Course Recommendation System for Company Job Placement Using Collaborative Filtering and Hybrid Model

Jaeheon Park[1][a], Suan Lee[2][b], Woncheol Lee[3][c] and Jinho Kim[4][d]
[1]Kangwon Joint Program of Software Convergence Course, Kangwon National University, Chuncheon, South Korea
[2]School of Computer Science, Semyung University, Jecheon, South Korea
[3]SeedsSoft Co., Ltd., Chuncheon, South Korea
[4]Dept. of Computer Science and Engineering, Kangwon National University, Chuncheon, South Korea

Keywords: Course Recommendation, Collaborative Filtering, Hybrid Model, Deep Learning.

Abstract: This study introduces a novel recommendation system aimed at enhancing university career counseling by adapting it to more accurately align with students' interests and career trajectories. Recognizing the challenges students face in selecting courses that complement their career goals, our research explores the efficacy of employing both collaborative filtering and a hybrid model approach in the development of this system. Uniquely, this system utilizes a company-course recommendation method, diverging from the traditional student-course paradigm, to generalize company-course relationships, thereby enhancing the system's recommendation precision. Through meticulous feature engineering, we improved the performance of the NeuMF model. Our experiments demonstrate that the proposed method outperforms other models by 10% to 79% based on the mAP metric, suggesting that the proposed model can effectively recommend courses for employment.

1 INTRODUCTION

University education plays a pivotal role in equipping students with major-related knowledge and guiding their career paths. Within the South Korean educational framework, from elementary through high school, students are mandated to enroll in school-prescribed courses. However, at the university level, while certain credits are mandatory, students have the autonomy to select their liberal arts and major-specific courses. This transition introduces complexity, as students must choose courses each semester that align with their interests and prospective career trajectories. The challenge is compounded by the vast array of available courses, making informed decision-making daunting, especially for courses previously unattended. This dilemma is particularly pronounced during open enrollment periods (Chung et al., 2015), where students curate their curriculum within set requirements, traditionally relying on advisors and senior students for guidance.

The reliance on experience-based decision-making through direct human interaction, a norm prior to the pandemic (Lee, 2020), has diminished as students increasingly turn to the internet and social media for information. This shift underscores the need for a revamped approach in university education guidance. Addressing this need requires universities to develop structured support systems (Baek et al., 2021) to alleviate the confusion and challenges faced by students in course selection.

This paper proposes the development of a course recommendation system aimed at assisting contemporary university students, who are adept at using the internet yet face difficulties accessing course information through traditional means. Our study compares and analyzes the efficacy of Neural Matrix Factorization (NeuMF) (He et al., 2017), a collaborative filtering technique (Schafer et al.,

[a] https://orcid.org/0009-0006-7984-5373
[b] https://orcid.org/0000-0002-3047-1167
[c] https://orcid.org/0009-0002-1029-9600
[d] https://orcid.org/0000-0003-1125-3938

Park, J., Lee, S., Lee, W. and Kim, J.
Course Recommendation System for Company Job Placement Using Collaborative Filtering and Hybrid Model.
DOI: 10.5220/0012792200003756

In *Proceedings of the 13th International Conference on Data Science, Technology and Applications (DATA 2024)*, pages 139-149
ISBN: 978-989-758-707-8; ISSN: 2184-285X

2007), and the Hybrid Model (Huang et al., 2019), a click-through rate (CTR) method (Yang et al., 2022), in implementing this recommendation system. We also explore feature engineering (Turner et al., 1999) to enhance the accuracy of the NeuMF model. Furthermore, our approach innovates by recommending courses not merely on the student-course axis but also in terms of aligning courses with potential career paths in specific companies, diverging from traditional recommendation practices. This system aims to broaden students' career perspectives, enabling them to make informed course selections that reflect their career ambitions and interests, thereby elevating the quality of university education and facilitating better job placement outcomes.

The contributions of this study are summarized as follows.

- **Introduction of a Company-Centric Recommendation Framework**: Shifted the focus from traditional student-centric models to a company-centric approach, aligning educational outcomes with the specific needs and preferences of potential employers.
- **Development and Enhancement of the NeuMF Model:** Employed and enhanced Neural Matrix Factorization (NeuMF) through feature engineering to capture non-linear user-item interactions and improve recommendation accuracy.
- **Comparison with Hybrid Models:** Conducted a comparative analysis between NeuMF and various hybrid models (DeepFM, xDeepFM, DCNv2, AFN+, and EulerNet) to evaluate their efficacy in course recommendation.

2 RECOMMENDER SYSTEM FRAMEWORK

This section explores the methodology employed in assembling the necessary datasets, integrating these datasets into a unified database, and subsequently training our recommendation model.

2.1 Data Collection & Integration

In the development of our course recommendation system aimed at enhancing employment prospects for students, we undertook a comprehensive data collection and processing approach, as illustrated in Figure 1.

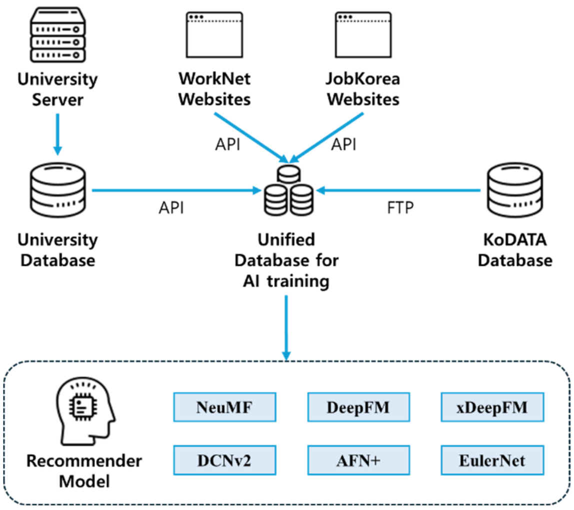

Figure 1: Overview of the proposed recommendation system framework.

