

FACULDADE DE ENGENHARIA DA UNIVERSIDADE DO PORTO

Alumni Multidimensional Job Analytics with AI agents

Carlos Eduardo Coelho Verissimo



Mestrado em Engenharia Informática e Computação

Supervisor: Prof. João Carlos Pascoal Faria

Second Supervisor: Eng. Rui de Sousa Neves

July 29, 2025

Alumni Multidimensional Job Analytics with AI agents

Carlos Eduardo Coelho Verissimo

Mestrado em Engenharia Informática e Computação

Approved in oral examination by the committee:

President: Prof. Gabriel de Sousa Torcato David

Referee: Prof. Luis Borges Gouveia

July 29, 2025

Resumo

Compreender e acompanhar as trajetórias profissionais dos alumni é cada vez mais importante para instituições académicas, estudantes e empregadores, mas os dados integrados e acionáveis continuam escassos. Esta dissertação apresenta o desenvolvimento da *Alumni-FEUP*, uma plataforma que responde a esta lacuna ao analisar os percursos de alumni dos cursos de Engenharia Informática, Engenharia de Software e outros cursos geridos pelo Departamento de Engenharia Informática da FEUP.

O sistema integra recolha de dados, enriquecimento baseado em IA e visualização interativa, oferecendo uma análise multidimensional dos destinos profissionais. Destaca-se pela utilização de agentes de IA dedicados à classificação de funções profissionais, estimativa de níveis de senioridade e resolução de dados geográficos ambíguos, recorrendo a modelos de linguagem de grande escala e pipelines com recuperação aumentada, assegurando adaptabilidade num mercado de trabalho em constante evolução. A plataforma suporta análise sobre várias dimensões, nomeadamente funções profissionais, senioridade, distribuição geográfica, percurso académico e evolução temporal das carreiras.

A validação combina abordagens quantitativas e qualitativas. As métricas quantitativas incluem a proporção de alumni abrangidos, a precisão das classificações de funções, senioridade e localização (validadas por feedback dos utilizadores e revisão manual), e taxas de adesão dos utilizadores à plataforma. A análise qualitativa baseia-se em feedback aberto dos utilizadores e respostas à escala de usabilidade do sistema (SUS) recolhidas numa implementação real. Os resultados demonstram um forte envolvimento dos utilizadores, contributos relevantes para o alinhamento curricular e validação da eficácia do pipeline de enriquecimento. O projeto é de acesso público, open source, e pretende servir de modelo para iniciativas semelhantes noutras instituições.

Abstract

Understanding and tracking alumni career trajectories is increasingly important for academic institutions, students, and employers, yet integrated, actionable data remains scarce. This thesis presents the development of *Alumni-FEUP*, a platform designed to bridge this gap by analyzing the career outcomes of alumni from Informatics Engineering, Software Engineering, and other degrees managed by the Informatics Engineering Department at FEUP.

The system integrates data collection, AI-driven enrichment, and interactive visualization to provide a multidimensional perspective on alumni outcomes. Key features include the use of dedicated AI agents to classify professional roles, estimate seniority levels, and resolve ambiguous location data, leveraging large language models and retrieval-augmented pipelines for adaptability in a dynamic job market. The platform supports analysis across several dimensions, such as professional roles, seniority, geographic distribution, academic background, and temporal career progression.

Validation combines both quantitative and qualitative approaches. Quantitative metrics include the proportion of alumni covered, automated classification accuracy for job roles, seniority, and location (validated by user feedback and manual review), and user engagement rates on the platform. Qualitative analysis draws on open-ended user feedback and System Usability Scale (SUS) responses from a live deployment. Results indicate strong engagement, meaningful insights for curriculum alignment, and validation of the enrichment pipeline's effectiveness. The project is publicly accessible, open source, and intended as a model for similar institutional initiatives.

UN Sustainable Development Goals

The United Nations Sustainable Development Goals (SDGs) provide a global framework to achieve a better and more sustainable future for all. It includes 17 goals to address the world's most pressing challenges, including poverty, inequality, climate change, environmental degradation, peace, and justice.

The Alumni-FEUP platform contributes to several of these goals, demonstrating its broader societal impact beyond academic research.

SDG 4 Ensure inclusive and equitable quality education and promote lifelong learning opportunities for all

SDG 8 Promote sustained, inclusive and sustainable economic growth, full and productive employment and decent work for all

SDG 9 Build resilient infrastructure, promote inclusive and sustainable industrialization and foster innovation

SDG 17 Strengthen the means of implementation and revitalize the Global Partnership for Sustainable Development

SDG	Target	Contribution	Performance Indicators and Metrics
4	4.4	Supports universities in analyzing alumni outcomes to refine curricula, improve educational quality, and validate program effectiveness.	Number of institutions using the platform; curriculum updates informed by data; % of alumni working in-field.
8	8.5	Enables mapping of alumni across 52 countries and 164 industries, supporting job market transparency, talent discovery, and workforce planning.	Geographic and industry diversity metrics; alumni employment rate; number of platform users exploring career data.
	8.6	Helps bridge the education-to-employment gap by surfacing trends in youth employment and early-career trajectories.	Increase in young alumni employment rate; reduction in mismatch between studies and career.
9	9.5	Tracks emerging job roles and skill requirements using AI agents, contributing to data-driven innovation in ICT and education infrastructure.	Number of roles classified by AI; rate of taxonomy updates; adoption of platform insights in institutional reports.
17	17.6	Open-source design and detailed documentation foster collaboration and reuse across academia and industry.	Number of external institutions adopting the platform; GitHub stars/forks; citations or reuse in other projects.

Acknowledgments

This thesis marks the end of a long and demanding chapter.

Balancing professional responsibilities with academic deadlines was a constant and demanding challenge, and completing this project would not have been possible without the support of the people around me.

To my whole family, thank you for your patience, love, and unwavering encouragement. Your support grounded me, even on the most overwhelming days.

To my friends, França, Nuno, and Tomás, thank you for the late-night talks and laughs, the distractions when I needed them the most, and for reminding me that I was never alone in this.

To my supervisor, Prof. João Faria, and co-supervisor, Eng. Rui Neves, I'm deeply grateful for your guidance, trust, and flexibility throughout this process. Thank you for giving me the space to explore what I believed mattered, and for pushing me to take this project further.

To those who listened, checked in, or simply crossed my path when I needed it the most,

Thank you :))

Carlos Veríssimo

Contents

1	Introduction	1
1.1	Context and Motivation	1
1.2	Problem	2
1.3	Objectives	2
1.4	Research questions and Hypothesis	3
1.5	Structure	4
2	Literature Review Methodology	5
2.1	Methodology	5
2.2	Identification Stage	6
2.2.1	Search Query Refinement	7
2.2.2	Final Search Queries and Results	8
2.3	Screening Stage	9
2.4	Eligibility Stage	10
2.4.1	Full Text Availability	10
2.4.2	Analysis of Eligible Documents	10
2.5	Inclusion Stage	11
2.6	Additional Resources	14
2.6.1	Standards for Job Title and Seniority Classification	14
2.6.2	Initiatives for Tracking Alumni Outcomes	18
2.7	Discussion and Summary	19
3	State of the Art Analysis	20
3.1	Job Title Classification	20
3.1.1	Overview of Methods	21
3.1.2	Standards and Frameworks	23
3.1.3	Challenges	25
3.1.4	Research Gaps	26
3.2	Alumni Tracking and Analytics Initiatives	28
3.2.1	Existing Platforms	28
3.2.2	Key Features and Innovation	29
3.2.3	Limitations of Current Approaches	30
3.2.4	Relevance to Research	31
3.3	Visualization and Data Integration of Multidimensional Data	32
3.3.1	Techniques and Tools	33
3.4	Critical Analysis and Synthesis	34
3.5	Conclusion	36

4	Alumni-FEUP: A Novel Approach to Alumni Career Tracking	37
4.1	Overview of the Solution	37
4.2	System Architecture	38
4.3	Key Components and Technologies	39
4.3.1	Frontend	39
4.3.2	Backend API	41
4.3.3	Database and Storage	42
4.3.4	Orchestration Service	43
4.4	Implementation Details	43
4.4.1	CSV Upload and Manual Form Submission	44
4.4.2	Enrichment Orchestration Overview	46
4.4.3	Alumni Profile Extraction	47
4.4.4	Enrichment Services	48
4.4.5	LinkedIn-Based Authentication	49
4.5	AI-Driven Agents	50
4.5.1	Overview of Agent Architecture	50
4.5.2	Seniority Agent	53
4.5.3	Location Agent	57
4.6	Job Classification Agent	60
4.6.1	Motivation and Role of the Agent	61
4.6.2	Leveraging the ESCO Taxonomy	61
4.6.3	Architecture and Workflow	63
4.7	Challenges and Design Decisions	67
5	Results and Discussion	68
5.1	Data Coverage and Composition	68
5.2	Associated Costs	72
5.3	Evaluation Methodology and Experimental Setup	73
5.4	Quantitative Results	74
5.4.1	User-Submitted Validation	74
5.4.2	LLM-based Confidence Scores	76
5.4.3	Manual Evaluation	79
5.4.4	Form Feedback on Classification Quality	83
5.4.5	Telemetry Data	84
5.5	Qualitative Results	84
5.5.1	Participant Overview	84
5.5.2	System Usability Scale (SUS) Feedback	85
5.5.3	Analytics Feedback	86
5.5.4	Open-Ended Feedback	88
5.5.5	Profile Experience Feedback	89
5.6	Discussion	90
6	Conclusions and Future Work	93
6.1	Summary of Contributions	93
6.2	Main Findings	94
6.3	Limitations and Open Issues	95
6.4	Future Work	98
6.5	Final Remarks	100

References	101
A Architecture	107
A.1 Job Classification Agent - Confidence Calculation	107
B Feedback Form and Results	109
B.1 Feedback Form	109
B.2 Results	121
B.2.1 Job Classification Confidence Scores	121
B.2.2 Real Misclassification Cases	122

List of Figures

1.1	Previous version of the platform, showing the original map-based visualization	2
2.1	PRISMA flow diagram	6
2.2	Document retrieval percentage, per database	10
2.3	Document distribution by category	11
2.4	Radford Global Job Leveling	15
2.5	European ICT Professional Role Profiles	17
2.6	SNOMED-CT Example - Orthopedic Surgeon	17
2.7	UBC Global Alumni Impact Map	18
4.1	High-level system architecture of Alumni-FEUP	38
4.2	Data flow in the Alumni-FEUP platform	39
4.3	Analytics page in the Alumni-FEUP platform	40
4.4	Alumni profile page in the Alumni-FEUP platform	41
4.5	Database schema of the Alumni-FEUP platform (simplified)	42
4.6	Admin tooling interface for uploading alumni CSV data	44
4.7	Manual submission form for alumni data	46
4.8	Shared agent architecture diagram	53
4.9	Edit seniority modal	55
4.10	Seniority agent architecture diagram	57
4.11	Location agent architecture	60
4.12	Logic used to determine valid leaf-level ESCO nodes	62
4.13	ESCO grouping analytics	63
4.14	Edit Classification modal showing alternatives, confidence, and LLM reasoning	65
4.15	End-to-end job classification workflow and architecture	66
5.1	Number of graduates per supported degree	69
5.2	LinkedIn profile availability across all graduates	69
5.3	Source distribution of the 1 353 LinkedIn URLs used for enrichment	70
5.4	LinkedIn profile availability by degree	70
5.5	Role validation modal displayed in user profiles	74
5.6	Distribution of participants by affiliation with FEUP	85
5.7	Distribution of System Usability responses	86

List of Tables

2.1	Search queries - first iteration	7
2.2	Final search queries	8
2.3	Document retrieval count, per database and search query	9
2.4	Inclusion and exclusion criteria	9
2.5	Document Retrieved for Data Visualization and Data Integration	11
2.6	Documents retrieved for Job Title Classification	12
2.7	Documents Retrieved for Alumni Tracking and Analytics	13
3.1	Job Title Classification Systems Comparison	34
3.2	Alumni Tracking Platforms Comparison	34
3.3	Data Visualization Approaches Comparison	35
4.1	Predefined seniority levels used by the agent	55
4.2	Comparison of role classification counts and ESCO node availability by level . . .	63
5.1	Summary of extracted alumni career data	71
5.2	Top 10 ESCO classifications represented in the dataset (all historical roles)	71
5.3	Top 10 ESCO classifications for alumni's current (latest) roles	72
5.4	Estimated cost breakdown of external services	73
5.5	Estimated OpenAI cost breakdown for the enrichment agents	73
5.6	User validation outcomes for AI-generated job title classifications	75
5.7	Validation accuracy by ESCO classification level	75
5.8	Validation outcomes for seniority classification across different levels	76
5.9	Interpretation of job classification confidence scores	77
5.10	Job classification confidence by ISCO Sub-Major Group (top 10 by count)	78
5.11	Distribution of LLM-based seniority classification confidence by level	78
5.12	Joint distribution of roles by seniority and confidence bucket	80
5.13	Five most and least common ESCO job titles in the manual evaluation sample . .	80
5.14	ESCO classification accuracy for the 10 most common roles in the evaluation sample	81
5.15	ESCO classification accuracy for the 10 least common roles in the evaluation sample	82
5.16	Job classification (ESCO) accuracy by seniority level in the manual evaluation subset	82
5.17	Seniority classification accuracy by level in the manual evaluation subset	83
5.18	Real Experience Score and Core Web Vitals	84
5.19	Distribution of participants by degree	85
5.20	Most commonly used filters on the analytics page	87
5.21	Perceived layout of the analytics page	89
5.22	Notable user suggestions and reported issues	89
5.23	User feedback on profile and classification experience (n=8)	90

B.1 Job classification confidence scores for all ISCO sub-major groups 121

List of Acronyms

AI	Artificial Intelligence
CNN	Convolutional Neural Network
DSSM	Deep Semantic Structure Algorithm
ESCO	European Skills, Competences, Qualifications, and Occupations
FEUP	Faculdade de Engenharia da Universidade do Porto
ISCO	International Standard Classification of Occupations
LEIC	Licenciatura em Engenharia Informática e Computação (pre-Bologna)
L.EIC	Licenciatura em Engenharia Informática e Computação (post-Bologna)
LLM	Large Language Model
MEI	Mestrado em Engenharia Informática
M.EIC	Mestrado em Engenharia Informática e Computação
MESW	Mestrado em Engenharia de Software
MIEIC	Mestrado Integrado em Engenharia Informática e Computação
ML	Machine Learning
NLP	Natural Language Processing
PRISMA	Preferred Reporting Items for Systematic Reviews and Meta-Analyses
PDMD	Programa Doutoral em Media Digitais
SNOMED	Systematized Nomenclature of Medicine
SOC	Standard Occupational Classification
SUS	System Usability Scale
SVM	Support Vector Machines
TF-IDF	Term Frequency-Inverse Document Frequency
WoS	Web of Science

Chapter 1

Introduction

1.1 Context and Motivation

Alumni networks play a crucial role in strengthening university communities worldwide, supporting lifelong connections, professional development, and institutional reputation. Increasingly, universities recognize the value of tracking alumni outcomes, both to improve academic programs and to demonstrate their broader societal impact. However, gathering and integrating this information remains a significant challenge for most institutions.

The proposal is motivated by the celebration of 30 years since the creation of the Bachelor of Informatics and Computing Engineering (LEIC) at FEUP in 1994. A platform, AlumniEI World¹, was developed to showcase the distribution of alumni from different programs (LEIC, MIEIC, L.EIC, M.EIC) on a world map, based on information extracted from their LinkedIn profiles and cross-checked with official data from FEUP's SIGARRA system.

Currently, the platform includes a timeline of historical events and a world map showing alumni distribution, but these components operate independently. Additionally, while statistical information about alumni's professional roles has been collected, it remains disconnected from both the spatial and temporal data. The challenge lies in integrating these disparate types of data (categorical, spatial, and temporal) into a cohesive and future-proof platform.

The platform is meant to celebrate a significant milestone for the program while offering a rich, interactive tool to visualize the career paths and geographical distribution of FEUP's alumni over time. By bridging the existing gaps between the timeline, map, and professional data, the proposal aims to provide a more holistic understanding of alumni career trends, including how their roles have evolved across industries, locations, and time.

This comprehensive approach will serve multiple stakeholders:

- **Academic Institutions:** Understanding graduate outcomes and program effectiveness
- **Current Students:** Exploring potential career paths and networking opportunities
- **Industry Partners:** Tracking technological trends and talent distribution

¹ Available at <https://eic30anos.fe.up.pt/alumniEIworld/>

- **Alumni:** Connecting with peers and discovering career opportunities

1.2 Problem

The core challenge is threefold:

First, there is no unified system for classifying and organizing the diverse job titles and seniority levels within the alumni dataset.

Second, the platform currently displays static maps and timelines of events without meaningful integration. The inability to correlate geographical movements with career progression and historical events limits the platform's utility for understanding alumni career trajectories.

Third, the data pipeline requires significant manual oversight, from initial data extraction to final visualization. The process of aligning LinkedIn data with SIGARRA's internal system lacks automation, making it difficult to maintain current information and scale the platform effectively.

These limitations restrict the development of an intuitive platform for exploring and understanding the career paths and professional landscapes of FEUP's alumni (see [Figure 1.1](#)).




Figure 1.1: Previous version of the platform, showing the original map-based visualization

1.3 Objectives

The main objectives of this dissertation are:

- To develop an innovative platform that supports interactive, multidimensional analysis and visualization of alumni career trajectories, leveraging concepts from online analytical processing (OLAP) to set new standards in the academic community
- To design an automated and resilient data ingestion and enrichment pipeline for processing, classifying, and updating alumni records from diverse data sources, including LinkedIn profiles, SIGARRA data, and other sources

- To create advanced AI-based systems for automated classification of job titles and seniority levels that can adapt to emerging trends
- To implement interactive visualizations that effectively integrate spatial, temporal, and categorical data for comprehensive exploration of alumni careers

These objectives aim to provide a robust foundation for educational institutions seeking to analyze and understand the professional outcomes of their alumni.

1.4 Research questions and Hypothesis

There are two main research domains that this thesis aims to address: the novel and innovative aspects of the platform, and the automated classification of job titles and seniority levels.

- **RQ1:** Does a similar platform exist? If so, how can this one differentiate itself to set a standard within the academic community and inspire initiatives across universities and industries?
- **RQ2:** How can the platform remain innovative and relevant within the rapidly evolving tech industry, where trends and technologies continuously change?
- **RQ3:** Is it feasible to build a system from scratch that accurately classifies job titles and seniority levels, while also being resilient to emerging and unforeseen trends?
- **RQ4:** Which interactive visualizations most effectively enhance user engagement with the alumni data?

In the second domain, significant effort will be devoted to investigating whether current job title classification methods can effectively adapt to the constantly evolving landscape of job titles and roles, as well as exploring potential improvements. An example of such an approach is the one proposed by Javed et al. [20], titled *Towards a Job Title Classification System*.

The following hypotheses are thus proposed to address the research questions raised above:

- **H1:** An approach combining LLMs with traditional machine learning can accurately classify job titles and seniority levels while maintaining adaptability to emerging roles
- **H2:** The classification methods can be integrated into a fully automated data pipeline requiring minimal human oversight
- **H3:** Alumni-FEUP can serve as a foundational platform for institutions to track career trajectories and foster enhanced networking

1.5 Structure

The remainder of this document is organized as follows:

- **Chapter 2** describes the methodology for the systematic literature review, including paper selection criteria, classification strategy, and review scope.
- **Chapter 3** analyzes the state of the art across job classification, alumni tracking platforms, and data visualization techniques, applying the systematic review methodology described in Chapter 2.
- **Chapter 4** details the system design and implementation, covering the data enrichment pipeline, backend infrastructure, AI agents, and frontend architecture.
- **Chapter 5** presents the evaluation methodology and results, including both quantitative and qualitative assessments of the platform.
- **Chapter 6** concludes the document with a summary of contributions, discussion of limitations, future work directions, and the platform's sustainable development impact.

Chapter 2

Literature Review Methodology

Understanding alumni career paths and job title classification is increasingly important for both educational institutions and employers seeking insights into workforce trends. Platforms like Alumni-FEUP aim to reveal where graduates work, what roles they fill, and how skills align with industry needs.

Alumni networks have long supported career development, but existing platforms often lack detailed analytics on professional trajectories. This chapter assesses whether current solutions already cover the proposed scope of Alumni-FEUP, which integrates richer data insights for universities and employers.

Recent advances in AI and machine learning have enabled automated classification of job titles from sources like LinkedIn, deepening our understanding of career progressions and graduate competencies. Similarly, student outcome tracking is now a priority for institutions demonstrating program effectiveness and alumni success.

This literature review follows a dual focus: first, surveying research on job title classification methods and frameworks; second, analyzing initiatives and platforms dedicated to tracking student outcomes, their methodologies, achievements, and challenges.

2.1 Methodology

This literature review uses the PRISMA (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) framework [38] to ensure a systematic, unbiased, and replicable approach. Given the rapid developments in machine learning, especially with large language models (LLMs), such rigor is essential.

Each PRISMA stage is detailed in the following sections, with the overall process summarized in [Figure 2.1](#).




Figure 2.1: PRISMA flow diagram

2.2 Identification Stage

The identification stage, the first step in the PRISMA framework, involves selecting relevant resources using targeted search queries across four databases: IEEE Xplore, SCOPUS, ScienceDirect, and the ACM Digital Library.

For job title classification, three queries were designed to capture papers, articles, and reports. For alumni tracking and analytics platforms, a broader query was used to include sources like websites, tools, and systems.

Initial queries aimed for broad coverage and were refined iteratively (see Table 2.1). The resulting document counts and their distribution across databases are shown in Table 2.3.

The first search query targets job title classification methods, accounting for both British and American spellings of "categorization." The second and third queries focus on interactive data visualization platforms and spatiotemporal data integration. The final query retrieves work on alumni tracking and analytics platforms.

ID	Search Query
SQ1	("job position") AND (classification OR categori?ation) AND ("machine learning")
SQ2	("interactive visuali?ation" OR dashboard) AND ("multidimensional data" OR "multivariate data") AND ("web-based" OR "business intelligence")
SQ3	("spatial data") AND ("temporal data" OR "time series" OR "historical data") AND ("data integration")
SQ4	("alumni tracking" OR "alumni analytics") AND (platform OR system OR tool)

Table 2.1: Search queries - first iteration

2.2.1 Search Query Refinement

Query refinement was performed on SCOPUS by searching within titles, abstracts, and keywords, mainly by analyzing result counts. Titles were reviewed as needed to ensure relevance. Further details on identification stage filters are provided in [subsection 2.2.2](#).

Refinement of SQ1

The initial SQ1 query returned only 13 results, too few for meaningful analysis. Expanding it to include "LLM," "language model," and alternate terms for "title" (such as "occupation" and "position") increased the results to 27. The final query used was:

("job position" OR "job title" OR "job occupation") AND (classification OR categori?ation) AND ("machine learning" OR LLM OR "language model").

Refinement of SQ2

The initial SQ2 query yielded 11 documents. It was refined by adding "data visualization" to the visualization terms and "online platform" to the platform terms, broadening the scope for multidimensional data tools. The refined query produced 26 results:

("interactive visuali?ation" OR dashboard OR "data visuali?ation") AND ("multidimensional data" OR "multivariate data") AND ("web-based" OR "online platform" OR "business intelligence").

Refinement of SQ3

SQ3 originally returned 19 documents. The query was expanded by including "geographic" and "location" in addition to "spatial" for the geographic aspect, increasing relevant results. The other sections were unchanged. The refined query returned 23 documents:

("spatial data" OR "geographic data" OR "location data") AND ("temporal data" OR "time series" OR "historical data") AND ("data integration").

Refinement of SQ4

The initial SQ4 query, targeting alumni tracking and analytics platforms, returned only 8 results on SCOPUS. To broaden coverage, keywords like "graduate," "career," and "employment" were

added, along with "application" and the wildcard operator ("*") to capture related terms. The refined query returned 17 documents.

2.2.2 Final Search Queries and Results

After the refinement process, the final search queries are as follows:

ID	Search Query
SQ1	("job title" OR "job occupation" OR "job position") AND ("classification" OR "categori?ation") AND ("machine learning" OR "LLM" OR "language model")
SQ2	("interactive visuali?ation" OR dashboard OR "data visuali?ation") AND ("multidimensional data" OR "multivariate data") AND ("web-based" OR "online platform" OR "business intelligence")
SQ3	("spatial data" OR "geographic data" OR "location data") AND ("temporal data" OR "time series" OR "historical data") AND ("data integration")
SQ4	("alumni tracking" OR "alumni analy*" OR "graduate tracking" OR "career tracking" OR "employment tracking") AND ("platform" OR "system" OR "tool" OR "application")

Table 2.2: Final search queries

Note that ScienceDirect does not support the "*" wildcard, so SQ4 used "alumni analytics" instead of "alumni analy*".

Querying and filtering options varied by database, but a standardized approach was used wherever possible.

