

AN EXPLAINABLE MULTI-TIERED CAREER RECOMMENDER SYSTEM USING SHAP-GUIDED NATURAL LANGUAGE INTERPRETATIONS

A Thesis Submitted to the
UNIVERSITY OF CALICUT
in partial fulfilment of the requirements for the award of the degree of
DOCTOR OF PHILOSOPHY IN COMPUTER SCIENCE
Under the Faculty of Science

By
NOUSHAD RAHIM M.

Under the guidance of
Prof. (Dr.) MOHAMED BASHEER K. P.
Principal
Amal College of Advanced Studies (Autonomous), Nilambur



P.G & RESEARCH DEPARTMENT OF COMPUTER SCIENCE
Sullamussalam Science College (Autonomous), Areekode - 673639
(Affiliated to the University of Calicut)
Malappuram Dist., Kerala, India



NOVEMBER 2025



**SULLAMUSSALAM
SCIENCE COLLEGE**
AREEKODE

Phone :0483-2850700

Mob :+919895811093

Website : www.sscollege.ac.in

E-mail : mail@sscollege.ac.in

Aided College | Affiliated to the University of Calicut | Re-accredited by NAAC at A Grade

Ref. HA-1077/2025-26/PhD

Date 19.05.2026.....

CERTIFICATE

This is to certify that the thesis entitled “AN EXPLAINABLE MULTI-TIERED CAREER RECOMMENDER SYSTEM USING SHAP-GUIDED NATURAL LANGUAGE INTERPRETATIONS” submitted by Mr. Noushad Rahim M, to the University of Calicut, for the partial fulfilment of the requirements for the award of the degree of Doctor of Philosophy (Ph.D) in Computer Science, is a bona fide research work done by Mr. Noushad Rahim M, under my supervision and guidance, in the PG & Research Department of Computer Science, Sullamussalam Science College, Areekode, Malappuram, Kerala. The content embodied in this thesis, in full or in parts, has not been submitted to any other University or Institute for the award of any degree.

The thesis is revised as per the modifications and recommendations reported by adjudicators. The attached soft copy is the same as that of the revised copy. The thesis is submitted as such, to the University of Calicut with reference to the letter number: No. 274679/RESEARCH-C-ASST-1/2025/Admn, Dated: 13.05.2026.

Prof. (Dr.) Mohamed Basheer K. P.

Research Guide, PG & Research Department of Computer Science
Sullamussalam Science College, Areekode

Prof. (Dr.) MOHAMED BASHEER K.P.
Principal
Amal College of Advanced Studies
(Autonomous)
Eranielmangal P.O., Nilambur - 679329

Ugrapuram Post, Areekode, Malappuram District, 673639



UNIVERSITY OF CALICUT

CERTIFICATE ON PLAGIARISM CHECK

1.	Name of the Research Scholar	NOUSHAD RAHIM M	
2.	Title of thesis / dissertation	AN EXPLAINABLE MULTI-TIERED CAREER RECOMMENDER SYSTEM USING SHAP-GUIDED NATURAL LANGUAGE INTERPRETATIONS	
3.	Name of the Supervisor	Prof. (Dr.) Mohamed Basheer K.P	
4.	Department/Institution	PG & Research Department of Computer Science, Sullamussalam Science College, Areekode, Malappuram, Kerala	
5.	Similar content (%) identified	Non Core	Core
		Introduction/ Theoretical overview/Review of literature/ Materials & Methods/ Methodology	Analysis/Result/Discussion / Summary/Conclusion/ Recommendations
		5	1
	Acceptable maximum limit (%)	10	10
6.	Software used	iThenticate	
7.	Date of verification	10.11.2025	

*Report on plagiarism check, specifying included/excluded items with % of similarity to be attached.

SEEMA M.V.
Assistant Librarian
University of Calicut
Malappuram - 673 635

Checked by (with name , designation & signature) Seema


Name and signature of the Researcher Noushad Rahim M

Name and signature of the Supervisor:
Prof. (Dr.) MOHAMED BASHEER K.P
PEN: 677975
Principal

The Doctoral Committee* has verified the report on plagiarism check with the contents of the thesis, as summarized above and appropriate measures have been taken to ensure originality of the Research accomplished herein.