Course History Data: Leveraging the university's server, we accessed and extracted students' historical course data. This dataset provides invaluable insights into the academic trajectories and preferences of students, serving as a foundational element for our recommendation algorithm.

Employment Information: Recognizing the importance of aligning course recommendations with market demand, we systematically collected employment-related data through APIs from the WorkNet (2024) and JobKorea (2024) websites. WorkNet is a South Korean employment information site operated by the Ministry of Employment and Labor and the Korea Employment Information Service, providing job search and recruitment information as well as career guidance since 1998. JobKorea is South Korea's leading job portal, offering real-time job listings, company analysis, and various employment services. This data collection process helped us identify job trends, skill requirements, and job openings across different industries.

Company Information: To enrich our dataset with a broader understanding of the employment environment, we collected data from the KoDATA (2024) site via FTP and integrated it with company information. KoDATA is Korea's premier comprehensive credit research and evaluation agency, leading the credit rating market with a database of about 12 million companies and a nationwide network. Established to promote credit lending and healthy credit commerce by providing reliable credit information, KoDATA's dataset includes company details such as name, type, size, form, status, and industry code in Korea.

Unified Database: Upon completing the data collection process, we integrated the disparate datasets into a coherent, unified database. This integrated database is designed to facilitate efficient data retrieval, manipulation, and analysis, serving as the backbone for training our recommendation model. The integration process involved normalizing data formats, resolving data inconsistencies, and establishing relational links between the datasets.

2.2 NeuMF

The Neural Matrix Factorization (NeuMF) (He et al., 2017) model represents an innovative approach in collaborative filtering by leveraging deep neural networks to encapsulate the intricate interrelations within data, specifically enabling the modeling of non-linear associations. At its core, NeuMF introduces a hybrid architecture that seamlessly integrates aspects of Generalized Matrix Factorization (GMF) and Multilayer Perceptron (MLP), each employing an embedding layer critical for transforming sparse input vectors into dense representations. These representations, denoted as $p_u = p^T x_u$ for GMF and $q_i = Q^T x_i$ for MLP, facilitate the mapping where p_u and q_i represent the embedding vectors for user u and item i, respectively.

The GMF component of the model enhances interaction scoring between users and items by conducting an element-wise multiplication of their latent vectors, further refined through an activation function and weight adjustments. This mechanism is augmented by incorporating a non-linear sigmoid function, offering a broader expressive capacity compared to traditional linear Matrix Factorization techniques.

On the other hand, the MLP component introduces a higher degree of flexibility and non-linearity by processing the concatenated user-item vector through multiple hidden layers. This process, symbolized as $h_u i = p_u \odot q_i$, where $\odot$ denotes element-wise multiplication, allows the MLP to adeptly navigate through the layers using various activation functions to determine the interaction score.

By amalgamating the linear characteristics of MF with the non-linear dynamism of MLP, the NeuMF model achieves a delicate balance, enhancing user-item interaction representation while preserving the essence of collaborative filtering—the modeling of user-item interplay. This blend, where MLP spans from representations $z_{ui}^{(1)} = f^{(1)}(p_u \oplus q_i)$ to $z_{ui}^{(L)} = f^{(L)}\left(z_{ui}^{(L-1)}\right)$, culminates in a combined GMF and MLP structure, represented as $y_{ui} = \sigma\left(a^T\left(h_{ui} \oplus z_{ui}^{(L)} + b\right)\right)$. Here, σ signifies the sigmoid activation function, a the weight vector, and b the bias, collectively refining the performance of collaborative filtering without compromising its fundamental principles.

2.3 DeepFM

The DeepFM (Guo et al., 2017) model stands at the confluence of deep learning and factorization techniques, uniquely structured to dissect and learn from the interplay among diverse attributes through an integrated framework of linear and non-linear methodologies. At its foundational level, the model initiates with an input layer responsible for generating embedding vectors that encapsulate the distinctive features of both users and items. This process transforms each attribute, initially represented as a one-hot encoded categorical variable, into a compact, low-dimensional vector via the embedding layer, formalized as $v_i = Embed(x_i)$, where v_i symbolizes the embedding vector for attribute x_i.

Central to the DeepFM model is the Factorization Machine (FM) component, adept at modelling both simple and complex feature interactions. By embracing linear and second-order quadratic interactions, the FM part efficiently uncovers patterns instrumental for refining recommendations, represented as $\mathrm{FM}_{out} = \sum_{i=1}^{n} v_i \cdot x_i$, where n indicates the count of features and the dot product is utilized to quantify interactions among them.

Complementing the FM's process, the Deep Neural Network (DNN) segment of DeepFM delves into the realm of nonlinear interactions. This nonlinear exploration is denoted as $\mathrm{DNN}_{out} = \sigma(W_{DNN} \cdot concat(v_1, \cdots, v_n) + b_{DNN})$, where σ embodies the activation function, and W_{DNN} and b_{DNN} correspond to the weights and biases intrinsic to the DNN layers, respectively.

