

Mind the Gap: Gender-based Differences in Occupational Embeddings

Olga Kononykhina^{1,2}, Anna-Carolina Haensch^{1,2}, Frauke Kreuter^{1,2,3},

¹Ludwig-Maximilians-Universität Munich,

²Munich Center for Machine Learning (MCML),

³University of Maryland, College Park, USA,

Correspondence: olga.kononykhina@lmu.de

Abstract

Large Language Models (LLMs) offer promising alternatives to traditional occupational coding approaches in survey research. Using a German dataset, we examine the extent to which LLM-based occupational coding differs by gender. Our findings reveal systematic disparities: gendered job titles (e.g., “Autor” vs. “Autorin”, meaning “male author” vs. “female author”) frequently result in diverging occupation codes, even when semantically identical. Across all models, 54%–82% of gendered inputs obtain different Top-5 suggestions. The practical impact, however, depends on the model. GPT includes the correct code most often (62%) but demonstrates female bias (up to +18 pp). IBM is less accurate (51%) but largely balanced. Alibaba, Gemini, and MiniLM achieve about 50% correct-code inclusion, and their small (< 10 pp) and direction-flipping gaps could indicate a sampling noise rather than gender bias. We discuss these findings in the context of fairness and reproducibility in NLP applications for social data.

1 Introduction

Occupational coding—the task of assigning standardized occupational categories to free-text job descriptions—is a cornerstone of labor market statistics, informing policy in areas such as employment, migration, and public health. This task is inherently challenging: individuals often describe their work in ambiguous or incomplete terms, and coders must map these descriptions to one of hundreds (or 1,300 in Germany) of possible categories. Historically a manual process, occupational coding has evolved with the rise of automatic solutions. More recently, large language models (LLMs) have been proposed as tools to further automate and advance this process by leveraging their semantic capabilities to match job titles with occupational codes.

This paper examines gender disparities in the coding suggestions made by LLM-based occupa-

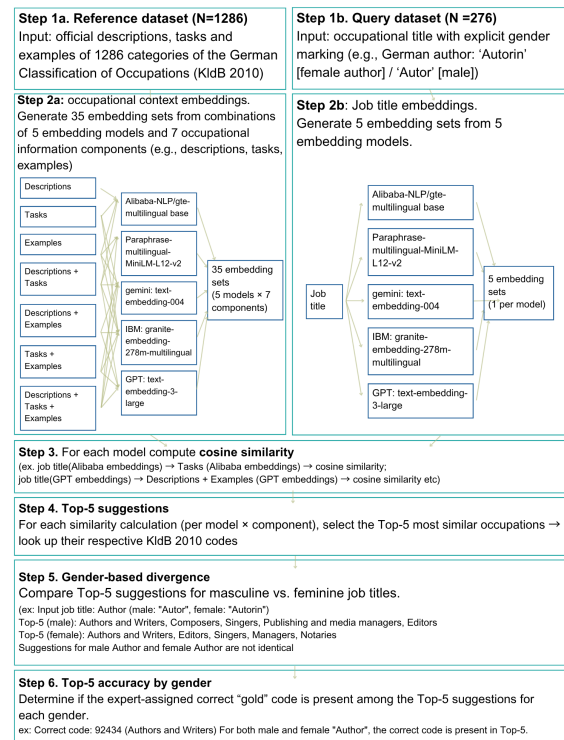


Figure 1: Research pipeline.

tional coding. Using German survey and the official German Classification of Occupations (KldB 2010), we analyze how often male and female forms of job titles receive divergent codes (see Figure 1 for the research pipeline). These differences are not only prevalent but occasionally substantial—pointing to potential downstream harms in labor statistics and policy.

2 Background

Occupational coding—the classification of free-text job titles into standardized categories—has long been recognized as susceptible to gender bias. In manual coding systems, biases can arise from historical taxonomy and human judgment. For example, earlier German occupation classifica-

tions documented that the occupational activities of men are covered more accurately than those of women, leading to misinterpretations in labor statistics (Matthes et al., 2008). Human coders might also inadvertently rely on gendered cues or stereotypes when interpreting ambiguous job titles, though systematic evidence is limited (Conk, 1981).

