

SkillGPT: a RESTful API service for skill extraction and standardization using a Large Language Model

Nan Li, Bo Kang, and Tijl De Bie

IDLAB - Department of Electronics and Information Systems (ELIS), Ghent University, Ghent 9000, Belgium {firstname.lastname}@ugent.be

Abstract. We present SkillGPT, a tool for skill extraction and standardization (SES) from free-style job descriptions and user profiles with an open-source Large Language Model (LLM) as backbone. Most previous methods for similar tasks either need supervision or rely on heavy data-preprocessing and feature engineering. Directly prompting the latest conversational LLM for standard skills, however, is slow, costly and inaccurate. In contrast, SkillGPT utilizes a LLM to perform its tasks in steps via summarization and vector similarity search, to balance speed with precision. The backbone LLM of SkillGPT is based on Llama, free for academic use and thus useful for exploratory research and prototype development. Hence, our cost-free SkillGPT gives users the convenience of conversational SES, efficiently and reliably.

Keywords: Skill extraction · Skill standardization · Information retrieval · Large language model.

1 Introduction

Automatic *skill extraction* aims to detect skills and qualities from unstructured job descriptions and user profiles without human involvement. *Skill standardization* further maps the free-styled skill annotations onto a standard taxonomy, such as the “European Skills, Competences, Qualifications and Occupations” (ESCO) standard [5]. Skill extraction and standardization (SES) benefit various downstream tasks, such as job recommendation and career path planning. The state of art SES methods are mostly supervised [12,13,7,1], relying on the costly and cumbersome process of human annotation. Weakly supervised techniques proposed such as by Zhang et al. [15] require heavy data preprocessing and feature engineering. BERT-based Large Language Models (LLMs) [3,9] were utilized recently [2,16,6], but merely as a refinement of the former static word embeddings.

Recent LLMs have created the possibility of solving SES without human annotation, heavy data-preprocessing or feature engineering. The latest generation (e.g. chatGPT) further shows the power and convenience of conversational interfaces, making code-less interactions possible for domain experts. However,

LLMs have some limitations that make them hard to apply directly for SES. Directly prompting a job posting for ESCO codes might well be answered with non-existent codes due to the *hallucination* problem [4,8]. To improve the accuracy, ESCO documents should be fed into the LLM as context. Yet, handling *lengthy prompts* of hundreds of pages is not straightforward. Even if the ESCO documents were somehow prompted (e.g. by chunk), model predictions would be *inefficient*, making it infeasible for online usage. Using powerful but *costly and close-sourced* models from OpenAI such as ChatGPT would hinder exploration and evaluation of possible solutions.

These considerations motivated us to design an SES tool to utilize the power of LLMs without sacrificing *accuracy, speed and cost-efficiency*. We call our solution **SkillGPT**: an API service for skill extraction and standardization from free-style text, that uses *fast* vector similarity search against precomputed ESCO embeddings, retrieves *with high precision* the matching codes given any prompted job description or user profile, and *costs little* (at least for academia) by using the latest open-sourced LLM Vicuna-13B [11]¹ (Sec. 2). SkillGPT has more merits: it is *multi-faceted* as it can extract multiple ESCO concept types, and it offers *multi-lingual* support for English, French and Dutch (Sec. 3). Our code is publicly available at <https://github.com/aida-ugent/SkillGPT>.

2 API Design

Fig. 1 is a high-level overview of SkillGPT’s components, with annotations of the specific models and tools² used in our demo.

There are two major phases of SkillGPT’s lifecycle. *First* is the system initialization for embedding and storing the ESCO taxonomy, where the ESCO entries are downloaded and converted into structured documents, then their vectorized representations are computed by the backbone LLM⁴ via a document embedding service², and finally stored in the vector database⁵.

Second is the main functionality of skill extraction and standardization. Extracting free-style textual skills from a job description or user resume is achieved by carefully prompting the LLM for summarization¹. The embedding (via ²) of the distilled text is used as a query to retrieve the top-k similar ESCO terms via vector similarity search³. All these are coordinated by an API gateway⁶.

