

# A Multi-Stage Framework with Taxonomy-Guided Reasoning for Occupation Classification Using Large Language Models

Palakorn Achananuparp and Ee-Peng Lim

Singapore Management University  
 palakorna@smu.edu.sg, eplim@smu.edu.sg

## Abstract

Automatically annotating job data with standardized occupations from taxonomies, known as occupation classification, is crucial for labor market analysis. However, this task is often hindered by data scarcity and the challenges of manual annotations. While large language models (LLMs) hold promise due to their extensive world knowledge and in-context learning capabilities, their effectiveness depends on their knowledge of occupational taxonomies, which remains unclear. In this study, we assess the ability of LLMs to generate precise taxonomic entities from taxonomy, highlighting their limitations, especially for smaller models. To address these challenges, we propose a multi-stage framework consisting of inference, retrieval, and reranking stages, which integrates taxonomy-guided reasoning examples to enhance performance by aligning outputs with taxonomic knowledge. Evaluations on a large-scale dataset show that our framework not only enhances occupation and skill classification tasks, but also provides a cost-effective alternative to frontier models like GPT-4o, significantly reducing computational costs while maintaining strong performance. This makes it a practical and scalable solution for occupation classification and related tasks across LLMs.

## 1 Introduction

The Standard Occupational Classification (SOC) systems and associated taxonomies, such as the O\*NET-SOC taxonomy (O\*NET-SOC), are essential to labor market analysis and workforce research (Hu et al. 2016; Guo et al. 2024; Kalhor et al. 2024). These systems provide a systematic framework for categorizing occupations, jobs, skills, and job relevant knowledge. By mapping labor-market data, such as job postings, resumes, or employment surveys, to standardized occupations (referred to as SOC codes and their corresponding SOC titles), researchers, policymakers, and businesses can gain useful insights into labor market trends and expectations. These insights are instrumental for shaping policies, optimizing workforce planning, and guiding industry/manpower development initiatives.

However, automatic occupation classification is inherently challenging due to the scarcity of large publicly available labeled datasets which require laborious efforts of la-

bor market experts. The lack of data resources further hinders the development and evaluation of robust supervised learning-based classification models. Given the sheer size and complexity of occupational taxonomies, non-expert annotators would require rigorous training to effectively interpret and apply the nuanced structures of these classifications, further complicating the issues.

Recent advancements in large language models (LLMs) have shown promises for addressing these challenges through prompt engineering (Li, Kang, and De Bie 2023; D’Oosterlinck et al. 2024). Despite their potential, critical research gaps remain unaddressed. First, the extent to which LLMs rely on their knowledge of occupational taxonomies for classification has not been explored. LLMs are pre-trained on broad datasets, which may include partial representations of occupational taxonomies. The quality and coverage of this taxonomic knowledge can also vary significantly across LLMs due to differences in their training data, architectures, and model sizes. Consequently, their classification effectiveness may be inconsistent. This raises important questions about the reliability and accuracy of LLM-based methods, particularly when dealing with less common or highly specific occupational categories, where the knowledge gaps are likely to be most pronounced.

Moreover, existing LLM-based methods often rely on simplistic prompting techniques, such as direct queries (Li, Kang, and De Bie 2023) or limited in-context examples (D’Oosterlinck et al. 2024), which fail to fully utilize the rich information embedded in taxonomies. Taxonomies like O\*NET-SOC contain detailed descriptions and hierarchical structures that can enhance classification accuracy. However, current methods fail to integrate this information effectively. While chain-of-thought (CoT) reasoning (Wei et al. 2022) is popular for general-purpose prompting technique, its adaptation to domain-specific tasks like occupation classification remains unclear. CoT decomposes problems into logical steps but does not inherently exploit taxonomic knowledge, limiting the performance of existing approaches.

To address these research gaps, we conduct a knowledge assessment study to evaluate the capabilities of various LLMs in generating accurate SOC titles and codes from the ONET-SOC taxonomy. Our findings reveal two major shortcomings: (1) LLMs, particularly smaller models, often struggle with generating precise SOC codes, particularly

for less common occupations, and (2) their classification accuracy varies significantly based on hierarchical granularity, indicating limited internalization of taxonomic knowledge. Building on these findings, we propose a multi-stage framework that integrates inference, retrieval, and reranking components. Specifically, we incorporate external knowledge from the O\*NET taxonomy to guide LLM’s reasoning through grounded in-context examples. Unlike computationally intensive fine-tuning or complex model modifications, our framework relies on efficient prompting and retrieval strategies, making it adaptable to a variety of LLMs without requiring additional training. By leveraging the detailed class information embedded in ONET-SOC, this framework enhances the accuracy and robustness of occupation classification tasks. Finally, we evaluate the framework on both occupation and skill classification tasks, demonstrating its effectiveness in mitigating the knowledge gaps and its adaptability across domains.

Our contributions are fourfold:

1. We are the first to systematically investigate LLM knowledge of the O\*NET-SOC taxonomy through a series of generation tasks, providing insights into their strengths and limitations.
2. We propose a novel framework that builds on retrieval-based prompting methods (Li, Kang, and De Bie 2023; D’Oosterlinck et al. 2024) by systematically integrating grounded rationals into in-context examples, aligning LLM reasoning and outputs with taxonomic information while minimizing prompt engineering effort.
3. We explicitly examine the relationship between LLMs’ taxonomic knowledge and their downstream classification performance. Our findings demonstrate that leveraging external taxonomic information throughout the classification pipeline, from inference to reranking, enhances alignment with the taxonomy and improves overall performance across different LLMs.
4. We demonstrate that our framework significantly reduces the computational costs while maintaining high classification accuracy. By enabling smaller, more efficient LLMs to achieve performance comparable to frontier models like GPT-4o, the framework provides a scalable and cost-effective alternative to GPT-4o’s direct prompting, making it practical for large-scale classification tasks.

## 2 Related Work

### Assessing LLM Knowledge

Existing research has investigated the ability of LLMs to memorize factual knowledge in the training data mainly through learning from high-frequency patterns, and their struggle with low-resource tasks (Tänzer, Ruder, and Rei 2022). Researchers have developed methods to systematically evaluate LLM knowledge of both common and long-tail facts in generic and specialized domains by creating tasks derived from knowledge graphs. (Sun et al. 2023; Luo et al. 2023; Kumar et al. 2024). In the context of taxonomic knowledge, recent work has examined LLMs’ understanding of hierarchical structures in specialized domains.

For example, Sun et al. (Sun et al. 2024) developed binary and multiple-choice questions to comprehensively test frontier LLMs’ comprehension of entities across ten taxonomies, spanning domains such as shopping, language, and medicine. Similarly, Soroush et al. (Soroush et al. 2024) assessed frontier LLMs’ capabilities in generating structured codes from medical taxonomies. Both studies revealed limitations in frontier LLMs’ ability to generate taxonomic entities from specialized domains.

Our study is the first to examine LLM knowledge of occupational taxonomies. Our task design is more comprehensive than previous work, encompassing those of Sun et al. (Sun et al. 2024) and Soroush et al. (Soroush et al. 2024). Additionally, we introduce a task variant that tests LLM performance when partial hints are provided. Moreover, our investigation focuses on efficient LLMs, which are practical for large-scale downstream applications such as our occupation classification.

### Occupation and Skill Classifications

Occupation classification has been explored in numerous studies using proprietary datasets and various methodologies to map job-related data to standardized occupations. Traditional approaches often involve multi-class text classification that rely heavily on feature engineering techniques, such as bag-of-words and word embeddings (Boselli et al. 2017), combined with standard machine learning models like SVM and random forest (Russ et al. 2016; Javed et al. 2015). Some studies focus on classifying occupations based on job titles (Boselli et al. 2017; Javed et al. 2015), while others use richer input data, such as full job postings (Li, Kang, and De Bie 2023).

Recent work has integrated LLMs in the occupation classification pipeline. For example, LLM4Jobs (Li, Kang, and De Bie 2023) employed a two-stage framework, utilizing Vicuna 33B to infer occupational labels from job postings and retrieve top-ranked candidates from the ESCO taxonomy. Similarly, LLMs have been applied to workforce-related tasks like multi-label skill classification, including generating synthetic training data for fine-tuning (Clavié and Soulié 2023; Decorte et al. 2023), inferring standardized skill labels (D’Oosterlinck et al. 2024), and reranking candidates (Clavié and Soulié 2023; D’Oosterlinck et al. 2024).

To our knowledge, this study is one of the few to evaluate occupation classification on large-scale job records derived from diverse, real-world online resume data, involving fine-grained SOC labels from the O\*NET-SOC taxonomy. Unlike fine-tuned systems (Zhang, Jensen, and Plank 2022; Clavié and Soulié 2023; Decorte et al. 2022), our framework is completely training-free and most comparable to LLM4Jobs (Li, Kang, and De Bie 2023) and Infer-Retrieve-Rank (IReRa) (D’Oosterlinck et al. 2024), where LLMs are utilized in conjunction with a taxonomy, along with retrieval and reranking modules in a classification pipeline. However, it differs from previous work in a few key aspects.

Firstly, our framework uniquely focuses on aligning LLM reasoning with taxonomic knowledge through taxonomy-guided in-context examples, which significantly enhances the overall accuracy. In contrast, LLM4Jobs and

SOC Code	Major	Minor	Broad	Detailed
15-2051.00	15	15-2	15-205	15-2051

Table 1: Hierarchical breakdown of the SOC code for the Data Scientists occupation (15-2051.00)

IReRa simply utilize zero-shot and general-purpose CoT prompting approaches, respectively. Secondly, we leverage rich external taxonomic information in the retrieval stage, a strategy similar to LLM4Jobs but not utilized in IReRa. This overcomes the limited taxonomic knowledge of many existing LLMs. Lastly, while LLM4Jobs focuses solely on retrieval after inference, our framework, like IReRa, introduces an additional reranking stage, where LLMs are used to further refine the list of retrieved candidates.

### 3 Knowledge Assessment Study

We begin by systematically examining the extent to which LLMs memorize and internalize occupational taxonomies through a series of generation tasks. For this investigation, we chose the O\*NET-SOC taxonomy – a widely adopted hierarchical classification system for standardized occupations – and evaluated seven LLMs, which are described in detail in subsequent sections.

#### The O\*NET-SOC Taxonomy

The O\*NET-SOC taxonomy (O\*NET-SOC) is a hierarchical system for organizing occupations within the Standard Occupational Classification (SOC) framework. Each occupation is represented by a unique SOC code, a title, a detailed description, and other associated information. The taxonomy comprises five levels of granularity, all encoded in the SOC codes: **major groups** (2-digit codes), **minor groups** (3-digit codes), **broad occupations** (5-digit codes), **detailed occupations** (6-digit codes), and **specific occupations** (8-digit codes). Table 1 illustrates its hierarchical structure using Data Scientists as an example. This study focuses on **O\*NET-SOC 2019**, the most recent version, which includes 1,016 specific occupations.

#### Choices of LLMs

We selected seven LLMs, comprising two frontier models: **GPT-4o** and **Gemini-1.5 Pro**, and five efficient models: **GPT-3.5 Turbo**, **Gemini-1.5 Flash**, **Claude Instant 1.2**, **Llama 3 8B Instruct**, and **Mistral Small**.

The frontier models were included to establish an upper bound on performance. These models represent the state-of-the-art in LLMs at the time of this study and are characterized by their significantly larger parameter sizes and more advanced architectures. In contrast, the efficient models were chosen for their balance between performance and computational cost. These models tend to have smaller parameter sizes or optimized architectures, enhancing their token processing speed and reducing inference costs. Their efficiency makes them particularly practical for large-scale downstream classification tasks such as occupation classification.

Mode	Direction		Granularity	Task
	Context	Output		
Complete Recall	Code	Title	8-digit	CT8
	Title	Code	8-digit 2-digit	TC8 TC2
Partial Recall	Code	Title	8-digit	CT8
	Title	Code	8-digit	TC8
Recognition	Code	Title	8-digit	CT8
	Title	Code	8-digit 2-digit	TC8 TC2

Table 2: Knowledge Assessment Tasks

#### Task Design

We devise a set of generation tasks to assess LLM knowledge by instructing LLMs to generate SOC titles and codes in order to assess both the ability to directly retrieve taxonomy information and the capacity to perform at varying levels of granularity within the taxonomy.

First, we structure the tasks into two modes of generation: **recall** and **recognition**. Recall tasks evaluate the model’s ability to freely generate titles or codes, while recognition tasks assess its ability in selecting the correct output from a list of candidates. Recall tasks are further divided into **complete** and **partial** recall modes. In complete recall, the model generates the full code or title without additional hints. In partial recall, the model is tasked with completing a partially provided code or title.

