

SlovKE: A Large-Scale Dataset and LLM Evaluation for Slovak Keyphrase Extraction

David Števaňák^{1,4} Marek Šuppa^{2,3,4}

¹University of Vienna, Austria

²Faculty of Mathematics, Physics and Informatics, Comenius University in Bratislava

³Cisco Systems, Inc. ⁴NaiveNeuron

Correspondence: marek@suppa.sk

Abstract

Keyphrase extraction for morphologically rich, low-resource languages remains understudied, largely due to the scarcity of suitable evaluation datasets. We address this gap for Slovak by constructing a dataset of 227,432 scientific abstracts with author-assigned keyphrases—scraped and systematically cleaned from the Slovak Central Register of Theses—representing a 25-fold increase over the largest prior Slovak resource and approaching the scale of established English benchmarks such as KP20K. Using this dataset, we benchmark three unsupervised baselines (YAKE, TextRank, KeyBERT with SlovakBERT embeddings) and evaluate KeyLLM, an LLM-based extraction method using GPT-3.5-turbo. Unsupervised baselines achieve at most 11.6% exact-match $F1@6$, with a large gap to partial matching (up to 51.5%), reflecting the difficulty of matching inflected surface forms to author-assigned keyphrases. KeyLLM narrows this exact–partial gap, producing keyphrases closer to the canonical forms assigned by authors, while manual evaluation on 100 documents ($\kappa = 0.61$) confirms that KeyLLM captures relevant concepts that automated exact matching underestimates. Our analysis identifies morphological mismatch as the dominant failure mode for statistical methods—a finding relevant to other inflected languages. The dataset (<https://huggingface.co/datasets/NaiveNeuron/SlovKE>) and evaluation code (<https://github.com/NaiveNeuron/SlovKE>) are publicly available.

Keywords: keyphrase extraction, Slovak, morphologically rich languages, large language models, benchmark dataset, evaluation metrics, Slavic NLP

1. Introduction

Keyphrase extraction—identifying the words or phrases that best represent a document’s topics (Hasan and Ng, 2014)—is essential for scientific literature discovery and classification (Pagiannopoulou and Tsoumakas, 2020). While recent advances have yielded strong results for English and other high-resource languages (Alzaidy et al., 2019; Basaldella et al., 2018; Nguyen and Kan, 2007), morphologically rich low-resource languages remain underexplored. In such languages, a single lemma can surface in dozens of inflected forms (Tsarfaty et al., 2013), creating a fundamental mismatch between the surface forms models extract and the canonical forms authors assign as keyphrases. This challenge is not unique to any single language—it affects the entire family of morphologically rich languages, including most Slavic, Finno-Ugric, and Turkic languages (Beliga and Martincic-Ipsic, 2016). Figure 1 illustrates the problem: author-assigned keyphrases in nominative form (e.g., *Rozvojový potenciál*) appear in the abstract only in inflected variants (e.g., genitive *rozvojového potenciálu*).

We study this problem through the lens of Slovak, a West Slavic language with limited keyphrase extraction research due to the lack of suitable corpora. Building on the work of Zelinka (2023), who col-

lected approximately 9,000 documents but faced limitations in data quality and scale, we construct a substantially larger and cleaner dataset and conduct a systematic evaluation of both established and LLM-based keyphrase extraction methods. We evaluate KeyLLM (Grootendorst, 2021), an LLM-based extraction approach, to determine whether generative models can better handle the morphological variability that hinders traditional extractive methods.

Our main contributions are:

- A rigorously cleaned dataset of 227,432 Slovak scientific abstracts with author-assigned keyphrases—a 25-fold increase over prior work—whose size and statistics are comparable to established English benchmarks such as KP20K (Meng et al., 2017).
- A controlled evaluation of statistical (YAKE), graph-based (TextRank), and embedding-based (KeyBERT) models on Slovak, revealing that exact-match $F1@6$ remains below 12% while partial-match $F1@6$ reaches 51.5%, exposing a ~ 40 -point gap attributable to morphological inflection.
- Evidence that KeyLLM substantially narrows this exact–partial gap by generating keyphrases in canonical form rather than extracting surface tokens, suggesting that gener-

Predmetom skúmania diplomovej práce je potenciál **vidieka**, územný rozvoj a možnosti **rozvoja vidieka**. Objektom skúmania je Žilinský samosprávny kraj a jeho historický **región** Dolná Orava. Cieľom práce je na základe teoretického rozpracovania problematiky a analýzy **rozvojového potenciálu** vymedziť možnosti **rozvoja vidieka**. V práci je sústredená pozornosť na analýzu **rozvojového potenciálu** Žilinského kraja a možností využitia jeho jednotlivých zložiek na zabezpečenie ekonomického, sociálneho a trvaloudržateľného **rozvoja**. Predpokladaný prínos práce vidíme vo formulovaní návrhov využitia **rozvojového potenciálu** kraja na zabezpečenie endogénneho **rozvoja vidieka**.

Keywords: **Vidiek**, **Rozvoj vidieka**, **Rozvojový potenciál**, Vidiecky **región**, Vidiecke sídlo, Nástroje rozvoja

Figure 1: Example abstract from the Test22K dataset (Slovak). Color coding indicates keyphrase occurrences: author-assigned keyphrases appear in the text in various inflected forms (**red**: exact match, **teal**: partial overlap, **green**: single-word match, **blue**: shared fragment). Note that surface forms in the abstract (e.g., *rozvojového potenciálu*, genitive) differ from the canonical keyphrase form (*Rozvojový potenciál*, nominative), illustrating the morphological mismatch challenge.

ative models are better suited to morphologically rich languages.

- A manual evaluation and error analysis ($\kappa=0.61$) identifying morphological mismatch as the dominant failure mode for extractive models and unmotivated adjective extraction as the primary weakness of KeyLLM.