Name & Signature of the HoD/HoI (Chairperson of the Doctoral Committee)
PRINCIPAL (i/c)
MUSLIM SCIENCE COLLEGE
AREEKODE, UGRAPURAM (P.O.)
MALAPPURAM (D.O. Pin: 673639)

*In case of languages like Malayalam, Tamil etc..on which no software is available for plagiarism check, a manual check shall be made by the Doctoral Committee, for which an additional certificate has to be attached.



DECLARATION

I, **Noushad Rahim M**, hereby declare that this thesis, entitled “**AN EXPLAINABLE MULTI-TIERED CAREER RECOMMENDER SYSTEM USING SHAP-GUIDED NATURAL LANGUAGE INTERPRETATIONS**”, is based on the original work done by me under the supervision of Prof. (Dr.) Mohamed Basheer K.P., Research Supervisor, PG & Research Department of Computer Science, Sullamussalam Science College (Autonomous), Areekode.

I confirm that,

- The work presented in this thesis has not been submitted previously for the award of any degree, either to this University or to any other university or institution.
- I have followed the guiding principles given by the University in organizing the thesis.
- Whenever I have used materials (theoretical analysis, data, figures, and text) from other sources, I have given due credit to them by citing them in the thesis and giving their particulars in the references.

Noushad Rahim M.

Areekode
19. 05. 2026

ACKNOWLEDGMENTS

I express my sincere gratitude to all who have guided, supported, and encouraged me throughout this research. Their insight and goodwill shaped the study from its initial conception to the final draft.

I owe my deepest thanks to Prof. (Dr.) Mohamed Basheer K. P., Research Supervisor and Principal, Amal College of Advanced Studies, Nilambur (Autonomous). His dedication set a clear standard for this work. He provided steady guidance on research design, asked precise questions at crucial moments, and helped refine the argument and structure of each chapter. He read multiple drafts with care, advised on methodological choices, and assisted in resolving practical difficulties during implementation. He also brought positive energy that sustained progress and raised expectations when the work risked settling for less.

I am equally grateful to Dr. Musthafa Farook P, Principal (I/C) and Chairperson of the Research Advisory Committee, Sullamussalam Science College, Areekode. As Principal, he created an enabling environment for research, including timely permissions, access to facilities, and effective coordination across departments. As Chairperson of the Committee, his suggestions were constructive and actionable, sharpening the scope, strengthening the evaluation plan, and improving the clarity of presentation.

I also thank Dr. Muhamed Ilyas P, former Principal, Sullamussalam Science College, Areekode, for his early encouragement and institutional support. I gratefully acknowledge Dr. Shameem Kappan, Associate Professor & Head, PG & Research Department, Sullamussalam Science College, Areekode, for effective coordination and timely guidance.

My sincere appreciation goes to Dr. Azeem C. M., State Co-ordinator, Career Guidance & Adolescent Counselling Cell, Directorate of General Education, Kerala. He played a pivotal role in securing the Memorandum of Understanding with the

Department of General Education, which enabled the use of the Kerala Differential Aptitude Test (KDAT) for this study.

I thank the members of the Research Advisory Committee for their oversight; faculty colleagues and administrative staff for assistance with scheduling, documentation, and logistics. I acknowledge with appreciation the expert counsellors who participated in the evaluations of the work. I am also grateful to the library staff for their day-to-day assistance.

I thank my parents, family, and friends for their constant encouragement. My heartfelt gratitude goes to my wife, Ashika Mumthas, and my daughter, Ayisha Gaib, whose steady support and patience sustained this work. To all who helped in any way, thank you.

Above all, praise be to the Almighty, whose mercy, guidance and grace sustained every step of this work.

Noushad Rahim M.

Dedicated

*To my kind-hearted father and ever-loving mother
(In grateful honour of their heavenly love)*

