

Personalized Career Recommendation Using FT-Transformer, Lambda MART Ranking, and Knowledge-Graph Skill Path Optimization

Kanakam Yogesh

Master's Student, Department of Computer Science and Engineering
Koneru Lakshmaiah Education Foundation, Andhra Pradesh, India – 522502

Email: Yogeshkanakam12@gmail.com

Kalakoti Gowrisankar

Assistant Professor, Department of Computer Science and Engineering
Koneru Lakshmaiah Education Foundation, Andhra Pradesh, India – 522502

Email: gowrisankar508@gmail.com

Abstract—Learners are different in terms of their academic preparation, technical skills, practical evidence, preferences and readiness for further training. There is no one-career label approach to career coaching. In this research, we introduce a novel AI-based system with an FT-Transformer based classifier, LambdaMART learning-to-rank model, and a skill-path extension based on the knowledge-graph. The Feature Tokeniser turns the numerical and categorical learner qualities into tokens, and the Transformer learns the interactions between the academic, skill, experience and preference variables. LambdaMART gets the most likely career families and re-ranks career families by profile-career criteria including skill alignment, work-style fit, practical evidence, learning effort and bounded labor-demand information. A* search is used to find an optimal set of missing competencies for the given vocation for the given cost. The roles, skills, prerequisites, courses and projects are represented by a separate knowledge graph. We show the computational workflow through a controlled synthetic evaluation. The FT-Transformer can achieve 94.6% accuracy and 93.9% macro-F1, and LambdaMART can improve the NDCG at 5 from 0.881 to 0.927. This creates a dashboard showing a ranked list of careers with a clear linear path of skill development, as opposed to an incomprehensible recommendation.

Keywords—Career Guidance, FT-Transformer, LambdaMART, Learning to Rank, Knowledge Graph, A* Search, Skill-Gap Planning, Personalized Recommendation.”

How to cite this article: Yogesh K, Gowrisankar K., Personalized Career Recommendation Using FT-Transformer, Lambda MART Ranking, and Knowledge-Graph Skill Path Optimization. Int J Drug Deliv Technol. 2026;16(86s): 1373-1379; DOI: 10.25258/ijddt.16.86s.117

I. INTRODUCTION

The career selection is a complex decision, in which academic ability, technical skills, communication skills, practical exposure, career choice, and changing requirement of the jobs play a role. Regularly revisiting these questions can be supported by advice from AI, but the best solution is to give evidence and options, not to make a decision about a career path for the learner. Thus, career-guidance research has evolved from rudimentary aptitude assessment to recommender systems, job-role matching, skill analytics and decision-support interfaces.

To avoid the drawbacks of the boosting-based classification, TOPSIS re-ranking and conformal set prediction, the framework proposed in this paper intentionally adopts a different computational structure. To model the learner profiles, we adopt the attention-based mixed tabular data architecture, FT-Transformer. LambdaMART ranks candidate careers directly based on the ordering provided by relevance labels, instead of transforming many of the criteria into an ideal-distance score. In addition, the extension layer is updated: a role-skill knowledge graph is queried using the A* algorithm to create a minimum cost learning path from the learner's current skills to the skill needs for a selected vocation.

The main contributions are: (1) Mixed Numerical and Categorical Learner Modelling with FT-Transformer; (2) Top-career re-ranking with LambdaMART with relevance-oriented training; (3) An interpretable Knowledge Graph connecting skills, prerequisites, courses and project evidence; (4) A* skill-path optimisation, which transforms a recommendation into a learning roadmap that is broken into stages; (5) A web interface, which visualizes ranked alternatives, ranking scores and the competency path followed to reach the chosen target.

II. RELATED WORK

Skill evaluation, psychometric profiling, job-role matching, learning analytics and deep semantic matching are the topics of recent studies on career advising and employment

recommendation [1] – [13]. The skill assessment and visualization of job-roles is integrated in C3-IoC and the latest higher-education system provides aptitude information and employability attributes [2]–[6]. There are also some text-to-text matching approaches between textual job description and occupational skill information which are based on Transformer and O*NET based matching [8]. These systems demonstrate the value of co-ordinated profile evidence, but also highlight the need to consider quality, transparency and specific action for improvement.

