Proceedings of
the 45th International Business Information Management Association Computer Science Conference

(IBIMA)

25-26 June 2025

ISBN: 979-8-9867719-7-7

ISSN: 2767-9640

Advancements in Artificial Intelligence and Computer Security in Modern Computing

Editor

Khalid S. Soliman

International Business Information Management Association (IBIMA)

Conference Chair

Khalid S. Soliman, International Business Information Management Association, USA

Conference Advisory Committee

Najib Saylani, Florida State College at Jacksonville, USA
Omar Al-Azzam, University of Minnesota Crookston, USA
Xiuzhen Feng, Beijing University of Technology, China
Hesham H. Ali, University of Nebraska at Omaha, USA
Magdy Serour, InContext Solutions, Australia
Ahmed Azam, DeVry University, USA

Program Committee

(it is IBIMA Policy to include a program committee member's name only after reviewing at least one submitted paper)

Katarzyna Skrzypek, University of Zielona Góra, Poland
Bahjat Fakieh, King Abdulaziz University, KSA
Gopikrishna Tatapudi, JNTU-KAKINADA, India
Muhammad Usman Tariq, Abu Dhabi University, UAE
Reem Almetrami, Zarka University, Jordan
Sławomir Radomski, University of Gdańsk, Poland
Bogdan Țigănoaia, National University of Science and Technology POLITEHNICA Bucharest, Romania
Anna Hamranova, University of Economics in Bratislava, Slovakia
Maciej Hojda, Wrocław University of Science and Technology, Poland
Grażyna Kowalewska, University of Warmia and Mazury, Poland
Cristian Bucur, Petroleum-Gas University of Ploiesti, Romania
Geerish Suddul, University of Technology, Mauritius (UTM), Mauritius
Clara Silveira, Instituto Politécnico da Guarda, Portugal
Jaroslaw Michalak, Military University of Technology, Poland
Mateusz Sergiel, Military University of Technology, Poland
Lorena Pocatilu (Batagan), Bucharest University of Economics, Romania
Łukasz Paśko, Rzeszów University of Technology, Poland
Samira Ellouze, University of Gabes, Tunisia
Grzegorz Popek, Wroclaw University of Science and Technology, Poland
Iwona Skalna, AGH University of Krakow, Poland
Mohammad Aizat Basir, Universiti Malaysia Terengganu (UMT), Malaysia
Gabriel Koman, University of Zilina, Slovakia
Alok Abhishek, University of Southern California, USA
Livia Sangeorzan, University Transilvania of Brasov, Romania
Diana Bratić, University of Zagreb, Croatia
Constantin Ilie, OVIDIUS University from Constanta, Romania
António Trigo, Instituto Universitário de Lisboa (ISCTE-IUL), ISTAR, Portugal
Sreekanth Ramakrishnan, Senior Software Engineer, Netflix, USA
Sebastian Meszyński, Nicolaus Copernicus University in Toruń, Poland
Zdzisław Sroczyński, Silesian University of Technology, Poland
Piotr Kohut, AGH University of Krakow, Poland
Łukasz Czerniszewski, MUT, Poland
Agnieszka Bier, Silesian University of Technology, Poland

Table of Content

Generating Learning Paths from Job Postings via Bayesian Networks

Florin STOICA
Lucian Blaga University of Sibiu, Sibiu, Romania, florin.stoica@ulbsibiu.ro
ORCID ID: 0000-0002-9073-0781

Dana SIMIAN
Lucian Blaga University of Sibiu, Sibiu, Romania, dana.simian@ulbsibiu.ro
ORCID ID: 0000-0002-5210-1810

Laura Florentina STOICA
Lucian Blaga University of Sibiu, Sibiu, Romania, laura.cacovean@ulbsibiu.ro
ORCID ID: 0000-0002-4758-6606

Elena-Cristina RAULEA
Lucian Blaga University of Sibiu, Sibiu, Romania, cristina.raulea@ulbsibiu.ro
ORCID ID: 0009-0001-5522-7128

Abstract

In today's dynamic job market, identifying the skills required by employers and helping learners acquire those skills efficiently is a significant challenge. In the context of the ENTEEF (Fostering Entrepreneurship through Freelancing) project, we address the problem of automatically generating structured learning paths using skills extracted from job offers. The ENTEEF project focuses on bridging the gap between education and industry by preparing students (and other target groups) with the competencies needed for freelancing careers. To this end, we propose a methodology that leverages Bayesian Networks (BN) to model relationships between skills extracted from ~30,000 job postings, and we use the network to derive recommended learning sequences. In our approach, each skill is a node in a BN and edges indicate dependency or co-occurrence relationships learned from job market data. We experiment with two weighting methods for the BN's arcs - mutual information score and normalized mutual information score - to quantify the strength of connections between skills. These weighting strategies, based on information theory, help rank possible skill progressions and guide the generation of optimal learning paths. Our method thus combines data-driven skill extraction with probabilistic modelling to produce learning pathways aligned with real-world job requirements.