The culmination of this model is a sophisticated predictive output, derived from the synergistic combination of FM and DNN outputs, denoted by $\hat{y} = \sigma(\mathrm{FM}_{out} + \mathrm{DNN}_{out})$, where a sigmoid function σ is applied to ensure the final prediction seamlessly integrates the linear and non-linear insights drawn from the user-item matrix. This hybrid approach enables DeepFM to achieve a nuanced understanding of feature interactions, significantly enhancing the precision of recommendations provided.

2.4 xDeepFM

The xDeepFM (Lian et al., 2018) model synergizes the foundational strengths of matrix factorization with the advanced representational capabilities of deep learning to enhance overall model performance. At the heart of xDeepFM lies the innovative Cross Network structure, designed to meticulously capture feature interactions along the cross dimension, thereby augmenting the model's representational efficacy. The Input Layer initiates this process by constructing embedding vectors for user and item attributes, transforming each attribute into a high-dimensional space for the model to effectively learn from. This transformation is succinctly encapsulated as $v_i = Embed(x_i)$, where v_i represents the embedding vector corresponding to attribute x_i. The Cross Network uniquely computes cross terms between input features, facilitating a nuanced representation of interactions and modelling complex, non-linear relationships between attributes. These interactions are mathematically expressed as $z_{k+1} = z_k + \sum_{i=1}^{n}\sum_{j=i+1}^{n}(v_i \odot v_j) \cdot w_{ij}^{(k)}$, where k denotes the order of the cross term, $\odot$ symbolizes element-wise multiplication, and $w_{ij}^{(k)}$ refers to the weight associated with the second cross term. Complementing this, the Deep Neural Network (DNN) component of xDeepFM delves into modeling the non-linear dynamics of input characteristics, employing an activation function denoted by σ. Within this framework, $h_{deep} = \sigma(W_{deep} \cdot v_{stack} + b_{deep})$ signifies the activation function output, v_{stack} represents the collective stack of all embedding vectors, and W_{deep} and b_{deep} are the respective weights and biases of the deep network.

The culmination of the xDeepFM model's operation is the integration of insights from both the Cross Network and DNN, yielding the final predictive output. This integration is denoted by $y = \sigma(z_{final} + h_{deep} \cdot W_{out} + b_{out})$, where z_{final} encapsulates the Cross Network's ultimate output, and W_{out} and b_{out} stand as the weights and biases of the output layer, ensuring a cohesive and powerful representation of user-item interactions through both linear and non-linear lenses.

2.5 DCNv2

The Deep & Cross Network version 2 (DCNv2) (Wang et al., 2021) is an advanced model tailored for enhancing ranking systems through deep learning, ingeniously integrating both a cross network for capturing explicit feature interactions and a stacked neural network for unveiling implicit feature interactions. At the foundation of DCNv2 lies the embedding layer, tasked with processing both categorical and dense inputs, encapsulated as $x_0 = Embedding(x)$, serving as the initial step in feature interaction exploration.

Within the cross network, situated at each cross layer, DCNv2 meticulously extracts explicit feature interactions across multiple layers. The process is mathematically represented as $x_{l+1} = x_0 \odot (W_l x_l + b_l) + x_l$, where x_l and x_{l+1} denote the input and output vectors of a given layer, respectively. These vectors undergo linear computations using layer-specific weights W_l and biases b_l, facilitating the extraction of feature interactions through iterative element-wise products of the output from the embedding layer x_0. Notably, the first cross layer incorporates a unique mechanism where the interaction information among features is encoded through the element-wise product of x_0 and its updated weight.

Complementing the explicit interaction mapping by the cross network, the Deep Network operates as a conventional feed-forward neural network. It leverages linear computations followed by activation functions, succinctly described as $h_{l+1} = f(W_l h_l + b_l)$, where h_l and b_l represent the weights and biases of the deep network, respectively.

DCNv2 distinguishes itself with two structural variants in combining the cross and deep networks: the Stacked Structure and the Parallel Structure. In the Stacked Structure, the initial input x_0 first traverses the cross network before being processed by the deep network. Conversely, the Parallel Structure concurrently feeds x_0 into both the cross and deep networks, with the final output layer synthesizing the outputs from both networks x_{L_c} from the cross network and h_{L_c} from the deep network) to form the final prediction. This prediction is articulated as $\hat{y}_i = \sigma\left(w_{log\ddot{y}}^{t} \cdot log\ddot{y}(x_{final})\right)$, where $w_{log\ddot{y}}$ is a dedicated weight vector for the logit, and σ denotes the sigmoid function, illustrating DCNv2's holistic approach to modeling feature interactions for improved ranking performance.

2.6 AFN+

Adaptive Feature Networks (AFN) (Cheng et al., 2020) models utilize a dynamic approach to learn the intersection of features of any order directly from the dataset, leveraging a logarithmic transformation layer. This innovative layer translates the influence of each feature within a feature combination into a learnable coefficient, enhancing the model's ability to understand complex relationships. The architecture of AFN is structured into several key components: an Input Layer and an Embedding Layer, a Logarithmic Transformation Layer for cross-feature learning, multiple Feed-forward Hidden Layers for integrating learned cross features, and an Ensemble mechanism that merges AFNs with Deep Neural Networks (DNNs) to create a comprehensive ensemble model.