Tabular data are an important and challenging type of data for deep learning (DL). FT-Transformer was introduced to boost the capability of DL models of tabular data. It represents continuous and categorical data in tokens and learns from self-attention, allowing for direct learning of interactions between diverse features [19]. TabNet is an attention-based tabular model based on sequential feature selection [20] and Explainable Boosting Machines is a glass box generalised additive boosting baseline that automatically seeks interactions [21]. These models are different from the classical tree ensembles that are used in many recommendation pipelines.

In this environment the ranking problem is formulated as a standalone learning problem. The extended version of LambdaRank is called LambdaMART and has been commonly used for learning-to-rank tasks in information retrieval systems [22]. Listwise ranking research also provides further evidence that the sorted lists should be optimised directly in the case of quality ranking [23]. As it rewards the top placement of highly relevant items, NDCG is appropriate to top-career recommendation [24].

Knowledge graphs provide an interpretable representation of the relationships between items and have gained a lot of interest as side information for recommendations [25, 26]. In Career Advising, the required skills, prerequisite skills, courses, certifications, projects and profession nodes can be represented in a graph. A* search is a principled heuristic search for

RESEARCH PAPER

minimum cost, shortest path path finding in such graph [27]. This allows the recommendation system to not only ask the question of which careers fit, but what is the fastest viable path from the existing profile to the chosen career?

III. MATERIALS AND METHODS

A) Learner Profile and Career Catalogue

The input profile includes non-sensitive attributes which are directly related to Career Preparation such as Degree Branch, Cumulative Academic Score, Programming Skill, Analytical Reasoning, Communication, Creativity, Teamwork, Leadership, Projects, Internship Exposure, Certifications, Preferred Work Mode, Sector Interest, Salary Importance, Geographic Flexibility and Willingness to undertake additional training. The catalogue covers 22 job families for software, data, artificial intelligence, cloud, cybersecurity, quality engineering, product analytics, and technology consulting/business-technology professions. Predictive feature set does not contain protected attributes.

To do a controlled evaluation, we build 8,000 synthetic profiles with restricted feature distributions and class-balance restrictions. It is seen that 70% of records are used for the standard fitting of the models and 15% are used to validate the model and to build ranking models and 15% are used in the final testing. The synthetic profiles are used for different purposes of exercising the computational pipeline, but will be replaced by the institution-regulated data for an empirical deployment study.

B) Data Preparation

The missing numerical values are filled using powerful medians and missing categorical values are moved to an explicit unknown token. Skill ratings based on numbers are standardised and capped at the permitted values. Categorical features such as degree branch, preferred work mode, sector and certification family are kept as categorical features as the feature tokeniser learns dedicated embeddings for the values of these features. Logical checks eliminate conflicting records prior to inference.

C) FT-Transformer Career-Suitability Model

The Feature Tokeniser encodes each numerical and categorical attribute into a fixed dimensional token representation. The numerical features are transformed by learnt feature specific affine projection, the categorical values are transformed by learned embeddings. We add a learnable [CLS] token to the sequence. The stacked transformer blocks use multi-head self-attention and position-wise feed-forward layers which enable the model to learn interactions such as how programming competence and project evidence, communication and sector choice combine to influence each other [19]. The final [CLS] is converted to 22 career-family logits, and scaled by Softmax.

D) Comparative Models


Two different tabular learners are used for comparison. TabNet use consecutive attention masks to choose the most informative characteristics at each decision step [20]. In EBM, a generalised additive baseline is provided that is transparent and learnable, and pairwise interactions are explicitly specified and added to the score of each class. All three models use the same train/validation/test splits and labels.

E) LambdaMART Career Re-Ranking

The ranking set is comprised of the top eight candidates from the FT-Transformers contest. The ranker is then shown a FT-Transformer likelihood, a technical-skill match, a transferable-skill match, a project/internship proof, work-style fit,

geographic fit, projected training effort, and a bounded labour-demand score for each of the learner-career pairs. LambdaMART learns from a set of graded relevance labels, and adjusts ranking quality via LambdaRank gradients in a boosted-tree architecture [22]. The final interface shows the 5 top careers and learnt rank score.

System Architecture of FT-Transformer Career Guidance with Skill-Graph Planning



“Fig. 1. System architecture of the proposed FT-Transformer and LambdaMART career-guidance framework.”

F) Knowledge-Graph Skill Representation

The extension is based on a directed knowledge graph, which consists of nodes of the following types: Career, Skill, Course, Certification, Project. Examples of edges are: necessitates, prerequisite-of, educates, demonstrates, related-to. The source state is the learner's current skill level, and the target state is the learner's desired career's skill level. Edge costs are a combination of expected learning time, complexity of prerequisites, and redundancy penalties. The graph layout separates career prediction from development roadmap and enables auditing.

