUSS		ISI-KG	
		F1	Pr	F1	Pr	F1	Pr	F1	Pr
WikidataTables Round1	CEA	0.069	0.24	-	-	-	-	-	-
	CTA	0.717	0.717	-	-	-	-	0.929	0.929
	CPA	0.677	0.734	-	-	-	-	0.898	0.988
WikidataTables Round2	CTA	0.194	0.279	-	-	-	-	0.956	0.956
	CPA	-	-	-	-	-	-	0.899	0.992
tBiodiv-Entity	TD	0.055	0.055	-	-	-	-	-	-
	CEA	0.926	0.926	-	-	-	-	-	-
	RA	0.002	0.002	-	-	-	-	-	-
tBiodiv-Relational	TD	0.780	0.780	-	-	-	-	-	-
	RA	0.719	0.758	-	-	-	-	-	-
	CEA	0.740	0.740	-	-	-	-	-	-
	CTA	0.648	0.648	-	-	-	-	-	-
	CPA	0.016	0.016	-	-	-	-	-	-
tBiomed-Entity	TD	0.029	0.029	-	-	-	-	-	-
	CEA	0.938	0.938	-	-	-	-	-	-
	RA	0.008	0.008	-	-	-	-	-	-
tBiomed-Relational	TD	0.621	0.621	-	-	-	-	-	-
	RA	0.411	0.411	-	-	-	-	-	-
	CEA	0.575	0.806	-	-	-	-	-	-
	CTA	0.749	0.749	-	-	-	-	-	-
	CPA	0.060	0.060	-	-	-	-	-	-
tBiodiv-Large-Relational	CEA	-	-	-	-	0.932	0.932	-	-
	CTA	-	-	0.741	0.741	0.615	0.615	-	-
tBiomed-Large-Relational	CEA	-	-	-	-	0.925	0.925	-	-
	CTA	-	-	0.867	0.867	-	-	-	-

Acknowledgments

We extend our heartfelt gratitude to all participants of this year’s challenge, as well as those from previous editions, for their invaluable feedback, active participation in discussions, and technical contributions, which have collectively shaped this challenge [33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 24, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 31, 26, 27, 28, 29, 30]. We also express our gratitude to the ISWC organizers and our sponsors for their support. Lastly, we acknowledge the significant role of the EasyChair conference management system and the CEUR-WS.org open-access publication service, which greatly simplified the organization of this challenge.

This document has been reviewed and refined with the support of AI tools. The authors assume full responsibility for the accuracy, integrity, and content of its text.

References

- [1] G. Fan, J. Wang, Y. Li, D. Zhang, R. J. Miller, Semantics-aware dataset discovery from data lakes with contextualized column-based representation learning, Proc. VLDB Endow. 16 (2023) 1726–1739. URL: <https://www.vldb.org/pvldb/vol16/p1726-fan.pdf>.
- [2] G. Fan, J. Wang, Y. Li, R. J. Miller, Table discovery in data lakes: State-of-the-art and future directions, in: S. Das, I. Pandis, K. S. Candan, S. Amer-Yahia (Eds.), Companion of the 2023