The process begins in the Input Layer, where both sparse categorical and numerical inputs are collected. These inputs are then transformed into a series of embedding vectors, $E = (e_1, e_2, \ldots, e_m)$, in the Embedding Layer. This is followed by the Logarithmic Transformation Layer, which employs specialized logarithmic neurons tailored to each vector. These neurons are designed to ascertain the significance (or order) of each attribute within the cross-attributes, $y_j = \exp\left(\sum_{i=1}^{m} w_{ij} \ln_{ei}\right)$, where w_{ij} is the weight allocated to the logarithmic neuron. The output of this layer is symbolized by $Y = (y_1, y_2, \ldots, y_n)$, indicating a nuanced understanding of feature interrelations.

Subsequent to the logarithmic transformation, the architecture advances to the Feed-forward Hidden Layers. Here, a sequence of fully connected layers is employed to amalgamate the cross features identified earlier. This operation is mathematically expressed as $z_l = ReLU(W_l z_{l-1} + b_l), l = 1, 2, \ldots, L$, where $z_0 = Y$, W_l symbolizing the weight matrix of the l-th layer, b_l representing the bias vector, and L marking the total number of hidden layers. The culmination of this process is the prediction layer $\hat{y} = w_p z_L + b_p$, where w_p and b_p stand for the prediction layer's weight vector and bias, respectively.

In the final stage, Ensemble with DNN, the AFNs are amalgamated in a fashion akin to DNNs, yielding an ensemble model denoted by $\hat{y}_{ensem} = w_1 \hat{y}_{AFN} + w_2 \hat{y}_{DNN} + b$. This ensemble model incorporates predictions from both AFN and DNN components ($\hat{y}_{AFN}$ and $\hat{y}_{DNN}$), balanced by corresponding weights w_1 and w_2, with b acting as the bias. This holistic approach not only enhances the model's predictive capabilities by combining the strengths of AFNs and DNNs but also offers a more nuanced understanding of feature interactions, setting a new benchmark in the field of machine learning for handling complex data relationships.

2.7 EulerNet

EulerNet (Tian et al., 2023) introduces an innovative approach for learning attribute interactions adaptively, capable of autonomously discerning interactions of any order directly from data. Distinguished by its unified architectural framework, EulerNet adeptly handles both explicit and implicit feature interactions by situating these interactions within a complex vector space, as inspired by Euler's formula. The structure of EulerNet encompasses an Embedding Layer for input embeddings, an Euler Interaction Layer for deciphering explicit feature interactions, and a predictive output segment specifically designed for Click-Through Rate (CTR) estimations, derived from the interaction layer's outcomes.

The core of EulerNet begins with the Embedding Layer, which translates input features into a series of feature embeddings $(e_1, e_2, \ldots, e_m)$. Progressing to the Euler Interaction Layer, it introduces a Complex Space Mapping technique to transition feature embeddings from the conventional real vector space to a complex vector space. Utilizing Euler's formula, it transforms these embeddings into complex numbers represented in polar coordinates, where $e_j = \mu_j \cos(e_j) + i\mu_j \sin(e_j)$, with μ_j being a modifiable coefficient parameter. The Generalized Multi-order Transformation, another pivotal component, facilitates the conversion of complex feature representations from Cartesian to polar coordinates, expressed through $\exp\left(\sum_{j=1}^{m} a_j \log(\lambda_j)\right)$ $\exp(i \sum_{j=1}^{m} a_j \theta_j)$, where $\lambda_j = \sqrt{r_j^2 + p_j^2}$ (inherently non-negative) and $\theta_j = atan2(p_j, r_j)$ denote the coefficient and phase vectors of the complex feature in polar form.

Within the framework of EulerNet, the Generalized Multi-order Transformation is pivotal, facilitating the delineation of interactions in the complex domain induced by the Euler transformation. This transformation is quantitatively expressed through $\Delta e = \sum_{k=1}^{n} l_k e^{i\Lambda k}$ and Λk, which respectively signify the explicit interaction and its complex representation within the Euler transform framework. The process extends to encompass Implicit Interaction Integration, capturing the nuances of implicit interactions. These interactions are mathematically represented by $r'k = ReLU(w_k r + b_k)$ and $p'k = ReLU(w_k p + b_k)$, where r and p denote the real and imaginary components, respectively, offering a comprehensive view of feature interrelations.

To amalgamate the explicit and implicit facets of feature interactions, EulerNet strategically combines these representations by aligning their real and imaginary components, encapsulated in the output vector $(o_k)_{k=1}^{n} = \left((\hat{r}_k + r'_k) + i(\hat{p}_k + p'_k)\right)_{k=1}^{n}$. This synthesis is instrumental in constructing a robust model that appreciates the depth of feature interactions.

In the phase dedicated to Click-Through Rate (CTR) predictions, EulerNet undertakes a linear regression analysis to deduce a scalar output, symbolized by $z = w^T\hat{r} + i(w^T\hat{p}) = z_{re} + iz_{im}$. This output, z, is dissected into its real z_{re} and imaginary z_{im} components, offering a dual perspective on the predictive analysis. Culminating the process, EulerNet intricately weaves together the explicit and implicit interaction insights, employing a sigmoid function, denoted as σ, to calculate the CTR value, represented by $\hat{y} = \sigma(z_{re} + z_{im})$. This approach not only exemplifies EulerNet's adeptness in harnessing complex feature interplay but also showcases its proficiency in translating these interactions into accurate predictive insights, underpinning its utility in the domain of CTR prediction.