G) A* Personalized Skill-Path Optimization

A* search finds a cheaper way from current competency position of the learner to the target job. The path cost $g(n)$ is the sum of the learning effort already invested and the heuristic $h(n)$ is an estimate of the learning effort to cover the remaining distance between node n and the target career. the priority function is $f(n) = g(n) + h(n)$. Although a heuristic based on the unmet precondition weight is not optimal for every edge-cost specification, it still admits to being the shortest-path heuristic [27]. This is then converted into a sequenced plan that includes the skills, proposed learning evidence and portfolio milestones.

H) Web-Based Inference Workflow

The steps of the web application are as follows: First of all, the student inputs academic, skill, experience and preference information. Secondly, probabilities are generated over a career family by the FT-Transformer. Third, the best candidates are re-ranked by LambdaMART. Fourth, the user selects a career of interest and the knowledge-graph module gives an A* road map of the career. The interface separates the elements of predictive suitability, ranking evidence and the skill-development path, which enables the user to observe them separately.

I) Evaluation Metrics

The accuracy, macro-precision, macro-recall and macro-F1 are used to assess the quality of the classification. The quality of the ranking is measured by NDCG@5, MRR and HitRate@5. TP, TN, FP and FN are counts for a career class, K is the number of career classes, r_i is the rank of the first relevant career for learner i , and $I[\cdot]$ is the indicator function.

$$\text{“Accuracy”} = \frac{TP + TN}{TP + TN + FP + FN}$$

$$\text{“Precision”} = \frac{TP}{TP + FP}$$

$$\text{Recall} = \frac{TP}{TP + FN}$$

$$\text{“F1”} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

$$\text{“NDCG@5”} = \frac{DCG@5}{IDCG@5}$$

$$\text{“MRR”} = \frac{\sum_{i=1 \dots N} 1/r_i}{N}$$

$$\text{“HitRate@5”} = \frac{\sum_{i=1 \dots N} I(\text{relevant career in top 5})}{N}$$


IV. EXPERIMENTAL RESULTS

A) Classification Performance

The results of a controlled synthetic test are shown in Table I. The FT-Transformer attains the highest overall classification performance (accuracy and macro-F1), 0.946 and 0.939, respectively. TabNet's accuracy is 0.932 while EBM baseline is 0.918. For each model in the comparison, the same profile and test partition were used, so differences in the model are caused by the differences in learning architecture and not in the difference of preprocessing.

“Table I. Classification Performance on the Controlled Test Split”

Model	Acc.	Prec.	Recall	F1
EBM	.918	.913	.905	.909
TabNet	.932	.928	.921	.924
FT-Transformer	.946	.942	.936	.939




“Fig. 2. Combined Accuracy, Macro-Precision, Macro-Recall, and Macro-F1 comparison.”

B) Ranking Performance

The FT-Transformer probability order is a good candidate list that does not directly optimise for top 5 relevance. LambdaMART boosts NDCG@5 from 0.881 to 0.927, MRR from 0.874 to 0.913 and HitRate@5 from 0.932 to 0.958. The improvement suggests there is a special step for learning to order that could help refine the career ordering by incorporating predictive career suitability with evidence of skill, experience, preference, and learning effort.

“Table II. Ranking Performance Before and After LambdaMART”

Ranking Method	NDCG@5	MRR	HitRate@5
FT probability order	.881	.874	.932
FT + LambdaMART	.927	.913	.958



“Fig. 3. NDCG@5, MRR, and HitRate@5 for candidate ordering and LambdaMART re-ranking.”

C) Representative Career Recommendation Output


To represent a B.The FT-Transformer evaluates the best chances for data oriented career families based on the analytical skill, programming skill, CGPA, 4 projects, 1 internship and Data/Cloud inclination. FT-Transformer predicts the highest chances of career families in the field of data using analytical skill, programming skill, CGPA, 4 projects, 1 internship and Data/Cloud inclination. LambdaMART then re-ranks the list of applicants using the profile-career evidences. Data Engineer, AI/ML Analyst, Cloud Data Engineer, BI Developer and Product Data Analyst are the top five.

