Gender Differences in ESCO Skill Acquisition: A Data-Driven Analysis in Greek Higher Education

Dimitrios Christos Kavargyris
School of Informatics, Aristotle University of Thessaloniki
Thessaloniki, Greece
dkavargy@csd.auth.gr

Konstantinos Demerdeslis
School of Informatics, Aristotle University of Thessaloniki
Thessaloniki, Greece

Lefteris Angelis
School of Informatics, Aristotle University of Thessaloniki
Thessaloniki, Greece

Abstract

Evidence suggests that male and female students may differ not only in disciplinary choices but also in the types of competencies they develop during their studies. This paper presents a data-driven analysis of gender-based differences in skill acquisition using a corpus of 1,400 undergraduate theses from four Greek universities (2009–2025). Leveraging the European Skills/Competences, Qualifications and Occupations (ESCO) taxonomy and the LLM-based ESCOX tool within the G-SKILL framework, as part of the EU Horizon-funded SKILLAB project, we extract and compare competencies across male and female students spanning all academic disciplines, from engineering and computer science to humanities, health sciences and social sciences. XGBoost-based classification (95% accuracy) and SHAP interpretability reveal that male students show stronger engagement with skills such as blockchain technologies, software configuration tools, and digital game development, while female students show higher representation in health-oriented and interpersonal domains.

Keywords

gender equality, IT education, ESCO, skill analysis, explainable AI (XAI), SKILLAB

1 Introduction

When a female engineering student chooses a thesis topic on user experience over network security, or a male medical student gravitates toward surgical technique over patient communication, these are not random choices, they reflect years of subtle signals about what each gender is "supposed" to be good at. Identifying where these patterns manifest at the skill level, and whether they are consistent across institutions and disciplines, is the first step toward dismantling them. Despite decades of policy efforts, women account for only approximately 17% of ICT graduates in the EU [1], and this structural gap may also shape how students develop and demonstrate research-related competencies during higher education [2]. Understanding which specific skills diverge across genders [3], across all academic disciplines, from engineering to humanities, is essential for designing inclusive curricula and targeted interventions.

This paper summarises key findings from the G-SKILL (Gendered Skill Insight for Learning and Labour) framework [4], developed as part of the EU Horizon-funded SKILLAB project [5]. G-SKILL combines machine learning (ML) and eXplainable AI (XAI) to quantify gender-based differences in skill acquisition from undergraduate theses, operationalising competencies via the European Skills/Competences, Qualifications and Occupations (ESCO) taxonomy. The contribution of this work lies not in claiming that one gender outperforms the other, but in providing an evidence-based framework through which Greek universities and policymakers can identify structural patterns in skill exposure and design more inclusive academic environments.

2 Methodology

A balanced corpus of 1,400 undergraduate theses was collected from four Greek universities: Aristotle University of Thessaloniki (AUTH)[1], National Technical University of Athens (NTUA)[2], National and Kapodistrian University of Athens (NKUA)[3], and University of Western Macedonia (UOWM)[4] (2009–Q2 2025), evenly split across gender (700F/700M) and discipline (STEM / Humanities). The breadth of disciplines is intentional: by sampling across all fields, the framework captures the full spectrum of skills that emerge in undergraduate research, including technical, biomedical, interpersonal and analytical competencies. Skills, defined as the competencies that students demonstrate through the writing and implementation of their undergraduate thesis, were extracted from titles and abstracts using ESCOX [6], an LLM-based tool mapping text to ESCO entries via cosine similarity (threshold = 0.5). This cosine similarity threshold is not selected arbitrarily, but is grounded in prior literature and validated through previous studies that applied the same tool to extract skills from academic papers [7], thesis texts [8], policy documents [9], online job advertisements [10], and, more generally, unstructured text.