Additionally, we focus on two main *context-to-output* generation directions: code-to-title (**CT**) and title-to-code (**TC**). In CT, we test the model’s ability to generate a title (output) from a given code (context), while in the TC direction, we test its ability to generate a code (output) from a given title (context). To analyze model performance across varying levels of granularity, each task direction is examined using both **8-digit** and **2-digit** codes. These combinations yield three task variants: **CT8** (title generation from 8-digit codes), **TC8** (8-digit code generation from titles), and **TC2** (2-digit code generation from titles). We exclude CT2 (title generation from 2-digit codes) as multiple titles can correspond to one 2-digit code.

The final task set, displayed in Table 2, includes CT8, TC8, and TC2 in the complete recall and recognition tasks and CT8 and TC8 in the partial recall tasks. TC2 is excluded from partial recall tasks as a single-digit hint provides insufficient context for meaningful evaluation.

#### Research Questions

Given these configurations, we pose the following research questions:

**RQ1:** How do frontier and efficient LLMs compare in their ability to generate occupational titles and codes from the O\*NET-SOC taxonomy?

**RQ2:** How does LLM performance vary across different levels of granularity within the O\*NET-SOC taxonomy?

**RQ3:** How do different task design elements, such as generation mode, direction, etc., influence LLM performance in generating titles and codes?

### Task Generation

For each occupation in the taxonomy, we generate recall and recognition tasks using prompt templates, shown in Figure 10 in the appendix. All tasks include a one-shot example to help models follow instructions accurately, which is especially helpful for smaller models.

For partial recall tasks, we provide the first half of the answer as a hint. For example, for the occupation *Data Scientists* (15-2051.00), the prompts for a TC8 partial recall task are: “*In the O\*NET-SOC 2019 taxonomy, the 8-digit code for the title Data Scientists is <blank>. Hint: It starts with 15-20.*”

For recognition tasks, we generate five randomly ordered options: the correct answer corresponding to the given occupation, two candidates randomly sampled from related occupations within the same major (2-digit code) groups, and two candidates randomly sampled from different major groups.

All experiments were conducted in October 2024. The models from OpenAI (GPT-4o and GPT-3.5 Turbo), Vertex AI (Gemini-1.5 Pro and Gemini-1.5 Flash), and Amazon Bedrock (Claude Instant 1.2, Llama 3 8B Instruct, and Mistral Small), were accessed via their respective APIs. The temperature was set to zero for all tasks and LLMs to ensure deterministic outputs.

### Metrics and Statistical Analysis

We evaluate performance using accuracy, defined as the proportion of generated answers that exactly match the corresponding values in O\*NET-SOC. The average accuracy score provides an overall metric measuring the model’s ability to generate taxonomic entities across tasks. A one-way ANOVA and linear regression were performed using Python’s SciPy (1.11.3) and Statsmodels (0.14) libraries. The data used in the analyses satisfy normality assumption.

### Knowledge Assessment Results

Table 3 presents the performance of the seven LLMs on the complete recall, partial recall, and recognition tasks. Overall, GPT-4o is the top performer, achieving an impressive average accuracy of 0.9211. Notably, GPT-3.5 Turbo follows as the second-best model with an average accuracy of 0.732, outperforming a more advanced model like Gemini-1.5 Pro. In contrast, Llama 3 8B Instruct has the lowest performance with an average accuracy of 0.2794. Most models perform better in recognition than recall tasks. Moreover, complete recall are generally more challenging than partial recall tasks.

**RQ1: How do frontier and efficient LLMs compare in their ability to generate occupational titles and codes from the O\*NET-SOC taxonomy?**

A one-way ANOVA was conducted to examine differences between frontier and efficient models. The results indicate a statistically significant difference between the performance scores of the frontier and efficient models ( $F =$

20.9,  $p < 0.001$ ). Specifically, the frontier models consistently achieve superior performance across all recall and recognition tasks, with GPT-4o setting a significant upper bound. The findings highlight the extensive knowledge and capabilities of frontier models in understanding and accurately generating titles and codes from the O\*NET-SOC taxonomy.

**RQ2: How does LLM performance vary across different levels of granularity within the O\*NET-SOC taxonomy?**

To analyze how LLM performance varies across different levels of granularity, a one-way ANOVA was performed to compare LLM accuracy on tasks involving 8-digit (specific occupations) and 2-digit (major occupational groups) SOC codes. The results demonstrate a statistically significant difference between the two levels ( $F = 12.667$ ,  $p < 0.001$ ), with LLMs performing significantly better on 2-digit tasks. The findings suggest that LLM performance varies with the granularity of the occupational codes. This disparity may be influenced by several factors, including the inherent difficulty of predicting finer-grained categories and potential limitations in the LLMs’ ability to fully leverage hierarchical information within the taxonomy.

**RQ3: How do different task design elements, such as generation mode, direction, etc., influence LLM performance in generating titles and codes?**

To explore how task design influences LLM performance in generating occupational titles and codes, we fit two Ordinary Least Squares (OLS) regression models with accuracy scores as the dependent variable.

**Recall vs. recognition:** In the first model, the independent variables included individual LLM models (e.g., GPT-4o), generation mode (recall or recognition), output type (title or code), code granularity (8-digit or 2-digit), and interaction terms to capture combined effects. An N-1 dummy coding was applied for categorical variables, with *Llama 3 8B Instruct*, *recall*, *code*, and *2-digit* as the reference categories for LLM, generation mode, output type, and code granularity, respectively. The model explained a significant proportion of variance in accuracy scores ( $N = 52$ ,  $R^2 = 0.759$ ).

Adjusting for other factors, significant main effects were found for LLM model, code granularity, and the generation mode and output type interaction. Firstly, the performance of individual LLMs showed significant differences. GPT-4o ( $\beta = 0.642$ ,  $p < 0.001$ ), Gemini-1.5 Pro ( $\beta = 0.44$ ,  $p < 0.001$ ), GPT-3.5 Turbo ( $\beta = 0.453$ ,  $p < 0.001$ ), and Claude Instant 1.2 ( $\beta = 0.24$ ,  $p < 0.01$ ) outperformed the baseline Llama 3 8B Instruct, whereas Gemini-1.5 Flash and Mistral Small did not show significant improvement over the baseline. This finding aligns with the results of RQ1, which demonstrated that frontier models significantly outperform efficient models across tasks.

Secondly, code granularity significantly affects accuracy. Tasks involving 8-digit codes were more challenging ( $\beta = -0.308$ ,  $p < 0.001$ ) compared to 2-digit codes. The finding is consistent with previous analyses in RQ2, confirming that tasks requiring greater taxonomic specificity are significantly more challenging.

Lastly, there is a significant interaction between recogni-

LLM	Complete Recall			Partial Recall		Recognition			Average
	CT8	TC8	TC2	CT8	TC8	CT8	TC8	TC2	
GPT-4o	<b>0.8612</b>	<b>0.9104</b>	<b>0.9754</b>	<b>0.9626</b>	<b>0.873</b>	<b>0.9242</b>	<b>0.939</b>	0.9232	<b>0.9211</b>
Gemini-1.5 Pro	0.2923	0.5453	0.8858	<u>0.8573</u>	0.6083	0.7067	<u>0.8907</u>	<b>0.9685</b>	0.7194
GPT-3.5 Turbo	<u>0.5551</u>	<u>0.6742</u>	<u>0.9183</u>	0.8041	<u>0.6693</u>	0.3878	0.8888	<u>0.9587</u>	<u>0.7320</u>
Gemini-1.5 Flash	0.002	0.1004	0.7333	0.5522	0.1614	0.2008	0.3622	0.9439	0.3820
Claude Instant 1.2	0.1663	0.3809	0.6663	0.6378	0.3661	0.373	0.6831	0.8829	0.5196
Llama 3 8B Instruct	0.0118	0.0768	0.247	0.3022	0.1919	0.3819	0.5846	0.439	0.2794
Mistral Small	0.0039	0.2185	0.7441	0.4734	0.248	0.1437	0.4813	0.8524	0.3957

Table 3: Performance of frontier and efficient LLMs on knowledge recall and recognition task variants. The best results are highlighted in bold and the second-best results are underlined.

tion mode and title generation ( $\beta = -0.277$ ,  $p < 0.05$ ), suggesting that recognizing valid titles is more challenging than recognizing codes. This disparity is likely due to titles tend to require nuanced linguistic understanding, whereas codes are more structured and deterministic, making them easier for LLMs to match accurately.

**Complete vs. partial recalls:** To examine the influence of recall modes, a second OLS regression model was fit using the same variables but replacing generation mode with recall mode (*complete recall* as the reference category). The analysis ( $N = 35$ ,  $R^2 = 0.921$ ) yielded consistent main effects for LLM model and code granularity. Additionally, significant effects were found for output type and the recall mode and output type interaction. Specifically, titles were more challenging to generate than codes in recall tasks ( $\beta = -0.072$ ,  $p < 0.05$ ). While there was no overall difference in accuracy between complete and partial recall, partial recall significantly improved performance for title generation compared to complete recall ( $\beta = 0.355$ ,  $p < 0.001$ ).

Altogether, these findings suggest that while current LLMs exhibit varying degrees of knowledge regarding the O\*NET-SOC taxonomy, their performance in tasks involving the taxonomy varies, with GPT-4o demonstrating the strongest performance. Factors such as task design elements, particularly code granularity and output type, should be considered when leveraging LLMs for downstream applications like occupation classification.

## 4 Framework

We propose a multi-stage framework for occupation classification, consisting of **inference**, **retrieval**, and **reranking** stages, shown in Figure 1. Let  $\mathbb{X}$  denote the input space (e.g., job title-company pairs),  $\mathbb{Y}$  the label space (e.g., SOC titles and codes), and  $\mathbb{T}$  the taxonomy (e.g., the O\*NET-SOC taxonomy) from which  $\mathbb{Y}$  are drawn. The framework maps an input  $x \in \mathbb{X}$  to corresponding labels  $y \in \mathbb{Y}$ , with  $\mathbb{Y} \subseteq \mathbb{T}$ .

The framework is specifically designed to address the limited domain-specific knowledge of LLMs by following these guiding principles: (i) leveraging the reasoning capabilities of LLMs to enhance their ability to infer and select correct class labels; and (ii) utilizing detailed class information from the taxonomy, such as class descriptions, to guide the reasoning process and enhance the retrieval of relevant candi-

date labels.

To align LLM reasoning, we introduce **taxonomy-guided reasoning example (TGRE)**, a simple method that incorporates *grounded rationales* into in-context examples with minimal prompt engineering. A grounded rationale is an explanation explicitly linked to the taxonomy, clarifying the relationship between the input and the corresponding label. In a typical classification task, an in-context example  $(x, y)$  consists of the task input  $x$  (e.g., a job title-company pair) and the corresponding ground-truth label  $y$  (e.g., an SOC title and code). Our method enriches these examples by injecting a grounded rationale  $r$  that explains the connection between  $x$  and  $y$ . We refer to this grounded example  $(x, r, y)$  as TGRE.

To construct TGRE, we define  $T(y)$  as a function that retrieves relevant information from the taxonomy corresponding to the ground-truth label  $y$ . Formally,  $T(y) = \{t_1, t_2, \dots, t_p\}$  where  $\{t_1, t_2, \dots, t_p\}$  represents the set of  $p$  taxonomic details associated with  $y$ , such as class descriptions, class instances, etc. In this study,  $T(y)$  returns class descriptions. The retrieved information  $T(y)$  is then used to create the grounded rationale  $r = f(T(y))$ , where  $f$  is a formatting function that translates  $T(y)$  into a natural language rationale. In this work, we adopt a simple approach where  $f(T(y))$  is the verbatim inclusion of the retrieved value  $T(y)$ . However, more sophisticated formatting functions  $f$ , such as those that rewrite or summarize taxonomic information, can be used to adapt  $T(y)$ . For each  $i$ -th example  $x_i$ , we create a grounded example  $(x_i, r_i, y_i)$ .

**Inference:** In the first stage, the input instance  $x \in \mathbb{X}$  is provided to the LLM along with a domain-specific prompt to infer a candidate class or label. The goal of the inference stage is to infer a candidate label  $y' \in \mathbb{Y}$  based on  $x$ . In addition, the prompt includes  $k$ -shot TGRE-based in-context examples  $\{(x_1, r_1, y_1), \dots, (x_k, r_k, y_k)\}$  to demonstrate the task and ensure the reasoning aligns closely with the taxonomy. Various methods, such as random selection, can be used to choose the examples from  $\mathbb{X}$ . The output of the LLM includes both a rationale  $r'$  explaining the inference and the inferred class label  $y'$ . For occupation classification,  $x$  is a job title-company pair,  $r$  is derived from the occupation description in O\*NET-SOC, and  $y'$  consists of SOC title and 8-digit code.