Keywords: Bayesian Networks, learning paths, mutual information.

Introduction

Aligning educational programs with the rapidly evolving demands of the job market is crucial for improving employability and lifelong learning. Learners often struggle to decide what to learn next to reach their career goals, especially in fields like freelancing where required competencies are multifaceted. The ENTEEF project, under the Erasmus+ programme, aims to address this challenge by improving entrepreneurship competences among students and other groups, preparing them to work as freelancers and promoting lifelong learning and micro-credentials (ENTEEF, 2025). ENTEEF's approach includes thorough analysis of the freelancer job market to identify key skills, a Competency Assessment Tool (CAT) to determine an individual's skill gaps, and a suite of 12 Massive Open Online Courses (MOOCs) to help fill those gaps. Through a competency gap test and targeted MOOCs, learners receive personalized learning paths that enable them to acquire the skills needed for successful freelancing careers.

However, manually constructing these learning paths from labour market data is labour-intensive. Job postings contain rich information about skills sought by employers, and mining this data can reveal which skills are most in-demand and how they relate to each other. If we can automatically infer a structured map of skills from job offer data, we can generate data-driven learning pathways guiding learners from their current competencies to those required by their target jobs. Recent research underscores the value of aligning learning content with job market needs: for example, the study (Carroll and Schlippe, 2023) achieved over 94% accuracy in identifying job-market skills within course materials using AI and found that showing such alignment to learners improved their motivation and provided valuable career insights. This exemplifies how connecting education to real job requirements can inspire and direct learners.

In this paper, we propose a method to generate recommended learning paths based on skills extracted from job offers. Our approach uses a Bayesian Network to model the relationships among skills, treating each skill as a node and learning a directed acyclic graph from data. The context for this work is the ENTEEF project's goal of guiding learners (prospective freelancers) to acquire competencies demanded by the market. By mining ~30,000 job postings for skills, we create a probabilistic model of how these skills co-occur or depend on one another. We then use information-theoretic measures

- specifically the mutual information score and the normalized mutual information score - to weight the connections (arcs) in the network. The weighted BN serves as the foundation for constructing learning paths: a higher weight on an edge suggests a stronger relation that can be interpreted as a recommended transition or prerequisite link between skills.

Therefore, this study is grounded in the hypothesis that "*A Bayesian Network learned from large-scale job postings, when weighted using mutual information measures, can effectively model skill relationships and generate learning paths that reflect realistic upskilling trajectories aligned with job market demands*". By extracting and structuring these relationships from collected data, we aim to provide a data-driven foundation for guiding learners in acquiring market-relevant competencies, particularly in freelancing contexts.

The paper is organized as follows. In the next section, we review related work. In the third section we detail our methodology including data processing, Bayesian Network learning, the two arc-weighting schemes and generating learning paths. The fourth section presents aspects related to the construction of models and their comparative evaluation. Fifth section offers a discussion of the findings, practical implications, and limitations of our study. Finally, are presented some conclusions and opportunities for future investigations in the last section.

Related work

There is a growing body of research focused on extracting skill requirements from job postings using Natural Language Processing (NLP). Skill extraction is considered a core task in computational job market analysis. Early approaches relied on keyword matching or statistical language models, but recent methods leverage deep learning. For instance, Zhang et al. (2022) introduce the SKILLSPAN dataset for skill extraction and demonstrate state-of-the-art results using BERT-based models on job posting text. Their work is among the first to apply advanced language models to identify both hard and soft skills from job ads, highlighting the feasibility of automatically obtaining structured skill data. Such techniques provide the initial input (a set of skills per job offer) for our problem.