2. Related Work

Keyphrase extraction has evolved from early statistical and graph-based methods (Hulth, 2003; Mihalcea and Tarau, 2004) to modern transformer-based approaches. Recent advances leveraging pre-trained language models (PLMs) have achieved state-of-the-art performance (Song et al., 2023), with BERT-based methods like KeyBERT (Grootendorst, 2020) and approaches utilizing domain-specific models like SciBERT (Park and Caragea, 2020) showing significant improvements. Supervised methods formulating keyphrase extraction as sequence labeling with BiLSTM-CRF architectures over contextualized embeddings (Sahrawat et al., 2020) have demonstrated strong performance, though they require substantial training data. More recently, large language models have been explored for keyphrase extraction (Grootendorst, 2021), though primarily for high-resource languages. Wu et al. (2022) demonstrated that PLMs exhibit significant performance degradation on low-resource languages, highlighting the challenges for languages like Slovak.

This performance gap has motivated a growing body of work on low-resource and multilingual keyphrase extraction. Gao et al. (2022) proposed retrieval-augmented multilingual keyphrase generation, while multilingual datasets like MAKED (Kulkarni et al., 2022), TermEval (Terryn et al., 2020), and the recent EUROPA dataset (Salaün

et al., 2024) covering 24 European languages support cross-lingual research. However, most multilingual work focuses on Romance and Germanic languages, with limited exploration of morphologically rich Slavic languages. Morphological complexity presents unique challenges (Tsarfaty et al., 2013), as languages like Slovak exhibit extensive inflection, with a single lemma having dozens of surface forms, complicating keyphrase matching and evaluation. Critically, no prior work has systematically quantified the impact of this inflectional variation on standard evaluation metrics—a gap our study addresses by reporting both exact and partial matching and analyzing the divergence between them.

Among Slavic languages, Slovak, Czech, and Polish are considered *under-resourced*, with only partial NLP infrastructure available (e.g., corpora, taggers, morphological analyzers) (Beliga and Martincic-Ipsic, 2016). Within the Slavic family, Czech has seen more development, with Docekal and Smrz (2022) testing a query-based method on 48,879 documents, though performance worsened on abstracts. Polish keyphrase extraction has also been studied (Giarelis et al., 2021; Pezik et al., 2022), with mixed results for deep learning approaches.

Slovak keyphrase extraction is still in its early stages. Varga et al. (2022) applied unsupervised methods to five court decisions, with TF-IDF achieving the best results. Zelinka (2023) made an important first step by collecting approximately 9,000 documents from the Slovak thesis registry; however, the dataset inherits noise typical of large-scale web harvesting—mixed-language content, inconsistent metadata, and variable keyphrase formatting—challenges that motivate the more extensive cleaning pipeline we describe in Section 4.2.

The main datasets for benchmarking keyphrase

extraction in English include KPTime (Gallina et al., 2019) (news articles) and, for scientific articles, KP20K (Meng et al., 2017), Inspec (Hulth, 2003), and SemEval (Kim et al., 2010). The KP20K dataset, with 530,000 training and 20,000 test documents, is most relevant to our work. While English benefits from hundreds of thousands of annotated documents, Slavic languages rely on datasets orders of magnitude smaller, making it difficult to determine whether low extraction performance reflects genuine model limitations or merely insufficient data.

3. Methodology

We select baselines that span three dominant paradigms—statistical (YAKE), graph-based (TextRank), and embedding-based (KeyBERT)—following the configuration of Zelinka (2023) to enable direct comparison with prior Slovak work. Crucially, all three are *extractive*: they return surface tokens from the input text, making them structurally vulnerable to the morphological mismatch between in-text inflected forms and author-assigned canonical keyphrases. This property is not specific to Slovak; it affects every morphologically rich language and makes these baselines a meaningful diagnostic tool for quantifying the impact of inflection on evaluation metrics. We contrast them with KeyLLM, a large-language-model-based approach that *generates* keyphrases and is therefore free to produce canonical forms regardless of the surface forms present in the text.

3.1. YAKE

YAKE (Campos et al., 2020) is an unsupervised keyword extraction algorithm that uses statistical methods to analyze document words. It scores keyphrases based on factors such as casing, term position, frequency, and relatedness to context. We use the yake library with Slovak stopwords and a uni- to bi-gram n-gram range, consistent with prior work.

3.2. TextRank

TextRank (Mihalcea and Tarau, 2004) is another unsupervised approach that employs a graph-based structure, where keyphrases form nodes in a graph, and their relevance is computed iteratively using a variant of the PageRank algorithm (Rogers, 2002). The pke library is used with custom adjustments to favor uni- and bi-grams.

3.3. KeyBERT

KeyBERT (Grootendorst, 2020) leverages BERT embeddings to extract keyphrases by measuring

their cosine similarity to the entire document. The approach assumes that a word effectively describing the document will have a vector representation close to that of the document itself. In our approach, we utilize the `kinit/slovakbert-sts-stsb` (Pikuliak et al., 2022), which has been fine-tuned for Slovak, and apply it to generate uni- and bi-gram keyphrases.

3.4. KeyLLM

KeyLLM builds upon KeyBERT and leverages Large Language Models (LLMs) such as GPT (Radford et al., 2018) for keyphrase extraction. Unlike KeyBERT, which relies on cosine similarity between BERT embeddings, KeyLLM uses prompted LLMs (GPT-3.5 (Ouyang et al., 2022) or GPT-4 (Achiam et al., 2023)) to directly generate keyphrases from the document text.

To reduce computational costs, KeyLLM offers an optional embedding-based clustering approach where documents are encoded using a sentence transformer (`all-MiniLM-L6-v2` for English or `kinit/slovakbert-sts-stsb` for Slovak). A similarity threshold parameter determines clustering: higher thresholds (0.90) create smaller clusters requiring more LLM calls, while lower thresholds (0.75) reduce costs by grouping diverse documents together. We analyze both versions—with clustering at various thresholds (0.75, 0.85, 0.90) and without embeddings—representing the first application to Slovak.

4. Dataset

4.1. Data Collection

Our data source is the Central Register of Theses and Dissertations in Slovakia (<https://opac.crzp.sk/>). It is a publicly available database of bachelor, diploma, rigorous, and similar types of theses from Slovak universities.