Figure 1: The proposed framework

**Retrieval:** From the output generated in the inference stage, we construct a query denoted as  $z'$  to retrieve  $m$  most similar candidates from an external taxonomy  $\mathbb{T}$ . Here,  $z'$  may represent  $y'$  (the inferred label),  $r'$  (the rationale), or  $(r', y')$ .  $x$  is not used in the query because it does not exist in the taxonomy  $\mathbb{T}$ . Let  $E(\cdot)$  denote an embedding function. To select the candidates from  $\mathbb{T}$ , we compute the cosine similarity between the embeddings of the query  $E(z')$  and taxonomic entities  $E(y_i)$  where  $\{y_1, y_2, \dots, y_{|\mathbb{T}|}\} \in \mathbb{T}$ . The taxonomic entities are ranked based on their cosine similarity scores, from highest to lowest, and the top- $m$  candidates  $\mathcal{Y}'_m$  are selected. Note that  $y'$  may already be in  $\mathcal{Y}'_m$ . For occupation classification,  $\mathcal{Y}'_m$  are SOC occupations, represented by titles and 8-digit codes, from O\*NET-SOC.

**Reranking:** In the final stage, the LLM refines the retrieval output by ranking up to  $n$  candidates ( $n < m$ ) that are most relevant to the input  $x$ , denoted as  $\mathcal{Y}'_n$  ( $\mathcal{Y}'_n \subseteq \mathcal{Y}'_m$ ). The LLM is prompted to rerank the retrieved candidates based on their contextual fit with  $x$ . The  $k$ -shot grounded in-context examples, similar to those used in the inference stage, are included in this prompt to demonstrate the task and ensure accurate alignment with the taxonomy. For occupation classification, this involves reranking the retrieved SOC titles and 8-digit codes to determine the most appropriate match for the job title-company pair.

**Key differences from IReRa:** While both IReRa (D'Oosterlinck et al. 2024) and our framework employ a three-stage pipeline (inference, retrieval, and reranking), we introduce the following key improvements. Unlike IReRa, which uses CoT reasoning, our framework employs TGRE to directly align rationales with taxonomic knowledge. Moreover, our retrieval stage uses rationale-augmented queries ( $z' = (r', y')$ ) incorporating both inferred labels and grounded rationales, whereas IReRa relies solely on label-based retrieval ( $z' = y'$ ). Finally, IReRa is a general framework for extreme multi-label classification, utilizing the DSPy programming model (Khattab et al. 2023). It features an iterative prompt optimization module aimed at enhancing accuracy but at the cost of significantly higher LLM calls and token usage. In contrast, our framework prioritizes efficiency by randomly selecting in-context examples, thereby minimizing computational overhead.

## 5 Occupation Classification Evaluation

To evaluate the effectiveness of the proposed framework, we conducted a multi-class occupation classification task. Specifically, the task involves assigning one of the pre-defined occupations in O\*NET-SOC to an input job title-company pair. This task is inherently challenging due to the ambiguity and diversity of job titles and the size and structure of the taxonomy.

### Dataset

We utilized US-based anonymized resume data gathered from various online professional networking platforms such as LinkedIn and Stack Overflow<sup>1</sup> by Lightcast, a labor market analytics company. From this resume data, we created a test set, called **Jobs12K**, by randomly sampling 11,920 job records. Each record consisted of a job title and its corresponding company name. Moreover, another 100 job records were randomly chosen as a validation set for selecting in-context examples. To ensure data quality, we excluded records that met the following criteria: (i) job titles containing non-alphanumeric or non-ASCII characters, and (ii) job titles exceeding seven words in length. After filtering, the average word counts for job records in the test set and validation set were 5.9 words (S.D. = 2.3) and 5.8 words (S.D. = 2.3), respectively. A detailed description of the Lightcast dataset, used for deriving Jobs12K, is included in Section B of the appendix.

### Experimental Setup

We tested various combinations of LLMs, reasoning methods, retrieval units, and ablated variants to evaluate the proposed framework.

**LLMs:** Each method was tested on five efficient LLMs: GPT-3.5 Turbo, Gemini-1.5 Flash, Claude Instant 1.2, Llama 3 8B Instruct, and Mistral Small. As previously mentioned, these LLMs were chosen for their cost effectiveness, which make them practical for large-scale classification tasks. GPT-3.5 Turbo was accessed in June 2024, while the remaining models were accessed in October 2024. The temperature was set to zero for all tasks and LLMs to reduce randomness and promote deterministic outputs.

**Reasoning methods:** We include two reasoning methods, **TGRE** and **CoT**, in both the inference and reranking

<sup>1</sup><https://kb.lightcast.io/en/articles/7153977-global-data-101>

stages. Since the Jobs12K dataset is unlabeled, we manually annotated a small sample ( $N < 10$ ) from its validation set to be used in constructing in-context examples. TGRE utilizes TGRE-based in-context examples, augmenting rationales with occupation descriptions from O\*NET-SOC. In contrast, CoT uses the chain-of-thought reasoning to generate rationales in the in-context examples without the taxonomy, serving as a competitive baseline. For ablated variants, indicated by single and double dashes, we exclude the reranking module (TGRE- and CoT-) or both the retrieval and reranking modules (TGRE-- and CoT--).

**Retrieval units:** Two retrieval units – **sentence** and **label** – are included. For the sentence-based retrieval unit, queries  $z'$  are rationales and inferred SOC titles and codes ( $r' \cup y'$ ) generated during the inference stage while documents  $y_i$  are SOC titles, 8-digit codes, and descriptions from O\*NET-SOC. For the label-based retrieval unit, queries  $z'$  are inferred titles and codes  $y'$  generated during the inference stage and documents  $y_i$  are SOC titles and codes. Sentence transformer models *multi-qa-mpnet-base-dot-v1* and *all-mpnet-base-v2* are used for generating embeddings for the sentence and label units, respectively. These models have been shown in our internal validation to produce optimal embeddings for the respective units. In the retrieval and reranking stages, we set both the number of top- $m$  retrieved candidates and top- $n$  selected candidates to 10.

**Other baselines:** Several method variants in the experiment represent approaches used in existing LLM-based classification methods. These include:

- **IReRa:** Our CoT with label-based retrieval variant is equivalent to IReRa (D’Oosterlinck et al. 2024) where chain-of-thought is utilized to infer class labels and rerank candidates. The top candidates is retrieved by computing the similarity between the inferred labels and the taxonomic class labels.
- **LLM4Jobs:** Our ablated sentence-based variants that omit the reranking module (e.g., TGRE- and CoT-) are generally comparable to LLM4Jobs (Li, Kang, and De Bie 2023). Unlike LLM4Jobs, which employs zero-shot prompting, our approach leverages in-context examples. The best ablated variant in the experiment is used to represent LLM4Jobs.

Moreover, we compare precision of the default SOC labels in Jobs12K provided by Lightcast (**LC**). These labels were inferred using proprietary algorithms.

**In-context examples:** For all variants, we include a one-shot in-context example. For TGRE, we randomly selected an input instance and its associated ground-truth SOC title and code from the validation set. Then, we retrieved the occupation description corresponding to the SOC title and code and included it into the in-context example as the grounded rationale. For CoT, we followed the protocol used in IReRa (D’Oosterlinck et al. 2024) to ensure a fair comparison by employing their optimization module to select the best  $k$ -shot in-context examples from the validation set (where  $k$  was empirically set to be 1). Then, we used the LLM to generate CoT rationales for the in-context examples.

All prompt templates and in-context examples are included in Section H of the appendix.

## Evaluation Metric

We used **Precision@ $K$**  ( $K = 1$ ), defined as the proportion of correctly classified labels among the top- $K$  predicted labels (SOC titles and 8-digit codes).  $K = 1$  is appropriate as occupation classification typically considers only the top prediction actionable. As  $K=1$  and each instance has a single ground-truth label, Precision@1 is equivalent to Recall@1, F1@1, and Accuracy@1.

Since ground-truth labels for the job records in Jobs12K were not available, we employed **GPT-4o**, which has been shown to be highly knowledgeable of the O\*NET-SOC taxonomy in Section 3, to identify the most relevant label. Specifically, we prompted GPT-4o to select the most appropriate SOC label from the pool of predicted labels generated by all tested methods for each job record. To mitigate positional bias, the lists of candidate labels were shuffled. GPT-4o was also instructed to explicitly indicate if none of the candidates were applicable. This evaluation was repeated for 3 runs, and majority voting was applied to select the final labels. In the case of ties, the final selection was made randomly. The prompt template for occupation selection is included in Figure 15 in Section H.

Consistent with previous findings (Liu et al. 2023; Sotana et al. 2023), our internal validation demonstrated that GPT-4o is a reliable evaluator for this task. Specifically, we conducted a comparison between human judgments and GPT-4o’s assessments on a small random sample of 100 job records. The human annotator, who is fairly knowledgeable of the O\*NET-SOC taxonomy, provided independent annotations, and the results showed 82% correlation between human and GPT-4o judgments. The strong agreement provides preliminary evidence that GPT-4o can serve as a reliable proxy for human evaluation in this task.

To compare results, one-way ANOVA with Tukey HSD post-hoc analyses were performed using Python’s SciPy (1.11.3) and Statsmodels (0.14) libraries. The data used in the analyses are normally distributed.

## Results

The results of occupation classification across methods and models are shown in Table 4. The best performance is achieved by **TGRE with sentence-based retrieval** on **GPT-3.5 Turbo**, attaining a Precision@1 of 0.8114, outperforming other baselines, including IReRa (0.7049), LLM4Jobs (0.6644 for CoT), and LC (0.634).

A one-way ANOVA indicated no significant differences between the LLMs ( $p > 0.05$ ). These results demonstrate the robustness of the framework, which consistently produces competitive performance across different LLMs despite their taxonomic knowledge gaps (as shown in Table 3). Even less knowledgeable models like Llama 3 8B Instruct and Mistral Small, achieved comparable Precision@1 scores using the framework.

Overall, methods using TGRE significantly outperform CoT ( $p < 0.05$ ), demonstrating the effectiveness of the grounded rationale approach. The post-hoc test shows that

Baseline Equivalence	Reasoning Method	Retrieval Unit	GPT	Gemini	Claude	Llama	Mistral
	TGRE	Sentence	<b>0.8114</b>	<b>0.7868</b>	<b>0.7616</b>	<b>0.7612</b>	<b>0.7687</b>
	CoT	Sentence	0.7685	0.6827	0.6944	0.6593	0.6506
IReRa	TGRE	Label	0.7694	0.6909	0.6841	0.5366	0.7202
	CoT	Label	0.7049	0.6618	0.6592	0.5491	0.5068
LLM4Jobs	TGRE-	Sentence	0.6644	0.6018	0.6084	0.5993	0.6466
LLM4Jobs	CoT-	Sentence	0.6384	0.508	0.4964	0.5261	0.3918
	TGRE-	Label	0.7015	0.5238	0.581	0.4969	0.5985
	CoT-	Label	0.5658	0.5005	0.403	0.3006	0.1356
	TGRE--	None	0.4875	0.0715	0.277	0.0597	0.0407
	CoT--	None	0.5002	0.0684	0.2886	0.0925	0.0857

Table 4: Precision@1 for occupation classification. The best results are highlighted in bold. The columns abbreviate the full names of the LLM models: GPT refers to GPT-3.5 Turbo, Gemini to Gemini-1.5 Flash, Claude to Claude Instant 1.2, Llama to Llama 3 8B Instruct, and Mistral to Mistral Small.

TGRE with sentence-based retrieval significantly outperforms CoT with label-based retrieval ( $p < 0.05$ ), with no significant differences between other method and retrieval unit combinations. Furthermore, the sentence-based retrieval unit outperforms the label-based unit in precision ( $p < 0.05$ ), highlighting its effectiveness in leveraging richer contextual information.

Next, there are significant differences in the performance of TGRE compared to its ablated variants, TGRE- and TGRE-- ( $p < 0.001$ ), and between CoT and its ablated variants, CoT- and CoT-- ( $p < 0.001$ ). These further underscore the importance of retrieval and reranking. The absence of reranking leads to a notable drop in performance. Removing both retrieval and reranking results in the steepest decline. Results for the inference-only variants (TGRE- and CoT-) are consistent with the taxonomic knowledge of each LLM, where more knowledgeable models greatly outperform less knowledgeable models.

