

# ADAPT-IA: Towards an Adaptive AI in Language Technologies applied to RIS3 Industrial sectors<sup>\*</sup>

Aitor Álvarez<sup>1,\*</sup>, Arantza del Pozo<sup>1</sup>, Montse Cuadros<sup>1</sup>, Thierry Etchegoyhen<sup>1</sup>, Juan Camilo Vásquez-Correa<sup>1</sup>, Ander González-Docasal<sup>1,2</sup>, Santiago Andrés Moreno-Acevedo<sup>1</sup>, Harritxu Gete<sup>1</sup>, Victor Ruiz<sup>1</sup>, Aingeru Bellido<sup>1</sup>, Alexander Platas<sup>1</sup>, Maia Agirre<sup>1</sup>, Ariane Méndez<sup>1</sup>, Javier Mikel Olaso<sup>1</sup>, Antonio Aparicio<sup>3</sup>, Asier López<sup>3</sup>, M. Inés Torres<sup>3,4</sup>, Begoña Arrate<sup>4</sup>, Joxean Zapiain<sup>4</sup>, Oscar Montserrat<sup>5</sup>, José María Echevarría<sup>5</sup>, José Ignacio Hormaetxe<sup>6</sup>, Pedro Fortea<sup>7</sup>, Sofia Olaso<sup>7</sup> and Jon Nuñez<sup>8</sup>

<sup>1</sup>Fundación Vicomtech, Basque Research and Technology Alliance (BRTA), Donostia – San Sebastián, Spain

<sup>2</sup>University of Zaragoza, Department of Electronics, Engineering and Communications

<sup>3</sup>University of the Basque Country - UPV/EHU

<sup>4</sup>Basque Centre for Terminology and Lexicography – UZEI

<sup>5</sup>Automotive Intelligence Centre - AIC

<sup>6</sup>Asociación Clúster de Energía del País Vasco - ACE

<sup>7</sup>Asociación de la Industria Ferroviaria Española - MAFEX

<sup>8</sup>Fundación de Investigación de la Máquina-Herramienta - INVEMA

## Abstract

Despite the fact that Language Technologies are increasingly benefiting Industry by accelerating and automating processes, Basque remains a largely overlooked language in these technologies, even though it is an official language of the Autonomous Community of the Basque Country and has a growing number of speakers. With the aim of promoting its use in industrial environments, ADAPT-IA was launched as a research and development project focused on Adaptive AI technologies applied to Language Technologies in multiple languages, including Basque, aiming to enhance their integration into industrial processes across key sectors in the Basque Country. We detail six use cases related to the development of advanced prototypes within the Automotive, Energy, Machine Tool, and Railway industries. These prototypes were developed through a collaborative effort involving the applied research center Vicomtech, the University of the Basque Country, the Basque Center for Terminology and Lexicography, and the corresponding industry clusters. This work demonstrates the potential of multilingual Language Technologies to address the specific needs of different industrial sectors.

## Keywords

Adaptive AI, Voice and text Assistants, Large Language Models, Adaptive Machine Translation, Language resources, MLOps

## 1. Introduction

The exponential growth of new technologies, driven by Artificial Intelligence (AI) and its increasing democratization in society, is becoming a pivotal factor in the digitization of Industry [1]. In recent years, Language Technologies (LT) has increasingly benefit the Industry across various areas. These advancements include many applications, such as the use of machine translation for translating technical manuals of industrial machinery [2], automatic classification of technical documents, manuals and reports [3, 4], data mining and extraction of key information from large volumes of technical texts [5],

SEPLN 2025: 41st International Conference of the Spanish Society for Natural Language Processing, Zaragoza, Spain, 23-26 September 2025.

\*Corresponding author.

<sup>†</sup>These authors contributed equally.

✉ aalvarez@vicomtech.org (A. Álvarez); adelpozo@vicomtech.org (A. d. Pozo); mcuadros@vicomtech.org (M. Cuadros); techegoyhen@vicomtech.org (T. Etchegoyhen)

🆔 0000-0002-7938-4486 (A. Álvarez); 0000-0003-3505-5514 (A. d. Pozo); 0000-0002-3620-1053 (M. Cuadros); 0000-0001-7253-1693 (T. Etchegoyhen)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

AI assistants to enhance operator interaction with intelligent systems [6], and speech processing technologies [7, 8], among others.