3 PROPOSED METHOD

Contrary to the traditional focus on predicting and recommending academic courses to students for upcoming semesters, as highlighted in previous research (Lee et al., 2021), this paper introduces a novel paradigm by exploring course recommendation from the perspective of potential employers. This innovative approach aims to align educational outcomes with the specific needs and preferences of companies where students aspire to work post-graduation. By analyzing data from graduates and their subsequent employment destinations, this study diverges from the prevalent student-centric models of course recommendation, proposing instead a company-centric framework.

This paradigm shift entails a significant realignment of the course recommendation process, with the primary criterion being the skill sets and knowledge areas valued by employers, rather than the academic or career interests of the students alone. Essentially, the study redefines the target user from the individual student to the company, advocating for a model where courses are recommended to companies based on the competencies they seek in their future employees. This approach underscores a strategic pivot towards enhancing the employability of graduates by directly addressing the demand-driven requirements of the job market, thereby fostering a more effective and pragmatic connection between educational institutions and the business sector.

3.1 Feature Engineering

This study utilizes a comprehensive dataset spanning two decades, from 2000 to 2020, encompassing student enrollment records. Given the sensitive nature of students' course records, these data are de-identified to ensure privacy. Central to our analysis is the incorporation of company information, a novel approach in our research. We use a dataset of 887,996 course records from students who are employed and have firm information. The dataset includes 16,695 students, 314 colleges, 16,995 courses, and 925 firms. It is important to acknowledge potential variances in students' course loads and interruptions in their academic journeys, which may introduce noise during model training. To align with our company-centric recommendation framework, the dataset underwent a transformation to focus on organizational attributes rather than individual student or course data.

In adapting to this new dataset configuration, we employ the Neural Matrix Factorization (NeuMF) technique, leveraging a user-item matrix devoid of specific features. This method involves creating a company-course matrix based on the 925 identified companies. For our hybrid model, which allows for the inclusion of additional features, companies are further categorized by incorporating college and gender information, forming a (company, gender, college) × subjects matrix. This stratification enhances the model's granularity, treating entities with differing gender and college affiliations as distinct, thereby refining our recommendation process.

Notably, the NeuMF framework, in its subsequent applications, does not incorporate explicit features. However, by classifying companies along with gender and college dimensions, we achieve a nuanced segmentation akin to that of the hybrid model. This approach effectively replicates feature engineering methods, albeit indirectly, for the NeuMF model, thereby aligning it more closely with the hybrid model's methodology.

In our experimental setup, the conventional student user base is redefined to reflect company-oriented course data, as detailed in Table 1. This transformation allows the hybrid model to categorize companies by college and gender, enriching the user profile with more descriptive attributes. Furthermore, to maintain consistency across models, the dataset is modified to attribute these additional characteristics to users, leveraging gender and college as categorical variables for feature engineering. This adjustment ensures that both NeuMF and hybrid models operate under a similar user definition framework.

Table 1: Kangwon National University Enrollment Dataset Description.

Model	NeuMF	Hybrid	NeuMF (with Feature Eng.)
User (Company)	925	9600	9600
Item (Course)	15,763	15,763	15,763
Average Item	287	27	27
Sparsity	0.9746	0.9962	0.9962

Regarding course data representation, the NeuMF model initially aggregates courses at the company level without additional attributes, resulting in a higher average number of courses per company. However, when employing the hybrid model's feature engineering approach, the average number of subjects per company drops significantly due to the division by college and gender, indicating a more focused and tailored recommendation system compared to the broader aggregation of the traditional NeuMF model.

4 EXPERIMENTS

4.1 Dataset & Hyperparameters

In this study, we analyze student enrollment records from Kangwon National University, covering a comprehensive span of 20 years from 2000 to 2020. The dataset incorporates essential student data such as student ID (anonymized), gender, college affiliation, graduation year, academic grade, course codes, employer details, and company identification codes. To adhere to privacy concerns, all personally identifiable information was removed, with unique ID values substituting for class numbers. This preprocessing step enabled us to forge a de-identified dataset that links student profiles with their respective course enrollments, facilitating the model's development and validation phases. The resultant dataset encompasses records from 16,695 students across 16,995 distinct courses.

For model training and evaluation, the dataset was partitioned into training and test sets, following a 70:30 split. This allocation ensures a robust framework for assessing the model's predictive accuracy while providing ample data for training. The model's hyperparameters were meticulously optimized to enhance performance, with Binary Cross Entropy (BCE) selected as the loss function, the Adam optimizer for gradient descent, a batch size of 128, a learning rate of 1e-4, and an extensive training duration of 500 epochs. This configuration was determined to offer the best compromise between computational efficiency and predictive accuracy, laying the groundwork for a rigorous evaluation of the proposed course recommendation system.

4.2 Evaluation Metrics

To assess our model's performance, we utilized a suite of evaluation metrics commonly employed in the analysis of recommendation systems. These metrics include Precision, Recall, F1-Score, and Mean Average Precision (mAP), each offering unique insights into the efficacy of our model across different dimensions of performance.

Precision: This metric calculates the ratio of correctly recommended courses (those actually taken by students) to the total number of courses recommended by the model. A higher Precision indicates that a greater proportion of the courses recommended by the model are relevant and useful to the students. Precision calculates the ratio of correctly recommended courses (true positives, TP) to the total number of courses recommended by the model (true positives TP + false positives FP).

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

Recall: Recall measures the fraction of courses actually taken by students that are successfully recommended by the model out of all possible relevant courses. It assesses the model's ability to capture all pertinent recommendations. Recall measures the fraction of courses actually taken by students (true positives, TP) that are successfully recommended by the model out of all possible relevant courses (true positives TP + false negatives FN).