All databases were filtered for documents published from 2016 onward, to focus on recent advances in the field and avoid overrepresenting very recent LLM literature. Filters were applied as follows:

- **ACM Digital Library:** 2016+ research articles - searching within all fields.
- **IEEE Xplore:** 2016+ documents - searching within all fields
- **ScienceDirect:** 2016+ research/review articles in Computer Science, Engineering, Decision Sciences.
- **SCOPUS:** 2016+ conference papers and articles in Computer Science, Engineering, Decision Sciences

Executing these queries across the four databases, with the described filters, retrieved 1,089 documents (including duplicates) and their distribution by database is shown in [Table 2.3](#) and [Figure 2.2](#).

Additionally, three documents were manually included in the total count; these consist of older articles not captured by the search queries but considered relevant to the topic, as well as documents discovered through references in the included papers. These will be discussed in more detail in [section 2.5](#).

Database	SQ1	SQ2	SQ3	SQ4	Total
ACM Digital Library	168	28	73	11	280
IEEE Xplore	49	57	9	6	121
ScienceDirect	205	90	281	19	595
SCOPUS	27	26	23	17	93
Total	449	201	386	53	1089

Table 2.3: Document retrieval count, per database and search query

The next section discusses the screening process, which includes the deduplication of records, as well as the clear definition of the inclusion and exclusion criteria.

2.3 Screening Stage

The screening stage, the second step in PRISMA, applies inclusion and exclusion criteria to all retrieved documents.

A total of 1,092 documents were imported into EndNote for screening. Due to the broad project scope and varied topics, defining the criteria was challenging.

Table 2.4 summarizes the eight screening phases used.

Criteria	Phase
Retrieved from online databases using search queries	1
Published in 2016 or later	2
Not a duplicate	3
Full text available	4
Published in English	5
Book or Conference Proceedings/Paper or Journal Article	6
High citation count or published in reputable sources	7
Title and/or abstract and/or keywords are pertinent to the domain of interest	8

Table 2.4: Inclusion and exclusion criteria

These criteria ensure relevance, quality, and accessibility. Documents had to be from targeted queries, published since 2016, unique, in English, and available in full text. Books, conference papers, and journal articles were prioritized for academic rigor, while citation count and reputable sources enhanced reliability. Relevance was checked via titles, abstracts, and keywords.

Duplicates were removed using partial title matching, reducing the initial set from 1,092 to 1,054 documents. Screening by the defined criteria (Table 2.4) excluded 945 documents by title in the first pass, leaving 109 for further abstract/keyword review.




Figure 2.2: Document retrieval percentage, per database

2.4 Eligibility Stage

The eligibility stage, the third step in PRISMA, involved in-depth review of the previously screened documents to assess relevance and full-text availability. Titles, abstracts, and keywords were evaluated against the criteria, resulting in 36 documents eligible for full-text screening. Reasons for exclusion are shown in the PRISMA flow diagram (Figure 2.1).

2.4.1 Full Text Availability

Using *EndNote*'s in-built tools for full-text searching, 8 PDF documents and 22 URLs were found out of the 36 documents. The remaining documents were manually retrieved from other sources, and a full-text article was linked to all 36 documents.

2.4.2 Analysis of Eligible Documents

A brief analysis of the documents' contents was then conducted to determine their relevance to the research questions, as well as to identify potential issues such as incorrect year of publication, mislabeled type and/or subject area.

At this stage, 1 document, published in 2014, was excluded. This document had been manually added from the reference list included in this dissertations' appendix, before the inclusion and exclusion criteria had been finalized.

Another article titled "Using Machine Learning for Labour Market Intelligence" [5] by Boselli et al. was marked as a duplicate of "Classifying online Job Advertisements through Machine Learning" [4] by the same authors.

Thus, the final number of documents eligible for inclusion in the analysis is 34.

2.5 Inclusion Stage

At this stage, full-text documents were analyzed in depth, focusing on their methods, results, and conclusions.

For each document, the following was annotated: authors and year, subject area, main research questions, and any machine learning or AI techniques used.

The annotated documents were split into three categories:

- **Category 1** - Job title classification
- **Category 2** - Alumni tracking and analytics platforms.
- **Category 3** - Visualization and data integration of spatiotemporal/multidimensional data.

Figure 2.3 shows the distribution of the documents across the three categories.




Figure 2.3: Document distribution by category

Details for each document are provided in Tables 2.5 through 2.7.

Purpose and Goals	Approach	Reference
Tackle challenges in Big Data analysis, focusing on measuring, displaying multidimensional data patterns for effective visualization	2D-Ratio algorithm, parallel axes, and clustering to manage complex, over-crowded data patterns	[57]

Table 2.5: Document Retrieved for Data Visualization and Data Integration

Purpose and Goals	Approach	Reference
Create a job recommendation system for efficient jobseeker-position matching	Machine learning classification algorithms	[45]
Improve job title classification based on descriptions	BERT-based model, multi-label classification	[26]
Develop system for resume classification and personality prediction	TF-IDF vectorization, Random Forest	[35]
Streamline candidate screening through resume analysis	NLP with cosine similarity	[24]
Enhance understanding of job demand across industries	SVM and Random Forest	[9]
Monitor real-time labor market for EU policies	Machine learning-based method	[4]
Enhance job title classification for online recruitment	CNN trained on 15M job posts	[54]
Address resume classification challenges	Feedforward NN with USE embeddings	[46]
Improve employer industry classification	SVM and Gradient Boosted Trees	[13]
Address information overload in recommendations	Deep Semantic Structure Algorithm	[28]
Extract structured labor market info using GenAI	Amazon Bedrock foundation models	[15]
Analyze Big Data profession skill requirements	Semi-automated ML with human judgment	[10]
Enhance skills matching between resumes and vacancies	Machine learning algorithms	[3]
Enhance semi-supervised classification system	Proximity-based classification	[20]
Improve job matching through title prediction	ML techniques for job title predictions	[40]
Build alumni database from LinkedIn data	Python web scraping of public profiles	[23]
Predict job titles from descriptions	Neural networks with pre-trained models	[50]
Develop two-stage job title identification	BERT + unsupervised ML	[39]
Create automatic text classification system	TF-IDF with ML and CNNs	[41]
Enable university alumni career tracking	LinkedIn data classification with SOC	[25]
Automate occupation coding for statistics	Comparative algorithm analysis	[44]
Standardize industry classifications	Data-mining methodology	[1]
Develop multilingual job market language model	Domain-adaptive pre-training, XLM-R	[58]

Table 2.6: Documents retrieved for Job Title Classification

Purpose and Goals	Approach	Reference
Examine how enhancing data quality affects alumni tracking dashboards	Conducted systematic review of data cleaning methodologies and implemented quality assessment frameworks for alumni tracking systems	[21]
Investigates the conversion rates of graduates to higher degrees and their employability at a Malaysian University	Employed data analysis techniques, including machine learning and time series analysis, to predict trends and evaluate outcomes.	[42]
Analyze the career trajectories of CEOs to identify unique characteristics contributing to their success.	Encoding job titles and companies with the SOC scheme. Needleman-Wunsch algorithm for optimal matching	[16]
Enhance alumni data utilization by developing an information analysis model that uses big data for alumni information mining and retrieval	Hierarchical clustering algorithm using ant colony optimization for distance calculations	[55]
Enhance alumni monitoring by integrating alumni data for improved accreditation and institutional development	Data warehouse using a normalized data store architecture, <i>Power BI</i> for visualization	[56]
Build alumni database from LinkedIn data	Python web scraping of public profiles	[23]
Implement a predictive model to assess career success through alumni tracking	Genetic algorithms to predict career success based on objective and subjective variables	[36]
Develop a web-based graduate tracer information system to streamline data collection and analysis for alumni	Rapid Application Development (RAD) model, HTML, Bootstrap, JavaScript, PHP, and MySQL	[30]
Enable universities to track alumni career progress efficiently using LinkedIn data	Developed an automated system to retrieve and classify LinkedIn job titles using SOC categories	[25]
The need for a robust online platform to enhance connectivity between colleges and their alumni	Development of a multifunctional application that encompasses many alumni-related features	[22]
Track alumni's post-graduate outcomes	Excel-based tracking tool to gather alumni data and improve engagement	[7]

Table 2.7: Documents Retrieved for Alumni Tracking and Analytics

Please note that two documents overlap between Category 1 (Job Title Classification) and Category 2 (Alumni Tracking):

- *Retrieving and Classifying LinkedIn Job Titles for Alumni Career Analysis* [25] by Li et al.
- *LinkedIn Alumni Profile Data Extraction* [23] by Lade et al.

Both studies focus on extracting and classifying LinkedIn profile data to analyze alumni career trajectories, demonstrating the natural intersection between job classification and alumni tracking domains.

2.6 Additional Resources

The following sections discuss various standards for job title and seniority classification from both Europe and the United States. Additionally, there is a summary of initiatives by universities and professional organizations aimed at tracking student and alumni outcomes.

2.6.1 Standards for Job Title and Seniority Classification

Standard Occupational Classification (SOC)

The Standard Occupational Classification (SOC) [51] is a U.S. federal standard for grouping workers into occupational categories. The 2018 SOC organizes jobs hierarchically:

- 23 major groups
- 98 minor groups
- 459 broad occupations
- 867 detailed occupations

SOC groups jobs by similar responsibilities, skills, education, and training. Each worker is classified into a detailed occupation based on their primary tasks.

The O*NET system [32] extends SOC by providing a richer database of occupational details and requirements.

O*NET-SOC

The 2019 O*NET-SOC, which incorporates the 2018 SOC taxonomy, includes 1,016 occupational titles, 923 of which represent O*NET data-level occupations. These are granular subdivisions of the SOC codes.

Skills Framework for the Information Age (SFIA)

SFIA [48] is a global benchmark for defining skills and competencies in ICT, software engineering, and digital transformation roles. It offers a structured way for organizations to assess and map professional skill sets.

Although not a job classification standard, SFIA is valuable for linking job roles to levels of responsibility and seniority, with seven levels ranging from Level 1 (Follow) to Level 7 (Set strategy, inspire, and mobilize).

Radford Global Job Leveling

Radford Global Job Leveling [2], developed by Aon Corporation, provides a consistent framework for categorizing job roles while allowing regional adaptation. Key features include:

- Clear job architecture (codes, titles, and functions) for structured career paths and salary frameworks
- Flexibility to accommodate local market demands and regulations
- Benchmarking to ensure competitive, equitable pay
- Dual career tracks, supporting both management and technical progression

Figure 2.4 shows the system is split between management and individual contributor tracks, supporting market alignment and clarity within organizations.

Figure 5

Sample Career Ladders within a Global Job Leveling System

Business Leadership	Executive	Management	7	Vice President			
	Grade 11		6	Sr. Director			
	Grade 10		5	Director	Professional	6	Principal
Management and Professional	Grade 9		4	Sr. Manager		5	Expert
	Grade 8		3	Manager		4	Advanced
	Grade 7		2	Sr. Supervisor		3	Career
	Grade 6		1	Supervisor		2	Developing
						1	Entry
Entry Level and Support	Grade 5				Support	5	Specialist
	Grade 4					4	Highly Skilled
	Grade 3					3	Senior
	Grade 2					2	Intermediate
						1	Entry

Figure 2.4: Radford Global Job Leveling

International Standard Classification of Occupations (ISCO)

The International Standard Classification of Occupations (ISCO) [19] is a global system for categorizing occupations based on job tasks. ISCO's hierarchical structure includes 10 major groups, subdivided into minor groups, unit groups, and detailed occupations. It is used internationally for comparing labor market statistics across countries.

While ISCO provides broad occupational categories, ESCO (European Skills, Competences, Qualifications and Occupations) builds on this framework by adding detailed skill and qualification

data, focused on the European context [12]. ISCO is ideal for cross-national labor market analysis, while ESCO is more granular and region-specific, emphasizing skills and competencies.

European Skills, Competences, Qualifications and Occupations (ESCO)

ESCO [12] is a classification system developed by the European Commission to categorize occupations, skills, qualifications, and competences in a way that is applicable across EU member states. Unlike ISCO, which focuses primarily on job titles, ESCO is more granular, detailing skills and competencies tied to specific job roles, with an emphasis on the European labor market. It supports multilingualism, offering content in 27 EU languages, which enhances its usability across diverse countries and sectors.

ESCO builds upon ISCO's hierarchical structure, offering deeper insights into how skills, qualifications, and occupations interrelate. It is particularly useful for creating a more dynamic labor market by facilitating cross-border mobility and alignment between education systems and job markets. Additionally, ESCO's detailed descriptions help match individuals' skills with the needs of employers, promoting labor market efficiency.

ESCO occupations are located at level five and lower, while ISCO-08, the latest version of ISCO, provides the top four levels of the occupational hierarchy.

As of the time of writing, the ESCOXML-R [58] model is the most promising approach to classifying job titles into this taxonomy and there are 3039 occupations in the ESCO classification system, available in 28 different languages.

This dissertation uses the ESCO taxonomy for job title classification.

European ICT Professional Role Profiles

The European ICT Professional Role Profiles [8], developed by CEN, define 30 ICT profiles grouped into seven categories: Process Improvement, Business, Technical, Design, Development, Service & Operation, and Support (see Figure 2.5).

Each profile includes mission, main tasks, deliverables, KPIs, and incorporates e-Competences from the European e-Competence Framework (e-CF). This makes the framework valuable for organizations managing ICT recruitment, educational institutions designing curricula, and professionals planning ICT careers.

By offering a European-focused perspective, these profiles complement broader classification systems and are especially relevant for analyzing technology-focused career paths.

Systematized Nomenclature of Medicine Clinical Terms (SNOMED-CT)

SNOMED CT [18] is a comprehensive medical terminology system, covering occupations as well as clinical findings, diagnoses, and procedures (see Figure 2.6).

While its scope goes beyond job roles, its sophisticated structure demonstrates how complex, domain-specific classification systems can be built, offering inspiration for specialized frameworks in fields like IT, even if its full complexity is not directly required here.




Figure 2.5: European ICT Professional Role Profiles




Figure 2.6: SNOMED-CT Example - Orthopedic Surgeon

2.6.2 Initiatives for Tracking Alumni Outcomes

University Initiatives

While many universities track student outcomes, few offer dedicated platforms for visualizing alumni data.

A basic search for “alumni map” yields many results, such as the Near East South Asia Center for Strategic Studies (NESA) [31] and the Escola Superior de Enfermagem de Coimbra (ESEnfC) [11]. However, the University of British Columbia (UBC) stands out for its comprehensive approach. UBC’s Global Alumni Impact Map (Figure 2.7), launched for its centenary, visually displays alumni locations worldwide [49].

The map allows filtering by role (e.g., students, alumni, staff) and lets users add themselves through an online form. Unlike the AlumniEI platform, which may require manual updates, UBC appears to automatically add new users once their information is submitted.




Figure 2.7: UBC Global Alumni Impact Map

Corporate Initiatives

Tracking alumni outcomes is less common in corporations. Corporations rarely track alumni outcomes, focusing instead on current employees. Corporate alumni networks, such as the McKinsey Alumni Center [27], Deloitte Alumni Network, and Microsoft’s AlumConnect, primarily support networking and career opportunities, but their directories are restricted to former employees and not publicly searchable.

Among all the tracked alumni initiatives, the University of British Columbia initiative most closely aligns with the scope of this project.

2.7 Discussion and Summary

In the inclusion stage of this systematic review, each document was thoroughly analyzed and annotated for authorship, publication year, type, subject area, research questions, and applied machine learning techniques. Records were grouped into three main categories: job title classification, alumni tracking and analytics, and visualization of spatiotemporal or multidimensional data.

Job title classification dominated the literature, reflecting a growing interest in labor market intelligence and recruitment. Twenty-three records explored diverse machine learning methods, demonstrating the range of approaches in this field: [[20], [4], [13], [9], [54], [41], [10], [1], [44], [50], [28], [24], [29], [39], [40], [45], [46], [58], [3], [15], [23], [26], [35]]. Two works spanned both job classification and alumni analytics, highlighting methodological overlaps [[25], [23]].

Alumni tracking also emerged as a key area, with 11 records examining platforms and methods for leveraging alumni data. Techniques included data warehousing and visualization tools like Power BI [[21], [42], [16], [55], [56], [23], [36], [30], [25], [22], [7]]. This focus aligns with the study's goal of understanding career trajectories and enhancing institutional value.

Although only one record fell strictly under visualization of spatiotemporal/multidimensional data [[57]], practical tools like D3.js, Looker, or Power BI are more commonly adopted in dashboards than described in academic literature. The emphasis in practice is on effective, interactive visualization, supporting better insight and decision-making.

This review also considered several classification frameworks: ESCO (the core European skills and occupations taxonomy used in this project), SOC and O*NET (US-centric occupational data), SFIA (ICT skills), Radford Global Job Leveling (career structure and pay equity), and the European ICT Professional Role Profiles (detailed European ICT roles). While SNOMED CT primarily serves healthcare, its structured approach to domain-specific classification provides a valuable perspective.

Among alumni tracking initiatives, UBC's Global Alumni Impact Map is the most noteworthy, though corporate alumni programs also exist.

This systematic review identified key works and frameworks spanning job classification, alumni tracking, and data visualization. The following chapter builds on these findings to inform the design and development of the Alumni-FEUP platform.

Chapter 3

State of the Art Analysis

The review builds upon the systematic review methodology presented in [chapter 2](#) and encompasses three main areas: job title classification methodologies, alumni tracking systems, and approaches to multidimensional data visualization.

Section [3.1](#) analyzes existing job title classification methods, covering traditional taxonomies and modern machine learning approaches, as well as international occupational classification standards and their applications.

Section [3.2](#) reviews current alumni tracking systems and analytics platforms, assessing their methodologies, technologies, and effectiveness in processing alumni career data.

Section [3.3](#) explores methods for visualizing and integrating complex multidimensional data, focusing on techniques for career trajectory analysis and educational outcomes.

Through this analysis, gaps in current research and practice are identified, providing direction for the novel contributions presented in subsequent chapters.

3.1 Job Title Classification

Job title classification, including the determination of seniority levels, is a critical challenge in alumni tracking and career pathway analysis for educational institutions.

The dynamic nature of the technology industry and the emergence of new roles demand more sophisticated, automated approaches.

Effective job title classification enables:

- Automated tracking of alumni career progression
- Identification of emerging career trends
- Understanding of skill evolution across seniority levels

Recent advances in machine learning, particularly the emergence of Large Language Models (LLMs), present new opportunities for addressing these challenges.

These models are promising for distinguishing between different seniority levels, even when companies use non-standard terminology. However, significant hurdles remain in standardization,

automation, and adaptation to emerging roles, especially in rapidly evolving sectors like technology.

3.1.1 Overview of Methods

Job title classification has evolved significantly over the past decade, transitioning from simple rule-based systems to sophisticated machine learning approaches.

The challenge of building resilient and accurate classification systems that can adapt to emerging trends has driven significant innovation in this field.

The two-stage classification approach, as demonstrated by Javed et al. [20], employs an SVM classifier to categorize jobs into broad SOC groups, followed by k-NN classifiers for fine-grained classification. This method allows for better identification of niche job titles while minimizing human oversight.

In the context of multilingual classification, recent studies have shown promising results. Tran et al. [50] demonstrated success with a Bi-GRU-LSTM-CNN model using pre-trained language models, while Zhang et al. [58] introduced ESCOXML-R, achieving state-of-the-art performance across multiple languages and datasets.

These advances in machine learning approaches suggest promising directions for improving classification accuracy and adaptability.

The evaluation of job title classification systems presents unique challenges due to the hierarchical nature of job categories. Key evaluation approaches include:

- **Traditional Metrics:** Accuracy, precision, recall, and F1-score provide baseline performance indicators.
- **Hierarchical Evaluation:** Metrics that account for partial matches in hierarchical classification schemes.

3.1.1.1 Traditional Machine Learning Approaches

Boselli et al. [4] demonstrated the effectiveness of classical machine learning techniques for classifying online job advertisements, utilizing features such as TF-IDF for text representation. Similarly, Patil and Shinde [35] explored the combination of K-Means clustering and Random Forest algorithms, though their application was limited by small dataset sizes.

Statistical methods remain relevant, particularly in scenarios where interpretability is crucial. Schierholz and Schonlau [44] conducted a comprehensive comparison of various machine learning approaches for occupation coding, finding that simpler algorithms like logistic regression and support vector machines (SVM) can perform competitively when appropriately tuned. These traditional approaches often benefit from their lower computational requirements and easier deployment in production environments.

3.1.1.2 Deep Learning Solutions

The advent of deep learning has revolutionized job title classification, introducing more sophisticated approaches capable of handling complex patterns and relationships. A notable contribution is the DeepCarotene system by Wang et al. [54], which introduced a multi-stream Convolutional Neural Network (CNN) architecture specifically designed for job title classification. This approach demonstrated superior performance in handling noisy, unbalanced datasets of job postings, processing over 15 million entries effectively.

Transformer-based models, particularly BERT and its variants, have shown remarkable results in recent years. Maitra et al. [26] presented an improved BERT model that achieved high accuracy in classifying job titles by leveraging contextual information from job descriptions. The study demonstrated that DistilBERT, a compressed version of BERT, could maintain high performance while reducing computational requirements. The multilingual capabilities of these models were further explored by Zhang et al. [58] with ESCOXML-R, which achieved state-of-the-art performance across multiple languages and datasets, particularly in handling cross-lingual classification challenges.

3.1.1.3 Hybrid and Novel Approaches

Recent research has increasingly focused on hybrid approaches that combine multiple techniques to leverage their respective strengths. Rahhal et al. [39] proposed a two-stage classification system that first identifies the broad category of a job title before making fine-grained classifications, significantly improving overall accuracy. This hierarchical approach has proven particularly effective in handling the complexity of modern job titles.

The integration of Large Language Models (LLMs) represents an emerging direction in job title classification. While Skondras et al. [46] demonstrated the potential of ChatGPT for dataset creation, its application was limited to data generation rather than direct classification. Similarly, Howison et al. [15] explored the use of generative AI for extracting structured data from job postings, though not specifically for title classification. These initial explorations suggest promising directions for future research, particularly in combining LLMs with traditional classification approaches.

3.1.1.4 Data Processing and Feature Engineering

A critical aspect of job title classification is data preprocessing and feature engineering. Several studies by Beristain et al. [3] and Li et al. [25] emphasize the importance of proper text normalization. The key preprocessing steps identified in the literature include:

- Standardization of abbreviations and acronyms (e.g., "Sr." to "Senior", "Dev" to "Developer")
- Removal of company-specific terminology and non-standard titles

- Handling of multilingual titles through translation or language-specific models
- Resolution of synonymous job titles (e.g., "Software Engineer" vs "Software Developer")
- Extraction of seniority indicators and role hierarchies

Lade et al. [23] specifically addressed the challenges of extracting and processing job titles from LinkedIn profiles, highlighting the importance of robust data cleaning pipelines. Their work emphasizes the need for automated validation mechanisms to ensure data quality while maintaining scalability (H2).

3.1.1.5 Evaluation Methodologies

The evaluation of job title classification systems presents unique challenges due to the hierarchical nature of job categories and the presence of overlapping roles. The literature identifies several key evaluation approaches:

- **Traditional Classification Metrics:** Standard measures including accuracy, precision, recall, and F1-score provide baseline performance indicators
- **Hierarchical Evaluation:** Metrics that account for partial matches in hierarchical classification schemes, particularly relevant for taxonomies like SOC and ISCO
- **Domain-Specific Assessment:** Specialized metrics that consider industry-standard taxonomies and role hierarchies
- **Cross-Lingual Performance:** Evaluation of classification accuracy across different languages and regional variations

Tran et al. [50] introduced a multi-label evaluation framework that better reflects the reality of jobs that span multiple categories or roles. This approach has proven particularly valuable for evaluating classification systems in dynamic industries where roles often combine multiple functions. The framework addresses the challenge of evaluating system resilience to emerging trends (RQ3) while maintaining classification accuracy across diverse job titles.

3.1.2 Standards and Frameworks

The International Standard Classification of Occupations (ISCO), maintained by the International Labour Organization (ILO), represents one of the primary frameworks. Boselli et al. [4] leveraged both ISCO and ESCO (European Skills, Competences, Qualifications and Occupations) taxonomies in their classification approach, demonstrating the complementary nature of these standards.

The Standard Occupational Classification (SOC) system appears frequently in the literature:

- Wang et al. [54] utilized SOC for their DeepCarotene system, processing over 15 million job postings

- Li et al. [25] employed it for its detailed occupation descriptors, though noting limitations in emerging IT roles
- Lalitha et al. [24] mapped raw job titles to SOC codes for their work on employer industry classification.
- Howison et al. [15] sampled their dataset against the SOC taxonomy, gathering 100 postings per occupational group

The O*NET-SOC framework, which builds upon SOC, sometimes referred to just as O*NET, has gained significant traction in recent studies:

- Skondras et al. [46] specifically chose O*NET-SOC taxonomy for their neural network classifier using USE embeddings.
- Javed et al. [20] at CareerBuilder implemented a two-stage classification system mapping their content into O*NET-SOC codes

The European e-Competence Framework (e-CF) and ESCO have shown particular utility in multilingual contexts:

- Beristain et al. [3] validated their clustering results against ESCO and ISCO
- Zhang et al. [58] leveraged ESCO for cross-lingual classification tasks

3.1.2.1 Framework Selection and Implementation

The choice of framework often depends on specific requirements:

- **Geographical Focus:** Studies targeting European markets tend to favor ESCO and e-CF [3, 58]
- **Industry Specificity:** Technology-focused research often combines O*NET with industry-specific taxonomies [25, 46]
- **Multilingual Needs:** ESCO proves particularly valuable for cross-lingual applications [26, 58]
- **Granularity Requirements:** SOC and ISCO serve well for broad categorization [20, 54]

3.1.2.2 Integration Challenges

Studies highlight several common challenges in implementing these frameworks:

- Mapping between different standards when combining multiple taxonomies
- Handling regional variations in job titles and responsibilities
- Accommodating emerging roles not yet included in standard frameworks

- Maintaining consistency across different organizational hierarchies

These challenges directly impact the development of resilient classification systems (RQ3) and the ability to automate the classification process (H2). The literature demonstrates that successful implementation often requires combining multiple frameworks while maintaining flexibility for emerging job titles and roles.