In summary, the findings demonstrate the effectiveness of the proposed framework, which combines LLM reasoning capabilities with taxonomy-guided examples, retrieval, and reranking to achieve robust performance across efficient LLMs in occupation classification tasks. Beyond performance improvements, the computational cost analysis in Appendix G indicates that the framework is highly cost-effective. While the framework incurs marginal processing time due to retrieval and reranking, our analysis shows that it substantially reduces API costs by 72% compared to direct GPT-4o inference, making it a more practical and scalable approach for large-scale classification tasks.

## 6 Skill Classification Evaluation

To demonstrate the versatility of the proposed framework, we extended it to a multi-label skill classification task. This task involves assigning one or more skills and competences from the ESCO taxonomy (ESCO) to a given job vacancy post. Similar to O\*NET-SOC, ESCO is the European standard classification system for Skills, Competences and Occupations. In this study, we employ the **ESCO v1.2** taxon-

omy, which contains 13,939 skills and competences (class labels).

By adapting the framework to this task, the input space  $\mathbb{X}$  consists of job vacancy posts, the label spaces  $\mathbb{Y}$  comprises skills and competences from the ESCO taxonomy  $\mathbb{T}$ . During the **inference** stage, the framework processes each job vacancy post ( $x$ ) to infer one or more skills or competences ( $y' \subseteq \mathbb{Y}$ ) relevant to the post. For TGRE, in-context examples are augmented with skill descriptions from ESCO. For CoT, the LLM generates chain-of-thought rationales for the in-context examples. In the **retrieval** stage, the inferred set of labels ( $y'$ ) and/or rationales ( $r'$ ) are used as queries to retrieve a set of candidate skills and competences from ESCO. Finally, the **reranking** stage refines the retrieved candidates using the LLM to evaluate their relevance to  $x$ .

### Dataset

We used a labeled dataset (Decorte et al. 2022; Zhang et al. 2022) of technology-sector job vacancy posts, where skill candidates for each post were manually annotated with corresponding ESCO skills and competences. To ensure reproducibility, we adopted the same test and validation splits as those used by D’Oosterlinck et al. (D’Oosterlinck et al. 2024). Duplicate skills and competences were removed from each job post in both sets to improve data quality. The test set includes 338 annotated job posts and 272 unique ground-truth labels, while the validation set, comprising 75 annotated posts and 67 unique ground-truth labels, was employed to supply in-context examples. On average, job posts in the test and validation sets contain 12.9 words (S.D. = 8.7) and 12 words (S.D. = 10.1), respectively. Each job post in the test set is associated with an average of 1.9 ground-truth labels (S.D. = 1.4), while those in the validation set have an average of 2 labels (S.D. = 1.3).

### Experimental Setup

We adopted the same experimental setup as described in Section 5, including LLM models and temperature setting, reasoning methods, retrieval units, and ablation variants.

The LLMs were also accessed during the same periods as described in the occupation classification evaluation. All prompt templates and in-context examples are provided in Section H of the appendix.

To handle the extensive label space in the ESCO taxonomy, we set the number of top candidates retrieved during the retrieval stage to  $m = 50$ , following D’Oosterlinck et al. (D’Oosterlinck et al. 2024). In the reranking stage, the number of top- $n$  candidates was set to  $n = 10$ . The number of in-context examples for both TGRE and CoT reasoning methods during inference and reranking was set to  $k = 2$ . For CoT, these in-context examples were selected by the optimization module in IReRa (D’Oosterlinck et al. 2024), whereas they were randomly chosen in TGRE-based variants.

To enable direct comparison with D’Oosterlinck et al. (D’Oosterlinck et al. 2024), we introduced additional GPT-4-based variants (**TGRE\*** and **CoT\***) for each top TGRE and CoT performer, using GPT-4o as the LLM in the reranking stage. These variants also serve as an upper bound for performance comparison. In addition, we report the performance of the best fine-tuned systems on the same dataset by Decorte et al. (Decorte et al. 2023) and Clavié et al. (Clavié and Soulié 2023), which were taken directly from their papers.

## Evaluation Metric

We employed **R-precision@K (RP@K)**, where  $K = 3, 5$ , and 10, to evaluate multi-label classification performance. Unlike **Precision@K**, **RP@K** measures the proportion of correctly classified labels among the top- $K$  predictions, normalized by the minimum of either the total number of predicted labels or the total number of ground-truth labels for a given job post. This metric, commonly used in skill classification tasks (Decorte et al. 2023; Clavié and Soulié 2023; D’Oosterlinck et al. 2024), accounts for label variability per instance. We performed the same statistical analyses as in occupation classification to assess significant differences in results.

## Results

Table 5 presents the results for **RP@10** across methods and models. Due to space limitations, **RP@3** and **RP@5** results are provided in Tables 8 and 9 in Section E of the appendix. The best performance is achieved by **TGRE\* with sentence-based retrieval**, attaining an **RP@10** of 0.7052 on **Gemini-1.5 Flash**, outperforming IReRa (0.6546)<sup>2</sup> and LLM4Jobs (0.6151). **CoT\*** also greatly outperforms the best CoT variant across all models. Consistent with previous findings (Clavié and Soulié 2023; D’Oosterlinck et al. 2024), these results reinforce the effectiveness of GPT-4 as a powerful reranker for skill classification.

Compared to fine-tuned systems, the best **TGRE\*** with sentence-based retrieval achieved an **RP@5** score of 0.6318,

<sup>2</sup>Note that the state-of-the-art **RP@10** of 0.7058 was achieved by the best IReRa variant as reported in their original experiment (D’Oosterlinck et al. 2024) though we were unable to independently replicate this result.

exceeding the **RP@5** scores reported by Decorte et al. (Decorte et al. 2023) (0.5462) and Clavié et al. (Clavié and Soulié 2023) (0.615). Meanwhile, the best TGRE with sentence-based retrieval achieved an **RP@5** score of 0.5681, outperforming Decorte et al.’s system, as shown in Tables 7 and 9 in the appendix.

In line with the findings from the occupation classification task, no significant differences were found between the performance of the various LLMs ( $p > 0.05$ ), indicating that all models are equally effective for this task. Next, TGRE significantly outperforms CoT across various LLMs and retrieval units ( $p < 0.001$ ), underscoring the benefit of incorporating taxonomy-guided reasoning. The post-hoc test indicates that most method and LLMs combinations do not significantly differ, with a few exceptions: TGRE on Claude Instant significantly outperforms CoT on Llama 3 ( $p < 0.05$ ), TGRE on Gemini-1.5 significantly outperforms CoT on Llama 3 ( $p < 0.01$ ), and TGRE on Gemini-1.5 significantly outperforms CoT on Mistral ( $p < 0.05$ ).

Additionally, the post-hoc analysis shows significant differences between most method and retrieval unit combinations ( $p < 0.01$ ), except for the comparison between TGRE and CoT when using label-based retrieval ( $p > 0.05$ ). These results highlight the robustness of TGRE across different domains, as demonstrated by its consistent superiority in both occupation and skill classification tasks.

Lastly, sentence-based retrieval units significantly outperform label-based units across most configurations ( $p < 0.05$ ), highlighting the importance of leveraging rich taxonomic context. For instance, TGRE with sentence-based retrieval on GPT-3.5 Turbo achieves an **RP@10** of 0.7045, while label-based retrieval achieves only 0.5744. Consistent with the occupation classification results, significant differences in performance were found between TGRE and its ablated variants ( $p < 0.001$ ), and between CoT and its ablated variants ( $p < 0.001$ ). These findings underscore the importance of both retrieval and reranking across domains.

Overall, these results validate the effectiveness and versatility of the proposed framework. By combining taxonomy-guided reasoning, retrieval, and reranking, the framework achieves strong and robust performance in both multi-class occupation classification and multi-label skill classification, demonstrating its adaptability across domains. Although we did not directly assess LLM knowledge of ESCO, the lower performance scores and task complexity in the skill classification may reflect a more limited knowledge of ESCO compared to O\*NET-SOC, underscoring the importance of the framework in bridging the internal knowledge gaps.

## 7 Limitations and Future Work

**Knowledge assessment:** This study focused exclusively on the ONET-SOC taxonomy to assess LLM knowledge, a choice justified by its widespread use. Expanding future studies to include other taxonomies, such as ESCO or ISCO, could improve generalizability.

**Framework:** While simplicity and adaptability are key strengths of the framework, there remain areas of improvement. First, some manual prompt engineering is still required for adapting the framework to other domains. Next,

Baseline Equivalence	Reasoning Method	Retrieval Unit	GPT	Gemini	Claude	Llama	Mistral
	TGRE*	Sentence	<b>0.7045</b>	<b>0.7052</b>	<b>0.6959</b>	<b>0.7023</b>	<b>0.6929</b>
	TGRE	Sentence	0.6331	0.6557	0.6202	0.5937	0.62
	CoT	Sentence	0.2466	0.2909	0.3968	0.2862	0.3105
	TGRE	Label	0.5744	0.6251	0.5811	0.4825	0.5624
IReRa	CoT*	Label	0.645	0.6546	0.6338	0.5704	0.6149
IReRa	CoT	Label	0.5993	0.5881	0.5381	0.4216	0.5113
LLM4Jobs	TGRE-	Sentence	0.5968	0.6031	0.6086	0.6094	0.5904
LLM4Jobs	CoT-	Sentence	0.181	0.1457	0.2637	0.1767	0.2067
	TGRE-	Label	0.5638	0.6151	0.5703	0.4705	0.5404
	CoT-	Label	0.5717	0.5224	0.4452	0.354	0.4506
	TGRE--	None	0.1527	0.1045	0.1164	0.0993	0.1008
	CoT--	None	0.0325	0.0592	0.028	0.0163	0.0466

Table 5: RP@10 for skill classification. The best results are highlighted in bold. The columns abbreviate the full names of the LLMs: GPT refers to GPT-3.5 Turbo, Gemini to Gemini-1.5 Flash, Claude to Claude Instant 1.2, Llama to Llama 3 8B Instruct, and Mistral to Mistral Small. For TGRE\* and CoT\*, we use GPT-4o for reranking.

different formatting function  $f$  could be designed for incorporating taxonomic knowledge into grounded rationales. Additionally, future work could explore ways to handle flat taxonomies or dictionaries, which lack hierarchical structures or rich class information, e.g., generating auxiliary data with LLMs to enrich such taxonomies or leveraging examples from similar taxonomies. Furthermore, we employed fixed  $k$ -shot in-context examples without dynamically selecting the optimal  $k$  examples or utilizing advanced selection strategies, such as  $k$ -nearest neighbor based methods (Liu et al. 2021). These strategies present promising directions for improvement. Similarly, future work could explore different choices for top- $m$  retrieval and top- $n$  reranking more thoroughly.

**Evaluation:** GPT-4 based evaluator used in occupation classification is not entirely free from bias. The manual validation of GPT-4o relied on a small human-annotated sample and scaling up these annotations could solidify findings. Similarly, the annotated SkillSpan dataset used for skill classification is relatively small. Using larger and more diverse datasets would improve the robustness of evaluation. For occupation classification, non-LLM or fine-tuned baselines were not included due to the lack of labeled training data. Future work could explore synthetic data generation to enable the fine-tuning of robust classification models.

## 8 Conclusion and Implications

In this paper, we investigate the capabilities of large language models (LLMs) in understanding the ONET-SOC taxonomy. The study begins with a knowledge assessment experiment to evaluate the performance of both frontier and efficient LLMs in generating precise occupational titles and codes. The results indicate that frontier models, especially GPT-4o, consistently outperform the efficient ones, demonstrating their in-depth taxonomic understanding. However, LLM performance varies across different levels of code granularity within the taxonomy, as evidenced by their lower

performance in generating specific occupations compared to major occupations.

To address these challenges, we propose a multi-stage framework for occupation classification that incorporates inference, retrieval, and reranking stages. Central to our approach is the taxonomy-guided reasoning examples that integrates grounded rationales from the taxonomy to improve classification accuracy. We apply the framework to both occupation and skill classification tasks, demonstrating its versatility and effectiveness. Our experiments show that the framework produces competitive performance across different LLMs, regardless of their underlying taxonomic knowledge.

This study offers several significant implications. Firstly, our analysis highlights the limitations of chain-of-thought (CoT) reasoning, particularly in tasks requiring domain-specific knowledge. While CoT is effective at general problem solving, our results show that it under-performs when deeper domain-specific understanding is necessary. Secondly, our findings highlight the limitations of LLM knowledge for occupation classification when directly inferring SOC labels. This underscores the crucial role of leveraging external taxonomic information to enhance the classification accuracy. By incorporating detailed class descriptions from taxonomies, LLMs generated better aligned queries, which in turn improved retrieval and reranking. This demonstrates how taxonomic knowledge can effectively address the gaps in LLMs’ understanding, particularly in specialized classification tasks. The effectiveness of this approach highlights the potential for extending the use of taxonomic or domain-specific knowledge to related classification problems. Lastly, the study contributes to computational social science (CSS) by addressing key challenges in automatic content analysis. The proposed framework offers an accessible and adaptable solution for CSS domains, which often face data scarcity, difficulties in data collection, and a reliance on domain-specific knowledge and taxonomies.