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

F1-Score: Serving as the harmonic mean of Precision and Recall, the F1-Score provides a single metric that balances both the precision of the recommendations and their completeness. It is particularly valuable when seeking a measure that accounts for both aspects of the model's performance, with higher values indicating a more balanced and effective recommendation system.

$$\text{F1-Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Mean Average Precision (mAP): Mean Average Precision (mAP) evaluates the quality of the entire ranking of recommendations produced by the model. It is the mean of the Average Precision (AP) scores for each query. This metric penalizes incorrect recommendations more severely when they appear higher in the list of recommendations, thereby emphasizing the importance of not only the accuracy but also the ranking of the proposed courses. Average Precision is calculated as the average of the precision values at the ranks where relevant items are found, accounting for the ranking order.

$$\mathrm{AP} = \frac{\sum_{k=1}^{n}(P(k) \times \mathrm{rel}(k))}{\text{number of relevant items}}$$

Where $P(k)$ is the precision at cut-off k and $\mathrm{rel}(k)$ is a binary indicator function that is 1 if the item at rank k is relevant, 0 otherwise.
Mean Average Precision (mAP) is then calculated as:

$$\mathrm{mAP} = \frac{1}{|Q|}\sum_{q \in Q} \mathrm{AP}(q)$$

Where Q is the set of all queries and $\mathrm{AP}(q)$ is the Average Precision for query q.

These metrics collectively enable a comprehensive evaluation of the recommendation model, scrutinizing its ability to not only identify relevant courses for recommendation but also prioritize them in a manner that aligns with actual student course selection patterns. Through this multidimensional analysis, we aim to ensure the development of a robust and effective course recommendation system that accurately reflects the needs and preferences of the target user base.

4.3 Evaluating Model Performance

In our research, we undertake a comparative analysis between the Neural Matrix Factorization (NeuMF) model and a Hybrid model to assess their efficacy within the context of a course recommendation system tailored for enterprises. The NeuMF model, grounded in deep learning methodologies, operates on the principle of utilizing implicit data derived from users' past behaviors to rank items, notably without the integration of explicit features. This approach distinguishes the NeuMF model as a pure collaborative filtering method, emphasizing the prediction of user preferences based on historical interaction patterns alone.

Conversely, the Hybrid model represents a more integrated approach, combining elements of collaborative filtering with the incorporation of additional data features to enhance recommendation accuracy. This juxtaposition of the NeuMF and Hybrid models in our study aims to elucidate the relative strengths and limitations of employing a solely behavior-based recommendation system against one that leverages a broader dataset encompassing explicit feature information.

By systematically comparing these models, our study contributes to a deeper understanding of the dynamics between collaborative filtering techniques and hybrid methodologies in the domain of enterprise-level course recommendations. This comparison not only serves to highlight the potential for improved recommendation precision and relevance through the inclusion of auxiliary data but also provides insights into the adaptability and scalability of these models in addressing the complex needs of enterprise environments.

In our experimental setup, to ensure consistency and control for variability, the division between the training and testing datasets was uniformly maintained across trials, with a split ratio consistently applied. To address the stochastic nature of negative sampling, we averaged the results over three iterations for each configuration. Typically, a student can sign up for 21 to 24 credits in a semester, and if they take a typical three-credit class, they take seven to eight courses each semester, so they take about 56 to 60 courses over the course of four years. However, there are also two-credit liberal arts courses, so we experimented with recommending courses in the range of 10 to 100. The models were assessed based on their ability to recommend the top 10, 25, 50, and 100 subjects, allowing for a comprehensive evaluation across varying scopes of recommendation breadth.

The performance of our models across varying recommendation sizes is summarized in four tables: Table 2 for the top 10, Table 3 for the top 25, Table 4 for the top 50, and Table 5 for the top 100 recommendations. Each table highlights key metrics such as Precision, Recall, F1-Score, and Mean Average Precision (mAP), offering insights into the models' effectiveness at different scales. In the results

Table 2: Metric results by model for the top 10.

Model	**P**	**R**	**F1**	**mAP**
DeepFM	0.16758	0.09438	0.12117	0.33515
xDeepFM	0.18488	0.10346	0.13654	0.34861
DCNv2	0.19331	0.10769	0.13857	0.35573
AFN+	0.19498	0.10946	0.14020	0.35256
EulerNet	0.19812	**0.11083**	**0.14214**	0.36322
NeuMF	**0.33286**	0.07552	0.12265	**0.50585**
NeuMF with FE	**0.36612**	**0.20352**	**0.26161**	**0.55461**

table, the highest-performing outcomes are highlighted with both bold and underline, the second-highest outcomes are emphasized in bold, and the third-highest outcomes are denoted by underline.

Table 3: Metric results by model for the top 25.

Model	P	R	F1	mAP
DeepFM	0.13389	0.18840	0.15979	0.28713
xDeepFM	0.14671	0.20444	0.17082	0.29901
DCNv2	0.15440	0.21497	0.17971	0.30423
AFN+	0.15333	0.21572	0.17925	0.30434
EulerNet	0.15697	**0.21944**	0.18302	0.30966
NeuMF	**0.27844**	0.14547	**0.19247**	**0.42953**
NeuMF with FE	**0.27112**	**0.37464**	**0.31458**	**0.47042**

Table 4: Metric results by model for the top 50.