Bayesian Networks have been widely used to model knowledge and learning in educational technology. BNs provide a principled way to represent probabilistic relationships among a set of variables (e.g. skills or concepts) and have been used extensively as student models in intelligent tutoring systems. By encoding dependencies between knowledge components, BNs can infer a learner's mastery level or suggest next learning steps. For example, the study presented in (Culbertson, 2016) reviews numerous assessment systems that employ BNs to diagnose student understanding. In the context of learning path generation, Shen et al. (2020) proposed using a Bayesian network-based association rule algorithm to discover optimal learning paths among microlearning units. Their study created navigation paths for learners by finding correlations among course units, which is conceptually similar to our goal of linking skills. These works validate the idea of using network structures to represent learning sequences.

Beyond Bayesian networks, researchers have explored various techniques for recommending or generating personalized learning paths. A recent systematic review by Rahayu et al. (2023) indicates that many learning path recommender systems rely on ontology or knowledge-based representations. In such systems, relationships between concepts (or skills) are explicitly modelled (e.g., prerequisites in a domain ontology), and algorithms then search these graphs for an optimal path tailored to the learner's profile. Our approach aligns with this knowledge-based trend, but instead of manually crafted ontologies it uses a data-driven BN learned from job data.

Methodology

Our methodology consists of four main steps: (1) Data Collection and Skill Extraction, (2) Bayesian Network Structure Learning, (3) Arc Weighting using Mutual Information, and (4) Learning Path Generation. Below we describe each step in detail, including key algorithms and code snippets to illustrate the implementation.

Dataset and Skill Extraction

We gathered a dataset of approximately 30,000 job offers, each accompanied by a list of skills required for that position. These job postings were collected from Upwork, a top-tier platform for freelancers and clients worldwide, using a custom web scraper developed in Python by members of the ENTEEF project team.

Ethical principles were consistently applied during the entire data collection process. The dataset was compiled exclusively from publicly accessible information, with all client-identifiable details removed to ensure privacy. Data selection was conducted without the use of filters, relying on Upwork's default sorting to retrieve the most recent and active job postings. The resulting dataset contains job postings featuring project titles, required skills, budget structures, client information, and geographic details related to the job offers.

The job postings were processed through an NLP pipeline to extract skill keywords, using techniques similar to those described in the literature (Rahayu et al., 2023).

We normalized skill names (e.g., handling synonyms and variations) to ensure consistency. By averaging the frequency of skill mentions in job postings across both continental and global levels, a ranking of 60 in-demand skills was generated. We will focus on these 60 most requested skills in the following. The result is a binary jobs × skills matrix M of size 30,000 × 60. An entry $M[i, j] = 1$ indicates that job posting i mentions skill j, and 0 otherwise. Each job posting thus provides a data point linking a combination of skills that employers expect together.

Learning a Bayesian Network Structure

Using the processed dataset, we learn a Bayesian Network structure that captures the probabilistic dependencies between skills. We treat each skill as a binary variable (present/absent in a job posting). Structure learning is performed with a Hill-Climbing Search (HCS) algorithm, a common greedy search method for BN structure discovery (Ankan and Textor, 2024; Dubois et al., 2008). The HCS algorithm starts with an empty network and iteratively adds, removes, or reverses edges to maximize a scoring function (such as the Bayesian Information Criterion or K2 score). In our implementation, we utilized the HillClimbSearch class from the pgmpy library in Python with a BIC scoring metric (Ankan and Textor, 2024).

This search yields a directed acyclic graph (DAG) G where nodes are skills and directed edges suggest a dependency (potentially interpreted as a prerequisite or strong association). For example, the algorithm might learn that an edge goes from skill A (e.g., *"HTML"*) to skill B (*"CSS"*), indicating that HTML knowledge often accompanies or precedes CSS in job requirements - a hint that learning HTML might be a prerequisite to learning CSS.

The learned BN structure encodes which skills tend to appear together in job descriptions and the directionality that best fits the data (note: the direction of an edge in a learned BN does not *strictly* imply pedagogical prerequisite, but we hypothesize it often aligns with a plausible learning order). The structure acts as a skill graph from which we can derive candidate learning paths.

Arc Weighting with Mutual Information

While the BN structure defines the qualitative relationships (which pairs of skills are connected), we next quantify the strength of each connection by computing weights for each arc. We use two different weighting schemes based on mutual information (MI):
(a) the raw mutual information score (RMI) and
(b) the normalized mutual information score (NMI).