## Ethical Statement

In this section, we discuss several ethical considerations relevant to our research to ensure the responsible use and development of automatic occupation classification systems.

**Model and taxonomic biases:** The LLMs utilized in this study inherently exhibit biases from their pretraining data, which can lead to inaccurate or unfair classifications. In addition, occupational taxonomies, such as O\*NET-SOC, may reflect inherent biases in how occupations are defined and categorized, particularly since the O\*NET-SOC taxonomy is US-centric. These biases can reinforce stereotypes or perpetuate inequalities, especially when certain occupations are associated with race, gender, or socioeconomic factors. Future work should focus on systematically addressing these biases to ensure equitable outcomes.

**Unintended consequences of automated systems:** Automated classification systems, including those developed in this study, may result in unintended consequences if misused or applied without human supervision. Even minor errors can lead to significant downstream effects. For example, a hiring platform relying on automated occupation classification systems might overestimate their accuracy. A misclassification could lead to a candidate’s resume being assigned an incorrect occupation, potentially impacting their employment opportunities. To mitigate such risks, these systems should be used to complement rather than fully replace human judgment.

**Access equity:** Our framework is designed to be lightweight and model-agnostic, allowing it to be implemented across a range of LLMs, including efficient models. This reduces barriers to adoption for researchers and organizations with limited computational resources. However, a gap remains between high-resource and low-resource organizations, particularly in accessing high-performance models. Bridging this gap will require broader access to high-performing LLMs and open-source alternatives.

**Environmental impact:** Although our framework is training-free and focuses on leveraging existing LLMs, model inference still requires GPU usage and contributes to the carbon footprint. By emphasizing the use of efficient LLMs, our research mitigates but does not fully eliminate this impact. To further minimize environmental impact, further efforts in optimizing LLM efficiency are needed.

## Acknowledgements

This research / project is supported by the Ministry of Education, Singapore, under its MOE Academic Research Fund Tier 2 programme (Proposal ID: T2EP20223-0047, Project ID: MOE-000615-00). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of the Ministry of Education, Singapore. The authors would also like to thank Yao Lu from the Data Science Institute (DSI) at Columbia University and the anonymous reviewers for their constructive feedback, which significantly improved the quality of this work.

## References

- Boselli, R.; Cesarini, M.; Mercorio, F.; and Mezzanzanica, M. 2017. Using machine learning for labour market intelligence. In *ECML PKDD 2017*, 330–342. Springer.
- Clavié, B.; and Soulié, G. 2023. Large language models as batteries-included zero-shot ESCO skills matchers. *arXiv preprint arXiv:2307.03539*.
- Decorte, J.-J.; Van Haute, J.; Deleu, J.; Develder, C.; and Demeester, T. 2022. Design of negative sampling strategies for distantly supervised skill extraction. *arXiv preprint arXiv:2209.05987*.
- Decorte, J.-J.; Verlinden, S.; Van Haute, J.; Deleu, J.; Develder, C.; and Demeester, T. 2023. Extreme multi-label skill extraction training using large language models. *arXiv preprint arXiv:2307.10778*.
- D’Oosterlinck, K.; Khattab, O.; Remy, F.; Demeester, T.; Develder, C.; and Potts, C. 2024. In-context learning for extreme multi-label classification. *arXiv preprint arXiv:2401.12178*.
- Guo, X.; Handzlik, D.; Jones, J. J.; and Skiena, S. S. 2024. The evolution of occupational identity in twitter biographies. In *Proceedings of ICWSM 2024*, volume 18, 502–514.
- Hu, T.; Xiao, H.; Luo, J.; and Nguyen, T.-v. T. 2016. What the language you tweet says about your occupation. In *Proceedings of ICWSM 2016*, volume 10, 181–190.
- Javed, F.; Luo, Q.; McNair, M.; Jacob, F.; Zhao, M.; and Kang, T. S. 2015. Carotene: A Job Title Classification System for the Online Recruitment Domain. In *Proceedings of IEEE BigDataService 2015*, 286–293.
- Kalhor, G.; Gardner, H.; Weber, I.; and Kashyap, R. 2024. Gender Gaps in Online Social Connectivity, Promotion and Relocation Reports on LinkedIn. In *Proceedings of ICWSM 2024*, volume 18, 800–812.
- Khattab, O.; Singhvi, A.; Maheshwari, P.; Zhang, Z.; Santhanam, K.; Vardhamanan, S.; Haq, S.; Sharma, A.; Joshi, T. T.; Moazam, H.; et al. 2023. Dspy: Compiling declarative language model calls into self-improving pipelines. *arXiv preprint arXiv:2310.03714*.
- Kumar, R.; Kim, Y.; Ravi, S.; Sun, H.; Faloutsos, C.; Salakhutdinov, R.; and Yoon, M. 2024. Automatic question-answer generation for long-tail knowledge. *arXiv preprint arXiv:2403.01382*.
- Li, N.; Kang, B.; and De Bie, T. 2023. LLM4Jobs: Unsupervised occupation extraction and standardization leveraging Large Language Models. *arXiv preprint arXiv:2309.09708*.
- Liu, J.; Shen, D.; Zhang, Y.; Dolan, B.; Carin, L.; and Chen, W. 2021. What Makes Good In-Context Examples for GPT-3? *arXiv preprint arXiv:2101.06804*.
- Liu, Y.; Iter, D.; Xu, Y.; Wang, S.; Xu, R.; and Zhu, C. 2023. G-Eval: NLG Evaluation using Gpt-4 with Better Human Alignment. In *Proceedings of EMNLP 2023*, 2511–2522.
- Luo, L.; Vu, T.; Phung, D.; and Haf, R. 2023. Systematic Assessment of Factual Knowledge in Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 13272–13286.

Russ, D. E.; Ho, K.-Y.; Colt, J. S.; Armenti, K. R.; Baris, D.; Chow, W.-H.; Davis, F.; Johnson, A.; Purdue, M. P.; Karagas, M. R.; et al. 2016. Computer-based coding of free-text job descriptions to efficiently identify occupations in epidemiological studies. *Occupational and environmental medicine*, 73(6): 417–424.

Soroush, A.; Glicksberg, B. S.; Zimlichman, E.; Barash, Y.; Freeman, R.; Charney, A. W.; Nadkarni, G. N.; and Klang, E. 2024. Large language models are poor medical coders—benchmarking of medical code querying. *NEJM AI*, 1(5): AIdbp2300040.

Sottana, A.; Liang, B.; Zou, K.; and Yuan, Z. 2023. Evaluation Metrics in the Era of GPT-4: Reliably Evaluating Large Language Models on Sequence to Sequence Tasks. In *Proceedings of EMNLP 2023*, 8776–8788.

Sun, K.; Xu, Y. E.; Zha, H.; Liu, Y.; and Dong, X. L. 2023. Head-to-tail: How knowledgeable are large language models (llm)? AKA will llms replace knowledge graphs? *arXiv preprint arXiv:2308.10168*.

Sun, Y.; Xin, H.; Sun, K.; Xu, Y. E.; Yang, X.; Dong, X. L.; Tang, N.; and Chen, L. 2024. Are Large Language Models a Good Replacement of Taxonomies? *Proc. VLDB Endow.*, 17(11): 2919–2932.

Tänzer, M.; Ruder, S.; and Rei, M. 2022. Memorisation versus Generalisation in Pre-trained Language Models. In *Proceedings of ACL 2022*, 7564–7578.

Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *NeurIPS*, 35: 24824–24837.

Zhang, M.; Jensen, K.; Sonniks, S.; and Plank, B. 2022. SkillSpan: Hard and Soft Skill Extraction from English Job Postings. In *Proceedings of NAACL 2022*, 4962–4984.

Zhang, M.; Jensen, K. N.; and Plank, B. 2022. Kompetenzer: Fine-grained skill classification in danish job postings via distant supervision and transfer learning. In *Proceedings of LREC 2022*, 436–447.

## Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
- (b) Do your main claims in the abstract and introduction accurately reflect the paper’s contributions and scope? **Yes**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
- (e) Did you describe the limitations of your work? **Yes**
- (f) Did you discuss any potential negative societal impacts of your work? **Yes, in the Ethical Considerations section**

- (g) Did you discuss any potential misuse of your work? **Yes, in the Ethical Considerations section.**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes, in the Ethical Considerations section.**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **Yes, the data used in ANOVA analyses meet the normality assumption.**
- (b) Have you provided justifications for all theoretical results? **Yes, ANOVA is a well-established method for hypothesis testing.**
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA**
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **Yes, we discuss the regression analysis in the knowledge assessment experiment.**
- (e) Did you address potential biases or limitations in your theoretical framework? **NA**
- (f) Have you related your theoretical results to the existing literature in social science? **Yes, the findings are related to prior research on LLM knowledge assessment and occupational classification**
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **Yes, the discussions are provided in the limitations and conclusion sections.**

3. Additionally, if you are including theoretical proofs...

- (a) Did you state the full set of assumptions of all theoretical results? **NA**
- (b) Did you include complete proofs of all theoretical results? **NA**

4. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **Yes, all prompt templates and in-context examples used in the study are included in the Appendix.**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes, in the experimental setup section of both occupation and skill classification evaluations.**
- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **NA**
- (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **Yes, in the experimental setup section.**

- (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes**
- (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? **Yes, in the Ethical Considerations.**
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity...**
- (a) If your work uses existing assets, did you cite the creators? **Yes, we use and cite both the original and derived datasets in the skill classification evaluation.**
- (b) Did you mention the license of the assets? **Yes, these are public datasets.**
- (c) Did you include any new assets in the supplemental material or as a URL? **Yes, we include our prompt templates in the appendix.**
- (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **NA**
- (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **Yes, we state that our online resume data are anonymized.**
- (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR? **NA**
- (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset? **NA**
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity...**
- (a) Did you include the full text of instructions given to participants and screenshots? **NA**
- (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **NA**
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **NA**
- (d) Did you discuss how data is stored, shared, and de-identified? **NA**

## A Code and Data Availability

The code and datasets are publicly available at <https://github.com/aeckpalakorn/TGRE-Classification>.

## B The Lightcast Dataset

The Lightcast dataset, obtained in October 2022, comprises approximately 141.6 million anonymized US-based resume profiles collected from various online professional networking platforms by Lightcast, a labor market analytics company. Each profile is enriched with occupation (SOC title and code), industry, and education details through machine learning-based standardization and processing<sup>3</sup>, and

<sup>3</sup><https://kb.lightcast.io/en/articles/7153977-global-data-101>

comprises job history, education history, and skills sections. To manage the dataset size, we partition profiles using a hashing-based method on their unique identifiers, resulting in 284 partitions with approximately 458,000 profiles each. These partitions, labeled “LC” followed by their index (e.g., LC0 to LC283), ensure a random and uniform distribution of the profile data.

To ensure data quality and authenticity, we performed data cleaning by applying these criteria: (i) each job record in the profile must include the job title, company name, city, state, country, and start date. Past job records also required an end date. Job records with end dates preceding start dates were excluded; (ii) job titles in the profile containing non-occupational terms (“student,” “intern,” or “owner”) as well as those indicating multiple roles within one title were excluded; and (iii) each education record in the profile must specify the degree obtained, start date, end date, and school name.

To construct the Jobs12K test set, we applied the following steps to sample 11,920 job records from the job histories of all profiles in the LC0 partition: (i) exclude job titles with non-alphanumeric, non-ASCII characters, or more than seven words; and (iii) For the 596 most common 6-digit SOC codes (those linked to at least 20 records), randomly sample 20 job records per code, resulting in 11,920 records.

Popular occupations used in the selection may not represent the true occupational distribution, as default SOC codes provided by Lightcast may lack precision at the 6-digit level. Instead, these codes indicate a preliminary signal of commonly represented occupations.

## C The Formatting Functions

In our proposed TGRE framework, the formatting function  $f$  is a component that converts taxonomic information retrieved for a specific class label  $y$  into a grounded rationale  $r$  that aligns the LLM reasoning process with the taxonomy. The rationale  $r$  serves as an intermediate explanation that clarifies the connection between the input  $x$  and the corresponding label  $y$ . For example, it explains how a job title-company pair is related to an SOC title and code.