The AI technology is currently experiencing a golden age. AI systems developed by leading technology companies, such as OpenAI, Deepseek, Mistral, Google, Meta, Claude, and Qwen, among others, are increasingly exhibiting remarkable generative capabilities in LT. Although the Large Language Models (LLM) integrated in these systems possess vast knowledge from being trained on extensive datasets -often accessible only to major industry players-, their performance and quality tend to decrease when applied to specific domains, languages, or data types that were not present during training. Despite their great potential, they still suffer from one of the main challenges of AI in general: models must be adapted with domain-specific data to perform optimally in specialized scenarios, industrial sectors, or low-resource minority languages such as Basque.

Within this context, the ADAPT-IA project aimed to advance research and development in Adaptive AI technologies applied to LT in several languages, including Basque, with the objective of facilitating their seamless integration into processes across various Basque Industrial sectors, such as automotive, machine tool, energy, and the railway industry, aligned with the objectives of RIS3 Euskadi's Research and Innovation Strategy for Smart Specialization.

We initially identified specific areas of interest to develop high-value use cases for each industrial sector using LT, in collaboration with the representative clusters participating in the consortium. The final selected use-cases were mainly focused on applications of (1) voice interaction between the user/operator and machines, (2) domain-expert and generic voice/text assistants for specialized and open-ended queries about machines, their operation and maintenance, (3) report classification system to ensure that unstructured inputs converge into a standardized text with consistent meaning, (4) voice control technology for managing and controlling a production line through spoken queries and data reporting, and (5) adaptive machine translation systems.

The Adaptive AI-related activities were accompanied by a significant effort to collect and generate in-domain data from large amounts of heterogeneous sources, including technical reports, manuals, specialized journals, technical glossaries, and more general digital repositories. Additionally, generative AI technology was also employed to produce synthetic data for training and adapting AI models. These data resources will be freely available to the community for research purposes.

The R&D activities were performed within the basic research project ADAPT-IA, partially supported by the Department of Economic Development of the Basque Government. The project started in April 2023 and finalised in December 2024, and was carried out by the following consortium: Vicomtech<sup>1</sup> (project coordinator), University of the Basque Country (UPV/EHU)<sup>2</sup>, the Basque Centre for Terminology and Lexicography (UZEI)<sup>3</sup>, INVEMA<sup>4</sup>, Cluster de Energía (ACE)<sup>5</sup>, Industria Ferroviaria Española (MAFEX)<sup>6</sup>, and Automotive Intelligence Center (AIC)<sup>7</sup>.

## 2. Main objectives

The main objective of the ADAPT-IA project was the research and development of Adaptive AI technologies applied to Language Technologies in multiple languages, with the aim of enhancing and facilitating their integration into industrial processes across various sectors in the Basque Country. The project also seeks to explore emerging methodologies that optimize the maintenance, adaptation, and continuous deployment of neural models in production, leveraging the growing potential of AI in combination with human knowledge and creativity. Figure 1 provides a simplified overview of the core concept of the project.