The raw mutual information is an information-theoretic measure of the dependency between two variables. In our context, it measures how much knowing one skill's presence in a job posting reduces uncertainty about the presence of another skill. Formally, for two skill variables X and Y, the raw mutual information $RMI(X; Y)$ is defined as:

$$RMI(X;Y) = \sum_{x\in\{0,1\}} \sum_{y\in\{0,1\}} P(X = x, Y = y)\, log \frac{P(X = x, Y = y)}{P(X = x)P(Y = y)}$$

where $P(X = x)$ is the probability of outcome x.

We treat the job dataset as empirical observations to estimate these probabilities for each pair of connected skills. The higher the RMI, the more information one skill gives about the other, meaning they strongly co-occur (either both present or both absent more often than chance). However, RMI is unbounded on the upper end and tends to grow with the entropy of the variables (for instance, very frequent or very rare skills can influence the magnitude of RMI).

To enable comparison across different skill pairs, we also compute the normalized mutual information (NMI), which scales the mutual information to a standardized range $[0,1]$.

The NMI is given by (Sklearn, 2025):

$$NMI(X;Y) = \frac{RMI(X;Y)}{\sqrt{H(X)H(Y)}},$$

where $H(\bullet)$ denotes entropy and is defined as follows:

$$H(X) = -\sum_{x\in\{0,1\}} P(X = x)\ log\, P(X = x)$$

Considering that the logarithm is in base 2, entropy is measured in bits.

This normalization accounts for the overall occurrence rates of skills X and Y, yielding 0 when skills are independent and 1 when knowing one perfectly predicts the other. Intuitively, NMI tells us the *fraction of maximum possible information* that X and Y share, thus providing a comparable "connection strength" on a zero-to-one scale.

We computed both metrics for every directed edge ($A \rightarrow B$) in the learned BN.

Using these metrics, we assign two sets of weights to the network's edges. For example, if skill A and skill B are connected in the BN, we might find $RMI(A; B) = 0.85$ bits and $NMI(A; B) = 0.42$. A higher weight (closer to 1 in NMI, or a larger RMI value) implies a stronger coupling of those skills in job ads.

A pruning threshold was applied (RMI < 0.009; NMI < 0.03) to eliminate weak dependencies from the initial network. The resulting reduced network retains only the stronger connections, for clearer path generation, and the summary statistics of their RMI/NMI scores are presented in Table 1 alongside those of the full network.

Table 1: Summary statistics of RMI/NMI weights in the Bayesian Networks

Statistic	**RMI – Full Network**	**RMI – Pruned Network** (threshold: 0.009)	**NMI – Full Network**	**NMI – Pruned Network** (threshold: 0.03)
Count	210	85	210	96
Mean	0.016	0.035	0.084	0.169
Median	0.006	0.027	0.024	0.136
Std Dev	0.022	0.023	0.112	0.117
Min	0	0.009	0	0.036
Max	0.102	0.102	0.492	0.492

Generating Learning Paths

The final step is to utilize the weighted Bayesian network to suggest learning paths. A learning path in this context is a sequence of skills $[S_1 \rightarrow S_2 \rightarrow \ldots \rightarrow S_k]$ that a learner could follow, where each transition $S_i \rightarrow S_{i+1}$ is an edge in the BN indicating a strong relationship. To generate a path for a given target job or role, we proceed as follows:

- **Identify Target Skills**: From the job role of interest (or job postings of that role), determine the set of key skills required. For example, a "Data Scientist" role might require {Python, Machine Learning, Data Visualization, Statistics}.
- **Subgraph Extraction**: Extract the subgraph of the BN that contains these target skills, and any skills directly or indirectly connected to them (this subgraph represents the domain of relevant competencies).
- **Path Search**: Within this subgraph, find paths that connect a learner's current skills to the target skills. We assume the learner's current skills (e.g., those they already possess or have mastered) are known via the competency assessment test (CAT). Starting from a current skill node, we perform a forward search through the network toward a target skill node. We prioritize moves along edges with higher weights (indicating stronger skill association). This can be implemented as a weighted graph search (e.g., Dijkstra's algorithm if interpreting 1-*weight* as a distance (for NMI method), or simply greedily following the highest-weight edge).

Example: If a learner knows HTML but needs to learn React (a JavaScript framework) for a job, the BN might contain a path HTML → CSS → JavaScript → React. Each arrow is supported by strong mutual information scores (indicating these skills frequently co-occur in jobs). The suggested learning path would then be to start with HTML (already known), then learn CSS, then JavaScript, and finally React, in that order. This path aligns with both job data (many web development postings list those skills together) and pedagogical logic (each skill builds on the previous).