Formally,  $f$  takes the retrieved taxonomic information  $T(y)$ , which consists of information of  $y$ , (e.g., occupation description), and transforms it into a rationale  $r = f(T(y))$ . The rationale is then embedded into in-context examples to guide the reasoning process of LLMs.

In this study, we use a simple implementation of  $f$  where the retrieved taxonomic information  $T(y)$  is incorporated verbatim in the rationale. This approach minimizes the framework’s complexity while achieving strong empirical results. For example:

Given an input  $x$  “records and evidence technician, cobb county police department” and its SOC label  $y$  “Police Identification and Records Officers (33-3021.02)” in the validation set, the proposed framework retrieves the matched taxonomic entity  $T(y)$  from the O\*NET-SOC taxonomy, comprising the followings:

- Occupation title: Police Identification and Records Officers

- Occupation code: 33-3021.02
- Description: Collect evidence at crime scenes, classify and identify fingerprints, and photograph evidence for use in criminal and civil cases.

The verbatim formatting function retrieves the description from  $T(y)$  and directly into the grounded rationale  $r$  as follows:

- Grounded rationale: Given the job title '**records and evidence technician**' and company '**cobb county police department**,' the individual is likely expected to **collect evidence at crime scenes, classify and identify fingerprints, and photograph evidence for use in criminal and civil cases**.

Here, the elements in bold are values derived from the task inputs and taxonomy that replace placeholders within the rationale template.

It is important to note that while we employ a simple verbatim approach in this study, LLMs can be leveraged to make more sophisticated formatting functions. For example, an LLM could rewrite the retrieved description for improved brevity and clarity. Moreover, it could incorporate other taxonomic information beyond descriptions, such as tasks and skills, to generate an enriched summary.

## D Optimization Module in IReRa

The optimization module in IReRa, as depicted in Figure 1 in D'Oosterlinck et al. (2024), employs bootstrapping to iteratively identify the most effective few-shot in-context examples from the validation set. This procedure involves systematically evaluating the performance of various combinations of examples and selecting those that achieve the best results based on a predefined metric. In our experiments, this optimization process identified a specific 1-shot in-context example ( $k = 1$ ) for the CoT-based variant in the occupation classification experiment. For the skill classification experiment, the optimization module selected a set of 2-shot in-context examples ( $k = 2$ ) for the CoT-based variant. While our TGRE variant also used  $k = 2$ , the examples were randomly chosen, distinguishing them from the optimized examples used by the CoT-based variant.

## E Supplemental Tables

This section provides supplemental tables supporting the results discussed in the main paper. Table 6 reports the Precision@1 scores for Lightcast's default SOC labels in the occupation classification task, evaluated by GPT-4o. Table 7 summarizes RP@5 and RP@10 results for two fine-tuned systems (Decorte et al. 2023; Clavié and Soulié 2023) in the skill classification evaluation. Additionally, Tables 8 and 9 present the RP@3 and RP@5 results for skill classification.

## F Supplemental Figures

This section presents supplemental figures supporting the results discussed in the main body of the paper.

Figures 2 and 3 illustrate the accuracy scores for the complete recall, partial recall, and recognition tasks of individual LLMs as well as frontier and efficient LLMs. Figures 4

Method	Precision@1
LC	0.634

Table 6: Precision@1 for Lightcast's default SOC labels in occupation classification evaluation.

Method	RP@5	RP@10
Decorte et al.	0.5462	N/A
Clavié et al.	0.615	0.6894

Table 7: RP@5 and RP@10 for two baseline fine-tuned systems in skill classification evaluation.

and 5 show the accuracy scores for tasks involving 8-digit and 2-digit codes for both individual LLMs and frontier/efficient LLMs. As highlighted in Section 3, GPT-4o emerges as the best overall performer, followed by GPT-3.5 Turbo and Gemini-1.5 Pro. Furthermore, frontier models consistently outperform efficient models across all tasks.

Figures 6 and 7 depict Precision@1 results for various reasoning methods and retrieval units in occupation classification, respectively. As discussed in Section 5, TGRE-based methods demonstrate superior performance compared to CoT-based methods, while sentence-based retrieval units outperform label-based retrieval units.

Likewise, Figures 8 and 9 present RP@10 results for various reasoning methods and retrieval units in skill classification. As highlighted in Section 6, TGRE-based methods achieve better performance than CoT-based methods. However, label-based retrieval units exhibit a higher average RP@10 compared to sentence-based retrieval units in most LLMs, influenced by some outliers (CoT with sentence-based retrieval variants). Nonetheless, a one-way ANOVA with post-hoc analysis indicates that sentence-based retrieval units significantly outperform label-based retrieval units ( $p < 0.05$ ).

## G Computational Cost Analysis

### Occupation Classification

We conducted a study using 100 randomly sampled records (job title-company name pairs) from the test set to assess the computational costs of TGRE compared to GPT-4o's direct prompting. GPT-3.5 Turbo was used as an underlying LLM for TGRE, whereas GPT-4o directly inferred SOC labels using CoT reasoning with the same CoT-based inference template (Figure 12). Both models were accessed via OpenAI APIs on March 18, 2025. Each test record was processed through a single API call, with no concurrent calls. The study was performed on a commodity machine equipped with an Intel Core i9-13900H 2600 MHz CPU (16 GB RAM) and an NVIDIA GeForce RTX 4060 GPU (16 GB VRAM).

We measured the following for each record to provide a comprehensive view of computational cost-effectiveness.

- **Average Tokens:** Computed as the total token usage (in-

Baseline Equivalence	Reasoning Method	Retrieval Unit	GPT	Gemini	Claude	Llama	Mistral
	TGRE*	Sentence	<b>0.5237</b>	<b>0.5685</b>	<b>0.5419</b>	<b>0.5311</b>	<b>0.5444</b>
	TGRE	Sentence	0.4423	0.4689	0.4428	0.4053	0.4561
	CoT	Sentence	0.177	0.2318	0.3018	0.2002	0.2229
IReRa	TGRE	Label	0.4034	0.4532	0.4181	0.3491	0.4186
IReRa	CoT*	Label	0.4132	0.4758	0.4339	0.4231	0.4502
IReRa	CoT	Label	0.4778	0.4053	0.3841	0.2816	0.3851
LLM4Jobs	TGRE-	Sentence	0.4201	0.4389	0.4398	0.4438	0.4142
LLM4Jibs	CoT-	Sentence	0.0868	0.0863	0.1677	0.0932	0.1179
	TGRE-	Label	0.4334	0.502	0.4679	0.3496	0.4127
	CoT-	Label	0.4157	0.3683	0.2411	0.2396	0.2791
	TGRE--	None	0.1543	0.1045	0.1134	0.0996	0.1011
	CoT--	None	0.0197	0.0316	0.0079	0.0015	0.0266

Table 8: RP@3 for skill classification. The best results are highlighted in bold. The columns abbreviate the full names of the LLMs: GPT refers to GPT-3.5 Turbo, Gemini to Gemini-1.5 Flash, Claude to Claude Instant 1.2, Llama to Llama 3 8B Instruct, and Mistral to Mistral Small. For TGRE\* and CoT\*, we use GPT-4o for reranking.

Baseline Equivalence	Reasoning Method	Retrieval Unit	GPT	Gemini	Claude	Llama	Mistral
	TGRE*	Sentence	<b>0.6075</b>	<b>0.6318</b>	<b>0.643</b>	<b>0.6107</b>	<b>0.6004</b>
	TGRE	Sentence	0.5252	0.5681	0.5258	0.4681	0.5226
	CoT	Sentence	0.1915	0.2677	0.352	0.2359	0.2587
IReRa	TGRE	Label	0.463	0.5338	0.4974	0.3989	0.4747
IReRa	CoT*	Label	0.4851	0.5675	0.524	0.4857	0.5323
IReRa	CoT	Label	0.5244	0.4832	0.4529	0.334	0.4407
LLM4Jobs	TGRE-	Sentence	0.5078	0.5163	0.5065	0.5033	0.4901
LLM4Jobs	CoT-	Sentence	0.1223	0.1072	0.2071	0.1168	0.1473
	TGRE-	Label	0.4857	0.5469	0.5131	0.4123	0.4704
	CoT-	Label	0.4929	0.4381	0.3213	0.2843	0.3484
	TGRE--	None	0.1523	0.1045	0.113	0.0994	0.1009
	CoT--	None	0.0301	0.0459	0.0136	0.0054	0.039

Table 9: RP@5 for skill classification. The best results are highlighted in bold. The columns abbreviate the full names of the LLMs: GPT refers to GPT-3.5 Turbo, Gemini to Gemini-1.5 Flash, Claude to Claude Instant 1.2, Llama to Llama 3 8B Instruct, and Mistral to Mistral Small. For TGRE\* and CoT\*, we use GPT-4o for reranking.

put tokens in the input context + output tokens generated by the model) summed across all records and divided by 100.

- **Average Cost** ( $\times 10^{-3}$  USD): Computed as the sum of costs per record, based on the API token pricing. As of March 2025<sup>4</sup>, GPT-3.5 Turbo costs \$0.5 per 1M input tokens and \$1.5 per 1M output tokens, while GPT-4o costs \$2.5 per 1M input tokens and \$10 per 1M output tokens.
- **Average Execution Time (seconds)**: Computed as the sum of API call latency (in seconds) for all records and divided by 100. For TGRE, this metric also includes embedding generation (0.0966 seconds) and retrieval time (0.0307 seconds).
- **Throughput (records per second)**: Computed as the in-

<sup>4</sup><https://openai.com/api/pricing/>

Metric	TGRE	GPT-4o	TGRE (Inference Only)
Tokens	986.5 $\pm$ 5.7	570.8 $\pm$ 15.4	262.8 $\pm$ 11.7
Cost	0.78 $\pm$ 0.04	2.76 $\pm$ 0.15	0.2 $\pm$ 0.02
Time	4.01 $\pm$ 0.57	2.88 $\pm$ 0.7	1.34 $\pm$ 0.18

Table 10: Computational cost comparison across models for occupation classification. The values represent the mean  $\pm$  standard deviation.

verse of execution time.

As shown in Table 10, TGRE processes 72.83% more tokens on average than GPT-4o. This is logical as TGRE involves both the inference and reranking stages, whereas




Figure 2: The performance of different LLMs on the complete recall, partial recall, and recognition tasks.




Figure 3: The performance of frontier and efficient LLMs on the complete recall, partial recall, and recognition tasks.

GPT-4o only processes tokens during the inference stage. However, when comparing only the inference stages, TGRE is significantly more token-efficient, processing 53.96% fewer tokens than GPT-4o.

TGRE also achieves a substantial cost reduction of 71.95% compared to GPT-4o. This demonstrates the cost-effectiveness when applied to smaller, more efficient LLMs. For large-scale classification tasks involving one million input instances, the total costs of TGRE with GPT-3.5 Turbo and GPT-4o amount to \$774.72 and \$2,761.78, respectively.

As for execution time, TGRE takes approximately one second longer per record (39.33%), compared to GPT-4o, due to the reranking step. Nevertheless, it remains highly efficient when comparing the execution time of their inference stages alone. Based on the average execution time, the estimated throughputs for TGRE and GPT-4o are approximately 0.26 records per second and 0.35 records per second, respectively. Additionally, TGRE incurs a one-time overhead for generating embeddings for O\*NET-SOC, which took about 455.65 seconds for all 1,016 SOC entities, averaging 0.45 seconds per entity.

In conclusion, The analysis shows that TGRE maintains a balance between computation cost and classification accuracy. While GPT-4o’s direct prompting is more streamlined and marginally faster, TGRE’s remarkable cost advantage makes it a practical and economical choice for large-scale classifications. The inference-only comparison further emphasizes its token efficiency and processing speed.

Metric	TGRE	GPT-4o	TGRE (Inference Only)
Tokens	2110.13 ± 31.06	288 ± 34.75	615.4 ± 56.09
Cost	1.34 ± 0.09	1.45 ± 0.33	0.46 ± 0.08
Time	4.12 ± 0.96	1.78 ± 0.62	1.97 ± 0.7

Table 11: Computational cost comparison across models for skill classification. The values represent the mean ± standard deviation.

### Skill Classification

We applied the same methodology to analyze computational costs for skill classification as was used for occupation classification. As shown in Table 11, TGRE’s token usage for skill classification is significantly higher compared to occupation classification tasks. This discrepancy is due to the inherent differences in task nature. Unlike the multi-class occupation classification, skill classification is a multi-label classification problem, requiring the inference and reranking of multiple labels per test instance. This results in higher token consumption and generation throughout the framework.

Interestingly, GPT-4o’s direct prompting exhibits greater token efficiency in skill classification compared to its performance in occupation classification tasks. Despite this, TGRE still achieves an approximate 7.45% cost reduction relative to GPT-4o, though the gap is considerably smaller than in the occupation classification scenario.