---

<sup>1</sup><https://www.vicomtech.org/en>

<sup>2</sup><https://www.ehu.eus/>

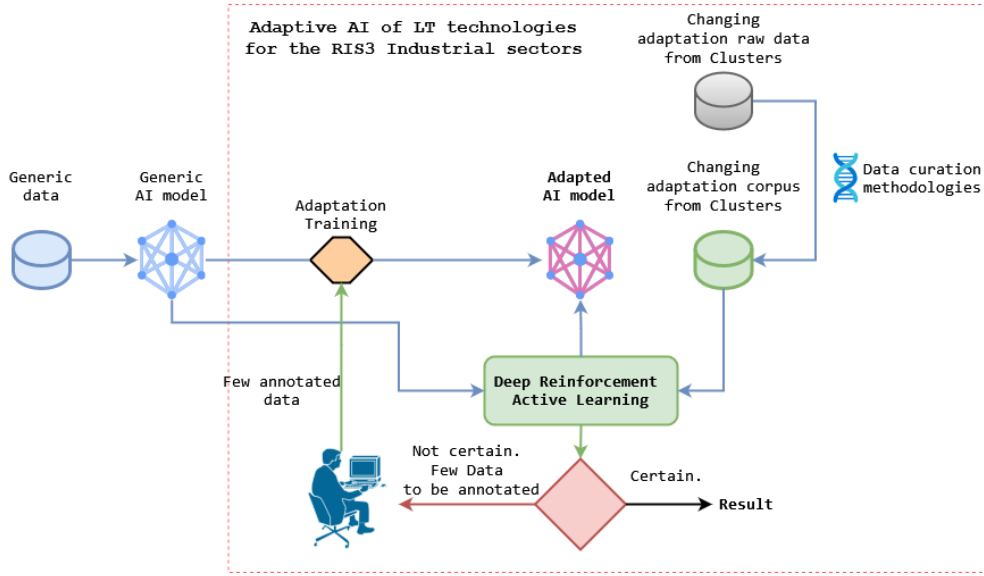
<sup>3</sup><https://uzei.eus/>

<sup>4</sup><https://www.afm.es/>

<sup>5</sup><https://www.clusterenergia.com/>

<sup>6</sup><https://mafex.es/>

<sup>7</sup><https://www.aicenter.eu/>



**Figure 1:** Main concept of the ADAPT-IA project.

More specifically, the project focuses on the following key technical objectives:

- R&D in Adaptive AI technologies applied to LTs in several languages, including paradigms such as Active Learning, Reinforcement Learning, and Transfer Learning, as well as their combination to enhance synergies and leverage the unique benefits of each approach.
- R&D in neural AI modeling for each language technology required in the project's priority use cases, with an emphasis on emerging neural architectures and a balance between performance and efficiency, ensuring the highest possible quality with minimal resource consumption.
- Generation of linguistic resources in several languages to enable the implementation of LT systems based on Adaptive AI, while promoting the use of the language and its terminological standardization in key industrial sectors of the Basque Country.
- Development of a prototype by each consortium member and sector, integrating the R&D activities in Adaptive AI applied to LTs.
- Validation of the prototypes with real users, including training and capacity-building processes.
- Scientific dissemination of the results in leading international forums, journals, and conferences.

The R&D activities, along with the development of prototypes for each sector and consortium member, required the compilation and generation of new data for training, adaptation, and evaluation purposes. These data serve as the foundation for building the systems developed for the following final use cases (UC):

- UC1: A Multilingual Speech-Based Driver Assistant
- UC2: Technical Norm Documentation Assistant
- UC3: Software Documentation Assistant
- UC4: PCB Assembling Assistant
- UC5: CNC Voice Assistant
- UC6: Adaptive Machine Translation

Table 1 presents a summary of the six developed use cases along with their main features.

### 3. Use cases and data resources

This section presents the use cases implemented in the ADAPT-IA project, along with the data resources compiled and generated for each. These use cases were defined through collaboration between the

**Table 1**

Overview of the use cases developed in the project and their key features

| UC  | Sector                                     | Main Technology | Maturity         | Test mode | Evaluation   |
|-----|--|-----------------|------------------|-----------|--|
| UC1 | Automotive                                 | Voice assistant | Prototype (TRL4) | Simulated | Qualitative with expert end-users  |
| UC2 | Machine Tool                               | Text assistant  | Prototype (TRL4) | Simulated | Qualitative with expert end-users  |
| UC3 | Energy                                     | Text assistant  | Prototype (TRL4) | Simulated | Qualitative with expert end-users  |
| UC4 | Energy                                     | Text assistant  | Prototype (TRL4) | Simulated | Qualitative with expert end-users  |
| UC5 | Machine Tool                               | Voice assistant | Prototype (TRL4) | Simulated | Qualitative with expert end-users  |
| UC6 | Automotive, Machine Tool, Energy, Railways | Adaptive MT     | System (TRL7)    | Simulated | Quantitative and qualitative with experts in terminology and translation |

technological partners and the industrial clusters, considering the needs and relevant applications identified by the latter.