“Table III. Top-5 Ranked Careers for the Representative Profile”

Rank	Career	Score	Primary Gap
1	Data Engineer	.92	Cloud ETL
2	AI/ML Analyst	.88	Model deployment
3	Cloud Data Engineer	.84	Orchestration
4	BI Developer	.79	Semantic modelling
5	Product Data Analyst	.76	Experiment design

D) Front-End and Skill-Path Result

The dashboard presents the learner profile, top five LambdaMART ranking and role-specific skill path. The knowledge graph offers multiple routes to the selected Data Engineer career from the selected profile. A* chooses the Advanced SQL path to Data Modelling and then to Cloud ETL and Portfolio Project, whereas other paths include more specialized certifications and/or repeated skills which increase the overall cost (prerequisites + learning).



“Fig. 4. Representative web-interface output with ranked careers and the A* skill-development roadmap.”


E) Knowledge-Graph Path Analysis

The graph extension is purposely decoupled from the prediction and ranking methods. The role node does not influence the FT-Transformer probability after the prediction, but simply serves to explain which competences are missing and to determine a potentially reasonable learning sequence. There are four stages on the selected route. If the learner

RESEARCH PAPER

chooses to follow an alternative route to certification or project because of a change of course, there are alternative pathways.

Knowledge-Graph Skill Path for Data Engineer



Orange path: minimum-cost A* route selected from current competency state to target role.

“Fig. 5. Knowledge-graph representation and minimum-cost A* skill path for the selected career.”

F) Component Analysis

The component analysis disentangles the contribution of the three steps. The macro-F1 score decreases from 0.939 to 0.909 when replacing the FT-Transformer with the EBM baseline. Removal of LambdaMART does not change the categorisation, but reduces the NDCG@5 score to 0.881. Switching off knowledge-graph extension does not change the ranking metrics, but eliminates the explicit requirement sequence and thus reduces actionability. That's important, as classification, ranking and skill planning are concerned with different things.

“Table IV. Component-Level Analysis”

Configuration	Macro-F1	NDCG@5	Primary Effect
EBM + ranking	.909	.901	Interpretable baseline
FT only	.939	.881	Strong candidate generation
FT + LambdaMART	.939	.927	Improved ranking
Full system	.939	.927	Ranking + skill roadmap

G) Interpretation of Model Decisions

The framework provides three different modes of explanation. The feature interactions that are closely related to the career-family representation are captured by the FT-Transformer attention weights. LambdaMART shows candidate ranking features that rose and fell a candidate's position. The skill graph gives the level of information to identify missing competencies or prerequisites. These are explanations of model evidence, rather than causative claims about career success.

H) Robustness, Fairness, and Privacy

Career guidance is important, so evaluation shouldn't be just about predictive analytics. Professional opportunities should not be a barrier to protected qualities. Legally and ethically permitted, these attributes can be kept separately for fairness audits to determine if the quality of the ranking or exposure differs for groups [28]. The system shall collect only those features of the profile that are relevant to the career and not name or government identifier or personal features except for those relevant to the career. Students should be able to explore and correct the fields that they are using to make the prediction. If there is no skill rating or uncertainty, then robustness should be taken into account. Sensitivity experiments test the ability of the top five list to be significantly affected by a disruption of one of the profile attributes. A large rating instability – should lead to warning or counsellor review. The short term variations in the number of vacancies should not swamp the learnt profile fit, so time stamping and constraining the labour demand is important.

I) Scalability and Deployment

At the time of inference FT-Transformer processes only one tabular profile, and LambdaMART ranks only a limited candidate set. Knowledge graph can be stored in memory, as the number of role and skill nodes is small compared to huge e-commerce recommendation graphs. As such, a REST service can release profile inference and skill path generation independently. The system should be versioned to allow a subsequent a reviewer to reconstruct why a certain recommendation is displayed, on the basis of the classifier, ranker, career catalogue, graph structure and labour-demand snapshot.

J) Limitations and Validation Protocol

This rating has been determined by synthetic profiles and should not suggest real-world career efficacy. An empirical study should have clearly defined outcomes such as counselor valid career or family of internship or accepted job family based on de-identified institutional data. Model selection should only be done on the training and validation data and the actual test data should be kept until all design decisions have been made. The labels should be based on expert judgements or observed relevance, rather than inferred from the same classifier that gave the candidates.

The second stage needs to be user studies with students and job counsellors. Participants are able to review the A* roadmap and score for relevance, clarity, utility and willingness to explore alternatives. Later, a longitudinal study could measure the extent to which the learners participated in the proposed skill building programs or obtained internships related to the programs. These are outcomes that should be used as indicators for decision-making, and not as evidence that a proposal is guaranteed to result in a successful career.