- International Conference on Management of Data, SIGMOD/PODS 2023, Seattle, WA, USA, June 18-23, 2023, ACM, 2023, pp. 69–75. URL: <https://doi.org/10.1145/3555041.3589409>. doi:10.1145/3555041.3589409.
- [3] A. Khatiwada, G. Fan, R. Shraga, Z. Chen, W. Gatterbauer, R. J. Miller, M. Riedewald, SANTOS: relationship-based semantic table union search, *Proc. ACM Manag. Data* 1 (2023) 9:1–9:25. URL: <https://doi.org/10.1145/3588689>. doi:10.1145/3588689.
- [4] P. Ouellette, A. Sciortino, F. Nargesian, B. G. Bashardoost, E. Zhu, K. Q. Pu, R. J. Miller, RONIN: data lake exploration, *Proc. VLDB Endow.* 14 (2021) 2863–2866. URL: <http://www.vldb.org/pvldb/vol14/p2863-nargesian.pdf>. doi:10.14778/3476311.3476364.
- [5] F. Nargesian, K. Q. Pu, B. G. Bashardoost, E. Zhu, R. J. Miller, Data lake organization, *IEEE Trans. Knowl. Data Eng.* 35 (2023) 237–250. URL: <https://doi.org/10.1109/TKDE.2021.3091101>. doi:10.1109/TKDE.2021.3091101.
- [6] A. Khatiwada, R. Shraga, W. Gatterbauer, R. J. Miller, Integrating data lake tables, *Proc. VLDB Endow.* 16 (2022) 932–945. URL: <https://www.vldb.org/pvldb/vol16/p932-khatiwada.pdf>.
- [7] A. Khatiwada, R. Shraga, R. J. Miller, DIALITE: discover, align and integrate open data tables, in: S. Das, I. Pandis, K. S. Candan, S. Amer-Yahia (Eds.), *Companion of the 2023 International Conference on Management of Data, SIGMOD/PODS 2023, Seattle, WA, USA, June 18-23, 2023*, ACM, 2023, pp. 187–190. URL: <https://doi.org/10.1145/3555041.3589732>. doi:10.1145/3555041.3589732.
- [8] S. Galhotra, U. Khurana, O. Hassanzadeh, K. Srinivas, H. Samulowitz, M. Qi, Automated feature enhancement for predictive modeling using external knowledge, in: P. Papapetrou, X. Cheng, Q. He (Eds.), *2019 International Conference on Data Mining Workshops, ICDM Workshops 2019, Beijing, China, November 8-11, 2019*, IEEE, 2019, pp. 1094–1097. URL: <https://doi.org/10.1109/ICDMW.2019.00161>. doi:10.1109/ICDMW.2019.00161.
- [9] M. palmonari, F. D’Adda, M. Cremaschi, E. Jimenez-Ruiz, Tutorial on Semantic Table Interpretation - TUTSTI @ ISWC2024, <https://unimib-datai.github.io/sti-website/tutorial/>, 2024.
- [10] Z. Zhang, Effective and efficient semantic table interpretation using tableminer+, *Semantic Web* 8 (2017) 921–957.
- [11] Z. Syed, T. Finin, V. Mulwad, , A. Joshi, Exploiting a Web of Semantic Data for Interpreting Tables, in: *Proceedings of the Second Web Science Conference*, 2010.
- [12] V. Mulwad, T. Finin, A. Joshi, Automatically Generating Government Linked Data from Tables, in: *Working notes of AAAI Fall Symposium on Open Government Knowledge: AI Opportunities and Challenges*, 2011.
- [13] V. Mulwad, T. Finin, Z. Syed, A. Joshi, T2LD: interpreting and representing tables as linked data, in: A. Polleres, H. Chen (Eds.), *Proceedings of the ISWC 2010 Posters & Demonstrations Track: Collected Abstracts, Shanghai, China, November 9, 2010*, volume 658 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2010. URL: <https://ceur-ws.org/Vol-658/paper489.pdf>.
- [14] P. Buche, J. Dibia-Barthélemy, L. Ibanescu, L. Soler, Fuzzy web data tables integration guided by an ontological and terminological resource, *IEEE Trans. Knowl. Data Eng.* 25 (2013) 805–819. URL: <https://doi.org/10.1109/TKDE.2011.245>. doi:10.1109/TKDE.2011.245.
- [15] G. Hignette, P. Buche, J. Dibia-Barthélemy, O. Haemmerlé, An ontology-driven annotation of data tables, in: M. Weske, M. Hacid, C. Godart (Eds.), *Web Information Systems Engineering - WISE 2007 Workshops, WISE 2007 International Workshops, Nancy, France, December 3, 2007*, *Proceedings*, volume 4832 of *Lecture Notes in Computer Science*, Springer, 2007, pp. 29–40. URL: https://doi.org/10.1007/978-3-540-77010-7_4. doi:10.1007/978-3-540-77010-7_4.
- [16] G. Hignette, P. Buche, J. Dibia-Barthélemy, O. Haemmerlé, Fuzzy annotation of web data tables driven by a domain ontology, in: L. Aroyo, P. Traverso, F. Ciravegna, P. Cimiano, T. Heath, E. Hyvönen, R. Mizoguchi, E. Oren, M. Sabou, E. Simperl (Eds.), *The Semantic Web: Research and Applications, 6th European Semantic Web Conference, ESWC 2009, Heraklion, Crete, Greece, May 31-June 4, 2009*, *Proceedings*, volume 5554 of *Lecture Notes in Computer Science*, Springer, 2009, pp. 638–653. URL: https://doi.org/10.1007/978-3-642-02121-3_47. doi:10.1007/978-3-642-02121-3_47.