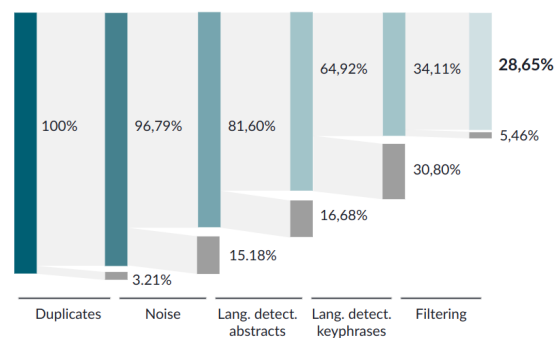


Figure 2: Filtering process and percentage of rows in the original dataset

The theses have been stored since 2010, which

provided us with a substantial number of available documents. At the time of our data collection (2024), there were a total of 794,527 available documents. Each thesis record in the register includes what the system labels a *primary abstract* (written in the thesis’s main language, typically Slovak) and a *secondary abstract* (a translation, typically into English), along with author-assigned keyphrases and metadata such as university, year of completion, and supervisor.

The website’s session limit of 20,000 documents required strategic filtering by year and alphabetical prefixes to access nearly all records through parallel processing. We successfully retrieved 793,722 documents, with the remaining 805 excluded due to non-alphabetical titles or unreadable text.

4.2. Data Cleaning

Starting from 793,722 scraped records, we applied a multi-stage cleaning pipeline (detailed in Appendix B) consisting of: (1) duplicate removal, prioritizing records with complete abstracts and keyphrases; (2) recovery of keyphrases appended to abstract text by universities lacking a dedicated keyphrase field; (3) removal of noisy metadata prepended to abstracts (e.g., author names, thesis type, page counts); (4) language verification using the `lingua` library (Lingua, 2021), which revealed that 20% of Slovak-labeled abstracts were actually in English; (5) keyphrase normalization using Stanza POS tagging (Qi et al., 2020) to split concatenated keyphrase lists and enforce a maximum length of four words; and (6) length-based filtering, retaining abstracts of 500–2000 characters with 4–15 keyphrases. The overall filtering process is summarized in Figure 2.

4.3. Dataset Statistics

After cleaning, we obtained 227,432 records, which we release as the **SlovKE** (Slovak Keyphrase Extraction) dataset. We split them randomly (with a fixed seed and without explicit stratification) into training (80%), validation (10%), and test (10%) sets. We additionally created a larger 20% test split (Test45K) for preliminary analysis; since no significant differences were found between the two test splits, we report results only for the 10% split (Test22K). The Test45K documents were not added back to the training set. As all evaluated models are unsupervised, the training split was not used for model training; it is included to support future supervised approaches.

The dataset statistics, compared to other scholarly datasets, are presented in Table 1. As shown, our dataset aligns well in size with larger keyphrase extraction datasets in English, even in terms of statistical properties. Notably, our absent-keyphrase

rate (37%) is comparable to that of KP20K (42.6%) and Inspec (44.3%). This indicates that the challenge of absent keyphrases in Slovak is not disproportionately harder than in English—contrary to what might be expected for a morphologically rich language—and that models and evaluation protocols developed for English absent-keyphrase generation can be meaningfully applied cross-lingually. As pointed out by Meng et al. (2017), a substantial absent-keyphrase rate is not undesirable; it can enhance model training and evaluation by testing the model’s ability to generate keyphrases that are not explicitly present in the text. The keyphrases predominantly consist of unigrams and bigrams.

It is worth noting that author-assigned keyphrases are inherently subjective: different authors may select different terms for similar content, and individual variation in keyphrase granularity, abstraction level, and terminology is well documented even in English datasets (Kim et al., 2010). In our dataset, this subjectivity is compounded by the fact that keyphrases are assigned by students with varying levels of academic experience, without standardized guidelines across universities. Rather than treating this as a limitation, we view it as representative of the realistic annotation conditions under which keyphrase extraction systems must operate, particularly in low-resource settings where expert re-annotation is impractical.

Additional statistics on abstract length distribution and token counts are provided in Appendix C.

5. Evaluation

5.1. Evaluation Metrics

We compare author-assigned keyphrases with those extracted by the models, first lemmatizing both sets to account for word variations. Two common matching techniques are used: *Exact matching*, where the extracted keyphrase must exactly match an author-assigned keyphrase, and *Partial matching*, where a match is counted if any fragment of the extracted keyphrase overlaps with a fragment of a golden keyphrase.

Evaluations are performed using $F1$ score at a fixed number k of keyphrases (denoted as $F1@k$), with recall capped at 1 when the number of matches exceeds the golden set (if $k > |\text{golden set}|$). The final $F1$ score is the micro-average across all test documents. We note that both metrics were designed primarily for English, where inflectional variation is limited. In morphologically rich languages, exact matching penalizes correct concepts expressed in a different case or number, while partial matching may over-reward coincidental token overlap. Reporting both metrics

Dataset (Test)	Inspec	KP20K	SemEval	Zelinka	SlovKE
Annotator	Expert	Author	Author+Reader	Author	Author
Document Type	Abstracts	Abstract	Papers (6-8p)	Abstracts	Abstracts
# Documents (Train)	1000	530K	144	-	182K
# Documents (Validation)	500	20K	-	-	22K
# Documents (Test)	500	20K	100	9K	22K
Avg. no. Words per Document	134.6	176	7,961	125.63	134.02
Avg. no. of Keyphrases	9.8	5.3	14.7	6.08	5
Avg. len of Keyphrase	2.3	2.6	2.2	1.68	1.68
% of absent Keyphrases	44.31	42.6	19.7	49.1	37

Table 1: Dataset comparison for Test sets

and analyzing their divergence therefore provides a more informative picture of model quality than either metric alone—a practice we recommend for any keyphrase evaluation on morphologically rich languages.

6. Results

6.1. Comparing Baseline Models

We compared models from three categories—statistical (YAKE), graph-based (TextRank), and pre-trained (KeyBERT)—using the same configuration as in prior work (e.g., $F1@6$, uni- and bi-gram extraction). Figure 3 provides an overview of the comparison with the Zelinka dataset.