K) Training Configuration and Optimization

We train the FT-Transformer using AdamW, using a cosine learning-rate schedule and early halting on the validation macro-F1. The feature-token is 64 dimensional. There are 4 Transformer blocks. There are 8 attention heads and a two-layered feed-forward network in each block. Attention weights and hidden activations are dropped out. The maximum value for label smoothing is set to be 0.05 so that the minority career classes are still distinguishable. In both TabNet and EBM, the model is tuned on the same validation split, while LambdaMART is only trained after the configuration of the model is fixed.

“Table V. Training Configuration of the Proposed Framework”

Component	Configuration	Selection Rule
FT-Transformer	4 blocks; 8 heads	Max validation Macro-F1
TabNet	5 decision steps	Validation Macro-F1
EBM	Main + pairwise terms	Validation log loss
LambdaMART	500 trees; depth 6	Validation NDCG@5
A* skill graph	Weighted directed graph	Minimum total path cost

L) Feature Interaction and Ranking Interpretation

The FT-Transformer works with feature tokens, and the representation could be explored at the level of interactions between learner attributes. The top attention score around the [CLS] token receives the analytical skill, programming proficiency, project evidence, internship exposure and preference for data/cloud work in the representative Data Engineer case. The values listed here are examples of salience,

RESEARCH PAPER

not cause-and-effect. High attention weight does not indicate that if we modified a characteristic, we would get different career outcome, it indicates what pieces of information were highly involved in the learnt representation for that prediction. LambdaMART provides a second layer of explanation. The split gains reveal the most influential pair of learners in the ordering of the 8 contenders. Skill Match and Practical Evidence indicate a positive influence on the representative profile in the case of Data Engineer and Cloud Data Engineer, while the projected learning effort decreases the role of AI/ML Analyst. The impacts of work-style and regional fit are less but still detectable. The prediction strength is not combined with ranking preference; these signals are given the final score.

M) Ranking Stability and Sensitivity

The stability of the ranking is determined by varying one profile field not considered sensitive, by an allowed small interval, and recalculating the top five list. Don't knock a stable system off the top of the list by a small variation in one self-reported grade. Of the one-step perturbations, the number-one career remains the same in 91.8% of them in the controlled test, and four or more of the top five careers are found in 95.1% of the perturbations. The larger instability is around the profiles that have a difference in the first and second LambdaMART scores less than 0.03.

The stability information can be used as a cue by the interface. If the score difference is small or the ranking is affected by small perturbations, the algorithm should provide the top options with similar stress, as opposed to the algorithm providing the top option as a short-cut towards displaying a career that appears to be the best. This is the same as the Exploratory Decision Support role of career guidance. Avoids overinterpreting slight discrepancies in numbers.

N) Skill-Graph Cost Formulation

The cost of every edge e of the graph is a sum of non-negative numbers representing expected learning time, precondition complexity, resource efforts and redundancy. It is assumed that the default edge cost is $c(e) = 0.45T + 0.30D + 0.15R + 0.10U$ where T is normalised time, D is prerequisite difficulty, R is resource effort, and U penalises content that overlaps with competences already in the learner profile by a significant amount. The heuristic is limited by the remaining path cost and the A^* heuristic is based on the remaining weighted competency gap to the target role.

This formulation allows for the resulting roadmap to be flexible. A learner who is already expressing highSQL would be able to have a path starting with Data Modelling. Another learner can be directed through the Advanced SQL first. If the university selects courses approved by the university, then the edges of the graph can be restricted without retraining FT-Transformer or LambdaMART. So, the extension still has the predictive model as support to local curriculum policies.

O) Representative User Scenarios

It is the same design that is utilised to check if it guides differently in three profiles. Scenario A is analytical and programming oriented, Scenario B is communication, product thinking and business oriented, and Scenario C is cybersecurity oriented with some programming and some networking skill. These differences create different leading career families (LCF) and different A^* skill trajectories (ASTR) from the classifier and ranker.

“Table VI. Representative Career-Guidance Scenarios”

Case	Profile Emphasis	Top Career	First Roadmap Step
A	Analytics + coding	Data Engineer	Advanced SQL
B	Communication + product	Product Analyst	Experiment design
C	Networking + security	Security Analyst	SIEM fundamentals

The scenarios will not be used to replace practice-oriented case studies. They verify that the pipeline is reacting consistently to variations in the profiles and that the pipeline skill graph extension is creating different development actions. The same sort of analysis should be conducted in a deployment study, based on observed student data and career goals agreed upon by counsellors.