3.1.3 Challenges

The field of job title classification faces several significant challenges that impact both research and practical implementations. These challenges are particularly evident in the context of rapidly evolving industries and multilingual environments.

3.1.3.1 Dynamic Nature of Job Titles

The continuous emergence of new roles and evolution of existing ones presents a fundamental challenge. Several studies highlight this issue:

- Li et al. [25] note how standard taxonomies struggle to keep pace with new technology roles
- Tran et al. [50] address the challenge of roles that span multiple categories or combine different functions
- Beristain et al. [3] highlight how company-specific terminology complicates standardization efforts

3.1.3.2 Data Quality and Standardization

Data-related challenges persist across different approaches:

- **Inconsistent Formatting:** Variations in how titles are written and formatted across different sources
- **Abbreviations and Acronyms:** Diverse use of shorthand and company-specific acronyms
- **Multilingual Aspects:** Handling job titles across different languages and regional variations
- **Seniority Indicators:** Inconsistent representation of career levels across organizations

3.1.3.3 Technical Implementation Challenges

The implementation of classification systems faces several technical hurdles:

- **Scalability:** Wang et al. [54] address the challenge of processing large-scale datasets efficiently
- **Real-time Processing:** The need for systems that can classify new titles as they emerge

- **Model Maintenance:** Keeping classification models updated with new patterns and titles
- **Integration:** Combining multiple data sources and classification approaches

3.1.3.4 Accuracy and Validation

Ensuring classification accuracy presents ongoing challenges:

- **Ground Truth:** Difficulty in establishing reliable training data for emerging roles
- **Cross-validation:** Limited availability of standardized benchmark datasets
- **Performance Metrics:** Challenges in evaluating hierarchical and multi-label classifications
- **Human Verification:** Balance between automation and manual validation

3.1.3.5 Organizational Differences

Company-specific challenges add another layer of complexity:

- **Title Inflation:** The practice of using inflated or non-standard titles such as "Ninja" or "Rockstar"
- **Hierarchy Variations:** Different organizational structures affecting title interpretation
- **Industry-Specific Terms:** Specialized terminology varying across sectors
- **Regional Differences:** Varying job title conventions across countries and cultures

These challenges directly impact the feasibility of building resilient classification systems (RQ3) and achieving automation (H2). The solution to be later presented in this thesis must address multiple challenges simultaneously while maintaining adaptability to new developments in the job market.

3.1.4 Research Gaps

Analysis of the current literature reveals several significant gaps in job title classification research and implementation. These gaps present opportunities for innovation and further study.

3.1.4.1 Methodological Gaps

Current approaches show limitations in several key areas:

- **LLM Integration:** While studies like Skondras et al. [46] explore LLMs for data generation, there is limited research on their direct application to classification tasks
- **Hybrid Approaches:** Few studies comprehensively combine traditional ML, deep learning, and LLM approaches

- **Seniority Classification:** Limited research on automated detection and standardization of seniority levels across different organizational structures
- **Cross-lingual Solutions:** Despite advances by Zhang et al. [58], gaps remain in handling regional variations and multiple languages simultaneously

3.1.4.2 Technical Gaps

Several technical challenges remain inadequately addressed:

- **Real-time Adaptation:** Current systems struggle to adapt to newly emerging job titles without retraining
- **Automated Pipeline:** Lack of end-to-end solutions that combine data collection, cleaning, and classification without human intervention
- **Validation Mechanisms:** Limited research on automated validation approaches that can ensure classification accuracy
- **Framework Integration:** Insufficient methods for automatically mapping between different classification standards

3.1.4.3 Data-Related Gaps

Significant gaps exist in data handling and processing:

- **Benchmark Datasets:** Lack of comprehensive, multilingual datasets for evaluating classification systems
- **Emerging Roles:** Limited data available for training models on newly created positions
- **Historical Tracking:** Insufficient methods for tracking job title evolution over time
- **Context Integration:** Limited use of additional context (company size, industry, location) in classification

3.1.4.4 Application-Specific Gaps

Several areas require further research for practical implementation:

- **Domain Adaptation:** Limited research on adapting general classification systems to specific industries or organizations
- **Career Progression:** Insufficient methods for analyzing and predicting career path trajectories
- **Skill Mapping:** Limited integration between job title classification and skill requirement analysis
- **Temporal Analysis:** Gaps in understanding how job titles and roles evolve over time

3.1.4.5 Future Research Directions

These gaps suggest several promising research directions:

- Development of hybrid systems combining traditional ML, deep learning, and LLMs (H1)
- Creation of fully automated classification pipelines with minimal human oversight (H2)
- Integration of multiple data sources and contexts for improved classification accuracy
- Design of adaptive systems capable of handling emerging job titles and roles (RQ3)
- Development of standardized evaluation frameworks for comparing different classification approaches

These research gaps highlight the potential for significant contributions to the field, particularly in addressing the challenges of automation, adaptability, and accuracy in job title classification systems.

3.2 Alumni Tracking and Analytics Initiatives

As highlighted in the literature search, while many universities maintain some form of alumni tracking, few have developed comprehensive platforms that integrate career progression data with interactive visualizations.

This analysis focuses on evaluating these existing systems across several key dimensions:

- Data collection and maintenance methodologies
- Integration with professional networks
- Privacy and data protection measures
- Visualization and interaction capabilities
- Real-time update mechanisms

3.2.1 Existing Platforms

Several universities and organizations have developed platforms for tracking and analyzing alumni data, as highlighted in the literature search:

- Universiti Sains Malaysia's Alumni Tracking Dashboard employs Power BI and machine learning techniques for analyzing graduate conversion rates and employment trends [16, 21].

- The California State University Long Beach's BUILD Snapshot system takes a more targeted approach, using personalized Excel files pre-populated with trainee data [7]. While less technologically sophisticated, this system achieves high response rates through its personalized approach.
- Universitas Muhammadiyah Yogyakarta's implementation uses a Normalized Data Store architecture for improved data integration [56], addressing common challenges in data consistency and accessibility.

3.2.2 Key Features and Innovation

Recent innovations in alumni tracking systems demonstrate significant technological and methodological advances across multiple dimensions. Data collection methods have evolved considerably, moving from traditional survey-based approaches to more sophisticated automated solutions. Particularly noteworthy is the automated extraction of professional data from LinkedIn [23, 25], though this approach requires careful consideration of API access and user permissions. Some platforms have implemented web-based graduate tracer systems with automated validation [30], while others maintain secured networks specifically for alumni data collection [22].

Analytics capabilities have seen substantial advancement, with several innovative approaches emerging:

- Time series analysis (ARIMA) combined with Support Vector Machines for forecasting graduate conversion rates [17]
- Genetic algorithms for predicting career success based on academic and professional factors [36]
- Machine learning approaches for analyzing career trajectories, including detailed analysis of CEO career paths [42]

Data quality management has become increasingly sophisticated, with institutions implementing various strategies to ensure data reliability. The success of these approaches is evident in cases like USM's implementation, which improved their contactable alumni base from 23,000 to 70,000 through systematic data cleaning processes [21]. Modern platforms employ normalized data store architecture [56] to maintain consistency across multiple data sources, while implementing robust validation mechanisms for survey responses and profile updates [7].

Privacy and security features have evolved in response to growing data protection concerns. Modern systems incorporate:

- Secured authentication and authorization mechanisms [22]
- Privacy-preserving data collection methods for professional network integration [23]
- Compliance frameworks for data protection regulations

Visualization and reporting capabilities have advanced significantly, with platforms now offering interactive dashboards powered by business intelligence tools like Power BI [16, 21]. These systems enable sophisticated geographic visualization of alumni distribution and career trajectory mapping. Some platforms have integrated predictive analytics visualization [36], providing insights into career success patterns and professional development trajectories.

The integration of external and internal systems presents unique challenges and opportunities. While platforms like LinkedIn offer rich professional data [25], accessing and integrating this information involves both technical and privacy considerations. The normalized data store approach demonstrated by [56] shows promise in maintaining data consistency across multiple sources, though implementation complexity increases with the number of data sources.

Career analysis capabilities have evolved substantially, moving beyond simple tracking to sophisticated prediction models. The genetic algorithm approach presented in [36] demonstrates the potential for predicting career success based on multiple factors, while the CEO career pattern analysis in [42] provides valuable insights into long-term career trajectory analysis. These advances suggest a trend toward more predictive and prescriptive alumni tracking systems.

Big data architectures have emerged as a crucial component in modern systems [55], enabling:

- Processing of large-scale alumni information
- Real-time data integration and analysis
- Scalable storage and retrieval mechanisms

These innovations collectively represent significant progress in alumni tracking capabilities, though challenges remain in achieving seamless integration and maintaining data quality at scale. The trend toward automated, data-driven approaches suggests a future where alumni tracking systems become increasingly sophisticated and capable of providing actionable insights for educational institutions.

3.2.3 Limitations of Current Approaches

Current alumni tracking platforms face several significant limitations that span technical, operational, and strategic dimensions. At the core of these challenges lies the persistent issue of data quality and maintenance. Even successful implementations, such as USM's dashboard, required extensive data cleaning efforts to improve alumni contactability [21]. The reliability of self-reported data remains questionable [36], while manual data entry processes continue to introduce inconsistencies and errors [30].

The most critical data-related challenges include:

- Incomplete or outdated information requiring extensive cleaning efforts [21]
- Difficulty in maintaining up-to-date alumni contact information [7]
- Challenges in verifying the accuracy of self-reported data [36]

Integration challenges present another significant barrier to effective alumni tracking. The BUILD program's experience highlights the difficulties in merging data from multiple sources [7], particularly when dealing with varying data formats and structures. The absence of standardized APIs for institutional databases further complicates automation efforts, while the normalized data store approach proposed by [56] addresses only part of the integration challenge.

Privacy and access restrictions pose increasingly complex challenges in the current regulatory environment. Systems relying on LinkedIn data face significant hurdles due to anti-scraping measures and API limitations [23]. The implementation of data protection regulations has complicated cross-border data sharing [22], while the need to balance data accessibility with privacy protection creates additional complexity [25]. These restrictions often limit the ability to automatically update alumni information and maintain current records.

Technical scalability emerges as a critical concern as alumni populations grow. The key technical limitations include:

- Performance degradation with increasing data volumes [55]
- Limited computational resources for complex analytics [16]
- Difficulty in handling multilingual data and international career paths

Analytics and visualization capabilities in existing platforms often fall short of modern requirements. Most systems focus on basic tracking rather than predictive analytics [42], lacking sophisticated tools for career trajectory analysis [36]. The inability to identify emerging career trends and visualize complex career paths [16] limits the strategic value these systems can provide to educational institutions.

Implementation barriers present additional challenges for institutions:

- High initial setup costs and resource requirements [56]
- Need for specialized technical expertise for system maintenance
- Challenges in achieving consistent user engagement [30]

These limitations underscore the need for more robust, automated, and privacy-conscious approaches to alumni tracking. The challenges in data integration and quality management particularly emphasize the importance of developing standardized approaches and automated validation mechanisms. Furthermore, the limited analytics capabilities of current systems suggest significant opportunities for innovation in career trajectory analysis and predictive modeling.

3.2.4 Relevance to Research

The analysis of existing alumni tracking platforms and their limitations directly informs the research questions and objectives of this thesis. The current state of the field reveals significant opportunities for innovation, particularly in automated data collection, integration, and analysis capabilities.

Regarding RQ1, the review reveals several key areas for innovation:

- Integration of multiple data sources beyond traditional survey methods
- Advanced analytics capabilities for career trajectory analysis
- Real-time data collection and processing mechanisms

The challenges identified in handling evolving job titles and career paths directly address RQ2. Current systems struggle to adapt to emerging roles and changing career patterns, particularly in the technology sector [25]. This limitation underscores the need for more adaptive and future-proof solutions that can automatically recognize and categorize new job titles and career trajectories.

The classification challenges faced by existing systems strongly align with RQ3. While some platforms attempt automated classification [42], they often lack:

- Sophisticated algorithms for handling non-standard job titles
- Methods for identifying career progression patterns
- Ability to adapt to emerging industry-specific terminology

The varying approaches to data visualization across platforms inform RQ4. Current visualization implementations, while functional, often fail to:

- Present complex career trajectories in an intuitive manner
- Enable interactive exploration of career patterns
- Provide meaningful comparisons across different career paths

Beyond the research questions, the review supports the hypotheses underlying this thesis. The prevalence of manual processes and limited automation in existing systems reinforces H2's assertion about the feasibility of building automated tracking systems. Similarly, the emerging use of machine learning techniques [16, 36] supports H1's proposition regarding the potential of LLMs for improving classification accuracy.

These findings suggest that while the field of alumni tracking has made significant progress, substantial opportunities remain for innovation. The gaps identified in current systems, particularly in areas of data integration, automated classification, and interactive visualization, align closely with the goals of this thesis to develop a more comprehensive and sophisticated platform for alumni tracking and analysis. The research directions proposed here address not only current limitations but also anticipate future needs in alumni career tracking and analysis.

3.3 Visualization and Data Integration of Multidimensional Data

The visualization of multidimensional data, particularly when combining spatial, temporal, and categorical information, presents unique challenges in data integration and representation. Zhang and Huang propose a 2D-Ratio method for visualizing multidimensional data patterns in large datasets, demonstrating key principles:

- Effective handling of multiple data dimensions simultaneously
- Reduction of visual clutter through dimension clustering
- Scalability for large datasets (validated with over 1 million records)

The practical implementation of alumni data visualization will likely leverage modern tools such as D3.js [6]. These established technologies offer several advantages:

- Built-in interactive features and user engagement capabilities
- Proven scalability for handling large datasets
- Regular updates and community support
- Integration capabilities with various data sources
- Customizable visualization options for different data types

3.3.1 Techniques and Tools

The implementation of effective data visualization for alumni tracking systems has leveraged various tools and technologies, ranging from specialized libraries to comprehensive business intelligence platforms. While several options exist, D3.js emerges as the most suitable choice for the Alumni-FEUP platform's requirements.

D3.js (Data-Driven Documents) stands out as the leading choice for several compelling reasons:

- **React Integration:** As the frontend is built with React, D3.js offers seamless integration through libraries like React-D3 or direct implementation
- **Customization:** Provides complete control over visualization design and interactivity
- **Community Support:** Extensive documentation, examples, and active developer community
- **Performance:** Efficient handling of DOM manipulations and animations
- **Flexibility:** Supports both geographical and temporal data visualization needs

While Power BI has shown success in alumni analytics dashboards, as demonstrated by USM's implementation [16], it is better suited for internal business intelligence rather than web-based interactive platforms.

Similarly, Looker, despite its robust features, is optimized for use with data warehouses like BigQuery, making it potentially over-engineered for the current architecture where a data warehouse implementation is not yet planned.

Other visualization tools considered include:

- **Plotly**: While powerful, lacks the fine-grained control offered by D3.js
- **Highcharts**: Commercial licensing could introduce unnecessary costs
- **deck.gl**: Specialized in geospatial visualization but may be excessive for our needs

The selection of D3.js aligns with both technical requirements and practical constraints:

- **Technical Fit**: Seamless integration with the existing React frontend
- **Cost Efficiency**: Open-source nature eliminates licensing concerns
- **Scalability**: Proven capability to handle large datasets efficiently
- **Development Control**: Full control over visualization implementation
- **Future-Proofing**: Active development and broad adoption ensure long-term viability

This choice enables the development of custom, interactive visualizations that can effectively represent the spatial (world map), temporal (timeline), and categorical (job titles) dimensions of alumni data while maintaining optimal performance and user experience.

3.4 Critical Analysis and Synthesis

The analysis uses capability matrices to compare approaches in each domain:

Capability	ESCOXLMR [58]	CEO Study [42]	BERT [26]	LinkedIn [25]	Two-Stage [39]	AlumniEI* (Proposed)
Multilingual Support	✓	✗	Partial	✓	✓	✓
Handles Emerging Titles	Partial	✗	✓	✓	✗	✓
Seniority Detection	✗	✓	Partial	✓	✗	✓
Two-tier Classification	✗	✗	✗	✗	✓	✓
Automated Classification	✓	✓	✓	✓	✓	✓
Industry-Specific Context	Partial	✓	Partial	✓	✓	✓

Table 3.1: Job Title Classification Systems Comparison

Capability	USM [21]	BUILD [7]	UMY [56]	ARIMA [16]	AlumniEI* (Proposed)
Automated Data Collection	✗	✗	✗	✓	✓
Career Path Analysis	✗	✓	✗	✓	✓
Privacy Compliance	✓	✓	✓	✓	✓
Predictive Analytics	✗	✗	✗	✓	Partial
Real-time Updates	✗	✗	✗	✗	✓
Open-Source Platform	✗	✗	✗	✗	✓

Table 3.2: Alumni Tracking Platforms Comparison

Capability	2D-Ratio [57]	Power BI [16]	AlumniEI* (Proposed)
Interactive Exploration	Partial	✓	✓
Geographic Visualization	✓	Partial	✓
Timeline Integration	✓	Partial	✓
Scalability	✓	✓	✓
Custom Visualizations	✓	✗	✓

Table 3.3: Data Visualization Approaches Comparison

Analysis of these matrices reveals several key patterns:

- **Job Classification:** While current approaches use traditional ML and deep learning models [26, 58], there's a notable gap in leveraging LLMs for classification tasks
- **Alumni Tracking:** Most systems lack automated data collection and real-time updates [7, 21, 56], focusing instead on periodic manual updates
- **Visualization:** While various tools exist, few systems successfully integrate spatial, temporal, and categorical data in an interactive manner [16, 57]

In job title classification, the most pressing challenges center around standardization and automation:

- Integration of multiple classification standards (SOC, ESCO, O*NET)
- Handling of emerging and non-standard job titles
- Automated seniority level detection

Current systems struggle particularly with rapidly evolving fields like technology, where new roles emerge frequently and traditional classification methods quickly become outdated.

The automated detection of seniority levels remains inconsistent across platforms, with most systems requiring manual intervention or relying on rigid rule-based approaches.

While privacy compliance is generally strong, several critical limitations persist:

- Lack of automated data collection capabilities
- Dependence on manual updates and validation
- Limited integration with professional networks

The integration with professional networks, though promising, remains constrained by API restrictions and privacy concerns. Real-time updates and validation of alumni information represent a particularly notable gap in current implementations.

In the visualization domain, key challenges include:

- Integration of spatial, temporal, and categorical data
- Scalability while maintaining performance
- Interactive exploration of complex career paths

3.5 Conclusion

The comprehensive analysis across job title classification, alumni tracking, and data visualization domains reveals both limitations and opportunities for innovation in addressing our research questions.

Regarding platform differentiation (RQ1), while initiatives like UBC's Alumni Map exist, current platforms lack integrated automated classification systems and real-time data validation, with limited interactive visualization capabilities.

The challenges of future-proofing (RQ2) and building resilient classification systems (RQ3) are evidenced in current systems' struggles with emerging job titles. Our analysis suggests combining:

- Integration of LLMs with traditional classification methods
- Two-tier classification architecture
- Automated validation and update mechanisms

The review of visualization approaches (RQ4) demonstrates the potential of D3.js for creating interactive, multidimensional visualizations that support:

- Integration of spatial and temporal data
- Interactive career trajectory exploration
- Scalable visualization architecture

These findings inform the development approach detailed in subsequent chapters.

Additionally, making Alumni-FEUP open-source enables community-driven improvements, cross-institutional collaboration, and standardization of alumni tracking methodologies, positioning the platform to set new standards in alumni career tracking and analysis.

Chapter 4

Alumni-FEUP: A Novel Approach to Alumni Career Tracking

This chapter details the design, architecture, and implementation of the Alumni-FEUP platform.

Building on the requirements and research questions outlined in [chapter 1](#), it presents the system’s structure, the technologies employed, and the key design decisions that shaped its development, with a special focus on the AI-driven agents that are at the core of the platform.

Each section progressively deepens the technical detail, beginning with a high-level overview and concluding with an analysis of the implementation trade-offs and challenges encountered.

4.1 Overview of the Solution

Alumni-FEUP is an interactive analytics platform developed to visualize and explore the career trajectories of FEUP alumni. It enables users to analyze alumni distribution across companies, countries, industries, degrees, job titles, and seniority levels, with a particular emphasis on temporal evolution and career progression patterns.

The platform is intended for a diverse set of stakeholders, including current students exploring career paths, academic staff evaluating graduate outcomes, and external organizations assessing institutional impact and talent trends.

The system is structured around three interconnected components: (1) a data ingestion and enrichment pipeline integrating institutional records (via SIGARRA) with public professional profiles (via LinkedIn), (2) a suite of AI-based classification agents responsible for transforming unstructured data into structured knowledge, and (3) a full-stack web application that presents this information through dynamic, interactive dashboards.

By including Large Language Models (LLMs) in the data enrichment pipeline, a novel approach to alumni insights is explored, especially in the area of job title classification, which had been previously dominated by traditional machine learning and data science approaches.

The implementation of the job classification agent is detailed in [section 4.6](#), as well as the reasoning behind this choice in [section 4.7](#).

Combining validated academic data with enriched labor market information allows the platform to offer a unified, data-driven view of post-graduation outcomes, directly addressing the research questions related to alumni tracking, skill taxonomy classification, and institutional footprint analysis.

4.2 System Architecture

The Alumni-FEUP platform follows a modular, service-oriented architecture composed of three primary components:

- A frontend web interface,
- A backend API server,
- A set of autonomous data orchestration and enrichment services.

This separation of concerns supports maintainability, scalability, and independent component evolution.




Figure 4.1: High-level system architecture of Alumni-FEUP

Figure 4.1 depicts the high-level system layout. The frontend includes interfaces for administrative tasks, such as CSV uploads, and user-facing visualizations. The backend API server handles authentication, token management, and all business logic related to alumni and user profile management.

For data enrichment and classification, the API server communicates with a dedicated orchestration layer composed of background agents. These agents are responsible for asynchronous tasks such as job title classification, seniority detection, and location resolution. They interface with third-party services to augment alumni profiles with structured, contextual information.

All processed and enriched data is persisted in a centralized relational database, from which the frontend retrieves information to power visual analytics and user queries.




Figure 4.2: Data flow in the Alumni-FEUP platform

Figure 4.2 illustrates the operational flow of data. The process begins with a CSV file upload by an administrator through the web interface. This file is parsed by the backend, and alumni and graduation records are inserted into the database.

Upon successful ingestion, the backend triggers asynchronous enrichment tasks through the orchestration service. These tasks include scraping LinkedIn profiles, resolving company and location metadata, and applying classification agents. Once complete, the structured data is stored and becomes available to end users via interactive dashboards.

4.3 Key Components and Technologies

The following subsections detail the key components and technologies used in the implementation of the platform, with a particular focus on the data processing and enrichment pipeline that happens in the AI powered orchestration service.

4.3.1 Frontend

The frontend of Alumni-FEUP is implemented using **Next.js**¹, leveraging its support for both static and server-side rendering. Static generation is used for general-purpose pages such as the landing page, feedback form, and about section, while dashboard views rely on dynamic rendering to fetch and display up-to-date alumni data from the backend.

Styling is handled using **Tailwind CSS**², allowing for highly customizable and responsive designs. Component structure and accessibility are supported by **ShadCN**³, which is used extensively across the application and extended where necessary with custom styles and behaviors tailored to the platform's needs.

¹<https://nextjs.org/>

²<https://tailwindcss.com/>

³<https://ui.shadcn.com/>

The visual analytics experience is powered by **D3.js**⁴, which renders all pie charts, trend lines, and interactive visualizations.

Users can filter alumni data through multiple dimensions, including hierarchical exploration of the ESCO taxonomy, filtering by faculty, degree, and graduation year, and geographic breakdowns by country and city - **Figure 4.3** demonstrates the platform's analytics page.




Figure 4.3: Analytics page in the Alumni-FEUP platform

These controls, combined with the responsiveness of the visual layer, make the platform adaptable to a wide range of exploratory and comparative queries.

On top of the visual analytics that include a set of interactive dashboards, each alumni profile is also available as a standalone LinkedIn-like page. Alumni can validate or update their profile data, and view their career progression over time - **Figure 4.4** showcases an example profile page.

⁴<https://d3js.org/>




Figure 4.4: Alumni profile page in the Alumni-FEUP platform

4.3.2 Backend API

The backend of Alumni-FEUP is implemented using **Nest.js**⁵, adopting a modular and declarative structure centered around RESTful controllers and service layers.