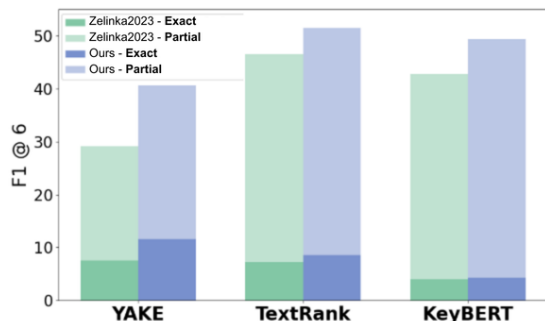


Figure 3: The $F1@6$ score for exact (stronger color) and partial (lighter color) matches is shown for the baseline models, comparing Zelinka and SlovKE.

The evaluations (Table 2; Table 3) showed slight improvements in exact matching, with YAKE potentially benefiting from the removal of noisy text during the data cleaning stage. This was because we cleaned the abstracts by removing words at the beginning of the text, such as *ABSTRACT* or *Bachelor Thesis*. Due to its underlying scoring mechanism, YAKE assigns high importance to such tokens. For partial matching, despite the notable improvement, the ranking of top-performing models remains consistent. Additional analysis across various values of k (Figure 4) shows that YAKE

Exact Matching			
Model	Metric	SlovKE	Zelinka
YAKE	Precision	10.4	6.8
	Recall	13.2	8.6
	$F1@6$	11.6	7.5
TextRank	Precision	7.7	6.5
	Recall	9.8	8.3
	$F1@6$	8.6	7.2
KeyBERT	Precision	4.0	3.8
	Recall	4.7	4.4
	$F1@6$	4.3	4.0

Table 2: Comparison of keyphrase extraction models using *Exact Matching*.

Partial Matching			
Model	Metric	SlovKE	Zelinka
YAKE	Precision	36.9	26.6
	Recall	46.0	33.2
	$F1@6$	40.6	29.1
TextRank	Precision	48.1	43.9
	Recall	56.3	50.6
	$F1@6$	51.5	46.5
KeyBERT	Precision	46.6	40.8
	Recall	53.7	46.4
	$F1@6$	49.4	42.8

Table 3: Comparison of keyphrase extraction models using *Partial Matching*.

achieves the highest $F1$ score in exact matching, while TextRank performs best in partial matching.

For all k , the gap between exact and partial matching is striking: YAKE’s exact $F1@6$ of 11.6 vs. partial $F1@6$ of 40.6 represents a 29-point spread, and TextRank shows an even larger gap of 43 points (8.6 vs. 51.5). We argue that this *exact-partial gap* is a diagnostic metric in its own right: it quantifies the degree to which morphological inflection distorts automated evaluation. In English,

where inflection is minimal, the two metrics are far closer. The magnitude of the gap we observe is therefore not merely a Slovak-specific finding—it is predictable for any language with rich nominal and verbal morphology (e.g., Czech, Polish, Finnish, Turkish, Hungarian) and suggests that standard exact-match $F1$ systematically underestimates extractive model performance in such languages. Ideally, we aim not only to achieve a higher $F1$ score for both exact and partial matching but also to narrow the gap between them, ensuring a more reliable evaluation.

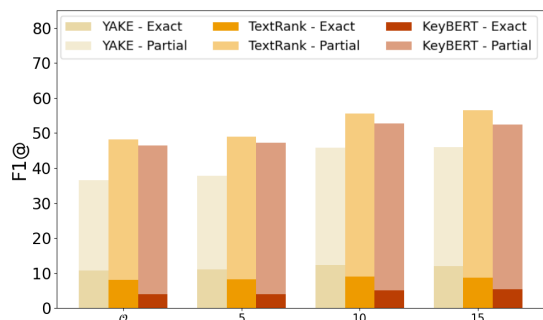


Figure 4: F1 score of exact and partial match for baseline models using different k values, where \emptyset denotes $k = |\text{golden set}|$.

6.2. KeyLLM Results

To extend our model evaluations, we employed GPT-3.5-turbo with prompts in both Slovak and English, yielding no significant differences, thus we present only results generated using an English prompt (see Appendix A).

As GPT-3.5-turbo is costly due to token processing, we initially used embeddings for cost-saving. We employed two Sentence Transformers: `kinit/slovakbert-sts-stsb` for Slovak, and `all-MiniLM-L6-v2` for English. This method involved clustering similar documents and selecting one representative for keyphrase extraction, reducing document processing. The model uses a threshold parameter: if the similarity score exceeds the threshold, documents are grouped together; otherwise, they are processed separately. A higher threshold (e.g., 0.9) results in more precise grouping with fewer documents per group, while a lower threshold (e.g., 0.7) reduces costs by clustering more diverse documents, though it may introduce noise.

However, this approach was ineffective with the recommended 0.75 threshold (Figure 5). Higher thresholds (0.85) improved results, especially for `all-MiniLM-L6-v2`, though both models performed similarly at 0.90, with results comparable to keyphrase extraction without embeddings. The

computational time and number of processed documents were also similar for the 0.90 threshold or the model without embeddings.




Figure 5: F1 score for KeyLLM with embeddings for thresholds 75, 85, 90, and without embeddings for two Sentence Transformers.

The results suggest that embeddings perform best when small clusters are formed (higher threshold), but for optimal keyphrase extraction, it is better to avoid them. Using the `all-MiniLM-L6-v2` transformer at thresholds of 0.80–0.85 yields significant improvements over baseline models for exact matching, making it well-suited for large datasets where some accuracy trade-off is acceptable.

Notably, KeyLLM’s exact-match $F1@6$ (~ 15.2) is substantially higher than the best extractive baseline (YAKE, 11.6), while its partial-match $F1@6$ (~ 49.1) is comparable. The resulting exact–partial gap for KeyLLM (~ 34 points) is narrower than that of the extractive models in relative terms: KeyLLM closes roughly 30% of the gap that YAKE exhibits at $F1@6$. This pattern provides direct evidence that generative models are more robust to morphological inflection than extractive methods, because they can produce canonical lemma forms rather than copying inflected surface tokens. We expect this advantage to generalize to other morphologically rich languages where the same extractive–generative asymmetry holds. However, partial matching showed no significant improvement over the best baselines, suggesting that KeyLLM’s gains are concentrated in form normalization rather than in identifying additional relevant concepts. We therefore conducted a further manual evaluation to disentangle these factors.