### 3.1. A Multilingual Speech-Based Driver Assistant

This prototype, developed collaboratively by Vicomtech, UPV/EHU, UZEI and AIC within the automotive domain, consists of a multilingual, speech-based driver assistance system operating in both Basque and English. Integrated into AIC’s industrial vehicle simulator, the assistant enables conversational access to real-time vehicle data, including speed, traffic conditions, tire pressure, and battery status.

To support the system, domain-specific corpora were created for both languages. For Basque, UZEI linguists generated sentence templates covering varied word orders, synonyms, omissions, and linguistic registers. These templates produced over 20 million sentence variations, which were reduced to a diverse and balanced set of 14K sentences using a cluster-based active learning approach [9]. An additional 612 naturally phrased sentences were collected from volunteers to enhance coverage of user intents.

For English, recent advancements in LLMs enabled efficient generation of artificial in-domain data. A small set of manually curated example queries served as seed data, reflecting typical user interactions with a car assistant. Based on these examples, additional utterances were generated using state-of-the-art LLMs (primarily ChatGPT-4o) resulting in a dataset of 1,700 unique and carefully validated sentences. Both corpora will be made publicly available upon project completion.

The assistant comprises three main components: an Automatic Speech Recognition (ASR) system, a Natural Language Processing (NLP) module, and a Text-to-Speech (TTS) unit. The ASR component was based on NVIDIA’s Parakeet models, which uses state-of-the-art recurrent neural network transducers (RNN-T) [10, 11]. Both the Basque and English systems employ the large fast-conformer architecture with 0.6 billion parameters [12]. The Basque model was trained from scratch on 1,258 hours of transcribed audio and fine-tuned with 12.8 hours of synthesized in-domain data generated using Vicomtech’s proprietary TTS system. The English model was adapted from NVIDIA’s pre-trained version and fine-tuned with 7 hours of synthesized in-domain 8,223 user queries.

The Natural Language Understanding (NLU) unit within the NLP component employs language-specific strategies due to the limited availability of robust Basque NLP models. For English, user queries are embedded using the 384-dimensional `all-MiniLM-L12-v2` model from Sentence-Transformers [13]. For Basque, embeddings are extracted using `Latxa 7B` [14], a language-specific LLM. In both cases, multilayer neural networks classify the embeddings into one or more user intents. A rule-based Dialog Manager governs system responses, offering enhanced controllability and robustness in domain-specific settings [15]. Finally, a template-based Natural Language Generation (NLG) module generates text responses, which are synthesized into speech via a proprietary Tacotron-2-based TTS systems [16].

### 3.2. Technical Norm Documentation Assistant

This prototype, developed by Vicomtech with INVEMA is designed to assist technicians with information related to Norms providing the documents where the information appears. The Retrieval-Augmented Generation (RAG) system developed by Vicomtech for INVEMA is a tailored chatbot designed for users

to make queries about the UNE regulations they own about tools and machines. This chatbot allows users to make queries about their regulations and obtain quick and accurate answers.

The RAG system was populated by segmenting the provided regulations using an optimized sentence-window chunking method. These segments were then vectorized with the BAAI/bge-m3c embedding model [17] and stored in a Qdrant vector database [18], managed by LlamaIndex [19] for storage and retrieval. The generative process directly integrated the retrieved contextual data with the user's query, using a locally hosted Phi-4 model [20], a small-scale LM suitable for deployment on systems with limited GPU memory.

A 20-question-answer test dataset, created by INVEMA professionals, was used to assess various chunking strategies and retrieval processes. This dataset allowed both human and LLM evaluation of different RAG approaches. The LLM evaluation used metrics derived from a RAG Triad strategy with the nuclia/REMi-v0 model [21] fine-tuned for RAG evaluation. These methodologies were integrated to deliver the final version of the RAG system.