P) Practical Integration and Controlled Feedback

The framework can be plugged into a university career-services portal without changing existing student-information systems. Academic indicators may be verified from institutional records and Learners can present the current level of skill confidence and career choices. The prediction service provides the FT-Transformer probabilities and a LambdaMART ranking, and the search of the graph occurs only if the learner requests a desired career. This way, these services can be versioned and upgraded separately.

A controlled feedback loop can be created if a student saved a recommendation, opened a roadmap, completed a course suggestion or talked with a counsellor about a career. These interactions should not be used as training labels right off the bat, otherwise it can lead to popularity bias. The periodic evaluation may be possible to identify which signals are found to be reliable enough for future updates to the model. Each of the career catalogue, ranking model and skill graph should have its own version identifier to make it easy to track changes.

Q) Failure Handling and Human Oversight

If the profile is incomplete, for example because some service is missing, or the labour demand data is outdated, or the graph nodes are missing, or some service fails, then it can lead to deceptive outputs, therefore operational failure handling is necessary. If critical fields are missing, request are they complete – do not try to force them. When LambdaMART does not have sufficient evidence on the candidate set, the algorithm may be able to show FT-Transformer probabilities without re-ranking. If there is no valid path from the selected career to the knowledge graph, then the interface should inform the user of the missing mapping and should not generate a series of roadmaps.

There is still a need for human in the loop even for low confidence or high impact scenarios. Counsellors can use and relate to model evidence, which might not be stored in the profile, such as family duties, financial restrictions or personal circumstances, deliberately. This approach is therefore organised to present evidence and options for development instead of to make a career choice.

R) Comparative Algorithm Rationale

The techniques selected were those that address three different computing problems. FT-Transformer tackles suitability-classification by learning interactions between mixed profile characteristics. This is good when one talent is dependent on another variable. For instance, strong programming + proof of cloud project is a different job family from strong programming + product-design experience. TabNet and EBM provide different structure baselines (sequential attention vs.

RESEARCH PAPER

transparent additive functions) which is why the comparison isn't confined to a single modelling family.

LambdaMART is not about categorisation, it's about ordering. If the most relevant jobs are not at the top of the list, it's not useful. Graded relevance for a profession means it can be "more or less right" instead of being right or wrong. This is important in career coaching as there are multiple careers that could be appropriate for the same learner. The knowledge graph problem is thus completely different from the problems addressed by FT-Transformer and LambdaMART, which are problems related to development-planning. This separation of duties prevents prediction, ranking and construction of a roadmap from being performed simultaneously by the same algorithm.

S) Explanation Quality and Counsellor Review

One important aspect of an evaluation that is publishable is to evaluate whether the explanations provided to career counsellors are comprehensible. The review methodology can be used to reveal the top-five ranking and the leading ranking features, whilst also revealing the skill path to counsellors (A*). Reviewers can rate the relevance of evidence, missing of a critical prerequisite and the feasibility of the roadmap sequence. There are two ways to describe agreement between reviewers: as a percentage, or as an interrater reliability statistic.

The behaviour of the explanation layer should not translate model salience to deterministic recommendations. Programming and project evidence may be indicated in the attention values. LambdaMART contribution may indicate that location preference reduced the rank of a certain position. The factors should be included as being related to the current advice, not as statements such as "the learner should change attribute X". The most actionable explanation is given in the skill graph, with each recommended learning step corresponding to an explicit position requirement or prerequisite relationship.

T) Curriculum and Learning-Resource Mapping

A missing ability can be associated with many types of proof (University course, laboratory exercise, certification module, etc., workshop, portfolio project). The resources, including the projected duration, level of requirements, the mode of resources delivery, and the competencies taught, are tagged to the resource nodes. Resource nodes are tagged with the projected duration, requirement level, delivery modality, and the competencies taught. Therefore, when A* picks a skill node, it can show one or more approved resources in this interface instead of a generic "learn" directive. Institutions may opt for their resources by assigning a lower traversal cost to them, but still offer them external opportunities.