The outcome is a recommended sequence of skills (and by extension, relevant MOOCs teaching those skills) personalized to the learner's goals and gaps. In ENTEEF's framework, once such a path is identified, the platform can present the learner with the specific MOOCs corresponding to each skill in the sequence. The Competency Assessment Tool ensures that the learner skips skills they have already mastered, focusing on the ones they lack, while the Bayesian network-derived structure ensures the order of learning is sensible and supported by real-world demand.

By using both raw mutual information and normalized mutual information for arc weighting, we can experiment with different path generation criteria. The raw mutual info score weighting might favour edges that involve generally highly demanded skills (since those contribute more bits of information), whereas normalized mutual info score might highlight niche but strongly linked skill pairs (by controlling for base frequency).

Model Development and Evaluation

The initial Bayesian Network structure obtained with HillClimbSearch is shown in Figure 1. By filtering out very weak connections, defined as edges with an NMI below a threshold, we obtained the pruned network presented in Figure 2.

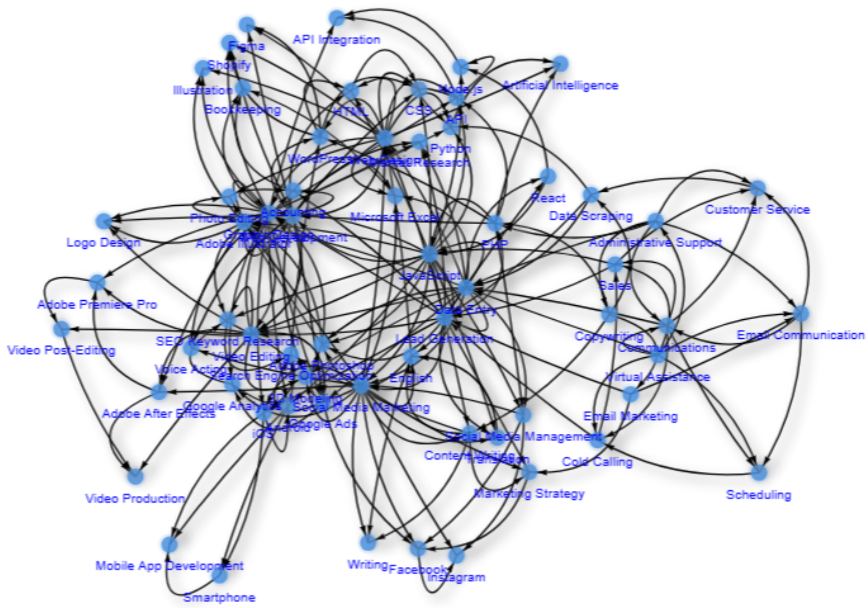

Fig 1. The initial learned Bayesian Network structure

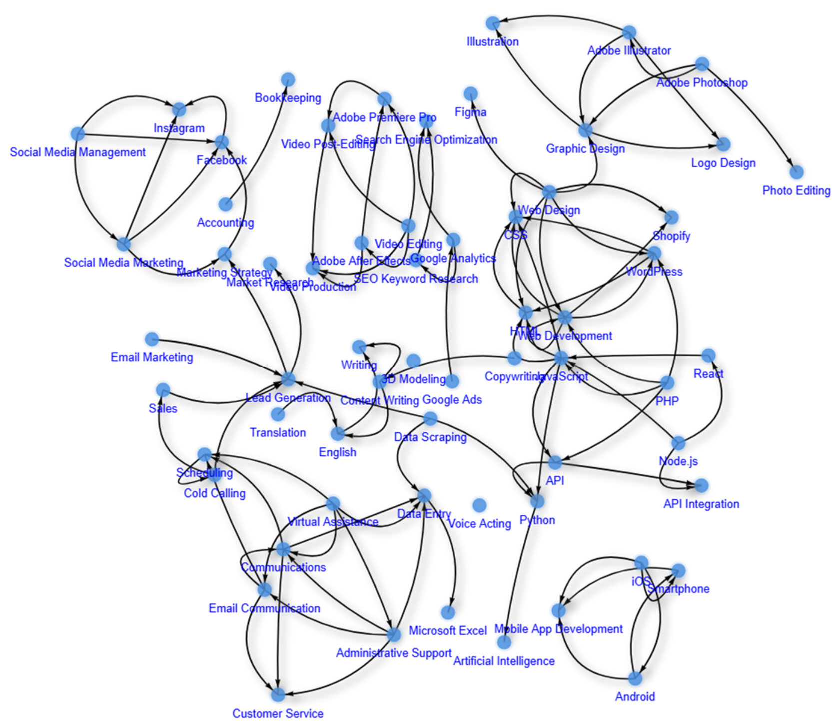