On top of providing the API endpoints needed for the frontend capabilities, the backend is also responsible for user authentication, session management, alumni record creation, and business logic for CSV ingestion and orchestration triggering.

CSV uploads are processed through a dedicated endpoint restricted to privileged users via **UseGuards**⁶.

Uploaded files undergo validation for type, size, and structural correctness, using a custom wrapper around the `csv-parse`⁷ library. Fields such as full name, graduation year, and LinkedIn URL are verified row-by-row, and errors are surfaced with descriptive exception messages.

For each valid record, the backend persists alumni and graduation data to the database and triggers the orchestration layer to initiate profile extraction.

The backend enforces **role-based access control** on privileged endpoints and integrates rate-limiting mechanisms to prevent abuse, particularly on routes that interact with external services (e.g., OTP requests or agent-triggering endpoints). Token-based authentication is used throughout, and middleware is employed for consistent error handling and request validation.

⁵<https://nestjs.com/>

⁶<https://docs.nestjs.com/guards>

⁷<https://www.npmjs.com/package/csv-parse>

4.3.3 Database and Storage

The platform uses **PostgreSQL** as the primary relational database, accessed via the **Prisma ORM**⁸. The schema is normalized and captures entities such as alumni, roles, graduations, locations, companies, and job classifications - **Figure 4.5** displays a simplified version of the database schema - the full schema can be consulted in the repository⁹.




Figure 4.5: Database schema of the Alumni-FEUP platform (simplified)

Strong referential integrity is maintained through foreign key constraints and cascade deletion rules for dependent records such as reviews and roles.

Fields critical for querying and filtering, such as the degree `conclusionYear`, the company `companyId`, and the seniority level `seniorityLevel`, are indexed for performance. Composite uniqueness constraints (e.g., on `alumni-course-year` tuples) prevent duplication and support integrity in analytics.

A lightweight **Redis**¹⁰ instance supports short-lived authentication mechanisms, including OTP code generation and email verification token handling. The service hashes all stored codes and enforces time-based expiration and retry limits.

⁸<https://www.prisma.io/>

⁹<https://github.com/carlosverissimo3001/alumni-feup/blob/main/api/prisma/schema.prisma>

¹⁰<https://redis.io/>

Cloudinary¹¹ is used as a media storage solution for profile images and company logos, providing efficient and secure blob storage with public URL delivery. This decouples image hosting from the core infrastructure while supporting dynamic frontend rendering.

4.3.4 Orchestration Service

The orchestration service acts as the central coordinator between the alumni ingestion pipeline and the downstream enrichment agents. Once a new alumni entry is created via the administrative CSV upload interface, this service triggers the necessary enrichment flows asynchronously.

Implemented as a single FastAPI application, it exposes dedicated endpoints to initiate extraction or enrichment tasks for each data dimension: job roles, companies, seniority levels, and locations. These endpoints are consumed by the NestJS backend after parsing and validating CSV rows.

The orchestration layer ensures that redundant fetches or classifications are avoided by checking for existing database entries or null results from external services. This guards the system from overfetching and reduces processing time.

Each enrichment task is isolated, allowing agents to execute in parallel. The orchestration logic ensures that tasks failing due to invalid data or unavailable third-party responses are gracefully skipped or retried depending on the scenario. To preserve system responsiveness and reduce redundant computation, the classification agent employs Redis-backed caching, particularly during the expensive embedding and reranking phases.

This modular and automated coordination of AI agents, each handling a different aspect of data enrichment, marks a novel approach in the context of academic alumni tracking platforms.

Rather than relying solely on deterministic rules or manual data cleaning, the orchestration service leverages LLM-backed agents to infer, enrich, and classify unstructured data at scale, offering a high degree of flexibility and adaptability.

The orchestration service thus abstracts the complexity of coordinating various AI modules and external scrapers, exposing a unified and robust interface to the main backend.

4.4 Implementation Details

This section presents a comprehensive walkthrough of the system's end-to-end implementation, detailing how raw data is transformed into structured and rich alumni insights.

Starting from the initial upload of a CSV file containing alumni records, the flow traverses through data validation, storage, enrichment, and ultimately presentation to the end-user, in many visualization formats.

Each stage is supported by modular components and asynchronous interactions designed to ensure scalability, reliability, and extensibility.

¹¹<https://cloudinary.com/>

To better illustrate the internal mechanics of the platform, the following technical subsections include flow diagrams, service breakdowns, and architectural patterns, with special attention given to the AI-based enrichment agents - one of the main innovative aspects of this project.

4.4.1 CSV Upload and Manual Form Submission

There are two ways to submit alumni data to the platform:

- Via a CSV file, through the admin tooling interface,
- Via a manual form submission, triggered by an end user.

Both methods are supported by the same backend service, and the data is similarly processed, with some minor differences outlined in the corresponding subsections.

CSV Upload

The data enrichment can begin with the upload of a CSV file via the platform's administrative interface (see [Figure 4.6](#)). This file is prepared by course administrators who collect LinkedIn profile links for recently graduated students. The CSV is expected to contain three key columns:

- The alumni's full name,
- The year in which the alumni graduated,
- The alumni's LinkedIn profile URL.

Although not used in the current implementation, a `student_id` field is often included in the source dataset for potential future enhancements.

Figure 4.6: Admin tooling interface for uploading alumni CSV data

Before submission, the administrator must also specify the associated faculty and course, which are sent along with the CSV file to the backend.

The frontend validates the file's size and ensures the correct file type before allowing the upload to proceed.

Once submitted, the backend performs an additional round of validation, including format verification, missing field checks, malformed row detection, and conclusion year parsing. Conclusion years are validated to fall within a reasonable interval (1950-2099), and improper formats (e.g., "Summer 2023") are parsed using regular expressions to extract the underlying year value. After parsing, each row is validated individually.

Duplicate detection is enforced based on LinkedIn profile URL, or, in its absence, by matching full name and graduation year - it is statistically unlikely that two alumni exist with the same name and graduation year, nonetheless, the system is designed to handle it.

Currently, student IDs are not stored or used for deduplication, as they are not needed for the platform's current scope. However, integrating official student identifiers would become important if the system were to support multiple universities or more complex data integrations in the future.

If a match is found, the insertion is aborted with an informative message, preserving data consistency and preventing duplicate enrichments.

New alumni are added to the system along with their graduation records. Metadata flags such as `hasSigarraMatch` and `wasReviewed` are automatically inferred from the context of the upload.

No data is committed to the database until the full file is parsed and validated. Errors are propagated back to the user interface for correction, ensuring that only clean and coherent datasets are accepted into the enrichment pipeline.

A detailed breakdown of the overall dataset coverage, including the total number of alumni, roles, and companies, as well as the methods used for LinkedIn URL collection, is presented in [section 5.1](#).

Manual Form Submission

Similarly, alumni can request to be added to the platform, via a manual form submission shown in [Figure 4.7](#).

While this is not expected to be a common use case, as all the alumni who have graduated from the supported degrees are already in the system, this feature was implemented in an early stage of the project, when it was planned to allow honorary members, to join the platform - these members are alumni from other degrees, people directly involved with the alumni community, who have not completed their studies, or other special cases, not handled by the automatic processes in place.

— Back

Join as Alumni

Email Verification ✓ Verified

Please verify your email to continue ⓘ

Email Address *

carlosverissimo3001@gmail.com Wait 25s

Full Name *

Bill Gates

LinkedIn URL *

https://www.linkedin.com/in/williamhgates

Faculty *

FEUP - Faculdade de Engenharia da Universidade do Porto

Completed Degrees *

L.EIC - Licenciatura em En... M.EIC - Mestrado em Eng...

Degree(s) Completion

Select the year you completed each degree

L.EIC - Licenciatura em Engenharia Informática e Computação (Bologna) 2023

M.EIC - Mestrado em Engenharia Informática e Computação 2025

Submit Application

Figure 4.7: Manual submission form for alumni data

To prevent fraudulent submission, the alumni is first asked to provide a personal email, which is then used to send a verification token, revealing the rest of the form, which, similarly to the CSV upload, includes the full name, graduation year, LinkedIn profile URL, as well as all the degrees they have graduated with, and the respective year of graduation.

4.4.2 Enrichment Orchestration Overview

Once alumni records are created and validated via CSV upload, the system immediately initiates the enrichment process. This is orchestrated by the backend API, which triggers the data extraction pipeline by forwarding alumni identifiers to the orchestration service.

This decoupling of the enrichment logic from the request lifecycle, ensures responsiveness and scalability and enables:

- Immediate triggering after ingestion, or delayed triggering via a separate profile update endpoint.

- Separation of concerns between data ingestion (NestJS) and enrichment (Python - FastAPI).
- Scalable background execution of multiple enrichment pipelines in parallel.

Critically, this orchestration acts as the entry point to the system's AI-based agents, bridging raw scraped data with the advanced classification logic described in later sections.

4.4.3 Alumni Profile Extraction

Once the alumni records are created in the backend, the orchestration service handles the extraction of the profile data from LinkedIn.

The service uses *Enrich Layer*¹², a third-party API specialized in scraping public LinkedIn data, to extract structured information from alumni profiles. This API operates under a pay-per-use model, and its cost implications are further detailed in [section 5.2](#).

The system ensures consistent and deduplicated queries by sanitizing each LinkedIn URL, removing prefixes (e.g., www), normalizing protocols to HTTPS, and trimming query parameters before use. If Enrich Layer returns no data, the task is aborted gracefully - this typically happens when the profile has been deleted.

The parsed response contains structured data on the alumni's professional experience, location, and profile image.

For each experience, the service checks whether the associated company already exists in the database (via sanitized LinkedIn URL matching).

If it does not, the company is inserted with basic information, and its enrichment is delegated to a background task for asynchronous processing - the company enrichment is further detailed in [subsection 4.4.4](#).

Next, each role is parsed and stored in the database, along with a corresponding raw entry containing fields such as the LinkedIn title and description (later used by the classification agent) - roles can be considered invalid if they don't meet certain criteria, such as having a start date, and an associated company.

The system also extracts the alumni's current location using available city, country, and country code fields. If a location cannot be inferred, the system triggers the location agent to resolve it using natural language reasoning - its implementation is described in more detail in [subsection 4.5.3](#).

If a valid profile picture is returned, it is uploaded to **Cloudinary**, to then be used in the frontend.

Finally, three critical asynchronous enrichment tasks are triggered:

- **Job classification**, which interprets the alumni's career trajectory based on parsed role data.
- **Role location resolution**, to resolve the geographical dimension of each professional entry.

¹²<https://enrichlayer.com/>

- **Company enrichment**, if needed, to update and verify metadata of newly discovered companies.

This parallelized architecture ensures that enrichment tasks are offloaded efficiently and do not block the main execution flow. It enables scalable and resilient data processing, making it feasible to handle large CSV uploads with minimal latency or overhead.

The service also implements a batch processing mechanism for updating multiple profiles concurrently, with configurable batch sizes and error handling.

Each failed extraction is logged with detailed error information, including the error type, message, and timestamp, which is stored in the alumni's metadata for debugging and monitoring purposes.

4.4.4 Enrichment Services

Once the alumni profile data is extracted, the system performs a series of targeted enrichment operations to complete and contextualize the dataset.

These services operate independently and asynchronously, each focusing on a specific aspect of the profile, such as companies, job roles, or geographic locations.

By decoupling the enrichment logic into dedicated modules, the platform ensures better maintainability, parallel processing, and extensibility for future improvements or third-party API integrations.

Company Enrichment

Once a role's company is detected and determined to be new, the system triggers an enrichment process to collect company metadata, using **Brightdata**¹³ as a third-party API.

This process is handled asynchronously via the company service, which is responsible for resolving company details such as industry, headquarters, founding year, logo, and more.

From there, the following updates are applied:

- **Basic Metadata:** Company name, website (cleaned and normalized), industry (mapped to internal categories), founding year, size, and organization type.
- **Visual Identity:** If a logo URL is provided, the system uploads it to Cloudinary, the platform's chosen blob storage provider, and stores the result.
- **External Links:** The system attempts to validate a matching *Levels.fyi*¹⁴ profile to later support salary benchmarking.
- **Location:** If headquarters and country code are available, the service invokes the Location Agent to resolve and associate a location object to the company.

¹³<https://brightdata.com/>

¹⁴<https://levels.fyi/>

Each company is associated with an industry category - these are sourced from a pre-defined list [37] that was created in advance in the database, and is based on LinkedIn's own industry classification.

To ensure scalability, enrichment is performed in batches of 10, using asynchronous tasks and exponential backoff polling - this is particularly useful when the admin requests the update of a large number of companies, something possible through the API documentation.

This modular procedure enables future integrations with other company intelligence APIs and isolates failures in any one external service, maintaining overall system robustness.

Role Enrichment

The role enrichment service is responsible for parsing the professional experience section of each alumni's LinkedIn profile. For each role, the system constructs two representations:

- A structured version (`Role`) containing metadata such as start and end dates, associated company ID, current status, and a placeholder for the location (initially set as remote).
- A raw version (`RoleRaw`) that preserves the original title, description, and location as free-text input for downstream classification and location enrichment.

This service is triggered during the alumni profile extraction process and once a role is parsed, it is persisted to the database and enriched asynchronously.

Job classification and location resolution are both delegated to dedicated AI agents using the structured and raw representations as input.

Location Enrichment

Location enrichment spans three entity types: alumni, companies, and roles. In all cases, a dedicated Location Agent resolves unstructured location data into normalized database entries.

For roles, the raw free-text location is first stored in the `RoleRaw` table and then processed in batch via `LocationService`, allowing for decoupled enrichment and auditability.

In contrast, for alumni and companies, the agent is called directly with inputs derived from third-party API responses (`Enrich Layer` or `Brightdata`), since these entities lack an intermediate raw storage layer.

This architectural divergence, favoring immediate resolution for some entities and deferred enrichment for others, was a deliberate design choice, which will be revisited in detail in [section 4.7](#).

4.4.5 LinkedIn-Based Authentication

The platform uses LinkedIn OAuth 2.0 with the `r_liteprofile` and `r_emailaddress` scopes to authenticate alumni. However, LinkedIn's Lite Profile API does not expose the user's vanity URL (i.e., the unique identifier following `https://www.linkedin.com/in/`), which is essential and the only fully reliable way to verify identity against the alumni dataset.

As a workaround, after successful login, the platform attempts to retrieve the user’s personal email from the API response and sends a verification token to that address.

If no email is returned, the user is prompted to manually provide one to complete the verification step. This introduces a potential security risk, which is discussed further in section 4.7, along with the rationale behind this design choice.

Only after email verification is the user asked to provide their LinkedIn profile URL. This is then matched against existing alumni records: if a match is found, the user is granted access to additional platform features; otherwise, access remains limited.

4.5 AI-Driven Agents

Large Language Model (LLM) agents are autonomous computational entities capable of combining natural language reasoning with structured tool use. When integrated into software pipelines, they can perform complex decision-making, data transformation, and enrichment tasks in a transparent and modular fashion.

In this work, LLM agents were leveraged to extract structured insights from unstructured alumni role data, tasks that would be prohibitively brittle or labor-intensive using traditional rule-based approaches.

The agent-based system was responsible for classifying job roles into the ESCO taxonomy, estimating the seniority of each role in its career context, and resolving ambiguous free-text locations into structured geographical identifiers.

This section presents the shared architecture that underpins all agents (subsection 4.5.1), followed by detailed descriptions of the seniority (subsection 4.5.2) and location (subsection 4.5.3) agents.

The more complex and central job classification agent is discussed in depth in section 4.6.

4.5.1 Overview of Agent Architecture

To orchestrate the data enrichment pipeline in a scalable and maintainable way, a modular agent-based architecture was adopted, grounded in the LangGraph extension of the LangChain framework¹⁵. Each enrichment task, be it job classification, seniority estimation, or location resolution, is implemented as a directed acyclic graph (DAG), composed of discrete nodes that represent atomic operations. These include database access, cache checks, language model interactions, and schema validations.

LangGraph and Tool-Driven Modularity

LangGraph enables the construction of reliable, auditable LLM agents by abstracting agent behavior into flow a flow of controlled nodes. Nodes can be deterministic (e.g., Redis lookup, DB write) or stochastic (e.g., LLM inference with a tool). LangChain’s `ToolNode` abstraction was

¹⁵<https://langchain-ai.github.io/langgraph/>

heavily relied on, which wraps model output into tool calls that must match a predefined JSON Schema. This ensures that every model output is structurally validated before being committed to the database or passed downstream.

For example, the job classification agent uses tools like `return_esco_choices` as well as `get_detailed_esco_classification`, while the location agent makes use of the tools: `return_geo_resolution` and `return_location_resolution`. By enforcing tool usage, consistent outputs are guaranteed even under the inherent stochasticity of LLMs.

Prompt Engineering and Schema Validation

Each tool call is paired with carefully constructed prompts¹⁶. These prompts are versioned to ensure traceability. When an LLM fails to produce valid output, either semantically or structurally, a retry policy using the `tenacity`¹⁷ library is applied. The retry mechanism leverages exponential backoff and enriches prompts with additional clarification steps on subsequent attempts. This robustifies the system against transient API issues, low-confidence model outputs, or schema mismatches.

Model Selection and Fallback Strategies

To balance performance, accuracy, and cost, two OpenAI models were deployed in production: `gpt-4.1-mini` as the default, and `gpt-4o-mini` as a fallback once token limits (e.g., 200k TPM (tokens per minute)) are reached. The cost implications of each of these models are further detailed in [section 5.2](#).

Each enrichment request tracks the model version used, ensuring auditability. Initial prototypes with Ollama [33], a framework for building and running LLMs locally, using models like `Mistral:7B`, `Qwen:8B`, and `Llama-3-Groq-8B-Tool-Use`, were found inadequate for production due to lower context lengths, performance limitations on constrained hardware (e.g., university-hosted VMs), and poor zero-shot alignment with the ESCO taxonomy.

Caching and Throughput Optimizations

Redis is used extensively across all agents to cache repeated queries, especially for high-cost operations like embedding generation and classification reranking. Cached entries are keyed by query hash and expire after 24 hours. This allows for cost control and improved latency when identical inputs recur, particularly useful when alumni hold common roles (e.g., *Software Engineer*).

To further manage throughput, especially with OpenAI's strict token-per-minute (TPM) limits, a custom `TokenRateLimiter` class was developed. This limiter tracks distributed token usage across processes via Redis, introduces safety buffers, and incorporates jitter-based sleep intervals to prevent request bursts.

This mechanism ensures that no single agent or user can overwhelm the rate cap, and enables the use of cost-effective high-throughput models in production.

¹⁶Visit <https://github.com/carlosverissimo3001/alumni-feup/blob/main/agents-api/app/utils/prompts.py> for the complete prompts.

¹⁷<https://tenacity.readthedocs.io/en/stable/>

Asynchronous Batch Processing

All agents support batch inputs, either directly (as with the seniority agent) or indirectly via parallel async tasks (as in the job classification agent). This was essential for processing hundreds of roles or locations concurrently while maintaining fault isolation. Batch processing also reduced overall latency by minimizing round-trips between the API layer and enrichment agents.

Deployment and Environment Constraints

Agents are deployed in a monolithic FastAPI service, running on a university-hosted virtual machine with limited resources. Given the compute constraints, local inference using larger open-source models was ruled out in favor of remote LLM APIs.

Rate limiting, caching, and tool-enforced minimal prompting allowed the system to operate reliably under these conditions, while keeping token costs and response times manageable - refer to [section 5.2](#) for more details about the underlying costs of running Alumni-FEUP.

Diagram Reference

A high-level diagram of the shared agent architecture, illustrating the LangGraph pipeline, tool invocation structure, retry loops, and Redis/cache interactions, is provided in [Figure 4.8](#).




Figure 4.8: Shared agent architecture diagram

4.5.2 Seniority Agent

The seniority agent estimates the level of experience reflected in each alumni role by analyzing their complete career trajectory rather than evaluating roles in isolation. This contextual approach helps the system disambiguate otherwise ambiguous roles. For instance, a *"Software Engineer"*

title could indicate an entry-level position or a mid-senior-level role depending on the individual's past experience, company context, and progression.

Career-Aware Batch Classification

Each alumni record is passed to the agent as a batch of roles, already enriched with ESCO classifications. These roles include metadata such as title, company, start and end dates, and a total experience indicator.

The roles are arranged chronologically and embedded into a structured prompt which is passed to a language model (either `gpt-4.1-mini` or, if necessary, `gpt-4o-mini`).

This setup allows the model to identify career progression patterns (e.g., from INTERN to ASSOCIATE to DIRECTOR) and estimate seniority in a way that simple keyword-based rules cannot replicate.

All role entries are included in the batch, even those already labeled, to provide maximum context.

Tool-Based Structured Output

The model is required to respond via a LangChain tool invocation (`return_seniority_choices`), which must match a strict JSON schema. This schema requires the model to output a list of results, each containing:

- `role_id`: a unique identifier
- `seniority`: one of the predefined enum levels
- `confidence`: a float score from 0.0 to 1.0
- `reasoning`: a concise explanation grounded in job title, experience, and context

Although the confidence score is derived solely from the model's internal reasoning and is not grounded in a deterministic metric, it is stored in the database to support future auditing and longitudinal analysis.

Users can review and modify the predicted seniority levels for their roles via the profile page, using the edit interface shown in [Figure 4.9](#).

This interface displays the model's original reasoning and confidence score, offering transparency into the classification process and helping users assess the reliability of the output.




Figure 4.9: Edit seniority modal

Supported Seniority Levels

The full list of supported seniority levels is shown in [Table 4.1](#). These levels are inspired by LinkedIn’s official seniority taxonomy, with some simplifications to reduce classification ambiguity.

Level	Description
INTERN	Initial professional experience
ENTRY-LEVEL	First full-time position
ASSOCIATE	Early career professional
MID-SENIOR	Experienced professional
DIRECTOR	Director of a department or team
C-LEVEL	Chief executive positions
EXECUTIVE	Senior management roles

Table 4.1: Predefined seniority levels used by the agent

Retries and Default Behavior

If the model fails to comply with the schema (e.g., due to malformed output or invalid enum values), a retry loop is triggered. This mechanism, powered by the `tenacity` library, enriches the prompt on each retry by appending additional instructions and formatting guidance. The model has up to three attempts to respond correctly.

In rare cases where the model continues to fail, the agent assigns the default value `ASSOCIATE`, reflecting a conservative midpoint in the seniority spectrum. These fallback assignments are flagged for potential future review.

Output Validation and Persistence

Once valid tool output is received, the agent updates each role in the database with the predicted seniority level, confidence score, and model reasoning. These entries are stored alongside the LLM version and can be audited or reclassified in future iterations.

Metadata is stored under a dedicated JSON field to support longitudinal analysis of classification confidence and model drift.

Why Not Rules-Based Logic

While a rules-based system could trivially detect titles like *Senior Manager* or *Intern*, such an approach lacks the nuance required for ambiguous or title-inflated roles. It also struggles with edge cases (e.g., *Founders*, *Consultant*, *Engineer II*). The use of LLMs allows the system to weigh experience, chronology, and linguistic subtleties, significantly reducing misclassifications compared to naive heuristics.

Diagram Reference

The seniority agent's graph-based flow, illustrating the batch input, prompt generation, tool call, retries, and DB update, is shown in [Figure 4.10](#).




Figure 4.10: Seniority agent architecture diagram

4.5.3 Location Agent

LinkedIn allows users to enter free-text location fields for roles, which are frequently informal, ambiguous, or inconsistent.

While alumni profiles and company headquarters typically include structured metadata such as city and country, role locations often contain expressions like *"Portugal and EMEA"* or *"Remote in LATAM"*, which cannot be resolved deterministically.

To address this, a dedicated location agent was implemented. It processes each input through a LangGraph-based pipeline that leverages LLM tool calls, structured schema validation, caching, and database lookups.

The goal is to transform arbitrary free-text into structured entities composed of city, country, and country code.

Resolution Pipeline

The agent operates in two phases. First, it uses a LLM tool (`return_geo_resolution`) to extract a city and *ISO 3166-1 alpha-2* country (two-letter identifier) from the raw input. Then, using `return_location_resolution`, it attempts to match this result against known entries in the location database.

If no match is found, the model is prompted to select or create the closest valid option from a pre-fetched candidate set. This approach ensures reusability of existing entries and avoids unnecessary duplication. If a country is identified but not the city, the agent is instructed to fallback to an abstract entry like `Other - Country`.

Fallback and Remote Logic

Currently, the agent uses a predefined entry labeled `REMOTE` (mapped to coordinates 0,0) whenever no structured location information can be resolved. However, this approach conflates genuinely remote roles with cases where the location data is simply unknown or missing. As the system relies on external data sources such as Enrich Layer, which do not explicitly differentiate between these scenarios, it is not currently possible to distinguish between *remote* and *unknown* locations.

A potential future improvement would involve tracking and labeling these cases separately to improve data accuracy, should more granular data become available.

Entity Integration

Once resolution is complete, the resulting object is linked to the relevant domain - alumni, role, or company:

- **Alumni:** If no structured location is found from LinkedIn scraping, the agent is invoked using city and country fields as hints.
- **Role:** Raw role locations (e.g., Sao Paulo Metropolitan Area) are stored in a dedicated table and resolved post-parsing.
- **Company:** When a new company is inserted, the headquarters and country code fields are used to trigger location resolution.