Model	P	R	F1	mAP
DeepFM	0.10655	0.29734	0.18250	0.24413
xDeepFM	0.11504	0.31834	0.16900	0.25338
DCNv2	0.12157	0.33680	0.17865	0.25865
AFN+	0.11806	0.33135	0.17409	0.26033
EulerNet	0.12179	**0.33814**	0.17907	0.26351
NeuMF	**0.23137**	0.21915	**0.22389**	**0.37450**
NeuMF with FE	**0.19051**	**0.51893**	**0.27870**	**0.41277**

Table 5: Metric results by model for the top 100.

Model	P	R	F1	mAP
DeepFM	0.07972	0.44053	0.13946	0.20595
xDeepFM	0.08438	0.46379	0.14278	0.21371
DCNv2	0.08877	0.48891	0.15097	0.21824
AFN+	0.08549	0.47821	0.14504	0.22110
EulerNet	0.08920	**0.49104**	0.15025	0.22320
NeuMF	**0.18336**	0.30793	**0.22985**	**0.32576**
NeuMF with FE	**0.12078**	**0.64714**	**0.20356**	**0.36793**

4.3.1 NeuMF vs. Hybrid

The outcomes of our experiments reveal distinct performance characteristics between the models. The Neural Matrix Factorization (NeuMF) model exhibited superior performance in terms of Precision, indicating its effectiveness in accurately recommending relevant subjects. Conversely, the Hybrid model demonstrated a stronger Recall metric, suggesting its proficiency in capturing a broader array of relevant recommendations. Notably, within the Hybrid category, the DeepFM model, despite registering the lowest Precision, surpassed the NeuMF model in Recall, highlighting the trade-offs between Precision and Recall across different model architectures.

Precision in our context is calculated by considering the actual number of recommendations made and affirming their relevance if they match the subjects in the test set. Given our shift in focus to company-oriented recommendations, the NeuMF model's lack of reliance on individual-specific features for training may contribute to its heightened Precision, as it operates on a more generalized basis. For Recall, which measures the proportion of relevant subjects in the test set that are successfully recommended, the performance is independent of the recommendation volume. Here, the NeuMF model tends to lag in Recall due to the larger denominator reflecting the grouped learning approach, whereas the Hybrid model, by incorporating user-specific features, demonstrates enhanced Recall owing to a more defined and smaller denominator.

The F1-Score, serving as the harmonic mean of Precision and Recall, offers a balanced measure of a model's performance, rewarding scenarios where Precision and Recall are closely matched. In the initial recommendation set of the top 10 subjects, the Hybrid model, particularly EulerNet, showcases superior performance due to its balanced Precision and Recall. However, as the recommendation list expands beyond the top 10, the NeuMF model's relative improvement in Recall, coupled with its strong Precision, positions it favorably in terms of F1-Score.

Regarding Mean Average Precision (mAP), which assesses the weighted order of recommendations, NeuMF continues to outperform, benefiting from its Precision advantage. This metric's emphasis on the sequencing of recommendations underscores NeuMF's ability to prioritize the most relevant subjects effectively, even when the evaluation incorporates the order of recommendations. This nuanced analysis of model performance across multiple metrics and recommendation depths provides valuable insights into the strengths and limitations of each approach within the specific context of company-oriented course recommendations.

4.3.2 NeuMF with Feature Engineering

In this study, we sought to augment the Recall performance of the Neural Matrix Factorization (NeuMF) model, which already exhibited commendable Precision, by incorporating feature engineering into the dataset. Initially, the model treated users as aggregated entities representing a single company. However, by applying feature engineering, we disaggregated these entities based on gender and college attributes, aiming to refine user segmentation and enhance the model's ability to capture user nuances.

This tailored approach led to a notable improvement in Precision, with an increase of approximately 3% to a Precision score of 0.03 in the top 10 recommendations compared to the baseline NeuMF model without feature engineering. Although a slight decrease in Precision was observed in subsequent recommendations, this variance can be attributed to the differentiated number of average items per user, as detailed in Table 1. The model incorporating feature engineering demonstrated increased interactivity, albeit with a lower average item count per user, potentially impacting performance in top_k recommendations due to a diminished item pool.

The enhancements in Mean Average Precision (mAP) across top_k recommendations underscore the enriched interaction dynamics post-feature engineering, surpassing the prior iteration of the NeuMF model. Moreover, the recalibrated approach yielded substantial gains in Recall, outperforming both the original NeuMF and the Hybrid model known for its strong Recall capabilities. This improvement is ascribed to the refined calculation of Recall, where the lower average item count and the expanded user base due to feature engineering contributed to a more favorable denominator, thereby elevating Recall performance.

Conclusively, the introduction of feature engineering not only preserved the Precision levels comparable to the NeuMF model's prior performance but also significantly enhanced interaction quality (as indicated by mAP) and Recall metrics, surpassing the Hybrid model's Recall efficacy. Additionally, the harmonized balance between Precision and Recall, reflected in the elevated F1-Score, attests to the comprehensive performance uplift achieved through this methodological refinement. This dual enhancement in both Precision and Recall dimensions underscores the effectiveness of integrating feature engineering into the NeuMF model, presenting a compelling case for its application in enhancing user-specific recommendation systems.

5 LIMITATIONS & FUTURE WORK

The NeuMF model, despite its strengths, exhibited limitations due to its inability to utilize features, necessitating a workaround by representing users as companies. This approach posed challenges in making personalized recommendations within a company, given the varied job roles and preferences across departments and genders. The feature engineering enhancements in this study, while improving performance, were constrained by the minimal use of features (college, gender, company) and were not applicable to continuous data with more extensive features.