Fig 2. The pruned network, using NMI with threshold 0.03

To compare the models produced by the two approaches (RMI and NMI), we extracted the subgraphs for each skill from the respective models, centering each subgraph around the corresponding skill.

In an empirical evaluation of the models generated for the most requested skills, it is observed that the two Bayesian network weighting methods (RMI and NMI) produce similar results. It can be seen in Figure 3 that for the Graphic Design skill, the subgraphs are identical. However, Figure 4 illustrates that for Web Development, the models are slightly different.

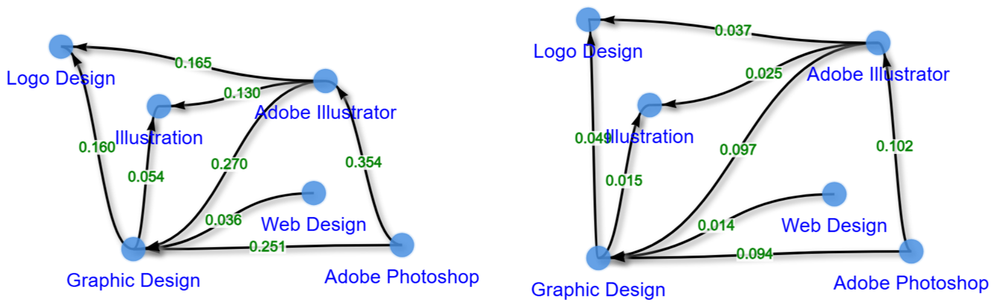

a) Arc weighting with NMI method   b) Arc weighting with RMI method

Fig 3. The identical extracted subgraphs for Graphic Design skill

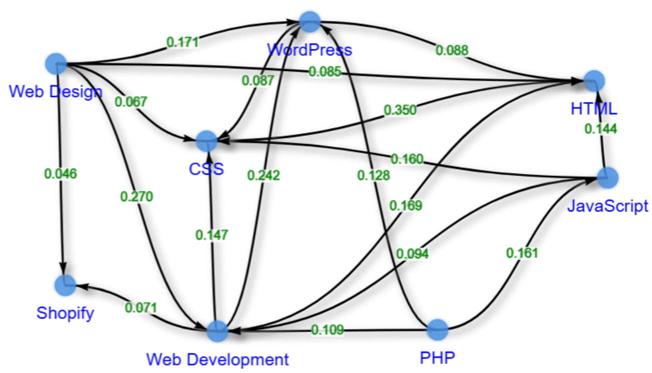

Fig 4. Extracted subgraph of the BN (NMI weights) for Web Development skill

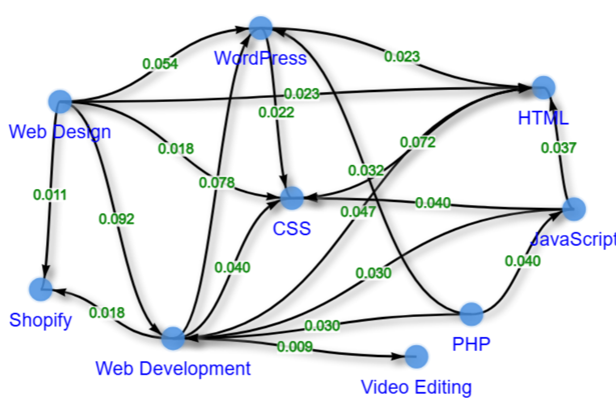

Fig 5. Extracted subgraph of the BN (RMI weights) for Web Development skill

For an automated comparative evaluation, we measured the similarity of all subgraphs using the Jaccard similarity metric focused on the node sets (node perspective).

Given a graph $G = (V, E)$ and two nodes *u* and *v*, the Jaccard similarity of their neighborhoods is defined as follows:

$$J(u,v) = \frac{|N(u) \cap N(v)|}{|N(u) \cup N(v)|}$$

where $N(u)$ is the set of nodes adjacent to u.