Once linked, the agent also triggers a background task to query OpenWeatherMap's API for geographical coordinates, enabling geospatial filtering and analysis.

Caching and Throughput Optimization

The agent is optimized for large-scale operation. Redis is used to cache recent resolutions, minimizing repeated calls for common inputs. Requests are processed in batches with sleep intervals between them to respect OpenAI rate limits, and failures are logged for future review.

Reliability and Validation

As with all other agents in the system, outputs are validated against strict JSON schemas. This ensures predictable formatting and prevents downstream failures. In cases of invalid output, the retry mechanism attempts to reformulate the prompt and recover up to three times before falling back to safe defaults.

Architecture Overview

Figure 4.11 shows a high-level view of the location agent's architecture and control flow. It demonstrates the LLM reasoning, database matching and fallback mechanisms.




Figure 4.11: Location agent architecture

4.6 Job Classification Agent

The job classification agent serves as the cornerstone of the enrichment pipeline. It transforms free-text role titles and descriptions into standardized occupational codes aligned with the ESCO taxonomy.

This process enables structured analysis of alumni career paths, aggregation across similar roles, and consistency in downstream visualizations and statistics.

In the subsections that follow, we describe the motivation behind this component, the ESCO taxonomy it builds upon, the multi-stage architecture powering its decisions, and its practical performance in real-world scenarios.

4.6.1 Motivation and Role of the Agent

The classification of job titles is a fundamental requirement for making sense of heterogeneous, unstructured career data. Alumni-entered roles on platforms like LinkedIn vary drastically in language, specificity, seniority, and format.

A single job function, such as software development, may be described in countless ways: *"Software Engineer"*, *"Developer"*, *"Backend Developer"*, *"Engineer II"*, or even *"Ninja Coder"*, for those who wish to be creative.

This lexical diversity poses a major challenge for any form of aggregation, analytics, or visualization. Without classification, comparing roles or tracking trends across alumni becomes unfeasible.

The job classification agent solves this problem by mapping each role to a standardized occupational category from the ESCO taxonomy.

It processes one role at a time, transforming raw titles and descriptions into structured, explainable, and auditable classifications.

These ESCO codes form the backbone of every downstream insight in the platform - from the global map to temporal evolution graphs.

Crucially, the agent removes the need for manually curated rules, hardcoded mappings, or keyword-based heuristics. It leverages vector embeddings, semantic similarity, reranking, and LLM-based reasoning to make high-confidence predictions that generalize across industries and job families.

As the most mature and central component of the AI enrichment pipeline, the job classification agent plays a pivotal role in enabling multidimensional analysis of alumni careers.

4.6.2 Leveraging the ESCO Taxonomy

As detailed in [section 2.6.1](#), the ESCO (European Skills, Competences, Qualifications and Occupations) taxonomy is a multilingual, skill-enriched occupational framework maintained by the European Commission.

ESCO extends the ISCO-08 standard with deeper granularity, offering over 3,000 occupation profiles across 28 languages. Each entry includes alternative labels, definitions, tasks, and related skills. ESCO's first four hierarchical levels mirror ISCO-08, ensuring interoperability, while additional skill-specific levels (Level 5 and beyond) enable nuanced occupational distinctions.

Classification Approach

In this work, roles are matched directly to ESCO leaf nodes, either terminal nodes or parent occupations that remain semantically precise (e.g., Financial Analyst, Software Developer). This balances classification specificity with interpretability.

Each ESCO occupation is stored in a dedicated table with precomputed embeddings using a multilingual model by OpenAI (`text-embedding-3-large` [34]). Embeddings are generated from key English-language fields (title, description, labels, tasks), but multilingual support is maintained for downstream compatibility.

During classification, only filtered leaf nodes are considered to ensure precise and interpretable matches, avoiding overly broad or generic assignments (see Figure 4.12).

Role Distribution and Hierarchy

A total of 3039 ESCO leaf nodes were used for classification. Most roles reside at Level 5 (6160 entries), followed by Level 6 (1832), with a small fraction at deeper levels 7 and 8.

Table 4.2 compares the distribution of classified alumni roles against the distribution of available ESCO leaf nodes by level, highlighting real-world usage patterns versus taxonomy structure.

This high granularity allows results to be grouped flexibly, from detailed ESCO roles to higher-level ISCO categories, supporting both fine-grained and aggregate analytics (see Figure 4.13).

Although deeper levels improve expressiveness, they also introduce additional classification complexity, as discussed later in section 5.4.




Figure 4.12: Logic used to determine valid leaf-level ESCO nodes




Figure 4.13: ESCO grouping analytics

Level	Number of roles	Percentage	Number of nodes	Percentage
Level 5	6160	75.86%	1795	59.07%
Level 6	1832	22.56%	1067	35.11%
Level 7	112	1.38%	136	4.47%
Level 8	16	0.19%	41	1.35%
Total	8120	100.00%	3039	100.00%

Table 4.2: Comparison of role classification counts and ESCO node availability by level

Future iterations may leverage multilingual ESCO entries directly and explore mitigation strategies for LLM bias toward common job roles during classification - limitations are further explored in [section 6.3](#).

4.6.3 Architecture and Workflow

The job classification pipeline operates as a background agent triggered automatically whenever alumni records are inserted or updated.

Classification operates at the role level, without context from the alumni's career trajectory or previous positions, ensuring modularity and reproducibility of each decision.

Trigger and Invocation

When a new alumni profile is created or updated, the backend system triggers a background job that offloads classification to the job classification agent. This ensures that even computationally expensive processing does not delay user-facing operations.

Historical classification data is overwritten in update scenarios, ensuring a single source of truth.

Workflow Overview

Figure 4.15 illustrates the full classification workflow. The process consists of three core stages:

1. **Retrieval:** The agent constructs an embedding query using the role's title and optional description - these are stored in the corresponding `role_raw` entry for this role. It retrieves the top 20 semantically similar ESCO leaf nodes from the `esco_classification` table using pgvector-based cosine similarity search. The embedding query is constructed using the `text-embedding-3-large` model by OpenAI.
2. **Reranking and Validation:** The retrieved candidates are reranked via a cross-encoder model (`ms-marco-MiniLM-L-6-v2`) to refine ordering. The top five are passed to the LLM, which selects the best three using a tool-based function call, along with reasoning and confidence scores.
3. **Persistence:** The best match is stored in the `job_classification` table, linked to the role via a foreign key. The full LLM output, including the reasoning and top-three choices, is preserved in a structured `metadata` JSON field for future auditing and transparency, allowing the role owner to inspect the classification and update it if necessary.

In order to enhance classification precision, the system applies a reranking step using a cross-encoder model (`ms-marco-MiniLM-L-6-v2`) on the top 20 ESCO matches retrieved via vector search. This step ensures that the final candidates presented to the LLM are semantically relevant and well-aligned, reducing the risk of irrelevant classifications due to embedding noise.

The LLM is then tasked with selecting the best match and returning structured metadata, including the confidence score, reasoning, and alternative candidates. This metadata is stored in the database and later surfaced in the user-facing application, allowing alumni to review and override their job classification, as shown in Figure 4.14.

This entire flow is built atop LangGraph's `StateGraph` abstraction, allowing composition of modular stages and clear retry logic for LLM calls. While the vector search and reranking steps are deterministic and executed synchronously, the GPT validation step includes exponential backoff and is the only component capable of triggering retries.

Update Your Job Classification

ESCO is the European multilingual classification of Skills, Competences, Qualifications and Occupations. Our AI has classified your role, but you can refine it below.

Current Classification

Software Developer2512.4

Confidence68%

LowMediumHigh

AI Classification Details

Classified by gpt-4.1-mini

Alternative Classifications:

Application Engineer71%

Integration Engineer69%

Reasoning: The role 'Software Engineer Intern' aligns best with 'Software Developer' as it directly reflects the core responsibilities of software engineering and development, making it the strongest match. 'Application Engineer' is also relevant as it involves designing and developing software applications, fitting well with the intern role in software engineering. 'Integration Engineer' is the third best match because it involves combining software components, which is a common task for software engineers, especially in a financial services context. Other candidates like 'Automation Engineer' and 'Cloud Engineer' are more specialized and less directly aligned with a general software engineering internship.

Show less ^

Select the specific job classification that best represents your role:

Software Developer

CancelUpdate Classification

Figure 4.14: Edit Classification modal showing alternatives, confidence, and LLM reasoning




Figure 4.15: End-to-end job classification workflow and architecture

Storage and Versioning

Each classification is stored in a dedicated `job_classification` table, which maintains the ESCO ID, confidence score, model used, and user override flags (e.g., if the user accepted or modified the suggestion), which allows for fine-grained metrics and future learning loops.

The ESCO hierarchy itself is static, only updated if new roles are introduced by the European Commission.

4.7 Challenges and Design Decisions

Throughout the development of the AI-driven classification pipeline, several design decisions were shaped by real-world constraints and tradeoffs. This section provides a high-level overview of the most relevant challenges faced, leaving deeper technical and ethical discussions for [section 6.3](#).

LinkedIn Enrichment and Legal Considerations

Direct access to the official LinkedIn API is not granted to third-party developers. As a result, the system relies on Enrich Layer, a third-party enrichment provider that uses cached LinkedIn profile data (typically less than 3 months old). While this mitigates scraping complexity, it raises questions around long-term sustainability and legal ambiguity - issues inherited from Enrich Layer's own data acquisition practices.

LLM Strategy: OpenAI vs Local Models

Initial experiments with local models such as `mistral:7b` via Ollama demonstrated that while self-hosting was cost-effective, the semantic reasoning and output structuring fell significantly short of production requirements. Furthermore, resource constraints on the deployment virtual machine ruled out larger local models. Ultimately, OpenAI's `gpt-4.1-mini` was selected for its performance, with total classification costs remaining under \$10.

ESCO Taxonomy Integration

One of the early challenges was identifying the correct subset of ESCO nodes to classify against. Naively restricting to terminal nodes would omit semantically important categories (e.g., *"Software Developer"*). The solution involved defining a curated set of leaf nodes that balance specificity with generality, as shown in [subsection 4.6.2](#).

Scalability and Rate Limiting

Classifying roles across thousands of alumni imposes load on both vector search and LLM endpoints. A token refilling strategy was implemented to throttle classification tasks and avoid rate limit violations, particularly during bulk alumni ingestion.

Retry Logic and Workflow Robustness

The only stage of the pipeline that supports retries is the GPT tool call, due to its higher latency and occasional transient failures. Vector search and reranking remain synchronous and deterministic, simplifying monitoring and debugging.

Further details, and other challenges are further explored in [section 6.3](#).

Chapter 5

Results and Discussion

This chapter begins with an overview of the dataset collected throughout the project ([section 5.1](#)), detailing the degrees tracked, LinkedIn coverage, and breadth of roles, companies, and classifications.

It then presents the evaluation of the platform’s enrichment agents: job classification and seniority estimation ([section 5.3](#)), responsible for transforming LinkedIn data into structured insights.

The evaluation combines quantitative and qualitative methods. Quantitatively, it includes alumni feedback, manual role review, and metrics like cosine similarity and model confidence. Qualitatively, it draws on feedback from a live focus group, structured forms, and System Usability Scale (SUS) [43] scoring.

The remaining sections cover the evaluation methodology, datasets, findings, and a discussion of limitations and open issues.

5.1 Data Coverage and Composition

The final dataset aggregates **3,042 graduates** drawn from the seven degrees currently uploaded to the platform: five in Informatics/Computer Engineering, the M.Sc. in Software Engineering (MESW), and the Ph.D. in Digital Media (PDMD), the latter two being added during this project.

[Figure 5.1](#) shows the distribution of alumni per degree; the long-standing MIEIC programme, active from 2006 through 2021, naturally dominates, but every programme is represented.

All graduates are loaded into the database; however, only those for whom a LinkedIn profile could be collected are visible in the platform interface.




Figure 5.1: Number of graduates per supported degree

LinkedIn coverage

Of the 3 042 alumni, **1 353** (44.5 %) have a valid LinkedIn profile that could be matched with their academic record; the remaining 1 689 lack an accessible URL (Figure 5.2).




Figure 5.2: LinkedIn profile availability across all graduates

The valid URLs originate from three distinct collection channels (Figure 5.3):

- **Previous platform dataset** - 953 links (70.4 %), inherited from the original AlumniEI World project and cross-referenced with the official LinkedIn alumni group;
- **Manual collection** - 271 links (20.0 %), added by the author through systematic cross-matching between institutional records and public LinkedIn profiles;

- **New degrees** - 129 links (9.5 %), obtained while onboarding MESW and PDMD during this project.




Figure 5.3: Source distribution of the 1 353 LinkedIn URLs used for enrichment

Figure 5.4 shows MESW and PDMD already surpass 75% LinkedIn coverage, with LEIC (pre-Bologna) at 73%, and older degrees like MEI and MIEIC at 68% and 48%.

This gap results from recent degrees being validated by administrators, while older ones used legacy datasets and manual matching. Most early data collection targeted long-standing programs (MIEIC, LEIC, MEI).

Newer degrees (M.EIC, L.EIC) have lower coverage because collection efforts were limited, and 2024 graduates were added during development, so their LinkedIn profiles were not yet collected.




Figure 5.4: LinkedIn profile availability by degree

Enriched entities

Processing the 1,353 LinkedIn profiles yielded the results summarized in [Table 5.1](#).

Metric	Value
Total roles (current)	8,120 (1,510)
Distinct companies	2,471
Unique ESCO classifications	640
Countries represented	52
Cities represented	279
Distinct industries	164

Table 5.1: Summary of extracted alumni career data

[Table 5.2](#) lists the ten most common historical ESCO roles, with current top roles in [Table 5.3](#). *Software Developer* leads by a large margin, but other technical roles like *Application Engineer* and *Software Architect* are also frequent. Consulting and academic positions (*ICT Consultant*, *Teaching Assistant*) further reflect the diversity of alumni career paths.

ESCO Classification	Code	Roles
Software Developer	2512.4	1,177
Application Engineer	2149.2	420
ICT Consultant	2511.10	248
Mobile Application Developer	2514.2.2	244
Web Developer	2513.5	220
Software Architect	2513.2	219
Integration Engineer	2511.17	160
University Teaching Assistant	2310.1.42	154
Project Manager	1219.6	145
Business Consultant	2421.2	142
...	...	...
Total		8,120

Table 5.2: Top 10 ESCO classifications represented in the dataset (all historical roles)

ESCO Classification	Code	Roles
Software Developer	2512.4	232
Application Engineer	2149.2	113
Software Architect	2512.3	56
Technical Director	2654.1.7	36
Chief Technology Officer	1330.3	32
Software Manager	1330.5.1	27
Enterprise Architect	2511.6	26
Data Engineer	2511.20	25
Assistant Lecturer	2310.1.5	25
Mobile Application Developer	2514.2.2	25
...	...	...
Total		1,510

Table 5.3: Top 10 ESCO classifications for alumni's current (latest) roles

Comparison of historical and current role distributions. Comparing Table 5.2 and Table 5.3 shows that while technical roles like *Software Developer* dominate both current and historical data, managerial roles (e.g., *Technical Director*, *CTO*) are more common among current positions, reflecting career progression. Early-career roles (e.g., *Teaching Assistant*, *Consultant*) are more frequent in historical data, illustrating initial career steps.

5.2 Associated Costs

The Alumni-FEUP platform relies on several third-party APIs. Most costs stem from data enrichment:

- **Enrich Layer:** \$81.18 for 1,353 LinkedIn profiles (using caching to minimize fees).
- **BrightData:** \$3.71 to enrich 2,471 company records.
- **OpenAI:** \$16.55 for all AI agent calls and embeddings across 8,119 roles.
- **Cloudinary, Vercel, Redis, PostgreSQL, and OpenWeatherMap** were used within free tiers or self-hosted at no cost.

Table 5.4 summarizes main expenses. Two enrichment cycles resulted in a total estimated cost of \$203. See Table 5.5 for OpenAI usage details.

Other models were tested but not used in production, so their costs are excluded.

Service	Purpose	Estimated Cost
Enrich Layer	LinkedIn profile scraping	\$81.18
BrightData	Company enrichment	\$3.71
OpenAI APIs	Classification, enrichment, embeddings	\$16.55
Total		\$101.44

Table 5.4: Estimated cost breakdown of external services

Model	Cost	Input Tokens (Millions)	Output Tokens (Millions)
gpt-4.1-mini	~\$13.66	~22.2	~4.4
gpt-4o-mini	~\$2.22	~11.8	~1.3
text-embedding-3-large	~\$0.16	~1.23	
Total	~\$16.55	~34.3	~5.7

Table 5.5: Estimated OpenAI cost breakdown for the enrichment agents

5.3 Evaluation Methodology and Experimental Setup

Evaluation combined quantitative and qualitative methods using production data (**1,353 alumni, 8,120 roles, 2,471 companies**).

Quantitatively, alumni validated AI classifications via a profile modal; results were logged for analysis. Due to low participation, the author also conducted a manual review of sampled roles by seniority and ESCO group.

Qualitatively, a feedback form (Appendix B.1) adapted from the System Usability Scale assessed usability, filtering, and visualizations. Users also gave free-text suggestions.

Finally, Vercel Analytics captured technical metrics like page load and interaction delays to supplement user feedback.




Figure 5.5: Role validation modal displayed in user profiles

5.4 Quantitative Results

This section quantifies the enrichment agents' performance in classifying job titles and estimating seniority.

Results are based on:

- Direct user validations submitted via the platform,
- Confidence scores associated with each AI-generated prediction,
- Manual review of a stratified sample conducted by the author,
- Structured feedback collected through the alumni form.

These combined sources provide a robust assessment in real and curated contexts. Telemetry data from Vercel Analytics was also considered for frontend responsiveness.

5.4.1 User-Submitted Validation

User-submitted validation directly captures alumni perceptions of the AI's classifications for their own profiles. This voluntary feedback offers valuable, expert insight into real-world classification accuracy.

Job Title Classification

A total of 37 job roles were evaluated by users through the platform's validation interface. These roles were automatically classified using the ESCO taxonomy and reviewed by alumni for accuracy.

Of the 37 roles, 29 (**78.4 %**) were confirmed as correct, while 8 (**21.6 %**) were marked as inaccurate (see [Table 5.6](#)). Notably, rejected cases often involved ambiguous titles such as "*Business Developer*" or "*Strategic Planning Manager*", likely due to vague LinkedIn descriptions or inherently fuzzy job titles.

Technical roles (e.g., *Software Developer*, *AI Engineer*) consistently mapped well to ESCO codes. In contrast, managerial and cross-functional positions showed more variability, reflecting the challenge of classifying less standardized roles.

Validation Outcome	Count	Proportion (%)
Accepted by the user	29	78.4
Rejected by the user	8	21.6
Total	37	100.0

Table 5.6: User validation outcomes for AI-generated job title classifications

Further analysis by ESCO hierarchy ([Table 5.7](#)) shows performance was higher for general Level 5 roles (87.5% accuracy) than for specific Level 6 roles (61.5%), highlighting increased error rates as classification becomes more granular. This aligns with expectations: subtle differences in job scope or organizational context make narrow taxonomic mapping harder.

ESCO Level	Validated Roles	Correct	Accuracy (%)
Level 5	24	21	87.5
Level 6	13	8	61.5
Total	37	29	78.4

Table 5.7: Validation accuracy by ESCO classification level

Overall, the observed **78.4% accuracy** reflects a baseline result, achieved with simple prompting and limited tuning. These findings reveal a trade-off between classification granularity and robustness, and point to clear avenues for future improvement, such as richer input features and better prompt engineering.

Seniority Classification

User validations assessed how well the AI estimated seniority levels. Of 36 reviewed roles, 30 (**83.3%**) were confirmed as correct and 6 (**16.7%**) as inaccurate.

The most represented seniority level among the reviewed roles was *Mid-Senior Level*, which also showed a strong validation rate of 87.5%. Other categories, such as *Intern*, *Director*, and *C-Level*, achieved full or near-full agreement from users (see [Table 5.8](#)). In contrast, the *Associate* level presented the lowest agreement rate, with only 71.4% of the roles validated as correct.

This likely reflects the label’s ambiguity, *Associate* means different things in different industries (e.g., transitional in tech, senior in finance), complicating both user perception and AI classification.

Seniority Level	Total Validated	Correct	Accuracy (%)
Intern	6	5	83.3
Entry-Level	4	3	75.0
Associate	7	5	71.4
Mid-Senior Level	11	9	81.8
Director	3	3	100.0
Executive	1	1	100.0
C-Level	4	4	100.0
Total	36	30	83.3

Table 5.8: Validation outcomes for seniority classification across different levels

These results indicate strong overall performance, but highlight challenges with ambiguous or overlapping levels, especially for *Associate*. The seniority agent currently does not consider industry-specific context, which may explain some misclassifications. Future work should incorporate such context and enhance transparency, as users often wanted clearer reasoning behind each classification (see [subsection 5.5.5](#)).

5.4.2 LLM-based Confidence Scores

The enrichment agents generate job title and seniority classifications, each with an associated confidence score stored in the database. Job classification confidence is based on embedding similarity, while seniority confidence reflects the LLM agent’s own certainty.

Job Classification Evaluation

Each role’s job classification is paired with a confidence score derived from vector-based semantic similarity between the role description and each ESCO occupation.

The pipeline first generates a high-dimensional embedding for each role using a state-of-the-art transformer model, and compares it to all ESCO occupation embeddings stored in the database. Similarity is computed using cosine distance, resulting in a raw similarity score between the role and each ESCO occupation.

Confidence Calculation: To improve interpretability and reliability, the raw similarity is post-processed into a confidence score that reflects both the absolute match strength and the relative distance to the best match. Specifically, the confidence for each candidate is computed as:

$$\text{Confidence} = \text{Scaling factor} \times \text{Absolute similarity} \times \text{Relative similarity}$$

Where:

- **Absolute similarity:** $(2.0 - d)/2.0$, with d being the cosine distance between the query and the ESCO occupation (range: 0–1).
- **Relative similarity:** $(2.0 - d)/(2.0 - d_{\text{best}})$, comparing each candidate to the best match (d_{best}).
- **Scaling:** Ensures reasonable distribution; final confidence is rounded to two decimals.

The full code listing is provided in Appendix A.1.

Interpretation: The confidence score is not a strict probability of correctness. Instead, it reflects the semantic similarity between the LinkedIn role (and description) and the closest ESCO occupation, serving as a proxy for overlap rather than true classification accuracy. Lower scores indicate weaker matches, helping prevent overconfidence in ambiguous cases (see Table 5.9).

Score Range	Interpretation
> 0.8	Very strong match
0.6–0.8	Plausible match
< 0.6	Review carefully

Table 5.9: Interpretation of job classification confidence scores

Distribution: Table 5.10 shows confidence scores grouped by ISCO sub-major categories, with most alumni found in ICT, science, engineering, and management roles, which matches their academic backgrounds.

The average confidence per group is usually between 67% and 71%. While this may seem moderate, the metric is intentionally conservative to reduce overconfidence, especially for ambiguous or overlapping roles.

In areas with many similar job titles, such as software or business, even a 0.6–0.8 score can represent a solid match.

This distribution highlights both the diversity of roles and the robust, careful design of the classification approach. More details are available in Appendix B.2.1.

ISCO Group	# Roles	Avg (%)	Min (%)	Max (%)
Information and Communications Technology Professionals	3,611	70.0	57.0	85.0
Science and Engineering Professionals	966	69.3	56.0	83.0
Administrative and Commercial Managers	777	70.3	54.0	85.0
Teaching Professionals	691	70.0	55.0	83.0
Production and Specialized Services Managers	555	70.2	57.0	84.0
Business and Administration Professionals	552	70.2	56.0	85.0
Legal, Social and Cultural Professionals	235	67.9	55.0	83.0
Chief Executives, Senior Officials and Legislators	141	67.0	56.0	82.0
Business and Administration Associate Professionals	125	67.0	53.0	79.0
Science and Engineering Associate Professionals	116	67.4	59.0	86.0
...	...	...	...	...
Total	8,120	69.7	53.0	86.0

Table 5.10: Job classification confidence by ISCO Sub-Major Group (top 10 by count)

Seniority Level Evaluation

The LLM-based confidence score reflects the model’s internal estimate of certainty when assigning a seniority level to a role, a non-deterministic process.

While these scores provide a useful indication of prediction reliability, they are not strictly calibrated, nor are they directly comparable to the embedding-based confidence used in job classification.

Accordingly, they are primarily reported to contextualize performance across different experience levels.

Seniority Level	# Roles	Avg. Confidence (%)	Min (%)	Max (%)
Mid-Senior Level	3,165	90.0	60.0	100.0
Associate	2,081	82.2	50.0	100.0
Entry Level	1,400	86.6	60.0	95.0
Intern	686	95.7	70.0	100.0
Director	431	91.8	70.0	100.0
C-Level	232	94.7	75.0	100.0
Executive	81	91.0	70.0	98.0

Table 5.11: Distribution of LLM-based seniority classification confidence by level

As shown in [Table 5.11](#), the LLM generally assigned high confidence to most roles, with averages above 0.8 for all levels. Notably, *Intern* and *C-Level* positions exhibited particularly high mean confidence, likely reflecting their more distinctive descriptions. Conversely, *Associate* roles

showed the lowest average and minimum confidence, aligning with known ambiguities and varied use of the term across industries.