6.3. Manual Evaluation

While standard keyphrase extraction evaluation relies on automatic matching against gold standard keyphrases (Kim et al., 2010), we additionally performed manual evaluation on a subset of 100 randomly selected documents to better understand the quality of our dataset and the practical utility of extracted keyphrases. This evaluation aimed to assess whether KeyLLM and the YAKE model

extract relevant keyphrases beyond exact matches with author-assigned keyphrases. We also tested YAKE with a larger n-gram range (1,3) to examine whether this adjustment improved results.

The manual evaluation was conducted to complement automated matching metrics and to better assess the semantic relevance of extracted keyphrases. A total of 100 documents were randomly sampled from the test set. Each document was evaluated with reference to the author-assigned keyphrases and the keyphrases extracted by the evaluated models.

For each document, annotators reviewed the abstract and the corresponding list of extracted keyphrases. A keyphrase was considered correct if it satisfied the following criteria:

1. **Topical relevance:** The keyphrase referred to a central concept explicitly discussed in the abstract.
2. **Conceptual adequacy:** The keyphrase represented a meaningful concept rather than a generic modifier (e.g., standalone adjectives were not considered sufficient).
3. **Non-redundancy:** When multiple surface variants of the same concept were extracted (e.g., inflected forms or reordered phrases), only the most representative form was retained.
4. **Linguistic validity:** Morphological variants (e.g., inflected forms) were accepted as correct if they clearly referred to the same concept as the author-assigned keyphrase.

Two evaluation settings were defined:

1. **Manual Exact**, which counted only keyphrases that matched or were equivalent to the author-assigned keyphrases (including acceptable morphological variants).
2. **Manual Extended**, which additionally included contextually relevant keyphrases present in the abstract but omitted from the author-assigned list (e.g., methodologies, named entities, or key concepts discussed but not listed as keyphrases).

The evaluation was primarily conducted by one of the authors, with a second author independently annotating a subset of 30 documents to assess inter-annotator agreement. Both annotators were native Slovak speakers with a background in NLP. Agreement was measured using Cohen's kappa coefficient, yielding $\kappa = 0.61$, which indicates substantial agreement according to standard interpretation guidelines (Landis and Koch, 1977). We note that keyphrase relevance judgments are inherently subjective—even on English benchmarks,

inter-annotator agreement for keyphrase annotation rarely exceeds $\kappa = 0.70$ (Kim et al., 2010)—and the morphological variability of Slovak introduces additional ambiguity in deciding whether an inflected form constitutes a match.

6.3.1. Key Observations and Failure Analysis

Manual vs. Automated Evaluation: Manual evaluations, particularly for exact matches, consistently outperformed automated methods. While automated algorithms struggled with lexical variation and word ordering, manual evaluation was able to recognize semantically related concepts expressed in different forms. For instance, in a thesis whose topic was obesity, the author-assigned keyphrases covered the core medical condition but omitted concepts that were explicitly discussed in the abstract, such as the psychological consequences of obesity or diagnostic procedures. KeyLLM extracted these as additional keyphrases (e.g., *psychologický problém*, *diagnostika*), which were judged relevant under our Manual Extended criterion despite not appearing in the author-assigned list. These are not matches for *obezita* but rather complementary keyphrases that improve topic coverage.

KeyLLM also outperformed YAKE in root-form handling and contextual understanding. It more reliably converted extracted keyphrases to their base forms and was able to connect words into coherent multi-word keyphrases (e.g., *linguistic characteristics* or *social characteristics*), even when the individual words did not appear contiguously in the text. This ability to produce lemmatized, canonical-form keyphrases without an explicit morphological analyzer is a practical advantage of generative LLMs that extends beyond Slovak: any language where authors assign keyphrases in citation form but text contains inflected forms would benefit similarly. Another notable strength of KeyLLM was its ability to capture distinct keyphrases without redundancy, including named entities such as company names, software tools, and medical systems. In contrast, YAKE frequently extracted semantically similar words or partial fragments of other keyphrases.

For YAKE, exact-match performance improved under manual evaluation, as annotators could recognize morphologically inflected forms as equivalent to author-assigned keyphrases. However, extended matching did not outperform partial automated matching, since manual evaluation retained only the most representative keyphrase per concept and excluded repetitive surface variations. When increasing the n-gram range to (1,3), YAKE's manual evaluation scores dropped significantly (Figure 6), as the model increasingly prioritized irrelevant trigrams over the document's main topics.




Figure 6: Comparison of precision, recall, and F1 score of automatically and manually evaluated YAKE model with n-gram range (1,3)

6.3.2. Failure Cases

Morphological mismatch was the most frequent failure mode observed for YAKE. Since YAKE extracts keyphrases strictly in the form in which they appear in the text, Slovak’s rich inflectional morphology regularly caused mismatches with author-assigned keyphrases even when the correct concept was identified (e.g., *aktívnej politiky trhu* instead of *aktívna politika trhu*). This also affected cases where a keyphrase appeared multiple times in different inflected forms or word orders, and YAKE selected a surface form that did not match the author’s version. This failure mode is inherent to any extractive method operating on a morphologically rich language: the same mismatch pattern would arise in Czech, Polish, Finnish, or Turkish, where a single lemma can surface in dozens of case- or agreement-inflected forms. The practical implication is that extractive keyphrase methods require a lemmatization or morphological normalization step to be fairly evaluated in such languages.

Wrong granularity was particularly prominent for YAKE with the (1,3) n-gram range. The model frequently prioritized irrelevant trigrams that were either surface-level variants of previously extracted keyphrases or generic constructions with little discriminative value (e.g., *analysis of trends*, *trends in data*). This failure mode is partially obscured by automated evaluation, where partial matching counts such variations as correct. Manual evaluation rejected these phrases as redundant, explaining the gap between manual and automated scores observed with larger n-gram ranges (Figure 6).