In Figure 6, only the skills with a node similarity score (Jaccard similarity) not equal to 1 are presented.

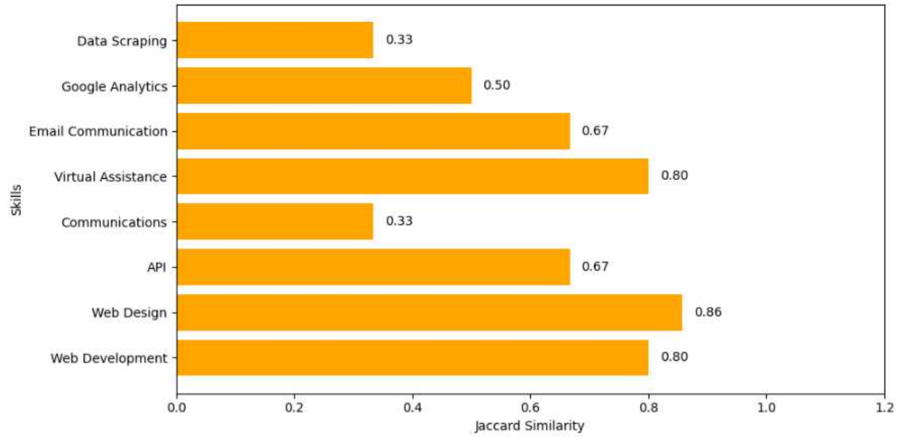

Fig 6. Skills with different Jaccard similarity

Discussion

Based on the ENTEEF project's goals and our data-driven modelling approach, we hypothesize that a Bayesian Network learned from large-scale job postings, when weighted using mutual information metrics, can effectively generate learning paths that reflect real-world skill progression and employer expectations. This hypothesis guides our methodological design and evaluation, aiming to contribute a scalable framework for personalized upskilling grounded in labor market data.

To validate this hypothesis, we provide:

1. *Empirical support from real job data*
   We constructed the Bayesian Network using a large dataset (~30,000 job postings) covering diverse freelancing roles. The structure learned reflects co-occurrence and dependency relationships between skills actually required in the market.
2. *Dual weighting and comparative evaluation*
   We employed two distinct arc-weighting methods (Raw Mutual Information and Normalized Mutual Information) to quantify the strength of skill relationships. We compared the resulting networks using Jaccard similarity and visual subgraph analysis for multiple key skills.
3. *Practical alignment with learning sequences*
   The learning paths generated align with both the BN structure and logical pedagogical progression, supporting the validity of the approach.
4. *Contribution to knowledge*
   Prior work in the learning path generation often relies on expert-defined ontologies or student learning data. Our contribution is a data-driven, scalable method that infers skill pathways from the job market itself, providing a novel bridge between labor analytics and educational technology.

The results of the Bayesian Network modeling, particularly the coherence of skill sequences and the similarity between RMI and NMI weighting methods, support our research hypothesis. Furthermore, the ability to extract clear subgraphs and measure node similarity confirms the structural consistency of our model, validating the hypothesis that data-driven BN structures can serve as effective frameworks for learning path generation.

This validation demonstrates the practical and theoretical soundness of our methodology and its contribution to knowledge in both learning path generation and data-driven curriculum design.

Next, we will present some limitations of our study:

- The learning paths are inferred from job data but not yet validated through actual learner outcomes (e.g., whether following a path leads to better job placement or skill mastery).
- The edges in the Bayesian Network are based on statistical dependencies, not necessarily *true pedagogical prerequisites*.

- The system has not yet been tested in real-world learning scenarios, so user experience and motivation effects are still to be studied.
- The pruning threshold for NMI was set empirically. Different thresholds may lead to different path suggestions, and this may affect consistency.

However, our approach has several important advantages and implications, including:

- Our method enables educational institutions to align learning paths with *real-world job market needs*, especially in fast-changing domains like freelancing and tech.
- Learners can follow targeted learning sequences based on actual labor demand, improving the efficiency and relevance of lifelong learning programs.
- Because the Bayesian Network is learned from job postings, the model can be updated regularly to reflect emerging skills and trends.
- While applied here to freelancing, this methodology can be adapted to any domain with sufficient job data, such as healthcare, cybersecurity, or digital marketing.