It is important to note that high confidence scores do not always equate to accurate predictions. To rigorously assess the practical effectiveness of the seniority classification agent, a structured manual evaluation will be performed.

This evaluation will systematically examine a representative sample of roles, stratified by seniority level and confidence score, allowing for a nuanced analysis of both correct and incorrect assignments. The methodology and detailed findings of this manual evaluation are presented in the following subsection.

5.4.3 Manual Evaluation

This subsection details the methodology for selecting roles to be manually evaluated and presents results for both job classification and seniority level predictions.

Roles were sampled to ensure stratified coverage by seniority, ESCO job code, and job classification confidence. This approach increases the representativeness of the manual evaluation subset and ensures that both frequent and rare categories, as well as a range of model certainty, are reflected in the assessment.

Sampling Methodology

To ensure a representative and reproducible sample, a stratified sampling procedure was employed:

- **Stratification:** Roles were grouped by *seniority level* and by their assigned *ESCO code*. Within each stratum, roles were further divided into confidence buckets (high, medium, low).
- **Random selection:** Roles were sampled within each (seniority, ESCO code, confidence bucket) group to ensure broad coverage across all strata, while maintaining a total sample size of 200. Due to uneven group sizes and practical constraints, some strata are more heavily represented than others, especially the most common seniority levels and job families. Rare combinations were included wherever possible. The final pool was randomly shuffled to avoid order bias and constitutes approximately **2.5%** of the total role count ($n = 8,120$).
- **Reproducibility:** Sampling was performed using deterministic hashing to ensure results can be exactly replicated on the current dataset.

Subset Composition

The resulting evaluation subset is broadly representative of the platform's alumni data, but slightly over-represents mid-senior and associate levels, which aligns with the workforce distribution of FEUP graduates.

Table 5.12 cross-tabulates both the seniority level and the confidence bucket, which is defined as follows:

- **High:** Confidence score greater than 0.8.
- **Medium:** Confidence score between 0.6 and 0.8.
- **Low:** Confidence score less than 0.6.

Table 5.13 shows the five most and least common ESCO job titles in the sample.

Seniority Level	High	Medium	Low	Total
MID_SENIOR_LEVEL	0	60	4	64
ASSOCIATE	6	50	1	57
ENTRY_LEVEL	3	33	0	36
DIRECTOR	1	16	0	17
INTERN	0	15	2	17
C_LEVEL	0	6	0	6
EXECUTIVE	0	3	0	3
Total	10	183	7	200

Table 5.12: Joint distribution of roles by seniority and confidence bucket

Most Common		Least Common	
ESCO Title	Count	ESCO Title	Count
Software Developer	9	Monitoring And Evaluation Officer	1
Web Developer	9	Database Developer	1
Mobile Application Developer	7	University Department Head	1
Application Engineer	6	Leisure Attendant	1
Integration Engineer	5	Employee Volunteering Programme Coordinator	1

Table 5.13: Five most and least common ESCO job titles in the manual evaluation sample

Taken together, these tables confirm that the manual evaluation sample covers a diverse set of roles, industries, and levels of model certainty. This supports a robust, fair, and unbiased assessment of the system's real-world performance.

Job Classification Evaluation

In a manual review of 200 roles, 140 (**70%**) were classified correctly, demonstrating strong overall reliability across a diverse, stratified sample. This aggregate figure reflects the system's overall reliability across a diverse, stratified sample.

Common roles (e.g., Web Developer, Software Developer) were classified with perfect accuracy (Table 5.14). Related titles (e.g., Mobile Application Developer, Application Engineer) also achieved high scores, though some ambiguity remains, especially for emerging or overlapping job

titles such as Integration Engineer and Cloud Software Developer, where taxonomy granularity or coverage limits accuracy.

A key challenge stems from ESCO itself: many real-world job titles map reasonably to multiple ESCO categories, depending on organizational context. For example, distinctions between Application Engineer and Software Developer are often blurred, resulting in some Software Engineer positions mapped as Application Engineer. Roles like Software Engineer (Cloud) were generally assigned to Cloud Software Developer, showing both taxonomy specificity and its limits. In edge cases, the best ESCO code may only partially reflect specialized job titles.

See Appendix B.2.2 for representative examples, including ambiguous mappings (e.g., IT Instructors assigned as Sports Instructor) due to missing domain-specific categories.

Rare roles (Table 5.15) showed a wider spread in accuracy: some (e.g., Membership Administrator) mapped well, while others (e.g., Database Developer) did not.

Classification accuracy by seniority (Table 5.16) was highest for C-level (83%) and associate roles (79%), but lower for directors (53%) and executives (33%). Other levels cluster between 65–72%. The system is most accurate for standardized titles and struggles with ambiguous or blended responsibilities, especially at senior levels.

Overall, the ESCO-based pipeline achieves solid accuracy for common tech roles but is limited by taxonomy design and the evolving nature of job titles. These results underscore the need for ongoing taxonomy updates, more diverse training data, and hybrid or multi-label approaches for edge cases.

ESCO Title	N	ESCO Correct (%)
Web Developer	9	100
Software Developer	9	100
Mobile Application Developer	7	86
Application Engineer	6	67
Business Consultant	5	100
Integration Engineer	5	40
Product Development Manager	5	60
Cloud Software Developer	4	25
Ict Consultant	4	100
Artificial Intelligence Engineer	4	50

Table 5.14: ESCO classification accuracy for the 10 most common roles in the evaluation sample

ESCO Title	N	ESCO Correct (%)
Membership Administrator	1	100
Financial Risk Manager	1	100
Monitoring And Evaluation Officer	1	100
University Department Head	1	100
Database Developer	1	0
Leisure Attendant	1	0
Employee Volunteering Programme Coordinator	1	0
Computer Network And Systems Engineer	1	100
Assistant Lecturer	1	0
Computer Network Technician	1	100

Table 5.15: ESCO classification accuracy for the 10 least common roles in the evaluation sample

Seniority Level	N	ESCO Correct (%)
ASSOCIATE	57	79
C_LEVEL	6	83
DIRECTOR	17	53
ENTRY_LEVEL	36	72
EXECUTIVE	3	33
INTERN	17	65
MID_SENIOR_LEVEL	64	67
Total	200	70

Table 5.16: Job classification (ESCO) accuracy by seniority level in the manual evaluation subset

Seniority Level Evaluation

180 out of 200 roles (**90%**) were assigned the correct seniority level according to manual review.

While the overall seniority classification accuracy is high, results vary across experience levels, as shown in [Table 5.17](#).

The model performed best for Mid-Senior Level (95%), Entry Level (92%), and Executive (100%), indicating strong capability in recognizing clear-cut cases at both extremes of the seniority spectrum.

Associate, Director, and Intern roles showed slightly lower accuracy (84–88%), likely reflecting the broader semantic range and overlap these titles have across industries.

Notably, classification errors were most common between adjacent seniority tiers.

For example, "Director" and "Head" positions were sometimes misclassified as C-level roles, and vice versa, due to variability in how organizations define these titles and the significant overlap in their real-world responsibilities.

Similarly, Director and Executive roles proved particularly challenging to distinguish, as the distinction between these senior positions is often blurred in practice, with organizations using these labels interchangeably or with differing scopes of authority.

These results suggest that, while the LLM-based agent reliably captures most seniority distinctions, certain ambiguous or organization-specific titles remain challenging, demonstrating the subjectivity and evolving nature of job hierarchies in professional settings.

Seniority Level	N	Correct (%)
ASSOCIATE	57	84
C_LEVEL	6	83
DIRECTOR	17	88
ENTRY_LEVEL	36	92
EXECUTIVE	3	100
INTERN	17	88
MID_SENIOR_LEVEL	64	95
Total	200	90

Table 5.17: Seniority classification accuracy by level in the manual evaluation subset

5.4.4 Form Feedback on Classification Quality

To assess the perceived accuracy of the AI-based enrichment agents, the feedback form included specific questions regarding the classification of job titles and seniority levels. These questions were only visible to respondents who successfully accessed their profiles, resulting in a focused subset of eight alumni, a very small sample size.

Overall, users expressed a generally positive impression of the classification quality. For job title classification, **75%** rated the system's accuracy as 4 or 5 out of 5, while **62.5%** gave similarly high marks for seniority classification. The majority also found the validation interface intuitive, and **63%** reported that updating or confirming the AI-generated classifications was easy.

Nonetheless, **63%** of participants noted that they would have benefited from additional context explaining the suggested classifications.

This suggests that while the model's outputs were often correct, a lack of transparency may reduce user confidence, especially for edge cases or less familiar role titles.

Taken together, these insights confirm that users generally trusted the enrichment pipeline's results, but also emphasize the importance of explainability and guided UX design in AI-driven systems.

These considerations will be further addressed in [section 6.3](#).

5.4.5 Telemetry Data

Platform telemetry provides meaningful insight into both user engagement and technical performance.

Since launch (approximately one month ago), the platform has seen **604 visitors** and **2,038 page views**, with a **56% bounce rate**.

The geographic distribution of traffic was predominantly from Portugal, followed by the United States, the Netherlands, and Ireland.

In terms of performance, the platform demonstrated excellent results on Core Web Vitals [14] and achieved a Real Experience Score (RES) [52] of **98**, a composite metric by Vercel that reflects overall user experience.

The First Contentful Paint (FCP) [53], a key responsiveness metric, was measured at **1.03 seconds**, indicating fast visual loading - these, and all other metrics, are shown in Table 5.18.

Metric	Value	Optimal Threshold
Real Experience Score (RES)	98	> 90
First Contentful Paint (FCP)	1.03s	≤ 1.8s
Largest Contentful Paint (LCP)	1.33s	≤ 2.5s
Interaction to Next Paint (INP)	184ms	≤ 200ms
Cumulative Layout Shift (CLS)	0.08	≤ 0.1
First Input Delay (FID)	14ms	≤ 100ms
Time to First Byte (TTFB)	0.36s	≤ 0.8s

Table 5.18: Real Experience Score and Core Web Vitals

These results suggest that the platform delivers a technically robust and responsive experience across a variety of connection types, an essential factor for sustained engagement.

5.5 Qualitative Results

The next sections summarize qualitative feedback from 21 users who completed the feedback form. While a short usability survey was included, the main focus was on user experience with the analytics dashboard and AI-generated profiles.

Both structured ratings and open-ended comments provided insights into strengths, limitations, and possible improvements. The full feedback form is available in Appendix B.1.

5.5.1 Participant Overview

The 21 feedback respondents represent a diverse alumni sample, with some holding multiple degrees.

Figure 5.6 shows most are alumni, and a few are both alumni and professors, while Table 5.19 details participants' academic backgrounds, with most from MIEIC, followed by pre-Bologna LEIC and MEIC. Other degrees, like ProDEI, MESW, MGI, and L.EIC, are also represented.




Figure 5.6: Distribution of participants by affiliation with FEUP

Degree	Number of Participants
MIEIC	11
LEIC (pre-Bologna)	4
ProDEI	4
M.EIC	3
MESW	1
MGI	1
L.EIC (Bologna)	1

Table 5.19: Distribution of participants by degree

5.5.2 System Usability Scale (SUS) Feedback

Platform usability was assessed using the System Usability Scale (SUS) with ten Likert-scale questions (see Figure 5.7). Responses were largely positive: most users felt confident, found the platform easy to use, and reported little need for technical support. Negative or neutral responses were rare, suggesting participants formed clear, favorable impressions.

Some users were unsure about frequent use, but feedback on ease of use, learnability, and feature integration was overwhelmingly positive. The mean SUS score was **72.7** (SD = 20.1), above the industry benchmark of 68, indicating the platform is highly usable. Scores ranged from 30.0 to 95.0, placing the platform solidly in the “B” range for usability.

Overall, most participants had a positive experience and no major usability issues were identified, providing a strong basis for evaluating analytics and profile features.




Figure 5.7: Distribution of System Usability responses

5.5.3 Analytics Feedback

The analytics dashboard is the main user-facing feature of Alumni-FEUP, enabling interactive exploration of alumni careers across multiple dimensions.

This section reviews user perceptions of its filters, visualizations, and analytical value. Feedback includes both quantitative (Likert-scale) and qualitative (open-ended) responses, offering a comprehensive view of how well the dashboard supports alumni data exploration and meets the interactive visualization goals outlined in **RQ4** (see [section 1.4](#)).

Filters Usability and Adequacy

Regarding the filters available on the analytics page, the feedback was overwhelmingly positive (see Questions 20 through 24 in [section B.1](#)).

Over 76 % of respondents agreed or strongly agreed that the filters were easy to use (Q20), while the remaining 24 % indicated a neutral stance. Notably, no users reported any negative experience (i.e., no disagree or strongly disagree responses).

In terms of adequacy, more than 71 % of participants felt the available filters were both useful and adequate for exploring the alumni data (Q21). Only a small minority (less than 10 %) expressed any reservations, while a handful of users felt neutral.

The most commonly used filters (Q22) were Role/Position (71 % of respondents), Country/City (62 %), and Company (57 %), as shown in Table 5.20. Seniority level, industry, and degree/conclusion year were also regularly used, while only a very small minority reported not using the filters at all. This pattern highlights that participants were primarily interested in exploring alumni by professional role, geographic location, and company affiliation.

Filter	Number of respondents (%)
Role/Position	15 (71%)
Country/City	13 (62%)
Company	12 (57%)
Seniority Level	8 (38%)
Industry	8 (38%)
Degree/Conclusion Year	6 (29%)
Did not use filters	2 (10%)

Table 5.20: Most commonly used filters on the analytics page

When asked about missing filters (Q23), the vast majority (86 %) indicated that the current selection was sufficient for their needs. Of the few who suggested improvements, most requests centered around greater flexibility in filtering by time-based criteria, such as year of entry, enrollment, or conclusion - it is assumed that the latter was a mistake by the respondent, as filtering by conclusion year has been available since the first day of the platform's deployment.

Additional suggestions included making it easier to apply past-data filters globally rather than limiting them to specific roles (Q24).

These findings suggest that the filtering system addresses most user requirements, while also pointing to possible enhancements for temporal and longitudinal analysis in future iterations.

Visualization Utility and Clarity

Overall, feedback on the visualizations (charts, tables, and timelines) was strongly positive (see Questions 25 through 28 in Appendix B.1).

More than 76 % of respondents agreed or strongly agreed that the visualizations were helpful (Q25), and no users reported strong negative experiences. The remainder were neutral or moderately positive, suggesting that the visualizations generally supported exploration and understanding of the alumni data.

Participants also reported a strong ability to drill down into the data meaningfully (Q26), with the majority indicating that they could find specific cohorts, titles, or categories as intended.

When asked whether any visualization felt unnecessary, confusing, or redundant, majority of respondents simply answered negatively (Q27). Among the specific and constructive feedback received, users mentioned:

- **Listing by faculty** was perceived as redundant by some respondents, since only FEUP is currently supported. This design decision was made by the author to ensure scalability in the event that the platform is extended to other faculties in the future.
- **Pie charts with more than 10 categories** were occasionally found to be misleading or unclear.
- **Desire for improved integration** between the map and analytics, particularly maintaining context when switching views. This was not directly actionable due to current platform constraints.

In terms of clarity (Q28), two-thirds of respondents agreed or strongly agreed that it was clear how to switch between different data views, while only a small minority found this aspect confusing.

Several minor usability issues, such as the table scrolling behavior, have since been addressed in later releases.

These findings are directly relevant to **RQ4** (see [section 1.4](#)), confirming that interactive charts and tables are highly valued and effective in enhancing user engagement with alumni data.

User feedback also highlights areas where additional improvements to interactivity and clarity could propel the platform to the next level.

5.5.4 Open-Ended Feedback

Participants were also invited to provide open-ended feedback (see Questions 29 through 33 in [Appendix B.1](#)) and to share notable use cases or suggestions for the analytics dashboard.

Page Usefulness and Layout

A strong majority found the analytics page helpful for answering specific questions about the alumni community (e.g., "Where are people who graduated with me working now?" or "Where are people from my hometown working?") (Q29). Most users could answer these questions effectively, with only one respondent expressing uncertainty.

Perceptions of the analytics page layout were overwhelmingly positive (Q31), with over 70 % describing it as well balanced. Only a few participants found it overwhelming or too minimal, and a single user reported a temporary scrolling issue, which has since been resolved.

Layout Perception	Number (%)
Well balanced	15 (71%)
Overwhelming (too much)	4 (19%)
Too minimal / lacking context	1 (5%)
Other (scrolling, too wide)	1 (5%)

Table 5.21: Perceived layout of the analytics page

Key Suggestions and Reported Issues

Open-ended feedback yielded a range of actionable insights and minor improvement requests (Q30, Q32, and Q33). Table 5.22 summarizes the most common themes. Notably, issue 5 (Show a list of persons matching selected filters) resulted from the earlier scrolling issue; since the list of persons appears below the tables, this usability issue has now been fixed.

Suggestion / Issue	Theme
1. Navigation breadcrumbs	Navigation
2. Make filters sharable/bookmarkable	Interactivity
3. Add overall AI confidence for experiences	Transparency
4. Ability to export results	Functionality
5. Show a list of persons matching selected filters	Data visibility
6. Pie chart is confusing for large categories	Visualization
7. Some scrolling/layout issues (now fixed)	Usability
8. Company missing when filtering by degree	Data completeness

Table 5.22: Notable user suggestions and reported issues

Illustrative Use Cases

Respondents highlighted several valuable use cases, including:

- Identifying where friends or alumni are currently working to aid job searches.
- Finding people in specific areas or companies.
- Combining filters (e.g., country and seniority) for targeted searches.

5.5.5 Profile Experience Feedback

Of the 21 survey participants, 8 alumni successfully logged in and accessed the profile validation features. Among those who accessed their profiles, the majority (63 %) reported that everything worked smoothly, while a quarter noted that some aspects were unclear and one respondent was unable to find or open their profile page (Q39 in Appendix B.1).

Regarding the AI-generated job classification, most users found it relatively easy to understand and update (Q40), though some expressed a need for additional guidance (Q43).

Accuracy was rated positively for both current role (Q41) and seniority classification (Q42), with the majority selecting 4 or 5 out of 5 for both metrics. However, a few users did report lower confidence in the system's accuracy, highlighting the importance of continued refinement.

Notably, 63 % of respondents indicated that more context or explanation for the suggested classifications would have been helpful (Q43), while the remainder either found the information clear or were unsure.

Feedback Item	Number (%)
Everything worked smoothly	5 (63%)
Some parts were unclear	2 (25%)
Unable to access profile	1 (13%)
Easy to update job classification (4 or 5 out of 5)	5 (62.5%)
High accuracy for job classification (4 or 5 out of 5)	6 (75%)
High accuracy for seniority classification (4 or 5 out of 5)	5 (62.5%)
Wanted more explanation/context	5 (63%)

Table 5.23: User feedback on profile and classification experience (n=8)

These findings indicate that, for most users who accessed the profile features, the platform provided a smooth and accurate experience. However, the repeated desire for more context and guidance around AI-driven classifications points to a clear opportunity for further enhancement, particularly in onboarding and explanation for non-technical stakeholders.

5.6 Discussion

The evaluation offers a comprehensive view of the Alumni-FEUP platform's effectiveness and its AI-driven enrichment pipeline, drawing on a large and diverse dataset (3,042 alumni, 8,120 roles, 2,471 companies).

Job Classification Evaluation

The ESCO-based job classification agent achieved an overall manual evaluation accuracy of 70%, representing a significant advancement over manual approaches for processing large-scale alumni datasets. The results reveal a clear pattern: canonical technical roles such as "Web Developer" and "Software Developer" achieved near-perfect accuracy (100%), while emerging or hybrid titles like "Cloud Software Developer" (25%) and "Integration Engineer" (40%) posed greater challenges.

This performance differential highlights the inherent tension between taxonomy comprehensiveness and the dynamic nature of job titles in the technology sector. The 70% accuracy rate,

while not perfect, demonstrates that automated classification can effectively handle the majority of roles while requiring human oversight primarily for edge cases and emerging positions.

The confidence scoring mechanism proved valuable for quality control, with scores typically ranging between 67% and 71% across different occupational groups. This conservative scoring approach helps identify cases requiring manual review while maintaining high precision for confident classifications.

Seniority Level Estimation

The LLM-driven seniority agent outperformed job classification with a 90% manual evaluation accuracy, demonstrating exceptional capability in understanding organizational hierarchies and career progression patterns. The highest reliability was observed for "Mid-Senior Level" (95%), "Entry Level" (92%), and "Executive" (100%) positions, which represent well-defined career stages with clear distinguishing characteristics.

The performance variation across seniority levels reveals important insights about organizational title conventions. "Associate" roles showed lower accuracy (84%), reflecting the semantic ambiguity of this term across different industries and organizational structures. Similarly, "Director" and "C-Level" positions occasionally suffered from classification errors due to the blurred boundaries between these senior management tiers in practice.

The high overall accuracy suggests that LLMs can effectively interpret the contextual cues present in LinkedIn role descriptions to make reliable seniority assessments, even when job titles themselves are ambiguous or organization-specific.

User Experience and Platform Engagement

The platform achieved a System Usability Scale (SUS) score of 72.7, placing it above the industry benchmark of 68 and indicating strong usability across the diverse alumni community. This positive user experience is particularly noteworthy given the platform's complexity and the variety of user backgrounds represented in the sample.

User feedback on interactive features was overwhelmingly positive, with over 76% of respondents finding visualizations helpful and 71% reporting that filters were adequate for their exploration needs. The most commonly used filters, Role/Position (71%), Country/City (62%), and Company (57%), align with the primary use cases identified during platform design, confirming that the interface effectively supports the intended user workflows.

The profile validation interface, while used by a smaller subset of users (8 out of 21 survey participants), received positive feedback regarding ease of use and classification accuracy. However, 63% of users expressed a desire for more context explaining AI-generated classifications, highlighting an important area for future enhancement in transparency and explainability.

Technical Performance and Scalability

Telemetry shows the platform performs well globally, with a Real Experience Score of 98 and a First Contentful Paint of 1.03 seconds, meeting web responsiveness standards. Since launch, it

has served 604 visitors and 2,038 page views across several countries, proving its reach. The 56% bounce rate is reasonable for an analytics tool.

Current performance easily handles 1,353 profiles and 8,120 roles, with analytics queries finishing in about three seconds. This provides a strong base for growth, though further optimization and scaling will be needed as data volume increases, especially since the platform is currently hosted on a single server.

Dataset Coverage and Representativeness

The comprehensive dataset compiled for this evaluation represents one of the largest and most diverse collections of alumni career data in the academic literature. With coverage spanning 52 countries, 164 industries, and 640 unique ESCO classifications, the dataset provides robust statistical power for analysis while capturing the global reach of FEUP's informatics graduates.

The 44.4% LinkedIn coverage rate (1,353 out of 3,042 alumni) represents a solid foundation for automated enrichment, though it highlights the ongoing challenge of maintaining current contact information for alumni communities. The coverage imbalance across degrees, with newer programs like MESW and PDMD showing higher rates due to curated datasets, suggests that systematic data collection efforts can significantly improve platform utility.

Methodological Insights

The evaluation methodology employed in this study represents a significant contribution to the field of AI-driven data enrichment evaluation. The combination of quantitative metrics (manual evaluation, confidence scores) with qualitative insights (user feedback, SUS scoring) provides a holistic assessment that captures both technical performance and user experience.

The stratified sampling approach for manual evaluation, covering diverse roles and seniority levels, ensures that results are representative of the full spectrum of professional contexts present in the alumni population. The integration of user validation creates a feedback loop that enables continuous improvement while providing real-world validation of system outputs.

The evaluation results demonstrate that AI-driven approaches can effectively process large-scale alumni datasets while maintaining acceptable accuracy levels and user satisfaction. The 70% job classification and 90% seniority estimation accuracy rates provide a strong foundation for automated alumni data enrichment, with clear pathways for improvement through enhanced prompting strategies, taxonomy updates, and user feedback integration into the underlying agent's model tuning process.

Chapter 6

Conclusions and Future Work

This final chapter closes the loop on the *Alumni-FEUP* project. It (i) recaps the artefacts delivered, (ii) highlights the key insights already emerging, (iii) acknowledges the work’s limitations, and (iv) outlines concrete next steps that can extend its impact.