With the shift toward automated coding, researchers have found that algorithms often perpetuate or even amplify existing gender biases. A large-scale study of English online biographies demonstrated significant bias in occupation classification: including gender indicators (like names or pronouns) skewed predictions and yielded different true positive rates for women vs. men in gender-imbalanced field (De-Arteaga et al., 2019). Even after removing explicit gender tokens, subtle proxies in text led to residual bias favoring the majority gender in a profession. Advanced large language models (LLMs) also reflect societal stereotypes: recent evaluations found LLMs three to six times more likely to assign a person an occupation stereotypical for their gender, often beyond actual labor force proportions (Kotek et al., 2023; Touileb et al., 2023; Kirk et al., 2021).

However, most bias studies focus on English and binary gender contexts (Bolukbasi et al., 2016; Van Der Wal et al., 2022; Savoldi et al., 2025), with less work on languages like German that feature gendered job titles. This highlights the need for further research on robust, bias-resistant coding methods and evaluation in diverse settings.

3 Data

Our empirical analysis draws on two data sources: the German classification of occupations and survey data. The primary reference taxonomy is the German Classification of Occupations 2010 (Klassifikation der Berufe 2010, KldB 2010; Bundesagentur für Arbeit, 2019), which defines 1,286 standardized occupational categories. Each includes a *description* (e.g., Authors and writers producing complex creative texts requiring advanced skills), typical *tasks* (e.g., Creating and writing literary, technical, and factual text), and *example* job titles (e.g., Authors, screenwriters) (see Table 2A in the Appendix for a full illustrative example). These form the basis for generating reference embeddings.

The query set consists of self-reported occupa-

tions from a computer-assisted telephone interview (CATI) survey conducted in Germany in 2019 by INFAS (Institute for Applied Social Science (INFAS), 2019). The representative sample includes 1,415 adults, of whom 1,379 reported either current or past employment. Respondents answered the question "What is/was the occupational task that you mainly perform/performed at your last job?". Open-ended responses (mostly job titles) were manually coded into the five-digit KldB 2010 scheme by professional coders. The process included two coding stages, and adjudication to ensure high-quality labels for evaluation. These professional codes serve as a "gold code" to measure accuracy of the models' suggestions.

A key linguistic feature of German is the use of grammatical gender in occupational titles, typically marked by a masculine base form (e.g., *Lehrer*) and a feminine suffix (e.g., *Lehrerin*). Traditionally, the masculine form has served as a *generisches Maskulinum* (generic masculine) meant to include all genders. For instance, *Lehrer* may refer to any group of teachers. However, research shows that such forms are not interpreted as truly neutral and often lead to male-biased mental representations (Glim et al., 2023; Braun et al., 1998).

This study examines whether embedding-based occupational coding systems reflect or mitigate the semantic and social distinctions introduced by gendered job titles. To assess this, we identified 276 jobs in the dataset that differ only by grammatical gender (e.g., *Autor* vs. *Autorin*, *Ingenieur* vs. *Ingenieurin*; see Table 1A in the Appendix). We then analyzed the similarity of each model's coding suggestions across gendered input.

4 Methodology

To assess the role of gender in embedding-based occupational coding, we evaluated five multilingual models on a set of gendered job title pairs. Given a single gendered job title (masculine or feminine), the system retrieves the five KldB-2010 occupation codes whose reference embeddings are most similar to that title. This ranked list of five codes is our *classification outcome*. We evaluate it with (i) *gender-based divergence*—whether the male and female forms of the same title receive different Top-5 suggestions—and (ii) *Top-5 accuracy*—whether the gold code appears in the Top-5 suggestions.

Embedding models are increasingly used in automatic text classification tasks with large label

Occupational Information Component	MiniLM	Alibaba	Gemini	GPT	IBM
Descriptions	54	72	77	64	67
Tasks	56	77	72	77	67
Examples	62	64	72	56	67
Descriptions + Tasks	72	54	79	62	64
Descriptions + Examples	54	56	59	64	64
Tasks + Examples	59	69	82	72	77
Descriptions + Tasks + Examples	72	56	77	64	72

Table 1: Gender-Based Divergence in Top-5 Job Classification.