In summary, our study demonstrates the feasibility and potential of using Bayesian Networks to derive learning paths directly from labor market data. While the approach offers strong implications for personalized and demand-driven education, its effectiveness depends on the quality of skill extraction and the interpretability of inferred dependencies. These limitations highlight the need for future validation through real learner outcomes and expert review.

Conclusions and Future Work

Our research presents a novel approach to generating learning paths by combining skill extraction from job postings with Bayesian Network modelling. In the ENTEEF project context, this approach automates the linkage between labour market requirements and educational content, paving the way for data-informed upskilling pathways. Early results indicate that mutual information-based weighting provides a meaningful measure of skill relatedness, and the resulting paths align with logical prerequisite structures.

As a direction for future work, we intend to explore the use of the NOTEARS (Non-combinatorial Optimization via Trace Exponential and Augmented lagRangian for Structure learning) algorithm (Zheng et al., 2018) as an alternative to traditional Bayesian Network structure learning. Unlike our current approach, which applies mutual information-based weighting after structure discovery, NOTEARS jointly learns both the structure and the edge weights through continuous optimization. This could offer advantages in terms of scalability, differentiability, and potentially finer-grained representations of skill dependencies. Comparing the learning paths derived from NOTEARS to those produced using NMI/RMI-weighted Bayesian Networks will help assess the relative strengths of these methods in modeling realistic upskilling pathways based on labor market data.

We will compare learning paths obtained under each weighting scheme in future evaluations to see which better guides learners to acquire comprehensive, job-relevant skill sets. Future work will involve validating these generated paths with expert educators and measuring learner outcomes when following the recommended paths, as well as integrating the approach into the ENTEEF platform for real-world use.

Acknowledgements

Co-funded by the European Union (Project no 2024-1-PL01-KA220-HED-000248152). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the Foundation for the Development of the Education System. Neither the European Union nor the entity providing the grant can be held responsible for them. We gratefully acknowledge the contributions of the ENTEEF team members: Bartłomiej Balsamski, Jakub Kanclerz, Mukhammad Andri Setiawan, Ahmad Fathan Hidayatullah and Ratih Dianingtyas Kurnia for their valuable work in the primary data collection and preparation phase of this research. Their efforts in gathering and organizing job posting data provided a critical foundation for the analysis and modelling presented in this study.

References

- Ankan, A. and Textor J. (2024), ‘pgmpy: A Python Toolkit for Bayesian Networks’. Journal of Machine Learning Research, 25(265), 1-8.

- Carroll, D. and Schlippe, T. (2023), ‘Connecting Learning Material and the Demand of the Job Market Using Artificial Intelligence’, Artificial Intelligence in Education Technologies: New Development and Innovative Practices, pp. 282-298.

- Culbertson, M. J. (2016). Bayesian Networks in Educational Assessment: The State of the Field. Applied Psychological Measurement, 40(1), 3–21.

- Dubois, D., Prade, H. and Smets, P. (2008), 'A definition of subjective possibility'. International Journal of Approximate Reasoning, 48(2), 352-364. Available: https://doi.org/10.1016/j.ijar.2007.01.005.

- ENTEEF Project Fostering Entrepreneurship through Freelancing (2025), Erasmus+ Programme Project Website. Available: https://enteef.uek.krakow.pl/.

- Rahayu, N.W., Ferdiana, R. and Kusumawardani, S.S. (2023), 'A systematic review of learning path recommender systems', Education and Information Technologies, 28 (6), 7437-7460. Available: https://doi.org/10.1007/s10639-022-11460-3.

- Shen, H., Liu T. and Zhang Y. (2020). 'Discovery of Learning Path Based on Bayesian Network Association Rule Algorithm', International Journal of Distance Education Technologies, 18(1), 117-130.

- Sklearn Metrics Normalized Mutual Info Score - Scikit-learn 1.7 documentation (2025). Available: https://scikit-learn.org/1.7/modules/generated/sklearn.metrics.normalized_mutual_info_score.html

- Zhang, M., Jensen K. N., Sonniks S. D. and Plank, B. (2022), 'SkillSpan: Hard and Soft Skill Extraction from English Job Postings', Proceedings of NAACL 2022 (Association for Computational Linguistics), 4962-4984.

- Zheng, X., Aragam, B., Ravikumar, P. and Xing, E. P. (2018), 'DAGs with NO TA¬E¬RS: Continuous optimization for structure learning'. Advances in Neural Information Processing Systems (NeurIPS 2018), 12 pages.