To perform an one-stop SES or to interact with any individual service, we provide two options: direct interaction with a RESTful API⁸, or via graphical user interface⁷ (e.g. Gradio App).

¹ Vicuna is based on Llama [14], reported to perform comparably to GPT-3, and fine-tuned using data from shareGPT for better alignment with human preferences. Our experiments showed that Vicuna-13B (the number of parameters is 13 billion) has sufficient capacity for SES. More importantly, as it is open-source we can to deploy our model on a local machine with minimal cost. Nonetheless, developers can replace Vicuna with any LLM they want, including paid ones.

² We choose the tools through careful considerations and experiments. Detailed justification of our design choices as well as ablation study will be presented in an extended manuscript.

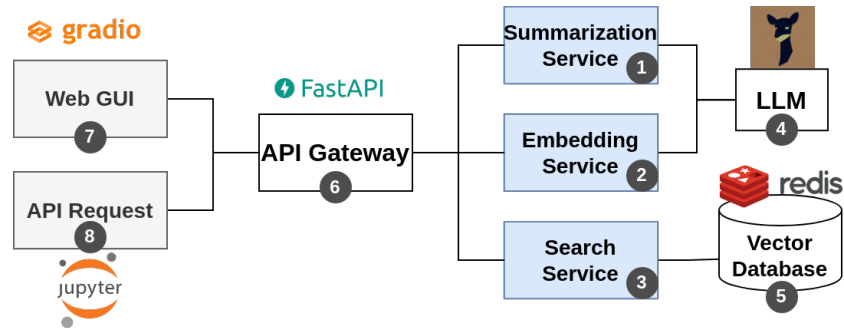


Fig. 1. Main components of SkillGPT. The specific tools and models are our careful choices for the current version, but our framework is flexible to accommodate other options. Note that this component diagram merely shows the design components, not the data flow.

SkillGPT
A RESTful API service for skill extraction and standardization from job descriptions and user profiles using large language model

Document

Description de la fonction **a**
Vous effectuez des audits financiers auprès de banques, de gestionnaires d'actifs et d'investisseurs institutionnels. Lors de la vérification et de la certification de produits financiers et de portefeuilles d'investissement, vous vous concentrez sur les critères et les politiques de durabilité. Vous renforcez notre gamme de nouveaux services financiers, par exemple des "Second Party Opinion". En tant que membre senior de l'équipe, vous soutenez vos collègues (Research Officers) et partagez vos connaissances sur le contenu. Vous coopérez avec eux dans des travaux de recherche sur la RSE et l'ISR, par exemple. Vous contribuez également à sensibiliser le public et à promouvoir le développement durable, grâce aux partenariats de Forum Ethibel. Par exemple, par le biais de l'étude annuelle ERSIS, via les médias...

Document type **b**
 job description user profile

Summary **d**
1. Audit financier: Vous êtes en mesure d'effectuer des audits financiers
2. Connaissances sur les critères et les politiques de durabilité: Vous renforcez notre gamme de nouveaux services financiers, par exemple des "Second Party Opinion". En tant que membre senior de l'équipe, vous soutenez vos collègues (Research Officers) et partagez vos connaissances sur le contenu. Vous coopérez avec eux dans des travaux de recherche sur la RSE et l'ISR, par exemple. Vous contribuez également à sensibiliser le public et à promouvoir le développement durable, grâce aux partenariats de Forum Ethibel. Par exemple, par le biais de l'étude annuelle ERSIS, via les médias...
3. Soutien à l'équipe: En tant que membre senior de l'équipe, vous soutenez vos collègues (Research Officers) et partagez vos connaissances sur le contenu. Vous coopérez avec eux dans des travaux de recherche sur la RSE et l'ISR, par exemple. Vous contribuez également à sensibiliser le public et à promouvoir le développement durable, grâce aux partenariats de Forum Ethibel. Par exemple, par le biais de l'étude annuelle ERSIS, via les médias...
4. Partenariats de Forum Ethibel: Vous contribuez à promouvoir le développement durable, grâce aux partenariats de Forum Ethibel. Par exemple, par le biais de l'étude annuelle ERSIS, via les médias...
5. Connaissance sur la RSE et l'ISR: Vous êtes capable de travailler sur des projets de recherche sur la RSE et l'ISR, par exemple. Vous contribuez également à sensibiliser le public et à promouvoir le développement durable, grâce aux partenariats de Forum Ethibel. Par exemple, par le biais de l'étude annuelle ERSIS, via les médias...