Shown is the percentage of job title pairs (male vs. female forms) where the model returned at least one different KldB classification in the Top-5 suggestions. Lower values indicate better gender consistency (ideal = 0%, where male and female forms receive fully identical suggestions).

spaces, such as the categorization of industries (Vidali et al., 2024; Milne et al., 2024), diseases (Nawab et al., 2024; Zhang et al., 2024a), or international trade (Chen et al., 2021). They provide a scalable way to retrieve a small set of relevant categories based on linguistic similarity prior to classification. In occupational coding, embeddings help to narrow the large number of fine-grained job categories by aligning free-text job descriptions with predefined classification labels (Achananuparp and Lim, 2025; Johary et al., 2025; Clavié and Soulié, 2023). We relied on the following models: MiniLM-L12-v2 (multilingual) (Reimers and Gurevych, 2019), Alibaba-NLP gte-multilingual-base (Zhang et al., 2024b), Gemini’s text-embedding-004 (Google, 2024), GPT text-embedding-3-large (OpenAI, 2024), and IBM’s granite-embedding-278m-multilingual (IBM-Research, 2024).

Our evaluation set originates from a CATI survey in which respondents named their occupation. We selected only those answers that met two criteria: (i) the job title is explicitly gendered in German (e.g., *Lehrer* ‘male teacher’, *Autorin* ‘female author’), and (ii) both the masculine and feminine form appeared in the sample and were professionally coded. Titles that were gender-neutral (e.g., *Babysitter*) or represented in only one grammatical gender (e.g., *Soldat* ‘male soldier’) were discarded.

This selection resulted in 276 gendered responses, covering 39 distinct job title pairs (*Lehrer* occurred = 22 times; *Lehrerin* - 30, *Autor* - 1 and *Autorin* - 1 (see Table 1A, Appendix)). Whenever a gender-marked title occurred more than once in the survey, we retained only the first occurrence of each form. Deduplication leaves N = 78 obser-

vations (39 masculine–feminine pairs). For each title we compute a contextualized embedding and compare it against a shared reference set of official job descriptions, tasks and examples from the KldB 2010 classification.

We apply the two indicators defined above — gender-based divergence and Top-5 accuracy — to every pair of masculine–feminine inputs. This setup allows us to test whether embeddings treat male and female occupational titles as semantically equivalent. Ideally, gendered inputs for the same occupation should result in identical suggestions and be classified with equal accuracy.

5 Results

Across all five embedding models, we observed systematic gender differences in Top-5 suggestions for otherwise identical job titles. The results reveal that current embedding approaches do not treat masculine and feminine occupational forms as semantically equivalent, despite their referential equivalence in context.

All five models exhibited gender-based divergence (Table 1), and most displayed at least some gender-related variation in accuracy (Figure 2).

The overall rate of **gender-based divergence** ranged from 54% to 82%, depending on the embedding model and the occupational information component from the reference dataset that was used for embeddings. For example, in one case, the term *Autor* (male form of "author") was matched to occupations such as *Komponist* (composer) and *Verlagskaufmann* (publishing manager), while *Autorin* (female form) yielded *Lektorin* (editor) and *Notarin* (notary) among its Top-5 suggestions. While both forms shared a common first suggestion, three

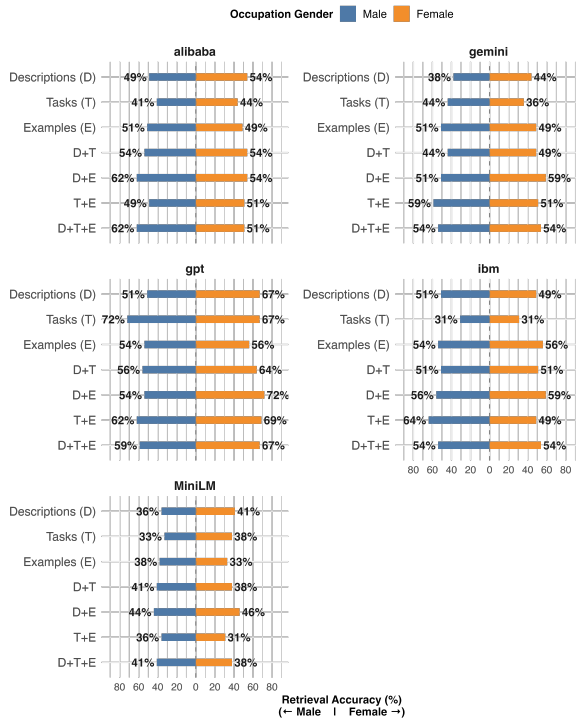


Figure 2: Top-5 Accuracy (%) - whether the gold code appears in the Top-5 suggestions