6.1 Summary of Contributions

1. **Open-source, production-ready platform.** A live, public deployment of *Alumni-FEUP*¹ was shipped and simultaneously released as open source², fulfilling the early-stage commitment to transparency and community reuse. The stack (Next.js, NestJS, PostgreSQL, Prisma, Redis) is semi-containerised and documented for simple installs.
2. **Triad of AI enrichment agents in production.** Three FastAPI services: *Job-Classification*, *Seniority*, and *Location* are orchestrated by the backend and invoked automatically after each LinkedIn profile extraction. Together they answer **RQ3** and lay the technical groundwork for **RQ2**.
3. **Novel hybrid ranking pipeline.** Each agent combines (i) dense embeddings, (ii) cross-encoder re-ranking, and (iii) a final GPT arbiter call that re-orders the top candidates and emits a JSON-validated choice. This cross-encoder → GPT rerank strategy reduces hallucinations while preserving LLM flexibility, a methodological contribution of its own.
4. **Comprehensive alumni dataset.** At the time of writing the platform tracks **3 041** alumni, **8 119** roles, and **2 471** companies (section 5.1). Of these, **1 351** alumni have an enriched profile, spanning 52 countries and 164 industries, satisfying the coverage goals behind **RQ1**.
5. **Human-in-the-loop validation interface.** A lightweight UI lets any alumnus up- or down-vote their automatically assigned role and seniority, feeding a feedback table that powers the quantitative evaluation described in section 5.4 and providing early evidence for **RQ4**.

¹<https://alumni-feup.com/analytics>

²<https://github.com/carlosverissimo3001/alumni-feup>

6. **Empirical evaluation via live focus group.** A moderated session (June 2025) served as a real-time demo and qualitative sanity check; insights collected in the 7-minute feedback form informed the discussion of strengths and pain points (section 5.4 and section 5.5).
7. **Reusable visual-analytics components.** All D3.js charts (pie charts and trend lines) were abstracted into a small shared library, enabling rapid reuse by other cohorts and addressing maintainability concerns highlighted in RQ2.
8. **Reproducibility and open science.** The public GitHub repository contains Dockerfiles, database schema, Mermaid-style architecture diagrams, ensuring the results can be replicated and extended by the academic community.

6.2 Main Findings

The evaluation of the Alumni-FEUP platform provides compelling evidence that addresses all four research questions and validates the proposed hypotheses. This section synthesizes the key findings and their implications for the field of alumni tracking and career analytics.

Research Question 1: Platform Differentiation and Innovation

The results demonstrate that Alumni-FEUP successfully differentiates itself from existing initiatives through its comprehensive integration of automated classification, interactive visualizations, and real-time data enrichment. The platform's ability to process over 8,000 roles from 3,000+ alumni with 70% job classification accuracy and 90% seniority estimation accuracy represents a significant advancement over manual approaches. The high user satisfaction scores (SUS score of 72.7) and positive feedback on interactive features confirm that the platform sets new standards for alumni tracking systems, supporting **H3** regarding its potential as a foundational platform for institutions.

Research Question 2: Adaptability to Evolving Tech Trends

The evaluation reveals both strengths and challenges in maintaining relevance within the rapidly evolving tech industry. The ESCO-based classification system shows strong performance for established roles (e.g., "Software Developer", "Web Developer") but struggles with emerging or hybrid titles (e.g., "Cloud Software Developer", "Integration Engineer"). This finding validates **H1** regarding the effectiveness of combining LLMs with traditional methods, while highlighting the need for continuous taxonomy updates and enhanced prompting strategies to handle evolving job titles. The platform's modular architecture and confidence scoring system provide the foundation for future adaptability.

Research Question 3: Feasibility of Automated Classification

The results strongly support **H2** and **H3**, demonstrating that automated classification is not only

feasible but highly effective. The 70% overall accuracy for job classification and 90% for seniority estimation, achieved through a combination of embedding-based similarity and LLM reasoning, significantly outperforms manual approaches while requiring minimal human oversight. The stratified evaluation design, covering diverse roles and seniority levels, confirms the system's robustness across different professional contexts. The confidence scoring mechanism provides transparency and enables quality control without extensive manual intervention.

Research Question 4: Interactive Visualization Effectiveness

User feedback overwhelmingly supports the effectiveness of interactive visualizations in enhancing engagement with alumni data. Over 76% of respondents found the visualizations helpful, with strong positive feedback on filtering capabilities, data exploration, and overall user experience. The telemetry data (RES score of 98, FCP of 1.03s) confirms that the platform delivers excellent technical performance, essential for sustained user engagement. The integration of spatial, temporal, and categorical data through interactive charts and maps successfully addresses the original problem of disconnected data components.

Hypothesis Validation

All three hypotheses have been validated through empirical evidence:

- **H1** is supported by the 70% job classification and 90% seniority estimation accuracy achieved through the hybrid LLM and traditional ML approaches.
- **H2** is confirmed by the successful integration of classification methods into a fully automated pipeline requiring minimal human oversight.
- **H3** is validated by the platform's comprehensive dataset coverage, high user satisfaction, and demonstrated ability to serve multiple stakeholders effectively.

Methodological Contributions

The evaluation methodology itself represents a contribution to the field, combining quantitative metrics (manual evaluation, confidence scores) with qualitative insights (user feedback, SUS scoring) to provide a holistic assessment of AI-driven alumni data enrichment. The stratified sampling approach ensures representative evaluation across diverse professional contexts, while the integration of user validation creates a feedback loop for continuous improvement.

6.3 Limitations and Open Issues

This section outlines the main limitations encountered during the development and evaluation of the Alumni-FEUP platform. While many of these were expected at this stage, they remain relevant considerations for future iterations and large-scale deployments.

Taxonomy Coverage Constraints

The ESCO classification system, while comprehensive, cannot capture every modern or emerging job title, particularly in fast-evolving tech sectors. This limitation is inherent to any standardized taxonomy and affects approximately 30% of classifications, where "nearest neighbor" assignments are made.

Evaluation Scope and Generalizability

The manual evaluation sample, while stratified and representative, covers only 2.5% of the total dataset. While this provides robust statistical insights, some rare role combinations may be underrepresented. Additionally, the evaluation focuses on FEUP alumni, limiting generalizability to other academic or industry populations with different title conventions and organizational structures.

User Engagement and Feedback Bias

User-submitted validations are voluntary and may be biased toward more engaged or technically skilled alumni. The relatively low participation rate limits the statistical power of user feedback analysis. The qualitative feedback, while valuable, may not capture the full spectrum of user needs and preferences.

Incomplete or Stale LinkedIn Data

The platform relies on the Enrich Layer API to retrieve publicly available LinkedIn data. While it offers a robust base endpoint, the default behavior prioritizes cached profiles, regardless of profile age (via the `use_cache=if-present` flag) - if the profile is not cached, Enrich Layer will attempt to source the profile externally.

To always fetch almost-live data, a higher-cost tier using `use_cache=if-recent` must be used, and the API will make a best effort to return a fresh profile no older than 29 days.

Given the credit limitations inherent to academic work, many profiles were returned from cache, resulting in the omission of recent roles and updates in several cases.

Furthermore, profiles that had been deleted or made private returned nullable values, reducing the overall completeness of the dataset.

Coverage Imbalance Across Degrees

Of the 3,042 alumni, only 1,353 had valid LinkedIn profiles enriched via Enrich Layer, giving about 45% coverage, a solid start for the platform.

MESW and PDMD are overrepresented due to curated datasets and fewer graduates, not higher engagement. In contrast, newer degrees like M.EIC and L.EIC have lower coverage due to limited collection so far. Coverage for these programs should improve as more LinkedIn profiles are matched to institutional records over time.

Technical Infrastructure Constraints

The platform currently depends on external APIs (e.g., OpenAI, Enrich Layer), which raises concerns about long-term sustainability and cost. The confidence scoring system is effective but not a true probability, scores should be viewed as relative measures.

Two backend limitations were noted:

- **API Rate Limits:** The enrichment layer sometimes hits OpenAI rate limits during large batch updates. While exponential backoff and retries maintain reliability, they introduce latency and complexity, and further optimization will be needed as the system grows.
- **Data Aggregation Performance:** Delivering analytics data to the frontend takes about three seconds for the current dataset (1,353 alumni, 8,120 roles). As data and usage increase, caching or selective loading will be necessary to keep the platform responsive.

Furthermore, the project experienced a significant disruption in data sourcing during development when Proxycurl, the initial LinkedIn data provider, abruptly shut down on the 4th of July 2025 following a lawsuit filed by Microsoft. The lawsuit alleged that Proxycurl operated "hundreds of thousands" of fake accounts to collect personal and professional information, despite the company's insistence that it only scraped publicly available data [47].

Enrich Layer, the new data provider, took over affected customers and honored existing credits while providing a matching API.

Frontend Limitations on Mobile Devices

Although the platform is responsive, some visual and usability issues persist on mobile devices, as reported by users. The interface is optimized for larger screens, especially in analytics, making the mobile experience less fluid and more limited. Further improvements are needed for better accessibility and readability on smaller devices.

Lack of Human-in-the-Loop Workflow

Currently, the enrichment process runs fully automatically, without human-in-the-loop (HITL) intervention for real-time corrections. While this approach offers consistency and scalability, it also means misclassifications, especially for ambiguous or emerging titles, may go unnoticed unless users manually report them.

Privacy and Ethical Considerations

Processing personal career data brings significant privacy and ethical responsibilities. While current best practices are followed, future work should include thorough privacy impact assessments and consider federated learning to reduce data centralization. As the platform grows, balancing data utility and privacy will remain a key priority.

The last-minute transition to Enrich Layer demonstrates the importance of having robust contingency plans and multiple data source options when building systems that rely on third-party data providers.

It also highlights how legal and regulatory challenges can significantly impact data collection strategies, even when operating within what appears to be a legal gray zone.

6.4 Future Work

Building on the evaluation and identified limitations, the following steps are proposed to further enhance the Alumni-FEUP platform and broaden its impact.

Enhanced Classification Pipeline

Future efforts should aim to:

- **Integrate Hybrid Taxonomies:** Extend ESCO with tech-focused or hybrid ontologies to better capture emerging and interdisciplinary roles.
- **Implement Multi-label Classification:** Allow roles to be assigned multiple titles to reflect real-world overlaps.
- **Enable Continuous Learning:** Use user feedback for ongoing model refinement and accuracy improvements.
- **Refine Prompting Strategies:** Develop smarter prompts for better handling of ambiguous roles and context-aware classification.

Scalability and Performance Optimization

To support growth and broader deployment:

- **Optimize Caching and Queries:** Improve caching and database efficiency for larger datasets.
- **Develop On-premise Alternatives:** Reduce reliance on external APIs, especially for LLM services.
- **Enable Distributed Processing:** Scale up processing for multiple concurrent enrichment tasks.
- **Automate Quality Monitoring:** Build automated checks for classification quality to reduce manual review.

Enhanced User Experience and Engagement

To increase impact and participation:

- **Mobile Optimization:** Redesign for smoother mobile use.
- **Improved Feedback Mechanisms:** Add better, possibly incentivized, feedback tools and more granular validation.

- **Personalization:** Offer personalized dashboards and recommendations based on user behavior.
- **Export and Integration:** Allow users and institutions to export data and integrate with other systems.

Multi-institutional Deployment

For broader adoption:

- **Modular Architecture:** Support multi-tenant setups with data isolation.
- **Standardized Data Models:** Create flexible models and APIs for varied institutional needs.
- **Cross-institutional Analytics:** Enable comparisons between institutions while maintaining privacy.
- **Community-driven Development:** Build a user community to share best practices and improvements.

Advanced Analytics and Insights

For deeper impact:

- **Career Path Analysis:** Identify and predict common career trajectories.
- **Skills Gap Analysis:** Connect program outcomes to industry needs.
- **Geographic Mobility Patterns:** Analyze alumni movement and location trends.
- **Industry Trend Analysis:** Track and respond to new roles and skill demands.

Privacy and Ethical Enhancements

To ensure trust and compliance:

- **Privacy Impact Assessment:** Regularly review and embed privacy-by-design.
- **Federated Learning:** Improve accuracy while minimizing data centralization.
- **Transparency and Explainability:** Clearly explain classifications and data usage.
- **User Control:** Empower users to manage their data and privacy choices.

Research and Academic Contributions

To foster community and knowledge sharing:

- **Open Dataset:** Share anonymized data for further research.
- **Methodological Publications:** Publish on methods and evaluation strategies.
- **Benchmark Development:** Establish benchmarks for alumni analytics.
- **Educational Resources:** Offer materials and workshops for institutional adoption.

6.5 Final Remarks

The Alumni-FEUP project marks a significant advance in alumni tracking and career analytics, showing the feasibility and value of AI-driven methods for understanding graduate outcomes. By combining automated classification, interactive visualizations, and real-time data enrichment, the platform demonstrates strong technical and practical results, achieving 70% accuracy in job classification, 90% in seniority estimation, and high user satisfaction.

While limitations exist, they highlight clear paths for future improvement. The platform's modular, open-source design supports ongoing development and broad adoption. Beyond its technical achievements, the project introduces new evaluation methodologies, user experience insights for AI systems, and a model for collaborative academic development.

Looking ahead, Alumni-FEUP can become a key tool for institutions, helping to analyze graduate outcomes, inform curricula, and foster academic-industry connections. The roadmap provides a clear direction for future impact. As alumni analytics evolves, this project demonstrates the power of combining advanced AI with thoughtful design and rigorous evaluation, laying a foundation for continued progress in the field.

References

- [1] A. A. Aletdinova and M. S. Murtazina. “The Collection and Processing Specifics of Online Data on Job Vacancies in the Russian Labor Market”. In: *2021 XV International Scientific-Technical Conference on Actual Problems Of Electronic Instrument Engineering (APEIE)*. 2021, pp. 1–5. ISBN: 2473-8573. DOI: 10.1109/APEIE52976.2021.9647682. URL: <https://ieeexplore.ieee.org/document/9647682/>.
- [2] Aon Corporation. *Aon Corporation - Radford Global Job Leveling*. Available at <https://humancapital.aon.com/aon.radford/media/files/articles/2018/radford-global-job-leveling-model-rises-to-the-challenge.pdf>, last accessed in November 2024. 2018.
- [3] S. I. Beristain, R. R. López Barbosa, and E. G. Barriocanal. “Improving Jobs-Resumes Classification: A Labor Market Intelligence Approach”. In: *International Journal of Information Technology and Decision Making* 23.4 (2024), pp. 1509–1525. ISSN: 02196220 (ISSN). DOI: 10.1142/S0219622023500013. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85145488159&doi=10.1142%2fS0219622023500013&partnerID=40&md5=5ed07efe3efeace099f7420a590c5118>.
- [4] R. Boselli et al. “Classifying online Job Advertisements through Machine Learning”. In: *Future Generation Computer Systems* 86 (2018), pp. 319–328. ISSN: 0167739X (ISSN). DOI: 10.1016/j.future.2018.03.035. URL: <https://www.sciencedirect.com/science/article/abs/pii/S0167739X17321830?via%3Dihub>.
- [5] R. Boselli et al. “Using Machine Learning for Labour Market Intelligence”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Ed. by M. Ceci et al. Vol. 10536 LNAI. Springer Verlag, 2017, pp. 330–342. ISBN: 03029743 (ISSN); 978-331971272-7 (ISBN). DOI: 10.1007/978-3-319-71273-4_27. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-85040225472&doi=10.1007%2f978-3-319-71273-4_27&partnerID=40&md5=ce88cada543a721580e5a28c7ce6718b%20https://link.springer.com/content/pdf/10.1007/978-3-319-71273-4_27.pdf.
- [6] Mike Bostock et al. *D3.js*. Available at <https://d3js.org/>, last accessed in December 2024.
- [7] N. Carrada Zuniga and A. Colbern. “The Annual BUILD Snapshot: Tracking Alumni Outcomes”. In: *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*. Ed. by C. Stephanidis et al. Vol. 13096 LNCS. Springer Science and Business Media Deutschland GmbH, 2021, pp. 345–360. DOI: 10.1007/978-3-030-90328-2_22. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-85119825561&doi=10.1007%2f978-3-030-90328-2_22&partnerID=40&md5=cfd11539be933b44154cca970b0db3a%2f

- 20https://link.springer.com/chapter/10.1007/978-3-030-90328-2_22.
- [8] CEN - European Committee for Standardization. *CEN - European ICT Professional Role Profiles*. Available at https://eufordigital.eu/wp-content/uploads/2019/10/EUROPEAN-ICT-PROF_ROLE-PROFILES-VERSION-2_PART-3_METHODODOLOGY.pdf, last accessed in November 2024. 2019.
 - [9] A. Chern et al. “Automatically Detecting Errors in Employer Industry Classification Using Job Postings”. In: *Data Science and Engineering* 3.3 (2018), pp. 221–231. ISSN: 23641185 (ISSN). DOI: 10.1007/s41019-018-0071-7. URL: <https://link.springer.com/content/pdf/10.1007/s41019-018-0071-7.pdf>.
 - [10] Andrea De Mauro et al. “Human resources for Big Data professions: A systematic classification of job roles and required skill sets”. In: *Information Processing & Management* 54.5 (2018), pp. 807–817. ISSN: 0306-4573. DOI: 10.1016/j.ipm.2017.05.004. URL: <https://www.sciencedirect.com/science/article/pii/S0306457317300018>.
 - [11] Escola Superior de Enfermagem de Coimbra. *Escola Superior de Enfermagem de Coimbra - Alumni Map*. Available at <https://alumni.esenfc.pt/en/map>, last accessed in November 2024. 2024.
 - [12] Eurostat. *Eurostat - ESCO*. Available at <https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Glossary:ESCO>, last accessed in December 2024. 2021.
 - [13] M. Goindani et al. “Employer Industry Classification Using Job Postings”. In: *IEEE International Conference on Data Mining Workshops, ICDMW*. Ed. by R. Gottumukkala et al. Vol. 2017-November. IEEE Computer Society, 2017, pp. 183–188. ISBN: 23759232 (ISSN); 978-153861480-8 (ISBN). DOI: 10.1109/ICDMW.2017.30. URL: <https://ieeexplore.ieee.org/document/8215661/>.
 - [14] Google. *Google Speed Insights*. Available at <https://developers.google.com/speed/docs/insights/v5/about>, last accessed in June 2025.
 - [15] Mark Howison et al. “Extracting Structured Labor Market Information from Job Postings with Generative AI”. In: *Digit. Gov.: Res. Pract.* (2024). DOI: 10.1145/3674847. URL: <https://doi.org/10.1145/3674847>.
 - [16] C. Y. Hung, E. Jeng, and L. C. Cheng. “Analysis of CEO career patterns using machine learning: taking US university graduates as an example”. In: *Data Technologies and Applications* (2024). ISSN: 25149288 (ISSN). DOI: 10.1108/DTA-04-2023-0132. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85200147756&doi=10.1108%2fDTA-04-2023-0132&partnerID=40&md5=d6765dee4e6e636abc889de09093302e>.
 - [17] Nguyen Quang Huy, Le Phuoc Thanh, and Doan Van Thang. “Research Arima Model Forecasts Student’s Graduation Rate”. In: *Proceedings of Second International Conference on Intelligent System*. Ed. by João Manuel R. S. Tavares et al. Singapore: Springer Nature Singapore, 2024, pp. 375–384. ISBN: 978-981-99-8976-8.
 - [18] International Health Terminology Standards Development Organisation. *International Health Terminology Standards Development Organisation - SNOMED CT*. Available at <https://www.snomed.org/>, last accessed in December 2024. 2024.

- [19] International Labour Organization. *International Labour Organization - ISCO*. Available at <https://ilostat.ilo.org/methods/concepts-and-definitions/classification-occupation/>, last accessed in December 2024. 2008.
- [20] F. Javed et al. "Towards a job title classification system". In: *arXiv preprint arXiv:1606.00917* (2016). URL: <https://arxiv.org/abs/1606.00917>.
- [21] A. I. M. Khairunizan and S. D. S. A. A. Feisal. "Alumni Tracking Dashboard: The Importance of Data Cleaning". In: *2023 7th International Conference on Business and Information Management (ICBIM)*. 2023, pp. 54–59. DOI: 10.1109/ICBIM59872.2023.10303146. URL: <https://ieeexplore.ieee.org/document/10303146/>.
- [22] P. Kumar, S. Swetha, and M. Sundari. "Secured Web-based Alumni Network and Information Systems". In: *Proceedings of the 7th International Conference on Intelligent Computing and Control Systems, ICICCS 2023*. Institute of Electrical and Electronics Engineers Inc., 2023, pp. 1427–1434. ISBN: 979-835039725-3 (ISBN). DOI: 10.1109/ICICCS56967.2023.10142761. URL: <https://ieeexplore.ieee.org/document/10142761/>.
- [23] S. Lade et al. "LinkedIn Alumni Profile Data Extraction". In: *2024 4th International Conference on Pervasive Computing and Social Networking (ICPCSN)*. 2024, pp. 174–178. DOI: 10.1109/ICPCSN62568.2024.00037. URL: <https://ieeexplore.ieee.org/document/10607617/>.
- [24] B. Lalitha et al. "Applicant Screening System Using NLP". In: *2023 International Conference on Innovative Data Communication Technologies and Application (ICIDCA)*. 2023, pp. 379–383. DOI: 10.1109/ICIDCA56705.2023.10099953. URL: <https://ieeexplore.ieee.org/document/10099953/>.
- [25] Lei Li et al. *Retrieving and Classifying LinkedIn Job Titles for Alumni Career Analysis*. Conference Paper. 2021. DOI: 10.1145/3450329.3476858. URL: <https://doi.org/10.1145/3450329.3476858>.
- [26] M. Maitra, S. Sinha, and T. Kierszenowicz. "An Improved BERT Model for Precise Job Title Classification Using Job Descriptions". In: *2024 IEEE 12th International Conference on Intelligent Systems (IS)*. 2024, pp. 1–6. ISBN: 2767-9802. DOI: 10.1109/IS61756.2024.10705204. URL: <https://ieeexplore.ieee.org/document/10705204/>.
- [27] McKinsey. *McKinsey Alumni*. Available at <https://www.mckinsey.com/alumni/about-us/about-the-alumni-program>, last accessed in November 2024.
- [28] Ravita Mishra and Sheetal Rathi. "Enhanced DSSM (deep semantic structure modelling) technique for job recommendation". In: *Journal of King Saud University - Computer and Information Sciences* 34.9 (2022), pp. 7790–7802. ISSN: 1319-1578. DOI: <https://doi.org/10.1016/j.jksuci.2021.07.018>. URL: <https://www.sciencedirect.com/science/article/pii/S1319157821001853>.
- [29] P. N et al. "Sentimental Analysis of Job Classification using Machine Learning Algorithms". In: *2023 8th International Conference on Communication and Electronics Systems (ICCES)*. 2023, pp. 1279–1284. DOI: 10.1109/ICCES57224.2023.10192717. URL: <https://ieeexplore.ieee.org/document/10192717/>.
- [30] K. J. E. Nabablit and E. S. Dajao. "Development of a Web-Based Graduate Tracer Information System with Data Analytics". In: *Lecture Notes in Networks and Systems*. Ed. by X. S. Yang et al. Vol. 694 LNNS. Springer Science and Business Media Deutschland GmbH, 2023, pp. 601–611. ISBN: 23673370 (ISSN); 978-981993090-6 (ISBN). DOI: 10.1007/

- 978-981-99-3091-3_50. URL: https://link.springer.com/chapter/10.1007/978-981-99-3091-3_50.
- [31] Near East South Asia Center. *Near East South Asia Center - Alumni Map*. Available at <https://nesa-center.org/alumni-map/>, last accessed in November 2024. 2024.
- [32] O*NET Resource Center. *O*NET - Occupational Information Network*. Available at <https://www.onetcenter.org/taxonomy.html>, last accessed in November 2024. 2024.
- [33] Ollama. *Ollama*. Available at <https://ollama.com/>, last accessed in June 2025.
- [34] OpenAI. *Text Embedding 3 Large*. Available at <https://platform.openai.com/docs/models/text-embedding-3-large>, last accessed in June 2025.
- [35] S. M. Patil and V. D. Shinde. "Analyzing the Effects of K-Means and Random Forest Algorithms on Resume Classification through Manipulation of Cluster Sizes and Datasets". In: *Proceedings of the 18th INDIACom; 2024 11th International Conference on Computing for Sustainable Global Development, INDIACom 2024*. Ed. by M. N. Hoda. Institute of Electrical and Electronics Engineers Inc., 2024, pp. 761–766. ISBN: 978-938054451-9 (ISBN). DOI: 10.23919/INDIACom61295.2024.10498631. URL: <https://ieeexplore.ieee.org/document/10498631/>.
- [36] R. Pico-Saltos et al. "Role of Alumni Program in the Prediction of Career Success in an Ecuadorian Public University". In: *Applied Sciences (Switzerland)* 12.19 (2022). ISSN: 20763417 (ISSN). DOI: 10.3390/app12199892. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85140030159&doi=10.3390%2fapp12199892&partnerID=40&md5=26cda6d9ff6b17d90eb3fff93cbb69099>.
- [37] Mollie Powell. *The Full List of Industries on LinkedIn: Unlocking Success through Strategic Selection*. Available at <https://www.linkedin.com/pulse/full-list-industries-linkedin-unlocking-success-mollie-powell-gajxf/>, last accessed in June 2025.
- [38] PRISMA. *PRISMA - guidelines for reporting systematic reviews*. Available at <https://www.prisma-statement.org/>, last accessed in October 2024. 2024.
- [39] I. Rahhal et al. "Two Stage Job Title Identification System for Online Job Advertisements". In: *IEEE Access* 11 (2023), pp. 19073–19092. ISSN: 21693536 (ISSN). DOI: 10.1109/ACCESS.2023.3247866. URL: <https://ieeexplore.ieee.org/ielx7/6287639/10005208/10050019.pdf?tp=&arnumber=10050019&isnumber=10005208&ref=>.
- [40] S. Rahman et al. "Job Title Prediction and Recommendation System for IT Professionals". In: *2023 5th International Conference on Sustainable Technologies for Industry 5.0, STI 2023*. Institute of Electrical and Electronics Engineers Inc. ISBN: 979-835039429-0 (ISBN). DOI: 10.1109/STI59863.2023.10464457. URL: <https://ieeexplore.ieee.org/document/10464457/>.
- [41] S. Ramraj, V. Sivakumar, and K. Ramnath G. "Real-Time Resume Classification System Using LinkedIn Profile Descriptions". In: *2020 International Conference on Computational Intelligence for Smart Power System and Sustainable Energy (CISPSSE)*. 2020, pp. 1–4. DOI: 10.1109/CISPSSE49931.2020.9212209. URL: <https://ieeexplore.ieee.org/document/9212209/>.