The INVEMA professionals can access this RAG through a REST API and a Gradio frontend interface [22]. The interface further grants access to documents referenced in the RAG output and allows user feedback for each system response.

### **3.3. Software Documentation Assistant**

Closely resembling the previous prototype, the RAG system developed by Vicomtech for a company of the ACE cluster is a domain-specific chatbot designed to assist users in navigating and understanding product documentation. This system leverages software manuals as its primary knowledge source, enabling efficient and context-aware access to technical information embedded within these documents.

To populate the RAG system, the provided manuals underwent segmentation using an optimized sentence-window chunking methodology. These chunks of the manuals are subsequently vectorized with the BAAI/bge-m3 [17] embedding model and stored in a Qdrant vector database, facilitated by LlamaIndex for both storage and retrieval operations. The generative process involved direct integration of the retrieved contextual data with the user's query, employing a locally hosted Mistral 7B Instruct v0.3 [23] model, a small-scale LM suitable for deployment on systems with limited GPU memory. A 24-question test dataset, developed by Optimotive software professionals, was employed to evaluate various chunking strategies and retrieval processes. This dataset enabled both human and LLM evaluation of different RAG approaches. With respect to the LLM evaluation, some metrics were derived from a RAG Triad strategy [24] with the Nuclia's REMi-v0 model fine-tuned for RAG evaluation. Through the integration of these methodologies, the final version of the RAG system was delivered.

The company professionals can access this RAG through a REST API and a Gradio frontend interface. The interface further grants access to documents referenced in the RAG output and allows user feedback for each system response.

### **3.4. PCB Assembling Assistant**

This prototype, developed by Vicomtech with a company from the ACE cluster, is designed to assist technicians in assembling electrical industrial components through a multimodal guidance system. The primary objective was to develop an intelligent assistant capable of guiding users through the assembly process step by step while allowing flexible navigation —both forward and backward— within the instructions. Additionally, the assistant provides supplementary information upon request, enhancing user support throughout the process.

The system was initially supplied with three separate manuals in Spanish, containing detailed component specifications, sequential assembly instructions, verification procedures, and testing report guidelines, supplemented with numerous interspersed images. However, the original manuals were not structured for direct and natural interaction, requiring manual edition and reorganization to break down excessively long assembly steps into a more user-friendly format.



In order to develop a system capable of assisting technicians through these guidelines, various LLMs were evaluated on their ability to process, interpret, and relay instructional content while maintaining accuracy, step order and displaying images when needed. The models tested include LLaMA 3.1-8B, Phi-4 14B, DeepSeek 7B, and Mistral 7B. Performance was assessed using a test set of approximately 25 queries, structured to include the previous step, the user's query, and the corresponding ground truth. The evaluation metrics applied include BLEU, ROUGE-L, Cosine Similarity, Exact Match, Levenshtein Distance, BERTScore, BLEURT, and G-EVAL.

Different prompting strategies were explored to optimize LLM performance for the procedural guidance task. Experiments included consolidating all manuals as context into a single prompt, splitting them into separate prompts, and varying the level of specificity with which instructions were provided to the model. Initial methods involved minimal instructions, specifying only that the assistant should guide users step by step. More advanced approaches incorporated memory mechanisms, instructing the model to retain the previous step and generate the next accordingly. To support visual guidance, image placeholders with corresponding identifiers were embedded in the system prompt, which were later replaced with actual images when the model referenced them in its step-by-step instructions. Results indicate that while more structured and detailed prompts can improve response accuracy, they also introduce higher latency, and do not always lead to better task execution. Further refinements are being explored to balance response quality and efficiency.

### **3.5. CNC Voice Assistant**

Developed by Vicomtech for a company of INVEMA, this prototype is an advanced solution designed to streamline communication between users and a Computer Numerical Control (CNC) system. It enables seamless interaction to efficiently obtain crucial parameter values, such as the number of machine axes and heads, their positions, maximum speeds, and the range of motion for different axes, among others.