- [17] E. Muñoz, A. Hogan, A. Mileo, Triplifying wikipedia’s tables, in: A. L. Gentile, Z. Zhang, C. d’Amato, H. Paulheim (Eds.), Proceedings of the First International Workshop on Linked Data for Information Extraction (LD4IE 2013) co-located with the 12th International Semantic Web Conference (ISWC 2013), Sydney, Australia, October 21, 2013, volume 1057 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2013. URL: https://ceur-ws.org/Vol-1057/MunozEtAl_LD4IE2013.pdf.
- [18] P. Venetis, A. Y. Halevy, J. Madhavan, M. Pasca, W. Shen, F. Wu, G. Miao, C. Wu, Recovering semantics of tables on the web, *Proc. VLDB Endow.* 4 (2011) 528–538. URL: <http://www.vldb.org/pvldb/vol4/p528-venetis.pdf>. doi:10.14778/2002938.2002939.
- [19] V. Efthymiou, S. Galhotra, O. Hassanzadeh, E. Jiménez-Ruiz, K. Srinivas, 1st international workshop on tabular data analysis (TaDA, in: R. Bordawekar, C. Cappiello, V. Efthymiou, L. Ehrlinger, V. Gadepally, S. Galhotra, S. Geisler, S. Groppe, L. Gruenwald, A. Y. Halevy, H. Harmouch, O. Hassanzadeh, I. F. Ilyas, E. Jiménez-Ruiz, S. Krishnan, T. Lahiri, G. Li, J. Lu, W. Maurer, U. F. Minhas, F. Naumann, M. T. Özsu, E. K. Rezig, K. Srinivas, M. Stonebraker, S. R. Valluri, M. Vidal, H. Wang, J. Wang, Y. Wu, X. Xue, M. Zaït, K. Zeng (Eds.), Joint Proceedings of Workshops at the 49th International Conference on Very Large Data Bases (VLDB 2023), Vancouver, Canada, August 28 - September 1, 2023, volume 3462 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023. URL: <https://ceur-ws.org/Vol-3462/TADA0.pdf>.
- [20] D. Vrandečić, M. Krötzsch, Wikidata: a free collaborative knowledge base, *Commun. ACM* 57 (2014) 78–85.
- [21] E. Jimenez-Ruiz, O. Hassanzadeh, V. Efthymiou, J. Chen, K. Srinivas, *SemTab 2019: Resources to Benchmark Tabular Data to Knowledge Graph Matching Systems*, in: *The Semantic Web: ESWC*, Springer International Publishing, 2020.
- [22] N. Abdelmageed, E. Jiménez-Ruiz, O. Hassanzadeh, B. König-Ries, KG2Tables: Your way to generate an STI benchmark for your domain, in: *Proceedings of the ISWC 2024 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice* co-located with 23rd International Semantic Web Conference (ISWC 2024), Hanover, Maryland, USA, November 11-15, 2024, volume 3828 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2024.
- [23] E. Jiménez-Ruiz, O. Hassanzadeh, V. Efthymiou, J. Chen, K. Srinivas, V. Cutrona, Results of SemTab 2020, in: *Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching* co-located with the 19th International Semantic Web Conference (ISWC 2020), 2020, pp. 1–8.
- [24] M. Marzocchi, M. Cremaschi, R. Pozzi, R. Avogadro, M. Palmonari., MammoTab: a giant and comprehensive dataset for Semantic Table Interpretation, in: *Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022*, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
- [25] E. A. Lobo, O. Hassanzadeh, N. Pham, N. Mihindikulasooriya, D. Subramanian, H. Samulowitz, Matching table metadata with business glossaries using large language models, in: *Proceedings of the 18th International Workshop on Ontology Matching* co-located with the 22nd International Semantic Web Conference (ISWC 2023), Athens, Greece, November 7, 2023, volume 3591 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023, pp. 25–36.
- [26] J. P. Bikim, C. Atezong, A. Jiomekong, A. Oelen, G. Rabby, J. D’Souza, S. Auer, Leveraging GPT Models For Semantic Table Annotation, in: *SemTab’24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024*, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.
- [27] V. Parmar, A. Algergawy, Wikidata-Driven CEA and CTA for Life Sciences Table Matching extending DREIFLUSS, in: *SemTab’24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024*, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.
- [28] W. Baazouzi, M. Kachroudi, S. Faiz, Kepler-aSI : Semantic Annotation for Tabular Data, in: *SemTab’24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024*, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.