- [42] D. M. Sanusi et al. “An Interactive Alumni Tracking Dashboard”. In: *2023 6th International Conference on Applied Computational Intelligence in Information Systems: Intelligent and Resilient Digital Innovations for Sustainable Living, ACIIS 2023 - Proceedings*. Ed. by S. H. S. Newaz et al. Institute of Electrical and Electronics Engineers Inc., 2023. ISBN: 979-835031742-8 (ISBN). DOI: 10.1109/ACIIS59385.2023.10367325. URL: <https://ieeexplore.ieee.org/document/10367325/>.
- [43] System Usability Scale. *System Usability Scale*. Available at https://en.wikipedia.org/wiki/System_usability_scale, last accessed in June 2025.
- [44] M. Schierholz and M. Schonlau. “Machine Learning for Occupation Coding - A Comparison Study”. In: *Journal of Survey Statistics and Methodology* 9.5 (2021), pp. 1013–1034. ISSN: 23250984 (ISSN). DOI: 10.1093/jssam/smaa023. URL: <https://academic.oup.com/jssam/article/9/5/1013/5952839>.
- [45] S. Sindhu et al. “A Comparative Study on Job Recommendation using Classification Algorithms”. In: *2023 7th International Conference on Computing Methodologies and Communication (ICCMC)*. 2023, pp. 645–652. DOI: 10.1109/ICCMC56507.2023.10083728. URL: <https://ieeexplore.ieee.org/document/10083728/>.
- [46] Panagiotis Skondras et al. “Efficient Resume Classification Through Rapid Dataset Creation Using ChatGPT”. In: *2023 14th International Conference on Information, Intelligence, Systems & Applications (IISA)*. Institute of Electrical and Electronics Engineers Inc., 2023, pp. 1–5. DOI: 10.1109/IISA59645.2023.10345870. URL: <https://ieeexplore.ieee.org/document/10345870/>.
- [47] StartupHub. *The #1 LinkedIn Scraping Startup Proxycurl Shuts Down*. Available at <https://www.startuphub.ai/ai-news/startup-news/2025/the-1-linkedin-scraping-startup-proxycurl-shuts-down/>, last accessed in July 2025.
- [48] The SFIA Foundation. *The SFIA Foundation - Skills Framework for the Information Age*. Available at <https://sfia-online.org/en>, last accessed in November 2024. 2024.
- [49] The University of British Columbia. *UBC - Global Alumni Map*. Available at <http://100.ubc.ca/impact-map/alumni/>, last accessed in November 2024. 2024.
- [50] H. T. Tran, H. H. P. Vo, and S. T. Luu. “Predicting Job Titles from Job Descriptions with Multi-label Text Classification”. In: *Proceedings - 2021 8th NAFOSTED Conference on Information and Computer Science, NICS 2021*. Ed. by V. N. Q. Bao and L. D. Manh. Institute of Electrical and Electronics Engineers Inc., 2021, pp. 513–518. ISBN: 978-166541001-4 (ISBN). DOI: 10.1109/NICS54270.2021.9701541. URL: <https://ieeexplore.ieee.org/document/9701541/>.
- [51] U.S Bureau of Labor Statistics. *U.S Bureau of Labor Statistics - Standard Occupational Classification (SOC)*. Available at <https://www.bls.gov/soc/>, last accessed in November 2024. 2024.
- [52] Vercel. *Real Experience Score*. Available at <https://vercel.com/docs/speed-insights/metrics#real-experience-score-res>, last accessed in June 2025.
- [53] Philip Walton. *First Contentful Paint*. Available at <https://web.dev/articles/fcp>, last accessed in June 2025.
- [54] J. Wang et al. “DeepCarotene -Job Title Classification with Multi-stream Convolutional Neural Network”. In: *2019 IEEE International Conference on Big Data (Big Data)*. 2019, pp. 1953–1961. DOI: 10.1109/BigData47090.2019.9005673. URL: <https://ieeexplore.ieee.org/document/9005673/>.

- [55] X. Wang. “Construction of Alumni Information Analysis Model Based on Big Data”. In: *Mathematical Problems in Engineering* 2022 (2022). ISSN: 1024123X (ISSN). DOI: 10.1155/2022/1587793. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85136569752&doi=10.1155%2f2022%2f1587793&partnerID=40&md5=0ebff319ca55c49ebf15f2238504a02d>.
- [56] Juanita Yusanti and Slamet Riyadi. “Dashboard-based Alumni Tracer Study Report Using Normalized Data Store Architecture”. In: *Proceedings of the 2019 8th International Conference on Software and Computer Applications*. Association for Computing Machinery Asroni, 2019, pp. 70–74. DOI: 10.1145/3316615.3316635. URL: <https://doi.org/10.1145/3316615.3316635>.
- [57] J. Zhang and M. L. Huang. “2D approach measuring multidimensional data pattern in big data visualization”. In: *2016 IEEE International Conference on Big Data Analysis (ICBDA)*. 2016, pp. 1–6. DOI: 10.1109/ICBDA.2016.7509823. URL: <https://ieeexplore.ieee.org/document/7509823/>.
- [58] M. Zhang, R. van der Goot, and B. Plank. “ESCOXLM-R: Multilingual Taxonomy-driven Pre-training for the Job Market Domain”. In: *Proceedings of the Annual Meeting of the Association for Computational Linguistics*. Vol. 1. Association for Computational Linguistics (ACL), 2023, pp. 11871–11890. ISBN: 0736587X (ISSN); 978-195942972-2 (ISBN). DOI: 10.18653/v1/2023.acl-long.662. URL: <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85174422231&doi=10.18653%2fv1%2f2023.acl-long.662&partnerID=40&md5=0f9d3de0b1295f8d64fb956a97cedc25>.

Appendix A

Architecture

A.1 Job Classification Agent - Confidence Calculation

The following code fragment implements the job classification confidence calculation as described in [subsection 5.4.3](#):

```
1      # Generate embedding for the query
2      # The query is the role title + description (if available)
3      query_embedding = generate_embedding(query)
4
5      computed_col = (EscoClassification.embedding.cosine_distance(query_embedding)).
6          label(
7              "similarity_score"
8          )
9
10     # The top k (20) most similar ESCO occupations are retrieved
11     search_query = (
12         select(
13             EscoClassification.id,
14             EscoClassification.code,
15             EscoClassification.title_en,
16             computed_col,
17         )
18         .where(EscoClassification.is_leaf.is_(True))
19         .order_by(computed_col)
20         .limit(top_k)
21     )
22
23     results = db.execute(search_query).fetchall()
24     # Vector distance of 0 means perfect match, hence min()
25     best_score = min(float(row.similarity_score) for row in results)
26
27     # Convert to confidence scores that balance absolute and relative similarity
```

```
27     # - Absolute similarity to ensure high confidence only for genuinely good
      matches
28     # - Relative similarity to preserve ranking
29     # - Scaling factor to ensure reasonable confidence distribution
30     results = [
31         EscoResult(
32             ...
33             confidence=round(
34                 # Start with absolute similarity (0 to 1)
35                 ((2.0 - float(row.similarity_score)) / 2.0)
36                 *
37                 # Multiply by relative similarity factor (penalizes matches much worse
                    than best)
38                 ((2.0 - float(row.similarity_score)) / (2.0 - best_score)),
39                 2,
40             ),
41         )
42     for row in results
43 ]
```

Here, `similarity_score` is the cosine distance between the role and ESCO occupation embeddings. The final confidence combines both absolute and relative similarity to enhance interpretability.

Appendix B

Feedback Form and Results

B.1 Feedback Form

The feedback form presented below was used to gather responses from platform users. Certain questions have been omitted, as they pertain to a parallel dissertation utilizing the same dataset but addressing features not covered in this work, such as the World Map and Reviews.

Alumni-FEUP Feedback

Welcome to the Alumni-FEUP platform

Thank you for testing the platform! This short survey (≈ 7 minutes) is part of the masters' thesis of Carlos Veríssimo (M.EIC) and José Pessoa (MESW) and will help validate the key features of Alumni-FEUP World.

Your feedback is incredibly valuable and appreciated.

[Access the Platform](#)

* Indicates required question

1. Email *

2. Do you consent to participate in this feedback survey as part of a research project? *

Mark only one oval.

☐ Yes

☐ No

Demographic Information

This section helps us understand who is interacting with the platform. Your answers will only be used for research purposes and will remain anonymous.

3. What is your relationship with the platform? *

If applicable, you can select more than one option

Tick all that apply.

- ☐ FEUP alumni
☐ Alumni from another institution
☐ Professor
☐ Faculty/administration staff
☐ Other: _____

4. (Optional) What was your degree? *

Mark only one oval.

- ☐ L.EIC - Licenciatura em Engenharia Informática e Computação (Bologna)
☐ LEIC - Licenciatura em Engenharia Informática e Computação (Pré-Bologna)
☐ M.EIC - Mestrado em Engenharia Informática e Computação
☐ MESW- Mestrado em Engenharia de Software
☐ MEI - Mestrado em Engenharia Informática
☐ MIEIC - Mestrado Integrado em Engenharia Informática e Computação
☐ Other: _____

System Usability Scale

For each of the questions below, please rate your level of agreement on a scale from 1 (Strongly disagree) to 5 (Strongly agree).

5. I think that I would like to use the application frequently. *

Mark only one oval.

	1	2	3	4	5	
Stro	☐	☐	☐	☐	☐	Strongly agree

6. I found the application unnecessarily complex. *

Mark only one oval.

	1	2	3	4	5	
Stro	☐	☐	☐	☐	☐	Strongly agree

7. I thought the application was easy to use. *

Mark only one oval.

	1	2	3	4	5	
Stro	☐	☐	☐	☐	☐	Strongly agree

8. I think that I would need the support of a technical person to be able to use the application. *

Mark only one oval.

	1	2	3	4	5	
Stro	☐	☐	☐	☐	☐	Strongly agree

9. I found that the various features (world map, filters, analytics, reviews) in the application were well integrated. *

Mark only one oval.

	1	2	3	4	5	
Stro	☐	☐	☐	☐	☐	Strongly agree

10. I thought there was too much inconsistency in this application. *

Mark only one oval.

1 2 3 4 5

Stro ☐ ☐ ☐ ☐ ☐ Strongly agree

11. I would imagine that most people would learn to use this application very quickly. *

Mark only one oval.

1 2 3 4 5

Stro ☐ ☐ ☐ ☐ ☐ Strongly agree

12. I found the application very cumbersome to use. *

Mark only one oval.

1 2 3 4 5

Stro ☐ ☐ ☐ ☐ ☐ Strongly agree

13. I felt very confident using the application. *

Mark only one oval.

1 2 3 4 5

Stro ☐ ☐ ☐ ☐ ☐ Strongly agree

14. I needed to learn a lot of things before I could get going with this application. *

Mark only one oval.

1 2 3 4 5


Stro ☐ ☐ ☐ ☐ ☐ Strongly agree

World Map

19. (Optional) Are there any additional visualization or filtering options that would be useful?

Analytics

This section focuses on the Analytics page, where users can filter and explore data on alumni roles, companies, locations, and more. We want to understand how intuitive and useful this experience was.



Filters

The questions below pertain to the adequacy, ease of use, and overall usefulness of the filters provided on the analytics page.

20. It was easy to use the filters on the Analytics page *

Mark only one oval.

1 2 3 4 5

Stro ☐ ☐ ☐ ☐ ☐ Strongly Agree

21. The filters for exploring the alumni data were useful and adequate *

Mark only one oval.

1 2 3 4 5

Stro ☐ ☐ ☐ ☐ ☐ Strongly Agree

22. Which filter(s) did you use the most? *

Tick all that apply.

- ☐ Company
- ☐ Country/City
- ☐ Role/Position
- ☐ Seniority level
- ☐ Industry
- ☐ Degree/Conclusion Year
- ☐ I didn't use the filters
- ☐ Other: _____

23. Did you feel like any important filters were missing? *

Mark only one oval.

- ☐ Yes
- ☐ No

24. (Optional) If you've answered **Yes** above, which filters would you have liked to see

Visualization

This section focuses on the different ways data is presented: charts, tables, trends, and how easily you could explore deeper levels of information (e.g., by course, ESCO classification, etc.).

25. Overall, I found the visualizations helpful (charts, tables, timelines)? *

Mark only one oval.

1 2 3 4 5

Strongly ☐ ☐ ☐ ☐ ☐ Agree

26. Were you able to drill down into the data in a meaningful way (e.g., finding specific cohorts, titles, or categories)? *

By “drill down,” we mean navigating through layers of data, such as going from a **country** to its **cities**, from **faculty** to related **courses**, from a **course** to its **cohorts**, or from **Level 1 ESCO groups** to more specific **job titles** in lower levels of the taxonomy.

Mark only one oval.

1 2 3 4 5

Not ☐ ☐ ☐ ☐ ☐ Absolutely

27. Did any visualization feel unnecessary, confusing, or redundant? *

28. It was clear to me how to switch between different views

Mark only one oval.

1 2 3 4 5

Strongly ☐ ☐ ☐ ☐ ☐ Agree

Other

This section covers general impressions of the analytics page, including how well everything fits together, if the information shown was relevant, and any additional suggestions you may have.

29. I found the page useful in answering specific questions (e.g. "Where are people who graduated with me working at", "where are people from my hometown working at", etc.) *

Mark only one oval.

1 2 3 4 5

Strongly ☐ ☐ ☐ ☐ ☐ Agree

30. (Optional) If there was any information you expected to find but couldn't, please describe here.

31. Was the layout of the analytics page overwhelming, balanced, or too simple?

Mark only one oval.

☐ Overwhelming (too much going on)

☐ Well balanced

☐ Too minimal / lacking context

☐ Other: _____

Qualitative Questions

This section is optional – but your input here really helps us improve.

32. (Optional) Was there a particular insight or use case that stood out to you?

33. (Optional) If you could add or change one thing in the analytics page, what would it be?

Personal Data & Interaction Flow

This section focuses on your experience interacting with your personal profile: including how easy it was to find, view, and update your information, and how consistent the experience felt when navigating between pages or features (e.g., analytics ↔ world). We're also interested in understanding how intuitive the classification editing was (job title, seniority, etc.).

Note: If you were not able to login with your LinkedIn account, please discard the questions related to updating and editing you information

Interaction

This section focuses on how well the platform connects data in the Analytics page with its representation on the Map, for example, clicking a country or city to view it spatially, using the "View on Map" button, or locating an individual alumni's position via the Alumni table

34. Did you use the "View on Map" button to explore alumni locations based on applied filters (course & graduation year)? *

Mark only one oval.

☐ Yes

☐ No

35. Did you try clicking on a country or city in the analytics view to locate it on the map? *

Mark only one oval.

☐ Yes

☐ No

36. Were you able to find and view the location of a specific alumni via their profile or entry in the Alumni table? *

Mark only one oval.

☐ Yes

☐ No

37. I feel that the connection between data and map is clear and consistent

Mark only one oval.

1 2 3 4 5

Strongly ☐ ☐ ☐ ☐ ☐ Strongly Agree

Profile Page

38. Were you logged in during your visit to the platform? *

Mark only one oval.

☐ Yes

☐ No *Skip to section 9 (End)*

Profile Page

39. Were you able to view and access your profile information, as well as your career timeline?

Mark only one oval.

☐ Yes, everything worked smoothly

☐ Yes, but some parts were unclear

☐ No, I couldn't find or open my profile

40. It was easy to understand and update my AI-generated job classification?

Mark only one oval.

1 2 3 4 5

Strongly ☐ ☐ ☐ ☐ ☐ Strongly Agree

41. How accurate was the AI-generated job classification for your current role?

The job title and seniority level you see on your profile were automatically generated using AI, based on the [ESCO taxonomy](#), a standardized European classification of skills, competences, and occupations.

ESCO doesn't include some commonly used titles like **"Software Engineer," "Backend Developer," etc...** As a result, your role may have been mapped to a broader or more formal equivalent (e.g., *Software Developer, Application Engineer*).

You can always review and refine the classification to better reflect your actual position.

Mark only one oval.

1 2 3 4 5

Correct ☐ ☐ ☐ ☐ ☐ Spot on

42. And what about your seniority level – how accurate was that classification?

Mark only one oval.

1 2 3 4 5

Not ☐ ☐ ☐ ☐ ☐ Very accurate

43. Would you have preferred more guidance or explanation for the suggested classifications?

Mark only one oval.

- ☐ Yes, more context would help
☐ No, the information was clear
☐ Not sure

B.2 Results

B.2.1 Job Classification Confidence Scores

ISCO Group	# Roles	Avg (%)	Min (%)	Max (%)
Information and Communications Technology Professionals	3,611	70.0	57.0	85.0
Science and Engineering Professionals	966	69.3	56.0	83.0
Administrative and Commercial Managers	777	70.3	54.0	85.0
Teaching Professionals	691	70.0	55.0	83.0
Production and Specialized Services Managers	555	70.2	57.0	84.0
Business and Administration Professionals	552	70.2	56.0	85.0
Legal, Social and Cultural Professionals	235	67.9	55.0	83.0
Chief Executives, Senior Officials and Legislators	141	67.0	56.0	82.0
Business and Administration Associate Professionals	125	67.0	53.0	79.0
Science and Engineering Associate Professionals	116	67.4	59.0	86.0
Legal, Social, Cultural and Related Associate Professionals	72	66.3	57.0	80.0
Sales Workers	44	69.1	59.0	83.0
General and Keyboard Clerks	42	68.2	59.0	76.0
Hospitality, Retail and Other Services Managers	33	69.6	57.0	76.0
Information and Communications Technicians	22	70.2	62.0	82.0
Drivers and Mobile Plant Operators	20	68.1	61.0	69.0
Health Professionals	20	69.9	59.0	81.0
Numerical and Material Recording Clerks	16	69.9	63.0	75.0
Customer Services Clerks	13	70.4	64.0	77.0
Personal Service Workers	11	71.5	62.0	84.0
Food Processing, Woodworking, Garment and Other Craft and Related Trades Workers	10	64.8	59.0	73.0
Market-oriented Skilled Agricultural Workers	8	68.6	64.0	71.0
Electrical and Electronics Trades Workers	8	68.3	63.0	75.0
Protective Services Workers	8	73.4	62.0	80.0
Market-Oriented Skilled Forestry, Fishery and Hunting Workers	6	69.5	68.0	73.0
Health Associate Professionals	5	67.8	64.0	74.0
Commissioned Armed forces Officers	3	73.3	69.0	77.0
Personal Care Workers	3	67.3	65.0	69.0
Metal, Machinery and Related Trades Workers	2	71.0	68.0	74.0
Handicraft and Printing Workers	1	56.0	56.0	56.0
Other Clerical Support Workers	1	61.0	61.0	61.0
Stationary Plant and Machine Operators	1	71.0	71.0	71.0
Non-commissioned Armed Forces Officers	1	83.0	83.0	83.0
Labourers in Mining, Construction, Manufacturing and Transport	1	70.0	70.0	70.0
Total	8,120	69.7	53.0	86.0

Table B.1: Job classification confidence scores for all ISCO sub-major groups

Table B.1 extends the results from Table 5.10 by including all the 34 sub-major groups found in the dataset.

B.2.2 Real Misclassification Cases

To provide concrete insight into the challenges encountered by the classification pipeline, the following paragraphs present a handpicked selection of real cases in which the predicted ESCO code differed from the expected classification. These examples illustrate common sources of ambiguity, including overlapping job titles, taxonomy coverage gaps, and limitations in distinguishing highly specialized or emerging roles.

Instructor (IT) → Sports Instructor

- **Original Title:** Instructor (IT)
- **Predicted ESCO:** Sports Instructor
- **Top 3 Choices:** Sports Instructor, Firefighter Instructor, Fitness Instructor

The system mapped an IT instructor to "Sports Instructor" due to the lack of an IT-specific instructor code in ESCO, resulting in a clear domain mismatch.

The LLM's reasoning acknowledges this, and states: "The top match is 'Sports Instructor' due to the highest confidence score, but it is not domain-specific to IT..."

Higher Education Lecturer (code 2310.1) would have been a possible choice for this role.

Tech Lead → Chief Technology Officer

- **Original Title:** Tech Lead
- **Predicted ESCO:** Chief Technology Officer
- **Top 3 Choices:** Chief Technology Officer, Food Technologist, Food Scientist

The system mapped a Tech Lead at a food delivery startup to Chief Technology Officer, primarily based on the candidate's leadership over technology strategy and development. However, the role description emphasizes hands-on engineering and project leadership, not necessarily C-level responsibilities, making CTO an overreach.

This demonstrates the inherent ambiguity of some of the roles in the dataset, where a clear-cut mapping is not always possible. For this role, a less overreaching choice would have been Software Manager (code 1330.5.1).

Chief of Engineering Staff → Marine Chief Engineer

- **Original Title:** Chief of Engineering Staff
- **Predicted ESCO:** Marine Chief Engineer

- **Top 3 Choices:** Marine Chief Engineer, Military Engineer, Engineering Assistant

This is a clear example where relying on embeddings and surface-level similarity can result in an implausible match. While Chief of Engineering Staff denotes a senior engineering leader (typically in a corporate or tech context), the model mapped it to Marine Chief Engineer, a role specific to maritime operations, simply because both contain the terms `Chief` and `Engineer` and indicate leadership. This highlights the risk of semantically similar, but contextually inappropriate, ESCO assignments.

A more context-aware system would have mapped this role to a generic Software Manager (code 1330.5.1), although this is far from a perfect match, denoting the taxonomy's limitations in less specialized roles.

Senior Full Stack Software Engineer (Tech Lead) → Cloud Software Developer

- **Original Title:** Senior Full Stack Software Engineer (Tech Lead)
- **Predicted ESCO:** Cloud Software Developer
- **Top 3 Choices:** Cloud Software Developer, Cloud DevOps Engineer, Integration Engineer

Although the system selected Cloud Software Developer to emphasize the use of modern cloud technologies, this role was primarily a senior full stack developer and tech lead working on logistics and supply chain software. A generic Software Developer code would have provided a simpler and more accurate classification, avoiding unnecessary specificity.

Senior Analyst / Full Stack Developer → Cloud Software Developer

- **Original Title:** Senior Analyst / Full Stack Developer
- **Predicted ESCO:** Cloud Software Developer
- **Top 3 Choices:** Cloud Software Developer, Cloud DevOps Engineer, Software Developer

The system assigned Cloud Software Developer, focusing on the cloud-native and DevOps elements in the description. However, the original role encompassed broader responsibilities, including analysis and full stack development for financial software. A standard Software Developer label would have more accurately captured the diversity of tasks in this position.

Chief Gaming Officer → Casino Gaming Manager

- **Original Title:** Chief Gaming Officer
- **Predicted ESCO:** Casino Gaming Manager
- **Top 3 Choices:** Casino Gaming Manager, Chief Product Officer, Chief Technology Officer

In this case, the system mapped "Chief Gaming Officer", a modern executive role often found in software or gaming companies, to "Casino Gaming Manager", a traditional ESCO occupation tied to the casino and gambling sector.

Since no role description was available on LinkedIn, the LLM relied solely on the title, which increases ambiguity, especially for creative or niche titles. This reflects a data quality issue, where the role did not contain enough information to make a more accurate mapping, which, possibly, would've been a C-level role.

Software Engineer → Application Engineer

- **Original Title:** Software Engineer
- **Predicted ESCO:** Application Engineer
- **Top 3 Choices:** Application Engineer, Software Developer, Embedded Systems Software Developer

Here, the system mapped the standard role Software Engineer to Application Engineer, prioritizing the application development aspect. However, Software Developer would have been a more appropriate and widely recognized match, as it is the canonical ESCO code for general software engineering tasks.

This example highlights the subtle ambiguity and overlap present in software job titles, and the need for careful disambiguation even in routine mappings.

Senior Backend Developer → E-Learning Developer

- **Original Title:** Senior Backend Developer
- **Predicted ESCO:** E-Learning Developer
- **Top 3 Choices:** E-Learning Developer, Corporate Training Manager, Enterprise Development Worker

The system mapped Senior Backend Developer to E-Learning Developer, primarily because the role included responsibilities for technical training and team mentoring. However, the core focus of the position was backend software development, not educational content creation. A standard Software Developer would have been more accurate.

This example demonstrates how auxiliary responsibilities (like training) can sometimes overly influence the classification, resulting in an ESCO title that misrepresents the actual job.