out of five recommendations differed, including assignments to distinct occupational major groups in KldB 2010.

In addition to the divergence ratio, we report **Top-5 accuracy by gender**—the proportion of masculine or feminine input titles whose gold KldB code appears among the five retrieved suggestions. Figure 2 reveals three patterns. (i) GPT shows a female bias: five of the seven reference configurations favour feminine titles, with the largest margin of +18 pp when Descriptions+Examples are used (72 % vs. 54 %). (ii) IBM is broadly gender-neutral except for the Tasks+Examples setting, where the masculine form is correct in 64 % of cases versus 49 % for the feminine form ($\Delta = 15$ pp). (iii) Alibaba, Gemini, and MiniLM display 50% accuracy, small (< 10 pp) and direction-flipping gaps whose sign depends on the reference subset. Such differences may reflect ranking variance rather than systematic bias.

Moreover adding more textual fields to the reference set (e.g. D \rightarrow D+T+E) does *not* systematically diminish gender differences. This suggests that lexical surface forms exert an influence on embedding similarity, even when more semantic context is introduced.

Taken together, these findings show that gender-

specific lexical variation in German occupational titles systematically affects LLM-based embedding outputs. This can have downstream consequences for fairness in automated occupational classification systems, and by extension, any research or policy relying on them.

6 Discussion

Our analysis demonstrates that embedding-based occupational coding behaves differently on gendered occupational titles in German. Across five state-of-the-art multilingual models and seven reference-set configurations, up to 82% of gendered pairs received divergent Top-5 suggestions. These differences involved distinct occupational codes that sometimes crossed major KldB groups. Such disparities highlight a critical limitation: current LLM-based coding approaches fail to generalize over morphological gender, treating formally different yet semantically identical titles as distinct occupations. This means that (if placed in the survey) two respondents who perform identical work but report it with different grammatical gender therefore would face different shortlists of suggested codes, raising an obvious fairness concern.

How harmful is the mismatch? That depends on whether the correct code still makes it into the list. GPT, for example, supplies the gold code for both forms in about 62% of cases on average and does so slightly *more* often for feminine titles (up to +18 pp). IBM has accuracy around 51% but it is almost balanced. For Alibaba, Gemini, and MiniLM the chance of seeing the gold code hovers around 50%. Coupled with the < 10 pp gender gaps that change sign across reference subsets - differences make it difficult to separate possible bias from sampling and retrieval noise. In short, **divergence is pervasive, but its practical impact varies by model.**

The stakes are high. In Germany, 1,300 standardized job categories inform labor and health statistics, and policy. Even minor classification differences can skew research on employment, wages, health and gender inequality. German adds complexity by making grammatical gender overt—most job titles appear in masculine and feminine forms (e.g., *Anästhesist* vs. *Anästhesistin*), with the generic masculine long dominant in records. While subtle linguistically, these differences are treated as semantically distinct by lan-

guage models, despite their functional equivalence.

To address this, future evaluation protocols incorporate controlled tests for gender consistency, particularly in morphologically rich languages. Survey infrastructures and coding systems should promote or accommodate gender-neutral occupational inputs, such as role-based terms (*Lehrkraft*) or inclusive forms (*Lehrer*in*), while also preparing models to interpret them reliably. Embedding models used in survey contexts may benefit from fine-tuning or contrastive alignment that enforces gender symmetry in professional roles.