- [29] M. Martorana, X. Pan, B. Kruit, T. Kuhn, J. van Ossenbruggen, Column Vocabulary Association (CVA): Semantic Interpretation of Dataless Tables, in: SemTab'24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.
- [30] N. Vandemoortele, B. Steenwinckel, S. V. Hoecke, F. Ongenae, Scalable Table-to-Knowledge Graph Matching from Metadata using LLMs, in: SemTab'24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.
- [31] B. Vu, C. Knoblock, F. Lin, Results of GRAMS+ at SemTab 2024, in: SemTab'24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.
- [32] D. Li Tin Yue, E. Jiménez-Ruiz, CitySTI 2024 System: Tabular Data to KG Matching using LLMs, in: SemTab'24: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2024, co-located with the 23rd International Semantic Web Conference (ISWC), 2024.
- [33] S. Yumusak, Knowledge graph matching with inter-service information transfer, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [34] G. Diallo, R. Azzi, AMALGAM: making tabular dataset explicit with knowledge graph, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [35] W. Baazouzi, M. Kachroudi, S. Faiz, Kepler-aSI : Kepler as a Semantic Interpreter, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [36] N. Abdelmageed, S. Schindler, JenTab: Matching Tabular Data to Knowledge Graphs, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [37] S. Tyagi, E. Jiménez-Ruiz, LexMa: Tabular Data to Knowledge Graph Matching using Lexical Techniques, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [38] R. Shigapov, P. Zumstein, J. Kamlah, L. Oberländer, J. Mechnich, I. Schumm, bbw: Matching CSV to Wikidata via Meta-lookup, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [39] M. Cremaschi, R. Avogadro, A. Barazzetti, D. Chierigato, MantisTable SE: an Efficient Approach for the Semantic Table Interpretation, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [40] V.-P. Huynh, J. Liu, Y. Chabot, T. Labbé, P. Monnin, , R. Troncy, DAGOBAN: Enhanced Scoring Algorithms for Scalable Annotations of Tabular Data, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [41] P. Nguyen, I. Yamada, N. Kertkeidkachorn, R. Ichise, H. Takeda, MTab4Wikidata at the SemTab 2020: Tabular Data Annotation with Wikidata, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [42] S. Chen, A. Karaoglu, C. Negreanu, T. Ma, J.-G. Yao, J. Williams, A. Gordon, C.-Y. Lin, Linking-Park: An integrated approach for Semantic Table Interpretation, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [43] D. Kim, H. Park, J. K. Lee, W. Kim, Generating conceptual subgraph from tabular data for