Six Machine Learning (ML) classifiers were evaluated under 10-fold cross-validation: Decision Tree, Random Forest, Logistic Regression, k-Nearest Neighbors, Gradient Boosting and **XGBoost**, an ensemble method that builds sequential decision trees, each correcting the errors of the previous, yielding robust and accurate predictions. SHapley Additive exPlanations (SHAP) [11] were then applied as XAI method to identify the skills most predictive of gender, moving beyond black-box accuracy toward a transparent understanding of which competencies drive the model's decisions. SHAP assigns each skill a value representing its marginal contribution to the gender prediction. In the binary classification task, male is encoded as the positive class (1) and female as the negative class (0). Accordingly, positive SHAP values indicate association with male authorship and negative values with female authorship.

[1] https://www.auth.gr
[2] https://www.ntua.gr
[3] https://www.uoa.gr
[4] https://www.uowm.gr

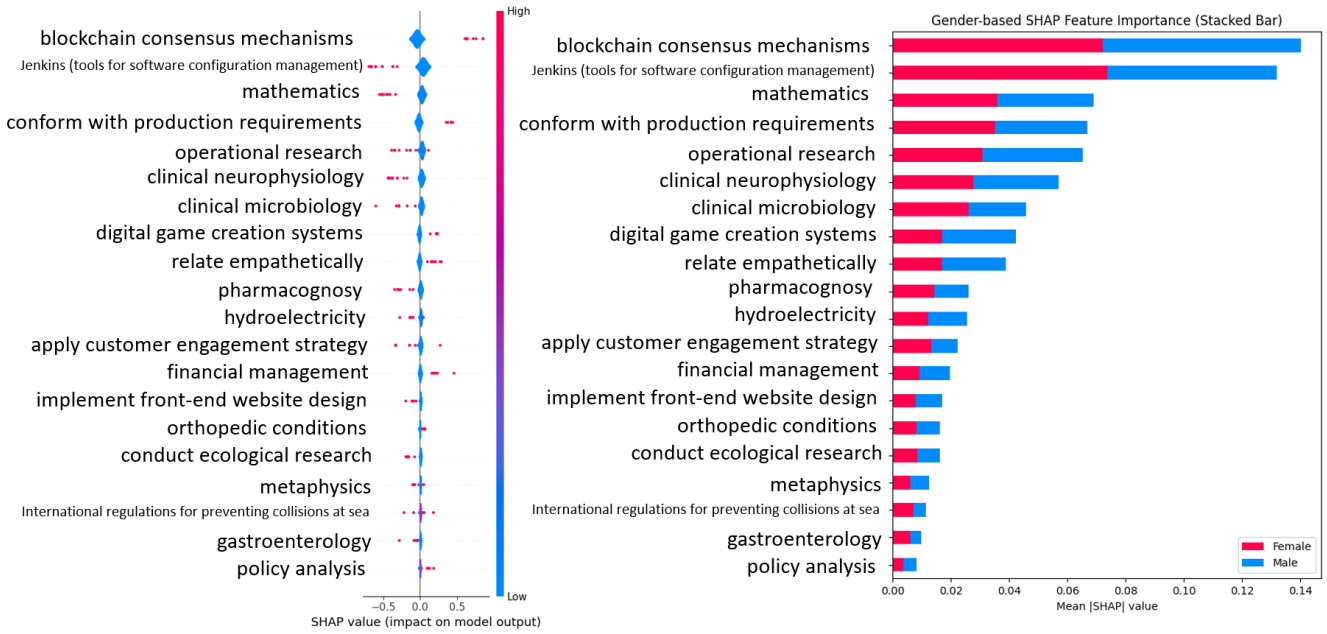

Figure 1: Gender-based SHAP feature importance across 20 ESCO skills (XGBoost, 10-fold CV, accuracy 0.95).

3 Key Findings

Figure 1 reveals two distinct gendered skill clusters. The **left panel** shows that *blockchain consensus mechanisms* and *digital game creation systems* carry the strongest positive SHAP values, pushing predictions toward *male* authorship. Conversely, *Jenkins (software configuration)*, *clinical neurophysiology*, *relate empathetically*, and *pharmacognosy* show negative SHAP values, driving predictions toward *female* authorship, a notable finding, as Jenkins is a technical tool typically associated with software engineering, suggesting that female students engaging with configuration and DevOps topics are underrepresented in the male-dominated narrative around such skills. The **right panel** confirms this split quantitatively: blockchain and digital game creation show blue-dominant (male) bars, while health-, interpersonal and software configuration skills show red-dominant (female) contributions.