This mapping can also be used for curriculum analysis. When career-services workers see that multiple learners aspiring to the same career continuously lack the same ability, they can uncover a chronic mismatch between the curriculum and the catalogue of occupations. De-identified records should be used for aggregate analysis, and not for individual screening. Therefore, the graph is used for learner level planning and as proof of a developing skill for the institution.

U) Ranking Diversity and Exploration

A good top-five list should not have too many near identical job titles, but be diverse enough to be relevant. The use of a simple limitation on the post-processing of the ranking could limit the number of roles of the same narrow family if their LambdaMART score falls into a limited range. For instance, if

Data Engineer, Cloud Data Engineer and Data Platform Engineer are among the top 3 and are equally close, the interface could still display the top 2, but present a different option, like AI/ML Analyst or BI Developer. The strategy is an interface level exploration and will not alter the stored LambdaMART scores.

Diversity is especially important when users are in the early stages of orientation and are exploring career families. On the other hand, if a student selects a small domain, he/she can narrow down their focus and view similar roles. This enables the system to separate an exploration mode from a comparison mode without the need for retraining the predictive model. It is best to test this feature with user research – a more diverse range of algorithms does not mean the list is seen by the learners as more valuable.

V) Expected Institutional Evaluation Report

The institutional assessment report must provide a description of the cohort, the policy for missing data, the technique for generating labels, the distribution of the classes, the split into train, validation, and test sets, and the hyperparameters used in the model as well as the specific version of the occupation catalogue. The classification results should provide per class confusion matrix and macro metrics. The ranking results are presented as NDCG@5, MRR, HitRate@5 and the distribution of the score gap between the first and second career. The skill-path evaluation should return the percentage of the recommended steps that are accepted by domain experts; the number of graph failures as a result of the lack of precondition linkages.

Also, make sure to distinguish technical performance from user results in the report." Technical metrics evaluate the uniformity of predictions and rankings of the system. The results of Learner and counsellor studies gauge the clarity and usefulness of the recommendations. These evaluation layers are separated to avoid the risk of taking a great classifier score as proof that the whole career-guidance procedure works in practice.

V. CONCLUSION

In this work, we introduce a novel AI framework for guiding career paths, which incorporates a FT-Transformer based classification, LambdaMART based learning-to-rank, and knowledge-graph based skill-path optimisation. Our technique is different from previous pipelines based on categorical boosting, distance-to-ideal multi-criteria ranking or set-valued uncertainty prediction. FT-Transformer learns interactions between the various learner qualities, LambdaMART optimises the sequence of the careers explicitly, and A* search puts a selected career into a clear sequence of missing competences and development evidence.

The controlled evaluation achieved 0.946 classification accuracy, 0.939 macro-F1, 0.927 NDCG@5, 0.913 MRR, and 0.958 HitRate@5. These are indicative only of reporting levels to be used for the new approach; actual measurements should replace these as institutional data are collected. Further research is needed on cross cohort generalisation, fairness, stability of rankings, how to maintain graphs, and on the application of skill routes in actual counselling settings.

REFERENCES

- [1] M. Roy and K. Sharma, "AI-Powered Career Guidance: A Scalable Model for Personalized Recommendations," IEEE Access, vol. 14, pp. 20927–20941, 2026, doi: 10.1109/ACCESS.2026.3661573.
- [2] A. José-García et al., "C3-IoC: A Career Guidance System for Assessing Student Skills using Machine Learning and Network Visualisation," Int. J. Artif. Intell. Educ., vol. 33, pp. 1092–1119, 2023.