Semantic redundancy was another recurring issue for YAKE. The model often occupied multiple extraction slots with surface variants of the same concept (e.g., extracting *rozvoj*, *rozvoj vidieka*, and *vidiecky rozvoj* as separate keyphrases), thereby reducing the effective coverage of distinct document topics.

Unmotivated adjective extraction represented the primary failure case for KeyLLM. The model frequently extracted standalone adjectives that lacked sufficient topical specificity when not paired with a noun. This behavior also explains the steeper decline in precision with increasing k : while highly relevant keyphrases are prioritized early, adjective-only extractions accumulate as more keyphrases are considered.

7. Conclusion

This work presents SlovKE, a 227,432-document Slovak keyphrase extraction dataset – a 25-fold increase over prior work – created through systematic data cleaning and quality control.

Our evaluation of baseline models reveals that data quality has a direct impact on keyphrase extraction: after cleaning, YAKE’s exact $F1@6$ improved from 7.5 to 11.6. More notably, the KeyLLM model (GPT-3.5-turbo) achieved an exact $F1@6$ of approximately 15.2, substantially narrowing the gap between exact and partial matching that characterizes all baseline models. This gap – with the best exact $F1@6$ at 11.6 versus the best partial $F1@6$ at 51.5 for baseline methods – reflects a fundamental evaluation challenge for morphologically rich languages, where surface-form variation makes strict matching unreliable and lenient matching overly generous. We argue that this exact–partial divergence is not merely a Slovak-specific artifact but a systematic issue that affects keyphrase evaluation in any language with productive inflectional morphology, including other Slavic, Finno-Ugric, and Turkic languages.

Manual evaluation on 100 documents ($\kappa = 0.61$) confirmed that KeyLLM captures semantically relevant keyphrases that automated metrics systematically miss – including named entities, methodological terms, and contextually central concepts absent from author-assigned gold standards. These findings suggest that standard automated evaluation underestimates the true quality of keyphrase extraction in morphologically rich languages and that developing morphology-aware evaluation protocols is an important direction for the field.

Beyond its immediate contribution to keyphrase extraction, SlovKE constitutes foundational infrastructure for Slovak NLP: at 227K documents with structured metadata, it is among the largest publicly available collections of Slovak scientific text and can support supervised keyphrase generation, document classification, cross-lingual transfer to typologically similar languages (Czech, Polish, Croatian), and the development of morphology-aware evaluation metrics. The dataset is available at <https://huggingface.co/datasets/NaiveNeuron/SlovKE> and the evalu-

ation code at <https://github.com/NaiveNeuron/SlovKE>.

8. Limitations

Dataset scope. While our dataset is the largest of its kind for Slovak, it draws exclusively from theses and dissertations registered in the Central Register, which may not capture the full diversity of Slovak academic writing. Extending coverage to journal articles, conference papers, and technical reports would broaden domain representation and is a natural next step for future dataset expansion.

Residual noise. Despite extensive cleaning, some records may contain formatting inconsistencies or uninformative keyphrases, as neither abstracts nor author-assigned keyphrases underwent expert review. Developing automated quality scoring (e.g., keyphrase informativeness classifiers) could help identify and filter such cases, and our dataset provides a large enough training signal to support such efforts.

Absence of keyphrase rankings. Our data source does not provide ranked keyphrases, precluding evaluation with rank-sensitive metrics such as MRR (Voorhees and Tice, 2000) or MAP (Liu, 2009). Crowdsourced or model-assisted ranking annotation over a subset of the dataset would enable these evaluations and constitutes a concrete direction for future work.

Scale of manual evaluation. Our manual evaluation covered 100 documents with two annotators ($\kappa = 0.61$). While this was sufficient to identify systematic patterns in model behavior, scaling this evaluation – potentially through crowdsourcing or active-learning-based annotation – would strengthen the generalizability of our findings across domains and document types.

Model optimization. The models evaluated in this work were not fine-tuned for Slovak. Given that our dataset provides over 180,000 training documents, supervised fine-tuning of sequence labeling or generative models represents the most direct path to improved performance. In particular, fine-tuning Slovak-specific language models (e.g., SlovakBERT) for keyphrase extraction could yield substantial gains and would directly leverage the resource we provide.

Unsupervised methods only. We evaluated only unsupervised and prompting-based approaches. Training supervised models on this dataset – including transformer-based sequence

taggers and keyphrase generation models – is an important next step that could reveal richer structure in the data and substantially improve extraction accuracy.

9. Data Availability

To support reproducibility and future research, we publicly release the SlovKE dataset under the Creative Commons Attribution 4.0 International License (CC BY 4.0) at <https://huggingface.co/datasets/NaiveNeuron/SlovKE>. The dataset comprises 227,432 cleaned Slovak scientific abstracts with author-assigned keyphrases, split into training (182,745 documents), validation (22,343), and test (22,344) sets. Each record contains the abstract text, keyphrases, and metadata (university, year, document type). All evaluation scripts and preprocessing code are available at <https://github.com/NaiveNeuron/SlovKE> under the MIT License. We encourage the community to use this dataset not only for keyphrase extraction but also for related tasks such as document classification, cross-lingual transfer to other Slavic languages, and the development of morphology-aware evaluation metrics.

Acknowledgements

This research was partially funded by the EU NextGenerationEU through the Recovery and Resilience Plan for Slovakia under the project No. 09I02-03-V01-00029.

10. Bibliographical References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Rabah Alzaidy, Cornelia Caragea, and C Lee Giles. 2019. Bi-lstm-crf sequence labeling for keyphrase extraction from scholarly documents. In *The world wide web conference*, pages 2551–2557.
- Marco Basaldella, Elisa Antolli, Giuseppe Serra, and Carlo Tasso. 2018. Bidirectional lstm recurrent neural network for keyphrase extraction. In *Italian Research Conference on Digital Libraries*, pages 180–187. Springer.