Our exploration of hybrid models—DeepFM, xDeepFM, DCNv2, AFN+, and EulerNet—utilized previously unexplored user and item information such as gender and college, marking a significant stride in offering students broader options. However, the hybrid models' potential was not fully realized due to limited metadata on students, restricting us to only three features and excluding critical information like detailed course attributes. The absence of comprehensive course features, such as major requirements and electives, limits our ability to guide students in course selection aligned with their career aspirations.

Future research aims to refine these recommendation systems by incorporating comprehensive course information, enabling personalized course recommendations based on students' career interests and academic history. This enhancement will facilitate a more detailed and personalized recommendation framework, better aligning educational outcomes with career aspirations.

6 CONCLUSIONS

In this study, we transitioned from traditional counseling methods to a student career recommendation system utilizing advanced recommendation algorithms. We employed the NeuMF model to recommend business-oriented subjects to students, leveraging a neural network to capture non-linear user-item interactions and recognizing more intricate patterns than linear models. Despite the challenges associated with feature utilization and user representation, our enhancements through feature engineering and exploration of hybrid models demonstrated significant strides in enriching recommendation systems.

Our investigation into incorporating nuanced features like gender and college in hybrid models highlighted both the progress and the limitations of our current approach. This study lays the groundwork for advanced recommendation systems that not only predict future dynamics of commercial districts but also assist students in navigating their educational and career paths more effectively. By addressing the identified limitations and focusing on comprehensive feature integration in future research, we aim to develop a more robust and effective course recommendation system tailored to students' career aspirations.

ACKNOWLEDGEMENTS

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [No.2022-0-00641, XVoice: Multi-Modal Voice Meta Learning]

REFERENCES

Baek, J., Kim, H., & Kwon, K. (2021). Artificial Intelligence-Based High School Course and University Major Recommendation System for Course-Related Career Exploration. *KIPS Transactions on Software and Data Engineering*, *10*(1), 35-44.

Bayer, I., He, X., Kanagal, B., & Rendle, S. (2017, April). A generic coordinate descent framework for learning from implicit feedback. In *Proceedings of the 26th international conference on world wide web* (pp. 1341-1350).

Cheng, W., Shen, Y., & Huang, L. (2020, April). Adaptive factorization network: Learning adaptive-order feature interactions. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 34, No. 04, pp. 3609-3616).

Chung, K. A., & Cho, J. H. (2015). Differences between major satisfaction, university life adjustment, self-esteem and career identity according to the time of department selection of students in the department of occupational therapy at one university. *Journal of the Korea Convergence Society*, *6*(5), 143-155.

Goutte, C., & Gaussier, E. (2005, March). A probabilistic interpretation of precision, recall and F-score, with implication for evaluation. In *European conference on information retrieval* (pp. 345-359). Berlin, Heidelberg: Springer Berlin Heidelberg.

Guo, H., Tang, R., Ye, Y., Li, Z., & He, X. (2017). DeepFM: a factorization-machine based neural network for CTR prediction. *arXiv preprint arXiv:1703.04247*.

He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T. S. (2017, April). Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web* (pp. 173-182).

Huang, Z., Yu, C., Ni, J., Liu, H., Zeng, C., & Tang, Y. (2019). An efficient hybrid recommendation model with deep neural networks. *IEEE Access*, *7*, 137900-137912.

Lee, J. K., Park, H., & Kim, W. (2021). Major Class Recommendation System based on Deep learning using Network Analysis. *Journal of Intelligence and Information Systems*, *27*(3), 95-112.

Lee, Y. (2020). A study on the Correlation of between Online Learning Patterns and Learning Effects in the Non-face-to-face Learning Environment. *Journal of the Korea Academia-Industrial cooperation Society*, *21*(8), 557-562.

Lian, J., Zhou, X., Zhang, F., Chen, Z., Xie, X., & Sun, G. (2018, July). xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining* (pp. 1754-1763).

Schafer, J. B., Frankowski, D., Herlocker, J., & Sen, S. (2007). Collaborative filtering recommender systems. In *The adaptive web: methods and strategies of web personalization* (pp. 291-324). Berlin, Heidelberg: Springer Berlin Heidelberg.

Tian, Z., Bai, T., Zhao, W. X., Wen, J. R., & Cao, Z. (2023, July). EulerNet: Adaptive Feature Interaction Learning via Euler's Formula for CTR Prediction. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 1376-1385).

Turner, C. R., Fuggetta, A., Lavazza, L., & Wolf, A. L. (1999). A conceptual basis for feature engineering. *Journal of Systems and Software*, *49*(1), 3-15.

Wang, R., Shivanna, R., Cheng, D., Jain, S., Lin, D., Hong, L., & Chi, E. (2021, April). Dcn v2: Improved deep & cross network and practical lessons for web-scale learning to rank systems. In *Proceedings of the web conference 2021* (pp. 1785-1797).

Yang, Y., & Zhai, P. (2022). Click-through rate prediction in online advertising: A literature review. *Information Processing & Management*, *59*(2), 102853.

WorkNet. (2024). National employment information site run by the Ministry of Labor and Employment and the Korea Employment Information Service. Retrieved May 18, 2024, from https://www.work.go.kr.

JobKorea. (2024). Korea's leading job platform. Retrieved May 18, 2024, from https://www.jobkorea.co.kr.

KoDATA, (2024). Korea Evaluation Data. Retrieved May 18, 2024, from http://www.kodata.co.kr/.