- knowledge graph matching, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2020, co-located with the 19th International Semantic Web Conference (ISWC), CEUR-WS.org, 2020.
- [44] N. Abdelmageed, S. Schindler, B. König-Ries, BiodivTab: Semantic Table Annotation Benchmark Construction, Analysis, and New Additions, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
- [45] N. Abdelmageed, S. Schindler, B. König-Ries, BiodivTab: A Tabular Benchmark based on Biodiversity Research Data, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [46] V.-P. Huynh, J. Liu, Y. Chabot, F. Deuzé, T. Labbé, P. Monnin, R. Troncy, DAGOBAN: Table and Graph Contexts For Efficient Semantic Annotation Of Tabular Data, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [47] W. Baazouzi, M. Kachroudi, S. Faiz, Kepler-aSI at SemTab 2021, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [48] N. Abdelmageed, S. Schindler, JenTab Meets SemTab 2021's New Challenges, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [49] B. Steenwinckel, F. D. Turck, F. Ongenae, MAGIC: Mining an Augmented Graph using INK, starting from a CSV, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [50] L. Yang, S. Shen, J. Ding, J. Jin, GBMTAB: A Graph-Based Method for Interpreting Semantic Table to Knowledge Graph, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [51] P. Nguyen, I. Yamada, N. Kertkeidkachorn, R. Ichise, H. Takeda, SemTab 2021: Tabular Data Annotation with MTab Tool, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [52] R. Avogadro, M. Cremaschi, MantisTable V: A novel and efficient approach to Semantic Table Interpretation, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2021, co-located with the 20th International Semantic Web Conference (ISWC), CEUR-WS.org, 2021.
- [53] I. Mazurek, B. Wiewel, B. Kruit, Wikary: A Dataset of N-ary Wikipedia Tables Matched to Qualified Wikidata Statements, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
- [54] K. Korini, R. Peeters, C. Bizer, SOTAB: the WDC schema.org table annotation benchmark, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), volume 3320 of *CEUR Workshop Proceedings*, 2022, pp. 14–19.
- [55] A. Jiomekong, C. Etoğa, B. Foko, M. Folefac, S. Kana, V. Tsague, M. Sow, G. Camara, A large scale corpus of food composition tables, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
- [56] A. Jiomekong, B. A. F. Tagne, Towards an Approach based on Knowledge Graph Refinement for Tabular Data to Knowledge Graph Matching, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International

- Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
- [57] W. Baazouzi, M. Kachroudi, S. Faiz, Yet Another Milestone for Kepler-aSI at SemTab 2022, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
 - [58] M. Cremaschi, R. Avogadro, D. Chierigato, s-elBat: a Semantic Interpretation Approach for Messy taBle-s, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
 - [59] N. Abdelmageed, S. Schindler, JenTab: Do CTA solutions affect the entire scores?, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
 - [60] X. Li, S. Wang, W. Zhou, G. Zhang, C. Jiang, T. Hong, P. Wang, KGCODE-Tab Results for SemTab 2022, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
 - [61] L. Mertens, A low-resource approach to SemTab 2022, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
 - [62] A. Sharma, S. Dalal, S. Jain, SemInt at SemTab 2022, in: Proceedings of the Semantic Web Challenge on Tabular Data to Knowledge Graph Matching, SemTab 2022, co-located with the 21st International Semantic Web Conference (ISWC), CEUR-WS.org, 2022.
 - [63] V.-P. Huynh, Y. Chabot, T. Labbé, J. Liu, R. Troncy., From Heuristics to Language Models: A Journey Through the Universe of Semantic Table Interpretation with DAGOBAN, in: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching (SemTab), CEUR-WS.org, 2022.
 - [64] W. Baazouzi, M. Kachroudi, S. Faiz, Kepler-aSI at SemTab 2023, in: SemTab'23: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2023, co-located with the 22nd International Semantic Web Conference (ISWC), 2023.
 - [65] S. Mehryar, R. Celebi, Semantic Annotation of Tabular Data for Machine-to-Machine Interoperability via Neuro-Symbolic Anchoring, in: SemTab'23: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2023, co-located with the 22nd International Semantic Web Conference (ISWC), 2023.
 - [66] V. Parmar, A. Algergawy, DREIFLUSS: A Minimalist Approach for Table Matching, in: SemTab'23: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2023, co-located with the 22nd International Semantic Web Conference (ISWC), 2023.
 - [67] B. Foko, A. Jiomekong, H. Tapamo, J. Buisson, S. Tiwari, Exploring Naive Bayes Classifiers for Tabular Data to Knowledge Graph Matching, in: SemTab'23: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2023, co-located with the 22nd International Semantic Web Conference (ISWC), 2023.
 - [68] E. G. Henriksen, A. M. Khorsid, E. Nielsen, A. M. Stück, A. S. Sørensen, O. Pelgrin, SemTex: A Hybrid Approach for Semantic Table Interpretation, in: SemTab'23: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2023, co-located with the 22nd International Semantic Web Conference (ISWC), 2023.
 - [69] I. Dasoulas, D. Yang, X. Duan, A. Dimou, TorchicTab: Semantic Table Annotation with Wikidata and Language Models, in: SemTab'23: Semantic Web Challenge on Tabular Data to Knowledge Graph Matching 2023, co-located with the 22nd International Semantic Web Conference (ISWC), 2023.