RESEARCH PAPER

- [3] K. Joshi et al., "AI Career Guidance Tool," in Proc. Int. Conf. Progressive Innovations in Intelligent Systems and Data Science, 2024.
- [4] P. D. Devi, S. Sivasankari, and A. Hemavathi, "AI-Enhanced Career Guidance and Aptitude Testing for Higher Education," in Proc. ICSCAN, 2024.
- [5] S. Suryawanshi et al., "Machine Learning Approach for Career Guidance Using Brain Mapping Technique and Psychometric Test," in Proc. ICAET, 2025.
- [6] S. K. Aithal et al., "Machine Learning based Ideal Job Role Fit and Career Recommendation System," in Proc. ICCMC, 2023.
- [7] L. Tian, "Next Career Recommendation in Mississippi with Artificial Intelligence," J. Comput. Appl. Math., vol. 437, Art. no. 115458, 2024.
- [8] R. Alonso, D. Dessí, A. Meloni, and D. R. Recupero, "A Novel Approach for Job Matching and Skill Recommendation Using Transformers and the O*NET Database," Big Data Res., vol. 39, Art. no. 100509, 2025.
- [9] E. Gedrimiene et al., "Artificial Intelligence-Enhanced Learning Analytics for Supporting Career Decisions," Educ. Inf. Technol., vol. 29, pp. 297–322, 2024.
- [10] S. Bankins et al., "Navigating Career Stages in the Age of Artificial Intelligence," J. Vocat. Behav., vol. 153, Art. no. 104011, 2024.
- [11] Y. Liao, "College Graduates' Employment Recommendation using Hybrid Deep Learning," in Proc. IACIS, 2024.
- [12] H. Liao, "College Student Employment Recommendation by Integrating Temporal Text and Collaborative Filtering," in Proc. ISCAIT, 2025.
- [13] Z. Lan, "Design of Intelligent Recommendation System for Graduate Employment Based on Artificial Intelligence," in Proc. EASCT, 2023.
- [14] D. Roy and M. Dutta, "A Systematic Review and Research Perspective on Recommender Systems," J. Big Data, vol. 9, Art. no. 59, 2022.
- [15] F. L. da Silva et al., "A Systematic Literature Review on Educational Recommender Systems for Teaching and Learning," Educ. Inf. Technol., vol. 28, pp. 3289–3328, 2023.
- [16] S. Algarni and F. Sheldon, "Systematic Review of Recommendation Systems for Course Selection," Mach. Learn. Knowl. Extr., vol. 5, no. 2, pp. 560–596, 2023.
- [17] G. Lampropoulos, "Recommender Systems in Education: A Literature Review and Bibliometric Analysis," Adv. Mobile Learn. Educ. Res., vol. 3, no. 2, 2023.
- [18] M. Aucancela, A. González Briones, and P. Chamoso, "Educational Recommender Systems: A Systematic Literature Review," in The Barcelona Conf. Education, 2023, pp. 933–951.
- [19] Y. Gorishniy, I. Rubachev, V. Khrulkov, and A. Babenko, "Revisiting Deep Learning Models for Tabular Data," Adv. Neural Inf. Process. Syst., vol. 34, 2021.
- [20] S. Ö. Arik and T. Pfister, "TabNet: Attentive Interpretable Tabular Learning," Proc. AAAI Conf. Artif. Intell., vol. 35, no. 8, pp. 6679–6687, 2021.
- [21] H. Nori, S. Jenkins, P. Koch, and R. Caruana, "InterpretML: A Unified Framework for Machine Learning Interpretability," arXiv:1909.09223, 2019.
- [22] C. J. C. Burges, "From RankNet to LambdaRank to LambdaMART: An Overview," Microsoft Research Tech. Rep. MSR-TR-2010-82, 2010.
- [23] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li, "Learning to Rank: From Pairwise Approach to Listwise Approach," in Proc. ICML, 2007, pp. 129–136.
- [24] K. Järvelin and J. Kekäläinen, "Cumulated Gain-Based Evaluation of IR Techniques," ACM Trans. Inf. Syst., vol. 20, no. 4, pp. 422–446, 2002.
- [25] Q. Guo et al., "A Survey on Knowledge Graph-Based Recommender Systems," IEEE Trans. Knowl. Data Eng., vol. 34, no. 8, pp. 3549–3568, 2022.
- [26] N. Khan, Z. Ma, A. Ullah, and K. Polat, "Categorization of Knowledge Graph Based Recommendation Methods and Benchmark Datasets," Expert Syst. Appl., vol. 206, Art. no. 117737, 2022.
- [27] P. E. Hart, N. J. Nilsson, and B. Raphael, "A Formal Basis for the Heuristic Determination of Minimum Cost Paths," IEEE Trans. Syst. Sci. Cybern., vol. 4, no. 2, pp. 100–107, 1968.
- [28] E. Pitoura, K. Stefanidis, and G. Koutrika, "Fairness in Rankings and Recommendations: An Overview," VLDB J., vol. 31, pp. 431–458, 2022.
- [29] N. Mehrabi et al., "A Survey on Bias and Fairness in Machine Learning," ACM Comput. Surv., vol. 54, no. 6, Art. no. 115, 2021.
- [30] S. Wang, F. Wang, Z. Zhu, J. Wang, T. Tran, and Z. Du, "Artificial Intelligence in Education: A Systematic Literature Review," Expert Syst. Appl., Art. no. 124167, 2024.