The system integrates multiple components. First, an ASR module transcribes spoken user input into text. This text is then processed by an LLM, which extracts user intent, relevant entities and a corresponding CNC string code. The extracted information is passed to a CNC interface, which acts as a bridge between the system and the CNC machine, retrieving the corresponding data from it. Finally, the retrieved information is presented to the user through a dedicated user interface. Additionally, the prototype includes simulation software provided to replicate CNC operations [25].

The LLM component is based on LLaMA 3.1-8B, which has been specifically instructed to map user input into intents, relevant entities, and CNC string code. For each possible user intent considered in the system, the LLM was provided with examples of both expected input and the corresponding output format and data. To ensure robustness, a validation mechanism was incorporated to verify that the LLM's responses conform to the expected JSON structure and permitted values. The ASR system, implemented using the Kaldi toolkit [26], utilizes a neural network-based architecture with iVector-based speaker adaptation to improve accuracy. For the acoustic and language models, we employed general-purpose pretrained models. Additionally, the language model was fine-tuned with task-specific sentences to enhance transcription accuracy. The CNC interface uses the Windows COM protocol to communicate with the CNC machine simulation software. By utilizing the API functions defined for the system, it retrieves the necessary data. Finally, the User Interface is a Gradio application that enables voice interaction. It includes several text fields where the relevant information is displayed to the user, providing an intuitive and straightforward way to present data retrieved from the CNC system.

### **3.6. Adaptive Machine Translation**

This use case covered all four industrial sectors in the project and had, as its main objective, the development of resources and systems for adaptive Machine Translation (MT), focusing on the Basque-Spanish translation pair in both directions. Three main aspects were tackled within the project: (i) resource creation, (ii) machine translation modelling, and (iii) system deployment. We briefly describe

each aspect in turn below.

Bilingual terminology is a critical component to reach quality translations in specific domains. For the four sectors of the project, namely Energy, Automotive, Railways and Machine Tool, no terminological resources were publicly available at the onset of the project, and part of the effort centred on building new term bases. For this purpose, a corpus was compiled for each industrial sector, which consisted of domain-specific private documents provided by ADAPT-IA participants from each sector and terminological reference resources such as dictionaries in Spanish and/or Basque. Bilingual terminologies were manually crafted and validated according to ISO standards, in particular ISO/TC 37, with around 600 new terminological records per sector. An additional objective was the development of sector-specific evaluation datasets to assess machine translation (MT) quality across four industrial domains. To this end, we selected Spanish-language documents representative of each sector, prioritizing high terminological density and ensuring they were suitable for open sharing with the scientific community. For each test set, our target was to extract a minimum of 50,000 words—corresponding to over 2,000 sentences—to enable robust and meaningful evaluation outcomes. An initial assessment of the shareable evaluation corpora from each sector indicated insufficient volumes of text including relevant terminology overall. We thus complemented the data from two sources. We first selected public Wikipedia entries dedicated to topics from each sector, as they contained significant volumes of relevant terms. Additionally, we generated synthetic text by querying both ChatGPT and DeepSeek to generate text involving a list of provided terms. The selected texts were professionally translated, and data were aligned at both paragraph and sentence levels. The newly created term bases and evaluation datasets form part of the ADAPTIA-MT suite, which will be shared with the scientific community under a CC-BY-NC-ND 4.0 license to foster research and development in adaptive MT for industrial sectors.

To exploit the resources created within the project, we developed several MT models, contrasting Neural Machine Translation (NMT) and LLM-based MT (LLMMT) approaches. The NMT models were adapted for terminologically-sound translation with variants of the approach in [27], where target language terms are injected in the source text wherever the input features matching terms. The models, based on an encoder-decoder Transformer architecture [28], were trained to respond to term injection, identified with specific tags, leading to significant improvements in translation accuracy in both translation directions. LLMMT models were based on Latxa [14], performing adaptation via few-shot in-context learning. We evaluated both approaches on the domain-specific test sets, on standard metrics. In terms of general translation quality, NMT models significantly outperformed LLMMT variants, whereas in terms of terminological accuracy the latter achieved better results. We performed additional human evaluations by experts in terminological translation for this language pair, which also indicated significant preference for NMT translations overall. Further details are provided in [29].