ESCO concept type **f**
 skills occupations skillGroups

ESCO concepts **i**

rank	preferredLabel	conceptType	code	altLabels
1	donner des conseils en matière de participation	Occupation	1213.8	sustainability analyst susta
2	assurer le développement du personnel	Occupation	2412.6.4	investor relation manager i
3	Élaborer des stratégies la passation de m	Occupation	2431.5	business continuity analyst
4	donner des conseils sur la gestion de ris	Occupation	2412.6.4	chargé de relations investisseurs chargée de
5	participer à la gestion de fonds	Occupation	1213.8	spécialiste du développement durable coordinat

ESCO concept type **g**
 skills occupations skillGroups

ESCO concepts **j**

rank	preferredLabel	conceptType	code	altLabels
1	responsable de la communication financière	Occupation	2412.6.4	chargé de relations investisseurs chargée de
2	responsable de la durabilité	Occupation	1213.8	spécialiste du développement durable coordinat
3	chargé de mission développement régional/ chargée de mission développement régional	Occupation	2428.12.14	chargé de mission développement régional charg

Document **e**
You carry out financial audits with banks, asset managers and institutional investors. When verifying and certifying financial products and invest services, for example "Second Party Opinion". As a senior member of the team, you support your colleagues (Research Officers) and share your cor public awareness and promote sustainable development, thanks to Forum Ethibel's partnerships. For example, through the annual ERSIS study.

Document type **c**
 job description user profile

Summary **h**
1. Financial Audit: The ability to carry out financial audits of banks, asset managers, and institutional investors to verify and certify financial prod
2. Sustainability knowledge: The expertise and understanding of sustainability criteria and policies to strengthen the range of new financial servi
3. Collaboration and Leadership: The ability to support and collaborate with colleagues (Research Officers) to share knowledge and expertise, an
4. Public Awareness and Promotion: The ability to raise public awareness and promote sustainable development through partnerships and medi
5. Technical Analysis and Reporting: The ability to analyze and report on financial data and research findings with accuracy and clarity, contribut

ESCO concept type **h**
 skills occupations skillGroups

ESCO concepts **j**

rank	preferredLabel	conceptType	code	altLabels
1	donner des conseils en matière de participation	Occupation	1213.8	sustainability analyst susta
2	assurer le développement du personnel	Occupation	2412.6.4	investor relation manager i
3	Élaborer des stratégies la passation de m	Occupation	2431.5	business continuity analyst
4	donner des conseils sur la gestion de ris	Occupation	2412.6.4	chargé de relations investisseurs chargée de
5	participer à la gestion de fonds	Occupation	1213.8	spécialiste du développement durable coordinat

Fig. 2. SkillGPT usage: multi-faceted and multi-lingual. The screenshots with blue background are from different sessions.

3 Use cases

The current version SkillGPT supports 2 document types (job description/user resume), 3 ESCO concept types (Skill/Occupation/Occupation group) and 3 languages (En/Fr/Nl), yielding **18** possible use cases. Due to space limitation, we only show the following condensed exemplary use cases, and kindly refer to our demo videos (<http://bokang.io/videos/SkillGPT.mp4>) for more. *First*, a user pastes a free-styled document in French (Fig. 2a), English (Fig. 2e), or Dutch into the top textbox, followed by choosing the type of the document between job description and user profile (Fig. 2b). *Then* the user clicks **Summarize** (Fig. 2c) to let the backbone LLM distill the skills contained in the document. The output skill list will be in the same language of the original input (Fig. 2d and g). *Next*, the user can choose which ESCO concept type (Fig. 2f) to standardize the free-style skill descriptions, and the corresponding most plausible ESCO terminologies will be returned (Fig. 2i). One *limitation* is that the extracted codes for the same content in different languages, although they mostly overlap, may differ somewhat (Fig. 2j). Our test trials also showed that SkillGPT seems to perform better in English documents than Dutch or French. The LLM probably models the English language better than other languages due to the non-uniform volume distribution of the training corpus. There is also some randomness in the generation procedure of skill summaries.