7 Conclusion

Our findings show consistent significant disparities: gendered job titles—such as *Autor* vs. *Autorin*—often lead to different occupation codes, despite having identical meanings. Our findings underscore the importance of grounding NLP innovations in language-specific sociolinguistic knowledge. Without rigorous attention to linguistic structure and social context, these tools risk perpetuating systemic biases—particularly in settings where semantic equivalence is masked by morphological variation. Addressing such challenges is crucial not only for the technical refinement of NLP systems, but for ensuring that their real-world applications advance rather than hinder equity.

Limitations

Our study offers a focused evaluation of gender-based divergence in embedding-based occupational coding using a representative German dataset. However, several limitations remain:

First, the analysis is restricted to a relatively small subset of gendered job titles (39 pairs). While these pairs are taken from the representative survey and mirror the titles an automated coder is most likely to encounter, a broader coverage of occupational terms—including less common or more ambiguous cases—will improve generalizability. We plan to extend our evaluation to a larger, more diverse set of occupations in future work.

Second, we focus exclusively on binary gender forms in German (e.g., *Lehrer* vs. *Lehrerin*), without including gender-neutral alternatives such as *Lehrkraft* or inclusive forms like *Lehrer*in*. Comparing how embeddings handle these alternatives would be a valuable extension, especially given their growing use in official communications and survey instruments.

Third, while our analysis uses the most detailed level of the German KldB 2010 classification system, we do not account for the hierarchical nature of occupational categories. Future work could investigate whether suggested categories systematically vary by skill level or specialization depending on gender, and whether gendered patterns emerge at higher aggregation levels within the hierarchy.

Fourth, our evaluation centers on semantic similarity retrieval from embedding spaces, which reflects only one mechanism of LLM-based classification. Other approaches—such as direct classification or few-shot prompting—may exhibit different patterns of gender sensitivity and merit separate analysis.

Fifth, we use cosine similarity as a proxy for human coding. An alternative would be an LLM-as-judge setup, where the model answers a binary prompt “Does title *t* belong to description *d*? yes/no”. This mirrors the human decision rule more closely but was beyond the present scope.

Finally, although we used multiple multilingual embedding models, our findings may not generalize to monolingual or fine-tuned models, particularly those explicitly designed for fairness or domain adaptation in occupational coding.

Bias Statement

In this paper, we study how German grammatical gender markers in job titles (ex. *Lehrer* vs. *Lehrerin* (male/female teacher)) shape the behavior of embedding-based occupation coders. When a model treats the two forms of the same job as semantically distinct, it produces representational harm: it implicitly endorses the idea that the work itself differs along gender lines, thereby imprinting occupational stereotypes. Because occupational codes feed official labor and epidemiological statistics, wage-gap analyses, such divergence can cascade into allocational harm. In other words, surface morphology and not actual job content may end up skewing policy, health, funding and public perception.

Our position is that the link between grammatical gender and occupational meaning is a relic of historical data-collection routines and modelling pipelines, not a reflection of today’s economic, social, or cultural realities. By auditing for those gender-conditioned divergences and mitigating them we can keep automated coders from reproducing or amplifying such harms.