To provide an operative environment, we adapted Vicomtech’s Batua MT ecosystem [30], which provides both text and document translation, the latter via private access. The environment was made accessible to all participants of the project and related companies within the clusters. The system leverages Vicomtech’s proprietary MT solution, Itzuli, which was adapted to support the methods and models for adaptive MT that were developed within the project. Participants could thus translate any content in their respective sectors, contrasting generic translation with terminology-enhanced MT in their domain. Both frontend and backend components of the Batua application were adapted to support terminology integration and selection, as well as model selection. Beyond the previously described controlled evaluations, to assess MT model quality, end-users performed subjective evaluations by translating and reviewing documents in their specific domains, with positive feedback overall.

## 4. Conclusions

In this work, we described ADAPT-IA, a research and development project focused on Adaptive AI technologies applied to multilingual Language Technologies, including Basque, with the goal of promoting their adoption and integration into industrial processes across multiple sectors in the Basque Country. We presented 6 use cases related to the development of advanced prototypes for the

Automotive, Energy, Machine Tool, and Railway sectors, implemented through collaboration between the applied research center Vicomtech, the University of the Basque Country, the Basque Center for Terminology and Lexicography, and the respective industry clusters.

This research has highlighted the importance of applying Language Technologies in these sectors, including the development of technologies for languages such as Basque. Future work will focus on the research and development of Small Language Models, with the aim of assessing their performance in these sectors, which would facilitate their use and deployment. Additionally, the research will be expanded to new sectors, with the goal of further developing new technologies and prototypes, as well as generating domain-specific linguistic resources and making them available to the research community.

## Acknowledgments

This work was partially supported by the Department of Economic Development and Competitiveness of the Basque Government (Spri Group) through funding for the ADAPT-IA project (KK-2023/00035).

## Declaration on Generative AI

During the preparation of this work, the authors used GPT-4 in order to: Grammar and spelling check. After using these tool, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

## References

- [1] A. Aldoseri, K. N. Al-Khalifa, A. M. Hamouda, Ai-powered innovation in digital transformation: Key pillars and industry impact, *Sustainability* 16 (2024) 1790.
- [2] C. Herold, H. Ney, On search strategies for document-level neural machine translation, in: A. Rogers, J. Boyd-Graber, N. Okazaki (Eds.), *Findings of the Association for Computational Linguistics: ACL 2023*, Association for Computational Linguistics, Toronto, Canada, 2023, pp. 12827–12836. URL: <https://aclanthology.org/2023.findings-acl.811/>. doi:10.18653/v1/2023.findings-acl.811.
- [3] S. Jiang, J. Hu, C. L. Magee, J. Luo, Deep learning for technical document classification, *IEEE Transactions on Engineering Management* 71 (2022) 1163–1179.
- [4] D. Song, A. Vold, K. Madan, F. Schilder, Multi-label legal document classification: A deep learning-based approach with label-attention and domain-specific pre-training, *Information Systems* 106 (2022) 101718.
- [5] A. Kumar, B. Starly, “fabner”: information extraction from manufacturing process science domain literature using named entity recognition, *Journal of Intelligent Manufacturing* 33 (2022) 2393–2407.
- [6] A. Borji, A categorical archive of chatgpt failures, *arXiv preprint arXiv:2302.03494* (2023).
- [7] H. Ludwig, T. Schmidt, M. Kühn, Voice user interfaces in manufacturing logistics: a literature review, *International Journal of Speech Technology* 26 (2023) 627–639.
- [8] T. Wang, P. Zheng, S. Li, L. Wang, Multimodal human–robot interaction for human-centric smart manufacturing: a survey, *Advanced Intelligent Systems* 6 (2024) 2300359.
- [9] S. A. Moreno-Acevedo, J. C. Vasquez-Correa, J. M. Martín-Doñas, A. Álvarez, Stream-based active learning for speech emotion recognition via hybrid data selection and continuous learning, in: *International Conference on Text, Speech, and Dialogue*, Springer, 2024, pp. 105–117.
- [10] V. Srivastav, S. Majumdar, N. Koluguri, A. Moumen, S. Gandhi, et al., Open automatic speech recognition leaderboard, [https://huggingface.co/spaces/hf-audio/open\\_asr\\_leaderboard](https://huggingface.co/spaces/hf-audio/open_asr_leaderboard), 2025.