4 Conclusion

We presented SkillGPT, a flexible and easy-to-use tool for skill extraction and standardization from free-styled job descriptions and user resumes. SkillGPT is *efficient*, *economical* and delivers often truthful and *plausible* results, thus meeting our design goals. However, SkillGPT in its current state has some *limitations*. First, treating the summarized text as a single document might cause certain subtle skills to be lost, since the synthesized embedding might be dominated by the most salient qualities. Second, many options for optimizing the performance of LLMs, such as fine-tuning and prompt engineering, have not properly been examined, due to time limitations and, more importantly, that how to harness LLMs is a booming field where new hacks and tricks are proposed everyday. In *future*, we plan to address the aforementioned limitations, perform both qualitative and quantitative evaluations on SES and various downstream tasks in e-recruitment recommendation [10], and optimize for smaller languages as well as support the full range of 25 European Union languages.

Acknowledgements The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) (ERC Grant Agreement no. 615517), and under the European Union’s Horizon 2020 research and innovation programme (ERC Grant Agreement no. 963924), from the Flemish Government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” programme, and from the FWO (project no. G0F9816N, 3G042220).

References

1. Bhola, A., Halder, K., Prasad, A., Kan, M.Y.: Retrieving skills from job descriptions: A language model based extreme multi-label classification framework. In: Proceedings of the 28th International Conference on Computational Linguistics. pp. 5832–5842 (2020)
2. Chernova, M.: Occupational skills extraction with FinBERT. Master’s thesis (2020)
3. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
4. Dziri, N., Milton, S., Yu, M., Zaiane, O., Reddy, S.: On the origin of hallucinations in conversational models: Is it the datasets or the models? arXiv preprint arXiv:2204.07931 (2022)
5. European Commission: European skills, competences, qualifications and occupations dataset v1.1.1 (2023), <https://esco.ec.europa.eu/en/use-esco/download>
6. Gnehm, A.S., Bühlmann, E., Clematide, S.: Evaluation of transfer learning and domain adaptation for analyzing german-speaking job advertisements. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference. pp. 3892–3901 (2022)
7. Green, T., Maynard, D., Lin, C.: Development of a benchmark corpus to support entity recognition in job descriptions. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference. pp. 1201–1208 (2022)
8. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P.: Survey of hallucination in natural language generation. *ACM Computing Surveys* **55**(12), 1–38 (2023)
9. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
10. Mashayekhi, Y., Li, N., Kang, B., Lijffijt, J., Bie, T.D.: A challenge-based survey of e-recruitment recommendation systems (2022)
11. Peng, B., Li, C., He, P., Galley, M., Gao, J.: Instruction tuning with gpt-4. arXiv preprint arXiv:2304.03277 (2023)
12. Sayfullina, L., Malmi, E., Kannala, J.: Learning representations for soft skill matching. In: Analysis of Images, Social Networks and Texts: 7th International Conference, AIST 2018, Moscow, Russia, July 5–7, 2018, Revised Selected Papers 7. pp. 141–152. Springer (2018)
13. Tamburri, D.A., Van Den Heuvel, W.J., Garriga, M.: Dataops for societal intelligence: a data pipeline for labor market skills extraction and matching. In: 2020 IEEE 21st International Conference on Information Reuse and Integration for Data Science (IRI). pp. 391–394. IEEE (2020)
14. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 (2023)
15. Zhang, M., Jensen, K.N., van der Goot, R., Plank, B.: Skill extraction from job postings using weak supervision. arXiv preprint arXiv:2209.08071 (2022)
16. Zhang, M., Jensen, K.N., Plank, B.: Kompetenzer: Fine-grained skill classification in danish job postings via distant supervision and transfer learning. arXiv preprint arXiv:2205.01381 (2022)