References

- Palakorn Achananuparp and Ee-Peng Lim. 2025. [A multi-stage framework with taxonomy-guided reasoning for occupation classification using large language models](#). *Preprint*, arXiv:2503.12989.
- Tolga Bolukbasi, Kai-Wei Chang, James Y Zou, Venkatesh Saligrama, and Adam Kalai. 2016. [Man is to computer programmer as woman is to home-maker? debiasing word embeddings](#). *arXiv preprint arXiv:1607.06520*.
- Friederike Braun, Anja Gottburgsen, Sabine Sczesny, and Dagmar Stahlberg. 1998. [Können geophysiker frauen sein? generische personenbezeichnungen im deutschen](#). *Zeitschrift für germanistische Linguistik*, 26(3):265–283.
- Bundesagentur für Arbeit. 2019. [Klassifikation der berufe 2010 \(kldb 2010\)](#). Online. Accessed: 2023-10-15.
- Hao Chen, Ben Van Rijnsoever, Marcel Molenhuis, Dennis van Dijk, Y. Tan, and B. D. Rukanova. 2021. [The use of machine learning to identify the correctness of hs code for the customs import declarations](#). In *2021 IEEE 8th International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10. IEEE.
- Benjamin Clavié and Guillaume Soulié. 2023. [Large language models as batteries-included zero-shot esco skills matchers](#). *Preprint*, arXiv:2307.03539.
- Margo A. Conk. 1981. [Accuracy, efficiency and bias: The interpretation of women’s work in the u.s. census of occupations, 1890–1940](#). *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 14(2):65–72.
- Maria De-Arteaga, Alexey Romanov, Hanna Wal-lach, Jennifer Chayes, Christian Borgs, Alexandra Chouldechova, Sahin Geyik, Krishnaram Kenthapadi, and Adam Tauman Kalai. 2019. [Bias in bios: A case study of semantic representation bias in a high-stakes setting](#). In *Proceedings of the Conference on Fairness, Accountability, and Transparency, FAT* ’19*, page 120–128, New York, NY, USA. Association for Computing Machinery.
- Sarah Glim, Anita Körner, Holden Härtl, and Ralf Rumber. 2023. [Early erp indices of gender-biased processing elicited by generic masculine role nouns and the feminine–masculine pair form](#). *Brain and Language*, 242:105290.
- Google. 2024. [Text embedding-004 model via gemini api](#). developers.googleblog.com/gemini-embedding. Accessed: 2025-06-07.
- IBM-Research. 2024. [granite-embedding-278m-multilingual](#). <https://huggingface.co/ibm-granite/granite-embedding-278m-multilingual>. Accessed: 2025-06-07.
- Institute for Applied Social Science (INFAS). 2019. [Standard monthly multi-topic cati survey 2019](#). Unpublished dataset. Nationally representative CATI survey conducted in Germany, April–June 2019. Administered by INFAS.
- Iman Johary, Raphael Romero, Alexandru C. Mara, and Tijl De Bie. 2025. [Jobhop: A large-scale dataset of career trajectories](#). *Preprint*, arXiv:2505.07653.
- Hannah Kirk, Yennie Jun, Haider Iqbal, Elias Benussi, Filippo Volpin, Frederic A. Dreyer, Aleksandar Shtedritski, and Yuki M. Asano. 2021. [Bias out-of-the-box: An empirical analysis of intersectional occupational biases in popular generative language models](#). *Preprint*, arXiv:2102.04130. Accepted to NeurIPS 2021.
- Hadas Kotek, Rikker Dockum, and David Sun. 2023. [Gender bias and stereotypes in large language models](#). In *Proceedings of The ACM Collective Intelligence Conference, CI ’23*, page 12–24. ACM.
- Britta Matthes, Carsten Burkert, and Wolfgang Bier-sack. 2008. [Berufssegmente: Eine empirisch fundierte neuabgrenzung vergleichbarer beruflicher einheiten](#). IAB-Discussion Paper 35/2008, Institut für Arbeitsmarkt- und Berufsforschung (IAB), Nürnberg.
- Cameron Milne, Yezzi Angi Lee, Taylor Wilson, and Hector Ferronato. 2024. [Ensemble retrieval strategies for an improved naics search engine in the u.s. census bureau](#). Zenodo. Global Consulting, 8115 Maple Lawn Blvd, Fulton, MD 20759.
- Khalid Nawab, Madalyn Fernbach, Sayuj Atreya, Samina Asfandiyar, Gulalai Khan, Riya Arora, Iqbal Hus-sain, Shadi Hijjawi, and Richard Schreiber. 2024. [Fine-tuning for accuracy: Evaluation of generative pretrained transformer \(gpt\) for automatic assignment of international classification of disease \(icd\) codes to clinical documentation](#). *Journal of Medical Artificial Intelligence*, 7:8. Published online: 21 June 2024.
- OpenAI. 2024. [Openai text-embedding-3-large model](#). <https://platform.openai.com/docs/guides/embeddings>. Accessed: 2025-06-07.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Beatrice Savoldi, Jasmijn Bastings, Luisa Bentivogli, and Eva Vanmassenhove. 2025. [A decade of gender bias in machine translation](#). *Patterns*, page 101257.
- Samia Touileb, Lilja Øvrelid, and Erik Velldal. 2023. [Measuring normative and descriptive biases in language models using census data](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2242–2248, Dubrovnik, Croatia. Association for Computational Linguistics.