- Slobodan Beliga and Sanda Martincic-Ipsic. 2016. [Network-enabled keyword extraction for under-resourced languages](#). In *Semantic Keyword-Based Search on Structured Data Sources - COST Action IC1302 Second International KEYSTONE Conference, IKC 2016, Cluj-Napoca, Romania, September 8-9, 2016, Revised Selected Papers*, volume 10151 of *Lecture Notes in Computer Science*, pages 124–135.
- Ricardo Campos, Vítor Mangaravite, Arian Pasquali, Alípio Jorge, Célia Nunes, and Adam Jatowt. 2020. Yake! keyword extraction from single documents using multiple local features. *Information Sciences*, 509:257–289.
- Martin Docekal and Pavel Smrz. 2022. Query-based keyphrase extraction from long documents. *arXiv preprint arXiv:2205.05391*.
- Ygor Gallina, Florian Boudin, and Beatrice Daille. 2019. Kptimes: A large-scale dataset for keyphrase generation on news documents. *arXiv preprint arXiv:1911.12559*.
- Yifan Gao, Qingyu Yin, Zheng Li, Rui Meng, Tong Zhao, Bing Yin, Irwin King, and Michael Lyu. 2022. [Retrieval-augmented multilingual keyphrase generation with retriever-generator iterative training](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1233–1246, Seattle, United States. Association for Computational Linguistics.
- Nikolaos Giarelis, Nikos Kanakaris, and Nikos Karacapilidis. 2021. A comparative assessment of state-of-the-art methods for multilingual unsupervised keyphrase extraction. In *IFIP International Conference on Artificial Intelligence Applications and Innovations*, pages 635–645. Springer.
- Maarten Grootendorst. 2020. [Keybert: Minimal keyword extraction with bert](#).
- Maarten Grootendorst. 2021. Keyllm: Keyphrase extraction using large language models. <https://maartengr.github.io/KeyBERT/guides/keyllm.html>. Accessed: 2024-03-04.
- Kazi Saidul Hasan and Vincent Ng. 2014. [Automatic keyphrase extraction: A survey of the state of the art](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1262–1273, Baltimore, Maryland. Association for Computational Linguistics.
- Anette Hulth. 2003. Improved automatic keyword extraction given more linguistic knowledge. In *Proceedings of the 2003 conference on Empirical methods in natural language processing*, pages 216–223.
- Su Nam Kim, Olena Medelyan, Min-Yen Kan, and Timothy Baldwin. 2010. [SemEval-2010 task 5 : Automatic keyphrase extraction from scientific articles](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*, pages 21–26, Uppsala, Sweden. Association for Computational Linguistics.
- Pruthwik Kulkarni, Krishnandu Singh, Manish Kulkarni, and Manish Shrivastava. 2022. [MAKED: A multi-lingual automatic keyword extraction dataset](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference (LREC 2022)*, pages 3614–3623, Marseille, France. European Language Resources Association.
- J. Richard Landis and Gary G. Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159–174.
- Lingua. 2021. Lingua: Language detection library. <https://github.com/lingua-project/lingua>. Accessed: 2024-03-04.
- Ling Liu. 2009. Encyclopedia of database systems.
- Rui Meng, Sanqiang Zhao, Shuguang Han, Daqing He, Peter Brusilovsky, and Yu Chi. 2017. Deep keyphrase generation. *arXiv preprint arXiv:1704.06879*.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, pages 404–411.
- Thuy Dung Nguyen and Min-Yen Kan. 2007. Keyphrase extraction in scientific publications. In *International conference on Asian digital libraries*, pages 317–326. Springer.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Eirini Papagiannopoulou and Grigorios Tsoumakas. 2020. A review of keyphrase extraction. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 10(2):e1339.
- Seoyeon Park and Cornelia Caragea. 2020. [Scientific keyphrase identification and classification by pre-trained language models intermediate task transfer learning](#). In *Proceedings of*

- the 28th International Conference on Computational Linguistics (COLING), pages 5409–5419, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Piotr Pezik, Agnieszka Mikołajczyk, Adam Wawrzyński, Bartłomiej Nitoń, and Maciej Ogrodniczuk. 2022. Keyword extraction from short texts with a text-to-text transfer transformer. In *Asian Conference on Intelligent Information and Database Systems*, pages 530–542. Springer.
- Matúš Pikuliak, Štefan Grivalský, Martin Konôpka, Miroslav Blšták, Martin Tamajka, Viktor Bachratý, Marian Simko, Pavol Balážik, Michal Trnka, and Filip Uhlárik. 2022. [SlovakBERT: Slovak masked language model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7156–7168, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D Manning. 2020. Stanza: A python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. Technical report, OpenAI.
- Ian Rogers. 2002. The Google PageRank algorithm and how it works. <https://www.cs.princeton.edu/~chazelle/courses/BIB/pagerank.htm>.
- Dhruva Sahrawat, Debanjan Mahata, Haimin Zhang, Mayank Kulkarni, Agniv Sharma, Rakesh Gosangi, Amanda Stent, Yaman Kumar, Rajiv Shah, and Roger Zimmermann. 2020. [Keyphrase extraction as sequence labeling using contextualized embeddings](#). In *Advances in Information Retrieval: 42nd European Conference on IR Research, ECIR 2020*, pages 328–335, Lisbon, Portugal. Springer.
- Olivier Salaün, Frédéric Piedboeuf, Guillaume Le Berre, David Alfonso-Hermelo, and Philippe Langlais. 2024. [EUROPA: A legal multilingual keyphrase generation dataset](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12718–12736, Bangkok, Thailand. Association for Computational Linguistics.
- Mingyang Song, Yi Feng, and Liping Jing. 2023. [A survey on recent advances in keyphrase extraction from pre-trained language models](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2108–2119, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ayla Rigouts Terryn, Véronique Hoste, and Els Lefever. 2020. [In no uncertain terms: A dataset for monolingual and multilingual automatic term extraction from comparable corpora](#). *Language Resources and Evaluation*, 54(2):385–418.
- Reut Tsarfaty, Djamé Seddah, Sandra Kübler, and Joakim Nivre. 2013. [Parsing morphologically rich languages: Introduction to the special issue](#). *Computational Linguistics*, 39(1):15–22.
- Dávid Varga, Simon Horvát, Zoltán Szoplák, Ľubomír Antoni, Stanislav Krajci, Peter Gurský, and Laura Bachnáková Rózenfeldová. 2022. [Keyphrase extraction from slovak court decisions](#). In *Proceedings of the 22nd Conference Information Technologies - Applications and Theory (ITAT 2022)*, volume 3226 of *CEUR Workshop Proceedings*, pages 142–150. CEUR-WS.org.
- Ellen M. Voorhees and Dawn M. Tice. 2000. [The TREC-8 question answering track](#). In *Proceedings of the Second International Conference on Language Resources and Evaluation (LREC'00)*, Athens, Greece. European Language Resources Association (ELRA).
- Di Wu, Wasi Uddin Ahmad, and Kai-Wei Chang. 2022. Pre-trained language models for keyphrase generation: A thorough empirical study. *arXiv preprint arXiv:2212.10233*.
- Andrej Zelinka. 2023. [Keyphrase extraction from slovak scientific documents](#). Bachelor's thesis, Comenius University, Bratislava, Bratislava, Slovakia. Available at <https://opac.crzp.sk/?fn=detailBiblioForm&sid=C03CB0FBA6588971DD3A2BD1B6BF>.