Figure 4: The performance of different LLMs on tasks at 8-digit (CT8 and TC8) and 2-digit codes (TC2).




Figure 5: The performance of frontier and efficient LLMs on tasks at 8-digit (CT8 and TC8) and 2-digit codes (TC2).




Figure 6: The performance of different reasoning methods on occupation classification.




Figure 7: The performance of different retrieval units on occupation classification.

More importantly, the inclusion of GPT-4o’s cost comparison is mainly for completeness. Unlike occupation classification, GPT-4o’s direct prompting approach in skill classification is severely hindered by a few major limitations. Firstly, GPT-4o exhibits a significantly limited knowledge of the ESCO taxonomy compared to ONET-SOC. A preliminary knowledge assessment study reveals GPT-4o’s accuracy in generating complete ESCO occupations to be only 0.010 – a stark contrast to its 0.9211 accuracy for ONET-SOC. This disparity strongly suggests that GPT-4o’s capability in ESCO skill inference is similarly limited. Secondly, When evaluated on the 100 random test instances used in

the computational cost analysis, GPT-4o’s direct prompting yields an RP@10 score of only 0.015. This demonstrates its inability to effectively predict ESCO skills without leveraging our framework.

## H Prompt Templates and In-context Examples

This appendix provides detailed prompt templates and in-context examples used across different stages of the study




Figure 8: The performance of different reasoning methods on skill classification.




Figure 9: The performance of different retrieval units on skill classification.

## Knowledge Assessment

In this section, we provide the prompt templates designed to evaluate the LLMs’ ability to generate SOC titles and codes based on the O\*NET-SOC taxonomy. Figure 10 displays recall and recognition prompt templates for the knowledge assessment tasks. Variables in the form  $\${var\_name}$ , highlighted in blue, are replaced with corresponding values from O\*NET-SOC.

## Occupation Classification

This section describes the prompt templates and in-context examples used in the inference and reranking stages of our framework for the occupation classification task. Figures 11 and 12 display TGRE-based and CoT-based prompt templates for occupation classification, respectively. Figures 13 and 14 display  $k$ -shot in-context examples for TGRE-based and CoT-based prompt templates for occupation classification, respectively. These in-context examples are integrated into the respective templates by replacing the placeholder  $\${examples}$ .

It is important to note that the TGRE prompt templates in Figure 11 do not specify the version of the O\*NET taxonomy, unlike the CoT prompt templates in Figure 12. During the prompt engineering, we found that specifying the ONET version did not affect task performance, as LLMs generally

lack the precision required to generate SOC titles and codes for specific O\*NET-SOC versions. As a result, we opted to omit the version number from TGRE prompts, while keeping them in the CoT prompts as the latter had already been used in prior experiments.

The output format for the CoT-based inference prompt template in Figure 12 is an adaptation from IReRa (D’Oosterlinck et al. 2024), originally designed for multi-label classification. To tailor this format to occupation classification, the output is truncated to retain only the top-1 predicted label.

## Skill Classification

Here, we present the prompt templates and in-context examples used in the inference and reranking stages of our framework for the skill classification task. Figures 16 and 17 display TGRE-based and CoT-based prompt templates for skill classification, respectively. Figures 18 - 23 display  $k$ -shot in-context examples for TGRE-based and CoT-based prompt templates for skill classification, respectively.

### Recall Prompt Template

Fill in the <blank>.  
---  
Output format:  
Generate either a canonical SOC title or code depending on the context.  
Encapsulate the answer in <answer></answer>.  
No other verbosity. Generate 0 if there is no suitable answer.  
---  
Example:  
In the O\*NET-SOC 2019 taxonomy, the `#{example_answer_field}` for the `#{example_query_field}` `#{example_query}` is <blank>.  
<answer>`#{example_answer}`</answer>  
---  
In the O\*NET-SOC 2019 taxonomy, the `#{answer_field}` for the `#{query_field}` `#{query}` is <blank>.

### Recognition Prompt Template

Select the best option to match entities in the O\*NET-SOC 2019 taxonomy.  
---  
Output format:  
Encapsulate the answer in <answer></answer>.  
No other verbosity. Generate 0 if there is no suitable answer.  
---  
Example:  
Which one of the options is the `#{example_answer_field}` of `#{example_query_field}` `#{example_query}`?  
Options:  
`#{ex_options}`  
<answer>`#{example_answer}`</answer>  
---  
Which one of the options is the `#{answer_field}` for the `#{query_field}` `#{query}`?  
Options:  
`#{options}`

Figure 10: Prompt templates for knowledge recall and recognition tasks. Variables in the form `#{var_name}`, highlighted in blue, are replaced with corresponding values from the O\*NET-SOC taxonomy.

### TGRE-based Inference Prompt Template

Given a job title and a company name, infer the job description and predict the most likely O\*NET-SOC occupational title and code. Only one SOC title and code per prediction. If insufficient context is given, specify "None" in the prediction. No other verbosity.

---

Follow the following format.

Reasoning: Given the job title "title" and company "company", the individual is likely responsible for...

Prediction: O\*NET-SOC occupations

---

Examples:

`#{examples}`

---

Your task:

Job title and company: `#{input}`

### TGRE-based Reranking Prompt Template

Given a job title and a company name, select up to 10 most applicable O\*NET-SOC occupational titles and codes to the job title and company name. If there is no applicable occupation, use the code "None (00-0000.00)." No other verbosity.

---

Follow the following format.

Reasoning: Given the job title "title" and company "company", the individual is likely responsible for...

Most likely occupations: A ranked list of occupations, one occupation per line separated by new line.

1. Occupation A

2. Occupation B

---

Examples:

`#{examples}`

---

Your task:

Job title and company: `#{input}`

Options:

`#{options}`

Figure 11: TGRE-based prompt templates for occupation classification. Variables in the form `#{var_name}`, highlighted in blue, are replaced with corresponding values from task inputs

### CoT-based Inference Prompt Template

Given a snippet of job title and company, predict O\*NET-SOC 2019 occupation titles and codes it most likely belongs to. Only one SOC title and code per prediction. If insufficient context is given, specify "None" in the prediction. No other verbosity.

---

Follow the following format.

Reasoning: Let's think step by step in order to predict the O\*NET 2019 occupation titles and codes it most likely belongs to. We ...

O\*NET-SOC 2019 Occupations: list of SOC titles and 8-digit codes separated by semi-colons, e.g., Title 1 (11-1101.00); Title 2 (22-2220.22)

---

Examples:

`#{examples}`

---

Your task:

Job title and company: `#{input}`

### CoT-based Reranking Prompt Template

Given a snippet of job title and company, select up to 10 O\*NET-SOC 2019 occupation titles and codes from the options most applicable to the job title and company name. Always include a numbered ranked list of O\*NET-SOC 2019 Occupations. If there is no applicable occupation, use the code "None (00-0000.00)." No other verbosity.

---

Follow the following format.

Job title and company: Job title, company

Options: A list of O\*NET-SOC occupational titles and codes to choose from

Reasoning: Let's think step by step in order to select up to 10 O\*NET-SOC 2019 occupation titles and codes from the options most applicable to the job title and company name. We ...

O\*NET-SOC 2019 Occupations: A ranked list of occupations, one occupation per line separated by new line.

1. Occupation A

2. Occupation B

...

---

Examples:

`#{examples}`

---

Your task:

Job title and company: `#{input}`

Options:

`#{options}`

Reasoning: Let's think step by step in order to

Figure 12: CoT-based prompt templates for occupation classification. Variables in the form `#{var_name}`, highlighted in blue, are replaced with corresponding values from task inputs.

### Examples for TGRE-based Inference Prompt

Job title and company: records and evidence technician, cobb county police department

Reasoning: Given the job title "records and evidence technician" and company "cobb county police department," the individual is likely expected to **collect evidence at crime scene, classify and identify fingerprints, and photograph evidence for use in criminal and civil cases.**

Prediction: Police Identification and Records Officers (33-3021.02)

### Examples for TGRE-based Reranking Prompt

Job title and company: records and evidence technician, cobb county police department

Options:

Detectives and Criminal Investigators (33-3021.00); Forensic Science Technicians (19-4092.00); Police Identification and Records Officers (33-3021.02); Fraud Examiners, Investigators and Analysts (13-2099.04); Private Detectives and Investigators (33-9021.00); Intelligence Analysts (33-3021.06); Retail Loss Prevention Specialists(33-9099.02); Transportation Security Screeners (33-9093.00); Animal Control Workers (33-9011.00); Appraisers and Assessors of Real Estate (13-2023.00)

Reasoning: Given the job title "records and evidence technician" and company "cobb county police department," the individual is likely expected to **collect evidence at crime scene, classify and identify fingerprints, and photograph evidence for use in criminal and civil cases.**

Most likely occupations:

1. Police Identification and Records Officers (33-3021.02)
2. Forensic Science Technicians (19-4092.00)
3. Detectives and Criminal Investigators (33-3021.00)
4. Fraud Examiners, Investigators and Analysts (13-2099.04)
5. Private Detectives and Investigators (33-9021.00)
6. Intelligence Analysts (33-3021.06)
7. Retail Loss Prevention Specialists (33-9099.02)
8. Transportation Security Screeners (33-9093.00)
9. Animal Control Workers (33-9011.00)
10. Appraisers and Assessors of Real Estate (13-2023.00)

Figure 13: In-context examples for TGRE-based prompts for occupation classification. Segments highlighted in blue are detailed description of the corresponding occupation retrieved from the taxonomy.

### Examples for CoT-based Inference Prompt

Job title and company: medical editor, medical editor and transcriptionist, zydoc medical transcription  
Reasoning: Let's think step by step in order to predict the O\*NET 2019 occupation titles and codes it most likely belongs to. We can start by breaking down the job title and company. "Medical editor" and "medical editor and transcriptionist" suggest that this person is responsible for editing and transcribing medical documents. The company, Zydoc Medical Transcription, confirms this assumption. Based on this information, we can predict that this job most likely belongs to the "Medical Transcriptionists" occupation (code 31-9094.00) in the O\*NET 2019 database. This occupation involves transcribing medical reports and documents. Other possible occupations could include "Medical Records and Health Information Technicians" (code 29-2071.00) or "Editors" (code 27-3041.00).  
O\*NET-SOC 2019 Occupations: Medical Transcriptionists (31-9094.00); Medical Records and Health Information Technicians (29-2071.00); Editors (27-3041.00)

### Examples for CoT-based Reranking Prompt

Job title and company: medical editor, medical editor and transcriptionist, zydoc medical transcription  
Options:  
Editors (27-3041.00); Medical Transcriptionists (31-9094.00); Medical Records Specialists (29-2072.00); Health Information Technologists and Medical Registrars (29-9021.00); Health Technologists and Technicians All Other (29-2099.00); Film and Video Editors (27-4032.00); Medical Equipment Repairers (49-9062.00); Emergency Medical Technicians (29-2042.00); Medical and Clinical Laboratory Technicians (29-2012.00); Health Informatics Specialists (15-1211.01)  
Reasoning: Let's think step by step in order to select up to 10 O\*NET-SOC 2019 occupation titles and codes from the options most applicable to the job title and company name. We can see that the job title includes "medical editor" and "medical transcriptionist", and the company name includes "medical transcription". This suggests that the job is likely related to medical editing and transcription. The options that best fit this description are Editors (27-3041.00) and Medical Transcriptionists (31-9094.00). Other options that could potentially fit are Health Information Technologists and Medical Registrars (29-9021.00) and Health Informatics Specialists (15-1211.01), as they also involve working with medical records and information. However, the job title does not mention any equipment repair or emergency medical services, so options such as Medical Equipment Repairers...  
O\*NET-SOC 2019 Occupations:  
1. Editors (27-3041.00)  
2. Medical Transcriptionists (31-9094.00)  
3. Health Information Technologists and Medical Registrars (29-9021.00)  
4. Health Informatics Specialists (15-1211.01)

Figure 14: In-context examples for CoT-based prompts generated by GPT-3.5 Turbo for occupation classification

### Zeroshot Occupation Selection Prompt

Given a job record (job title and company), select one or more O\*NET-SOC labels (titles and codes) from the options which are most applicable to the job record.  
---  
Follow this format:  
Answer: List of applicable O\*NET-SOC labels separated by semicolons. Each selected label should contain both title and code verbatim. If multiple labels are equally applicable, separate them with semicolons in your answer. If there is no applicable label, select or indicate "None (00-0000.00)." No other verbosity.  
---  
Your task:  
Job record: `#{input}`  
Options:  
`#{options}`

Figure 15: Zero-shot occupation selection prompt templates for occupation classification. Variables in the form `#{var_name}`, highlighted in blue, are replaced with corresponding values from task inputs.