- Oskar Van Der Wal, Jaap Jumelet, Katrin Schulz, and Willem Zuidema. 2022. [The birth of bias: A case study on the evolution of gender bias in an English language model](#). In *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, pages 75–75, Seattle, Washington. Association for Computational Linguistics.
- A. Vidali, N. Jean, and G. Le Pera. 2024. [Unlocking nace classification embeddings with openai for enhanced analysis and processing](#). *Preprint*, arXiv:2409.11524. CC BY-NC-ND 4.0 license.
- Mingyuan Zhang, Jiayi Wang, and et al. 2024a. Retrieval-augmented icd coding: A two-stage system beats vanilla llms by 94 *arXiv preprint arXiv:2401.12345*.
- Xin Zhang, Yanzhao Zhang, Dingkun Long, Wen Xie, Ziqi Dai, Jialong Tang, Huan Lin, Baosong Yang, Pengjun Xie, Fei Huang, and 1 others. 2024b. [mgte: Generalized long-context text representation and reranking models for multilingual text retrieval](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1393–1412.

A Appendix

	English job title	Male German job title	Female German job title	Male titles N	Female titles N
1	Department Head	Abteilungsleiter	Abteilungsleiterin	3	1
2	Employee	Angestellter	Angestellte	4	5
3	Public Sector Employee	Angestellter im öffentlichen Dienst	Angestellte im öffentlichen Dienst	1	1
4	Doctor	Arzt	Ärztin	9	6
5	Author	Autor	Autorin	1	1
6	Bank Clerk	Bankkaufmann	Bankkauffrau	5	6
7	Construction Manager	Bauleiter	Bauleiterin	3	1
8	Civil Servant	Beamter	Beamtin	14	5
9	Consultant	Berater	Beraterin	1	1
10	Accountant	Buchhalter	Buchhalterin	1	3
11	Bookseller	Buchhändler	Buchhändlerin	1	1
12	Office Administrator	Bürokaufmann	Bürokauffrau	1	4
13	Retail Salesperson	Einzelhandelskaufmann	Einzelhandelskauffrau	1	2
14	Electrician	Elektriker	Elektrikerin	5	1
15	Childcare Worker	Erzieher	Erzieherin	5	11
16	Tax Officer	Finanzbeamter	Finanzbeamtin	2	1
17	Janitor / Caretaker	Hausmeister	Hausmeisterin	3	2
18	Engineer	Ingenieur	Ingenieurin	7	1
19	Legal Expert	Jurist	Juristin	1	3
20	Clerical Assistant	Kaufmännischer Angestellter	Kaufmännische Angestellte	4	9
21	Nurse	Krankenpfleger	Krankenpflegerin	5	2
22	Warehouse Worker	Lagerist	Lageristin	1	1
23	Teacher	Lehrer	Lehrerin	22	30
24	Educator	Pädagoge	Pädagogin	1	2
25	Nursing Assistant	Pflegehelfer	Pflegehelferin	1	1
26	Police Officer	Polizeibeamter	Polizeibeamtin	3	1
27	Lawyer	Rechtsanwalt	Rechtsanwältin	3	1
28	Administrative Clerk	Sachbearbeiter	Sachbearbeiterin	7	5
29	School Principal	Schulleiter	Schulleiterin	1	3
30	Social Worker	Sozialpädagoge	Sozialpädagogin	2	4
31	Social Insurance Clerk	Sozialversicherungsfachangestellter	Sozialversicherungsfachangestellte	2	1
32	Taxi Driver	Taxifahrer	Taxifahrerin	3	1
33	Salesperson	Verkäufer	Verkäuferin	2	8
34	Insurance Clerk	Versicherungskaufmann	Versicherungskauffrau	2	1
35	Administrative Assistant	Verwaltungsangestellter	Verwaltungsangestellte	2	5
36	Administrative Officer	Verwaltungsbeamter	Verwaltungsbeamtin	4	2
37	Administrative Specialist	Verwaltungsfachangestellter	Verwaltungsfachangestellte	2	2
38	Dentist	Zahnarzt	Zahnärztin	2	1
39	Dental Technician	Zahntechniker	Zahntechnikerin	1	2