Table 1 reports the largest marginal differences ($p < 0.05$, Chi-square test). Crucially, this is not a gap in analytical ability. Mathematics appears more frequently in female theses (53.5% vs. 44.8%), likely because female students in health and social sciences disciplines explicitly employ statistical analysis, as reflected in skills such as pharmacognosy and policy analysis. In contrast, male students in engineering and computing apply mathematical reasoning implicitly through code, where formal mathematical terminology rarely surfaces in thesis abstracts. Yet specific IT domains, game development, blockchain, front-end engineering, remain male-dominated. The gap is one of exposure and channelling, not capability. At the institutional level, NTUA and NKUA exhibit greater variance in gender-linked skill distributions, while UOWM shows more homogeneous profiles, reflecting differences in disciplinary focus and curriculum structure across Greek universities.

Table 1: Selected IT/STEM skills with significant gender differences ($p < 0.05$).

Skill	F%	M%	Gap
Digital game creation systems	31.7	48.4	+16.7 M
Blockchain consensus mechanisms	39.0	48.0	+9.0 M
Implement front-end website design	46.1	54.0	+7.9 M
Mathematics	53.5	44.8	+8.7 F
Clinical neurophysiology	57.7	38.0	+19.7 F
Relate empathetically	55.6	36.0	+19.6 F
Conduct ecological research	56.6	40.8	+15.8 F
Policy analysis	54.5	47.2	+7.3 F

4 Implications for Gender Equality in Greek Universities

Curriculum design. Female students are underrepresented in specific IT domains, game development, blockchain, front-end engineering, not in analytical skills broadly. Introducing these topics in inclusive, application-driven ways (e.g. linking blockchain to policy or healthcare contexts) can lower perceived barriers. Institutions should also audit their thesis supervision practices, ensuring female students receive equal encouragement to pursue technically intensive research topics regardless of disciplinary norms.

Rethinking the mathematics narrative. The finding that mathematics appears more frequently in female theses [12] challenges the persistent stereotype that women avoid quantitative reasoning. Universities should actively communicate this evidence to counteract discouraging narratives that steer female students away from technical fields before they even begin.

Institutional awareness. The variance observed across universities, with NTUA and NKUA showing greater gender-linked skill divergence and UOWM more homogeneous profiles, suggests that institutional culture and curriculum structure play a significant role. Universities with higher divergence should prioritise targeted interventions, while those with more balanced profiles can serve as models for best practice.

Mentoring. Male students show lower engagement with collaborative and communication skills increasingly valued in tech workplaces. Cross-gender mentoring programmes addressing both sides of the gap contribute to more balanced IT teams [13]. Pairing students across skill clusters, technically strong male students with socially and analytically strong female peers, can foster mutual competency development and challenge entrenched stereotypes on both sides.

Data-driven monitoring. The G-SKILL framework offers a replicable tool to continuously audit skill gaps by gender and institution, supporting progress toward the EU Gender Equality Strategy 2020–2025. Deploying such monitoring at scale across Erasmus+ partner universities would enable the first systematic, cross-national benchmark of gendered IT skill supply in European higher education. The aggregated results could further inform national funding bodies and employers, aligning graduate skill profiles with the demand for diverse and inclusive tech workforces.

5 Conclusion

Using 1,400 ESCO-aligned undergraduate theses from four Greek universities, this paper demonstrates a structural IT skill gap between genders rooted in exposure and channelling, not ability. The G-SKILL framework provides a transparent, data-driven mechanism to identify and address these gaps, contributing to gender-equal participation in computing. Within STEM, preliminary observations suggest that sub-disciplinary culture further shapes these patterns: Electrical and Computer Engineering departments, for instance, show stronger male-associated technical skill clusters compared to Computer Science departments where skills overlap more evenly across genders, a distinction that warrants dedicated investigation in future work. Future work will extend the analysis across broader European university systems to examine whether these patterns persist across different cultural and educational contexts, enabling cross-country comparisons within the European Higher Education Area. It should be noted that while the corpus is balanced across broad disciplinary categories, within-category departmental composition may still vary by gender.