A. KeyLLM Prompt

Prompt: I have the following document:
[DOCUMENT]
Based on the information above, extract 15 keywords that best describe the topic of the text. Make sure to only extract keywords that appear in the text. Provide the extracted keywords as a list separated by a semicolon. Sort the keywords by their relevance to the text from the most relevant to the least relevant. Do not add any additional text such as "Keywords:" etc.

B. Data Cleaning Details

We identified and removed duplicates, prioritizing documents with complete abstracts and keyphrases. Our final dataset mainly consists of Bachelor and Diploma theses.

We observed a significant absence of primary abstracts and keyphrases, especially in the first five years (Figure 7), while missing secondary abstracts and keyphrases persisted across later years.




Figure 7: Percentage of missing primary abstracts, secondary abstracts, and keyphrases per year

As shown in Table 4, some top-contributing universities exhibit a high proportion of missing keyphrases. We discovered that some universities lack a designated keyphrase field, prompting students to append keyphrases at the end of their abstracts. We attempted to recover these missing keyphrases by identifying and extracting lists appended to abstracts.

We removed overly short abstracts and keyphrases while also cleaning noisy entries. Many abstracts contained irrelevant metadata prepended to the text, such as:

University	Miss. PA	Miss. SA	Miss. K
UK BA	8.65	15.84	99.94
STU BA	10.39	100.00	0.01
SPU Nitra	0.04	80.71	7.91
KU RK	56.55	60.65	38.29
VŠD Sládk.	0.83	19.33	86.10

Table 4: Top contributing universities with the percentage of missing or short primary abstracts (PA), secondary abstracts (SA), or keyphrases (K). Percentages exceeding 50% are highlighted in bold.

ABSTRACT Peter, Doe: Analysis of Various Techniques ... [Bachelor thesis]. Comenius University in Bratislava. Faculty of Philosophy. Thesis Supervisor: Mgr. John Smith, PhD. Year of Defense: 2019. Number of Pages: 80.

This thesis explores various techniques for successfully writing an abstract...

The extra content, typically starting with the author's name or "ABSTRACT" and ending with the publication year, was irrelevant and could obscure the abstract. We identified these patterns and trimmed the abstracts accordingly.

We identified abstracts with appended keyphrases (matching 'kľúčové slová' or variations), used them to fill missing values, and removed them from the abstract text to prevent model bias.

We focused on Slovak keyphrase extraction and found that some abstracts labeled as Slovak were actually written in other languages. Using the lingua language detection library, we found that 20% of Slovak-labeled abstracts (Figure 8) were actually in English. We implemented multilingual detection for abstracts with mixed languages, filtering out those with over 40% non-Slovak content.




Figure 8: Languages detected in primary abstracts labeled as Slovak

Keyphrases were often submitted incorrectly, either as long blocks of text or mixed-language lists. We filtered out keyphrases longer than

300 characters and removed records with insufficient keyphrases or multilingual versions. For keyphrases containing multiple items, we applied multilanguage detection with a 20% non-Slovak language threshold to exclude English phrases while retaining relevant terminology.

Keyphrase fields often contained lists separated by commas, dots, dashes, or semicolons. We used the Stanza tokenizer to identify words and POS tags, split keyphrases into individual items, removed non-alphabetical characters, and deleted records if any keyphrase exceeded four words. This improved keyphrase length consistency, as shown in Figure 9.




Figure 9: Distribution of the number of keyphrases before and after keyphrase cleaning

Finally, we filtered abstracts by length (500–2000 characters) and keyphrase count (4–15), retaining documents sufficiently descriptive for model evaluation.

C. Additional Dataset Statistics

Figure 10 shows the distribution of abstract lengths before final filtering, and Figure 11 summarizes the number of sentences and tokens per abstract.




Figure 10: Distribution of abstract lengths before final filtering

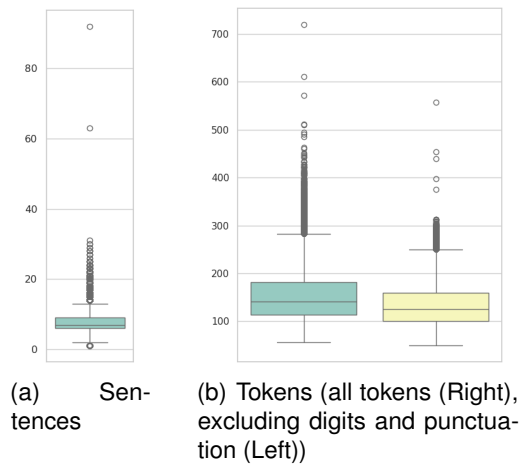


Figure 11: Box plots for the mean number of sentences and tokens in abstracts