Table 1A: Gendered German Job Title Pairs from Survey Responses (with English Translations). Based on open-ended responses from a survey of 1,379 adults in Germany, we identified 39 occupations that appeared in both masculine and feminine grammatical forms (e.g., *Lehrer / Lehrerin* for “teacher”). These job titles were reported directly by respondents (columns: *Male German job title* and *Female German job title*). Some titles were mentioned by multiple respondents (e.g., *Lehrer* = 22, *Lehrerin* = 30). For the analysis, only the first occurrence of each gendered form was retained, resulting in 78 unique observations. The table lists translated *English job title*, the respondents answers - *male and female German forms*, and the number of times each gendered form was mentioned in the survey (columns “*Male titles N*” and “*Female titles N*”)

Occupational Information Component	Associated Text (translated from german into english)
Descriptions	All authors and writers whose work is highly complex and requires a correspondingly high level of knowledge and skill. Members of these professions write screenplays for feature films, documentaries or short film reports or write speeches, novels, short stories, poems, plays and other non-journalistic texts for publication or presentation
Tasks	Conceive and write novels, short stories, poems, plays or radio plays Prepare speech manuscripts, for example for company events such as presentations or annual press conferences or for private events such as weddings or birthdays Write scripts for film and television productions, developing the content, plot and characters of a story Elaborate dialogues, describe locations, provide detailed information about spatial and temporal sequences, props, sounds, music, lighting and moods write brochures, manuals and similar technical publications research factual content and obtain other necessary information select materials for publication and make contact with publishers or literary agencies
Examples	Author Screenwriter Speechwriter Writer
Descriptions + Tasks	All authors and writers whose work is highly complex and requires a correspondingly high level of knowledge and skill. Members of these professions write screenplays for feature films, documentaries or short film reports or write speeches, novels, short stories, poems, plays and other non-journalistic texts for publication or presentation. Conceive and write novels, short stories, poems, plays or radio plays Prepare speech manuscripts, for example for company events such as presentations or annual press conferences or for private events such as weddings or birthdays Write scripts for film and television productions, developing the content, plot and characters of a story Elaborate dialogues, describe locations, provide detailed information about spatial and temporal sequences, props, sounds, music, lighting and moods write brochures, manuals and similar technical publications research factual content and obtain other necessary information select materials for publication and make contact with publishers or literary agencies
Descriptions + Examples	All authors and writers whose work is highly complex and requires a correspondingly high level of knowledge and skill. Members of these professions write screenplays for feature films, documentaries or short film reports or write speeches, novels, short stories, poems, plays and other non-journalistic texts for publication or presentation. Author Screenwriter Speechwriter Writer
Tasks + Examples	Conceive and write novels, short stories, poems, plays or radio plays Prepare speech manuscripts, for example for company events such as presentations or annual press conferences or for private events such as weddings or birthdays Write scripts for film and television productions, developing the content, plot and characters of a story Elaborate dialogues, describe locations, provide detailed information about spatial and temporal sequences, props, sounds, music, lighting and moods write brochures, manuals and similar technical publications research factual content and obtain other necessary information select materials for publication and make contact with publishers or literary agencies. Author Screenwriter Speechwriter Writer
Descriptions + Tasks + Examples	All authors and writers whose work is highly complex and requires a correspondingly high level of knowledge and skill. Members of these professions write screenplays for feature films, documentaries or short film reports or write speeches, novels, short stories, poems, plays and other non-journalistic texts for publication or presentation. Conceive and write novels, short stories, poems, plays or radio plays Prepare speech manuscripts, for example for company events such as presentations or annual press conferences or for private events such as weddings or birthdays Write scripts for film and television productions, developing the content, plot and characters of a story Elaborate dialogues, describe locations, provide detailed information about spatial and temporal sequences, props, sounds, music, lighting and moods write brochures, manuals and similar technical publications research factual content and obtain other necessary information select materials for publication and make contact with publishers or literary agencies. Author Screenwriter Speechwriter Writer

Table 2A: Illustrative example of KldB Code 92434 (authors, writers) components for embedding construction (one of 1 286 codes in the reference dataset)