Furthermore, extending the corpus to postgraduate and doctoral theses would allow examination of whether gendered skill patterns persist or converge with increasing academic specialisation. Additionally, while ESCOX leverages an LLM for skill extraction, potential gender bias inherent in such models is partially mitigated by constraining the output to the fixed, standardised ESCO taxonomy, ensuring that skill identification is anchored to predefined competency definitions rather than free-form model generation. Beyond disciplinary stratification, two further limitations should be noted: first, the binary gender classification reflects institutional metadata and does not capture the full spectrum of gender identities; second, as the skill extraction relies on an LLM-based tool, residual model biases may influence which competencies surface, though constraining output to the fixed ESCO taxonomy partially mitigates this risk.

Acknowledgments

This work is funded by the SKILLAB project, EU Horizon Europe Framework, Grant Agreement No. 101132663.

References

[1] European Institute for Gender Equality (EIGE). 2018. *Gender Equality Index 2018: ICT Graduates*. https://eige.europa.eu/

[2] F. Siddiq and R. Scherer. 2019. Is there a gender gap? A meta-analysis of the gender differences in students' ICT literacy. *Educational Research Review*, 27, 205–217.

[3] W. J. He and W. C. Wong. 2021. Gender differences in the distribution of creativity scores: Domain-specific patterns in divergent thinking and creative problem solving. *Frontiers in Psychology*, 12, 626911.

[4] Demerdelis, K. G. (2025). Skill analysis from university theses in relation to the author's gender (Undergraduate Thesis, School of Informatics, Aristotle University of Thessaloniki, Greece.)

[5] M. Aluas, L. Angelis, I. Arapakis, *et al.* 2024. SKILLAB: Skills matter. In *Proc. 50th Euromicro Conf. SEAA*, IEEE, pp. 491–498.

[6] D. C. Kavargyris, K. Georgiou, E. Papaioannou, K. Petrakis, N. Mittas, and L. Angelis. 2025. ESCOX: A tool for skill and occupation extraction using LLMs. *Software Impacts*, 25, 100772.

[7] D. C. Kavargyris, N. Mittas, and L. Angelis. 2026. Growing skills from code via SciESCO: A tri-phasic bibliometric-driven framework of scientific software development. *Journal of Systems and Software*, 112864.

[8] C. C. Chilmi, D. C. Kavargyris, N. Mittas, and L. Angelis. 2025. ESCOClassifyAI: A semi-automated software solution for ESCO-skill-based document classification. In *Proceedings of the 16th International Conference on Information, Intelligence, Systems & Applications (IISA)*, 1–8. IEEE.

[9] D. C. Kavargyris, P. Gkorezis, N. Mittas, and L. Angelis. Predicting skills in digital publication policies: A linear discriminant analysis and machine learning classification approach.

[10] V. Ntaoulas, K. Georgiou, N. Mittas, and L. Angelis. 2025. H-TURF: Detecting optimal green software engineering skillsets using TURF analysis and hierarchical cumulative voting. In *Proceedings of the Euromicro Conference on Software Engineering and Advanced Applications*, 121–136. Springer Nature Switzerland.

[11] S. M. Lundberg and S. Lee. 2017. A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems (NeurIPS)*, 30.

[12] F. Ulusoy, M. Sadak, L. Incikabi, and M. Pektas. 2025. Exploring gender differences in mathematical creativity: linking problem posing and problem solving. *Psychology in the Schools*, 62(10), 4327–4343.

[13] A. Alblas. 2022. Globally distributed technology projects thrive on experience working together apart: Impacts of team balance, task balance, and task heterogeneity. *Journal of Operations Management*, 68(6–7), 728–754.