- [11] J. C. Vázquez-Correa, A. Álvarez, H. Arzelus, S. A. Moreno Acevedo, A. González-Docasal, J. M. Martín-Doñas, The vicomtech speech transcription systems for the albayzín 2024 bilingual basque-spanish speech to text (bbs-s2t) challenge, in: Proc. IberSPEECH 2024, 2024, pp. 305–309.
- [12] Nvidia, Parakeet rnnt 0.6b, 2025. <https://huggingface.co/nvidia/parakeet-rnnt-0.6b> [Accessed: 21 May 2025].
- [13] W. Wang, F. Wei, L. Dong, H. Bao, N. Yang, M. Zhou, Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers, 2020. [arXiv:2002.10957](https://arxiv.org/abs/2002.10957).
- [14] J. Etxaniz, O. Sainz, N. Miguel, I. Aldabe, G. Rigau, E. Agirre, A. Ormazabal, M. Artetxe, A. Soroa, Latxa: An open language model and evaluation suite for basque, in: Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 14952–14972.
- [15] A. Vázquez, A. López Zorrilla, J. M. Olaso, M. I. Torres, Dialogue management and language generation for a robust conversational virtual coach: Validation and user study, *Sensors* 23 (2023) 1423.
- [16] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, Y. Wu, Natural tts synthesis by conditioning wavenet on mel spectrogram predictions, in: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2018, pp. 4779–4783. doi:10.1109/ICASSP.2018.8461368.
- [17] J. Chen, S. Xiao, P. Zhang, K. Luo, D. Lian, Z. Liu, Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation, 2024. [arXiv:2402.03216](https://arxiv.org/abs/2402.03216).
- [18] Qdrant, 2025. <https://qdrant.tech/> [Accessed: 21 May 2025].
- [19] Llamaindex, 2025. <https://www.llamaindex.ai/> [Accessed: 21 May 2025].
- [20] Microsoft, phi-4, 2025. <https://huggingface.co/microsoft/phi-4> [Accessed: 21 May 2025].
- [21] Nuclia, Nuclia-eval, 2025. <https://github.com/nuclia/nuclia-eval> [Accessed: 21 May 2025].
- [22] Gradio, 2025. <https://www.gradio.app/> [Accessed: 21 May 2025].
- [23] MistralAI, Mistral-7b-instruct-v0.3, 2025. <https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3> [Accessed: 21 May 2025].
- [24] TruLens, The rag triad, 2022. [https://www.trulens.org/getting\\_started/core\\_concepts/rag\\_triad/](https://www.trulens.org/getting_started/core_concepts/rag_triad/) [Accessed: 21 May 2025].
- [25] F. Automation, 8060/65/70 free simulator, 2022. <https://www.fagorautomation.com/es/documentacion#/descarga/8060-65-70-free-simulator-2/896496> [Accessed: 21 May 2025].
- [26] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz, J. Silovsky, G. Stemmer, K. Vesely, The kaldi speech recognition toolkit, in: IEEE 2011 Workshop on Automatic Speech Recognition and Understanding, IEEE Signal Processing Society, 2011. IEEE Catalog No.: CFP11SRW-USB.
- [27] G. Dinu, P. Mathur, M. Federico, Y. Al-Onaizan, Training neural machine translation to apply terminology constraints, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 3063–3068. URL: <https://aclanthology.org/P19-1294/>. doi:10.18653/v1/P19-1294.
- [28] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in Neural Information Processing Systems* 30 (2017).
- [29] T. Etchegoyhen, H. Gete, B. Arrate, J. Zapirain, V. Ruiz, Machine translation in industrial domains: Resources and evaluations, in: Proc. SEPLN 2025, 2025.
- [30] Vicomtech, Batua.eus itzultzailea, 2025. <https://www.batua.eus> [Accessed: 21 May 2025].