### TGRE-based Inference Prompt Template

Given a snippet from a job vacancy, generate rationale in order to identify all ESCO skills explicitly mentioned. In the rationale, encapsulate ESCO concept and description in a square bracket.

---

Follow the following format.

Reasoning: The job vacancy snippet explicitly mentions...Therefore, we can infer the following ESCO skills and knowledge...[ESCO concept: description]...

---

Example #1:

`#{example_1}`

---

Example #2:

`#{example_2}`

---

Your task:

Text snippet: `#{input}`

### TGRE-based Reranking Prompt Template

Given a snippet from a job vacancy and a list of ESCO skills and knowledge as options, first generate a rationale then select verbatim up to 10 options most applicable to skills and knowledge explicitly mentioned in the snippet. No other verbosity.

---

Follow this out format.

Reasoning: The job vacancy snippet explicitly mentions...These phrases give us clues in determining the most relevant options by aligning them with the specific requirements in the snippet...

Therefore, we select the following 10 most applicable options: A ranked list of ESCO skills and knowledge, one skill per line separated by new line.

1. Skill A

2. Skill B

...

---

Example #1:

`#{example_1}`

---

Example #2:

`#{example_2}`

---

Your task:

Text snippet: `#{input}`

Options:

`#{options}`

Figure 16: TGRE-based prompt templates for skill classification. Variables in the form `#{var_name}`, highlighted in blue, are replaced with corresponding values from task inputs.

### CoT-based Inference Prompt Template

```
Given a snippet from a job vacancy, identify all the ESCO job skills mentioned. Always return skills.
---
Follow the following format.
Vacancy: Job vacancy snippet
Reasoning: Let's think step by step in order to find the skills. We ...
Skills: list of comma-separated ESCO skills
---
${example_1}
---
${example_2}
---
Your task:
Vacancy: ${input}
Reasoning: Let's think step by step in order to
```

### CoT-based Reranking Prompt Template

```
Given a snippet from a job vacancy, generate a chain-of-thought rationale and select the 10 most
applicable ESCO skills from the options that are directly expressed in the snippet.
---
Follow the following format.
Reasoning: Start the rationale with "Let's think step by step in order to identify the most applicable
skills from the snippet. We..."
Skills: Ranked list of ESCO skills separated by semicolon
---
Example #1:
${example_1}
---
Example #2:
${example_2}
---
Your task:
Text snippet: ${input}
Options:
${options}
```

Figure 17: CoT-based prompt templates for skill classification. Variables in the form  $\${var\_name}$ , highlighted in blue, are replaced with corresponding values from task inputs.

### Example 1 for TGRE-based Inference Prompt

Text snippet: You'll be required to apply your depth of knowledge and expertise to all aspects of the software development lifecycle as well as partner continuously with your many stakeholders on a daily basis to stay focused on common goals.

Reasoning: The job vacancy snippet explicitly mentions "software development lifecycle," "partner continuously with your many stakeholders," and "stay focused on common goals." Therefore, we can infer the following ESCO skills and knowledge:

- "Software development lifecycle" likely refers to the ESCO knowledge concept [systems development life-cycle: [The sequence of steps, such as planning, creating, testing and deploying and the models for the development and life-cycle management of a system](#)].
- "Partner continuously with your many stakeholders" likely refers to the ESCO skill/competence concept [manage relationships with stakeholders: [Create and maintain solid internal and external relations with stakeholders at operational level based on mutual trust and credibility in order to achieve organisational goals. Ensure organisational strategies incorporate strong stakeholder management and identify and prioritise strategic stakeholder relationships](#)].
- "Stay focused on common goals" likely refers to the ESCO skill/competence concept [identify with the company's goals: [Act for the benefit of the company and for the achievement of its targets](#)].

### Example 2 for TGRE-based Inference Prompt

Text snippet: You are proficient in Python and English

Reasoning: The job vacancy snippet explicitly mentions "Python" and "English." Therefore, we can infer the following ESCO skills and knowledge:

- "Python" likely refers to the ESCO knowledge concept [Python (computer programming): [The techniques and principles of software development, such as analysis, algorithms, coding, testing and compiling of programming paradigms in Python](#)].
- "English" likely refers to the ESCO knowledge concept [English: [The English language](#)].

Figure 18: In-context examples for TGRE-based inference prompts for skill classification. Segments highlighted in blue are detailed description of the corresponding occupation retrieved from the taxonomy.

### Example 1 for TGRE-based Reranking Prompt

Text snippet: You'll be required to apply your depth of knowledge and expertise to all aspects of the software development lifecycle as well as partner continuously with your many stakeholders on a daily basis to stay focused on common goals.

Options:

Agile development; DevOps; Iterative development; Prototyping development; Rapid application development; Spiral development; Waterfall development; align efforts towards business development; collaborate on international energy projects; collaborate with coaching team; collaborate with designers; content development processes; cooperate to resolve information issues; cooperate with colleagues; cooperate with education professionals; create a work atmosphere of continuous improvement; create social alliances; develop production line; develop therapeutic relationships; engage with stakeholders; ensure cross-department cooperation; ensure cross-department cooperation; establish collaborative relations; focus on service; identify customer objectives; identify with the company's goals; incremental development; levels of software testing; liaise with cultural partners; liaise with engineers; maintain relationship with customers; maintain relationship with suppliers; manage personal progression; manage relationships with stakeholders; motivate employees; negotiate with stakeholders; negotiate with stakeholders; procurement lifecycle; product life-cycle; project commissioning; promote the principles of democracy and rule of law; research your target community; social alliances; solution deployment; strive for company growth; systems development life-cycle; talent management strategies; teamwork principles

Reasoning: The job vacancy snippet explicitly mentions "software development lifecycle," "partner continuously with your many stakeholders," and "stay focused on common goals." These phrases give us clues in determining the most relevant options by aligning them with the specific requirements in the snippet.

- The term "software development lifecycle" in the snippet directly corresponds to [systems development life-cycle: [The sequence of steps, such as planning, creating, testing and deploying and the models for the development and life-cycle management of a system](#)], making it the most relevant option.

- The term "partner continuously with your many stakeholders" directly points to [manage relationships with stakeholders: [Create and maintain solid internal and external relations with stakeholders at operational level based on mutual trust and credibility in order to achieve organisational goals. Ensure organisational strategies incorporate strong stakeholder management and identify and prioritise strategic stakeholder relationships](#)].

- The term "stay focused on common goals" suggests the importance of aligning individual efforts with the company's objectives, highlighting the need for employees to [identify with the company's goals: [Act for the benefit of the company and for the achievement of its targets](#)].

Therefore, we select the following 10 most applicable options:

1. systems development life-cycle
2. manage relationships with stakeholders
3. identify with the company's goals
4. engage with stakeholders
5. negotiate with stakeholders
6. maintain relationship with customers
7. Waterfall development
8. Iterative development
9. Agile development
10. Rapid application development

Figure 19: In-context example 1 for TGRE-based reranking prompts for skill classification. Segments highlighted in blue are detailed description of the corresponding occupation retrieved from the taxonomy.

### Example 2 for TGRE-based Reranking Prompt

Text snippet: You are proficient in Python and English

Options:

APL; Ancient Greek; Arabic; C++; COBOL; English; Erlang; French; German; Greek; Groovy; ICT system programming; Japanese; Java (computer programming); Latin; MATLAB; ML (computer programming); Objective-C; Pascal (computer programming); Python (computer programming); Ruby (computer programming); Russian; Scala; TypeScript; Yiddish; academic English; aviation English; comprehend the material to be translated; computer programming; computer science; computer technology; develop digital educational materials; e-learning; interpret languages in conferences; keywords in digital content; manage business knowledge; manage localisation; query languages; teach ESOL language class; teach ESOL literacy class; teach languages; teach linguistics; technical terminology; technology education; translate foreign language; translate spoken language; understand spoken English; understand written English; use localisation tools; write English

Reasoning: The job vacancy snippet explicitly mentions "Python" and "English." These phrases give us clues in determining the most relevant options by aligning them with the specific requirements in the snippet.

- The term "Python" directly corresponds to [Python (computer programming): [The techniques and principles of software development, such as analysis, algorithms, coding, testing and compiling of programming paradigms in Python](#)], making it the most relevant option.

- The term "English" directly corresponds to [English: [The English language.](#)] and it can also be relevant to "academic English," "understand spoken English," "understand written English," and "write English," depending on the context in which English proficiency is required.

Therefore, we select the following 10 most applicable options:

1. Python (computer programming)
2. English
3. understand written English
4. understand spoken English
5. write English
6. academic English
7. ICT system programming
8. computer programming
9. technical terminology
10. technology education

Figure 20: In-context example 2 for TGRE-based reranking prompts for skill classification. Segments highlighted in blue are detailed description of the corresponding occupation retrieved from the taxonomy.

### Example 1 for CoT-based Inference Prompt

Vacancy: performed code reviews for your teammates

Reasoning: Let's think step by step in order to find the skills. We need to find the skills that are mentioned in the vacancy.

Skills: Code review, teamwork

### Example 2 for CoT-based Inference Prompt

Vacancy: You are a good communicator and enjoy working in a team

Reasoning: Let's think step by step in order to find the skills. We need to find the skills that are mentioned in the vacancy.

Skills: Communication, Teamwork

Figure 21: In-context examples for CoT-based inference prompts GPT-3.5 Turbo for skill classification

### Example 1 for CoT-based Reranking Prompt

Text snippet: performed code reviews for your teammates

Options:

work in fitness teams; collaborate in company's daily operations; risk management; work in assembly line teams; lead a team; set up effective working relationships with other sports players; craftsmanship; encourage teambuilding; capacity building; work as a team in a hazardous environment; peer group methods; team building; build team spirit; supervise crew; motivate others; organise supporters; plan the work of teams and individuals; social alliances; teamwork principles; conduct ICT code review; military drill; lead a team in forestry services; work in teams; stimulate creativity in the team; manage the security team; coordinate engineering teams; plan teamwork; coaching techniques; manage a team; perform inspection analysis; assemble an artistic team; encourage teams for continuous improvement; personnel management; work in a construction team; motivate supporters; collaborate with coaching team; create social alliances; establish collaborative relations; coach staff for running the performance; evaluate performance of organisational collaborators; communication; report to the team leader; coordinate communication within a team; work with supporting team in community arts programme; contribute to the programmer's reflection process; cooperate with colleagues; ensure positive relations within the artistic team; write inspection reports; facilitate teamwork between students; shape organisational teams based on competencies

Reasoning: Let's think step by step in order to identify the most applicable skills from the snippet. When the snippet mentions "performed code reviews for your teammates," it highlights the need for skills related to understanding and assessing code quality, providing constructive feedback, and collaborating effectively within a technical team context.

Skills: conduct ICT code review; teamwork principles; communication; facilitate teamwork between students; coordinate communication within a team; evaluate performance of organisational collaborators

Figure 22: In-context example 1 for CoT-based prompts generated by GPT-3.5 Turbo for skill classification

### Example 2 for CoT-based Reranking Prompt

Text snippet: You are a good communicator and enjoy working in a team

Options:

use communication devices; inter-professional communication in the psychosocial field; encourage teambuilding; stimulate creativity in the team; disseminate internal communications; coordinate communication within a team; lead a team; use therapeutic communication techniques; communication; team building; promote good habits to avoid communication disorders; disseminate messages to people; communication related to hearing impairment; teamwork principles; communicate with media; work in teams; manage a team; digital communication and collaboration; communication studies; maintain internal communication systems; handle communications in the food processing industry; plan teamwork; body language; circulate information; cognitive psychology; plan the work of teams and individuals; cooperate with colleagues; teach communication methods; encourage teams for continuous improvement; communication disorders; communicate by use of interpretation services; communication principles; use communication equipment; use communication techniques; communication sciences; facilitate teamwork between students; teach communication to clients; apply technical communication skills; communicate by telephone; advise on communication strategies; promote organisational communication; use different communication channels; electronic communication; use a complex communication system; communicate in healthcare; promote communication between parties; develop communications strategies; develop inclusive communication material; psychology; build team spirit

Reasoning: Let's think step by step in order to identify the most applicable skills from the snippet. We need to focus on skills directly related to communication and teamwork. The snippet mentions being a "good communicator" and "enjoy working in a team," so we will prioritize skills that align with these two aspects.

Skills: communication; teamwork principles; work in teams; use communication techniques; promote communication between parties; coordinate communication within a team; build team spirit; cooperate with colleagues; lead a team; manage a team

Figure 23: In-context example 2 for CoT-based prompts generated by GPT-3.5 Turbo for skill classification