AN EXPLAINABLE MULTI-TIERED CAREER RECOMMENDER SYSTEM USING SHAP-GUIDED NATURAL LANGUAGE INTERPRETATIONS

Noushad Rahim M.
Research Scholar

Prof. (Dr.) Mohamed Basheer K. P.
Research Guide

ABSTRACT

Career guidance plays a vital role in helping students make informed educational and occupational choices. However, the shortage of trained career counsellors in countries such as India limits access to personalised guidance, often resulting in inappropriate course selection, reduced employability, and increased career switching, leading to the underutilisation of human resource potential and reduced returns on public investment in education. To address these challenges, this thesis proposes an explainable multi-tiered career recommender system that combines machine learning, explainable artificial intelligence, and natural language generation to deliver scalable, transparent, and personalised career recommendations. The primary contribution of this research is a novel hierarchical and multi-tiered recommendation architecture designed to generate diverse yet relevant career recommendations aligned with a candidate's aptitude profile. Unlike conventional recommender systems that often produce narrow recommendation lists, the proposed architecture organises careers into aptitude-based clusters and integrates multiple machine learning models within a content-based filtering framework, enabling recommendations from diverse career domains while preserving their suitability to the candidate's abilities. A second major contribution is a hybrid explainability framework that integrates the local and global explainability capabilities of SHAP (SHapley Additive exPlanations) to provide both candidate-level and career-level explanations. These explanations are automatically translated into natural language interpretations using NLG techniques, making recommendation outcomes more transparent and understandable. The recommendation framework is supported by a first-of-its-kind aptitude-based career dataset developed using career ability profiles derived from the Occupational Information Network (O*NET) database and candidate aptitude profiles generated using the Kerala Differential Aptitude Tests (KDAT), thereby addressing the traditional data sparsity and cold-start limitations encountered in career recommender systems. The thesis further introduces a scalable semi-supervised career expansion mechanism that enables the incorporation of emerging careers into the recommendation framework without retraining the underlying recommendation model. Experimental and expert evaluations demonstrated the robustness, effectiveness, and acceptability of both the recommendations and the generated explanations. Overall, this research presents novel contributions in both recommendation and explainability architectures and demonstrates how explainable machine learning can be effectively integrated into career counselling to deliver scalable, diverse, interpretable, and trustworthy recommendations.

Keywords: Career counselling; recommender systems; explainable artificial intelligence; SHAP; natural language explanations; scalable recommender systems.



Prof. (Dr.) MOHAMED BASHEER K.P
PEN: 677975
Principal
Amal College of Advanced Studies
Eranielangad P.O., Nilambur - 679329

AN EXPLAINABLE MULTI-TIERED CAREER RECOMMENDER SYSTEM USING SHAP-GUIDED NATURAL LANGUAGE INTERPRETATIONS

സംഗ്രഹം

തൊഴിൽ മാർഗ്ഗനിർദ്ദേശക - വിദ്യാർത്ഥി അനുപാതം വളരെ കുറവുള്ള ഇന്ത്യ പോലുള്ള സാമൂഹിക സാഹചര്യങ്ങളിൽ ഓരോ വിദ്യാർത്ഥിക്കും അവരുടെ അഭിരുചിയും താല്പര്യവും പരിഗണിച്ച് വ്യക്തിഗതമായ തൊഴിൽ മാർഗ്ഗനിർദ്ദേശം നൽകാൻ നിർമ്മിതബുദ്ധി ഉൾക്കൊള്ളുന്ന സാങ്കേതിക വിദ്യകൾ പ്രയോജനപ്പെടുത്തുക വഴി ഉന്നത വിദ്യാഭ്യാസ സ്ഥാപനങ്ങളിലെ കൊഴിഞ്ഞുപോക്ക് തടയാനും, വ്യവസായ മേഖലയിലെ തൊഴിൽ മാറ്റത്തിന്റെ തോത് കുറയ്ക്കാനും വിദ്യാഭ്യാസ മേഖലയിലെ സർക്കാരിന്റെ നിക്ഷേപം ഫലപ്രദമായി വിനിയോഗിക്കാനും സാധിക്കും. ഈ സാധ്യതകളെ മുൻനിർത്തി, നൂതനമായ ഘടനയോടുകൂടിയതും സുതാര്യത ഉറപ്പാക്കാൻ, സാധാരണക്കാരന് മനസ്സിലാകുന്ന സ്വാഭാവിക ഭാഷയിൽ കാര്യങ്ങൾ വിശദീകരിക്കാൻ ശേഷിയുള്ള ഒരു തൊഴിൽ ശുപാർശ സംവിധാനം ഈ ഗവേഷണം മുന്നോട്ടുവെക്കുന്നു. വിവരത്തിന്റെ അപര്യാപ്തത, പ്രാരംഭ ശുപാർശ പ്രതിസന്ധി തുടങ്ങിയ ശുപാർശ സംവിധാനങ്ങളിലെ പ്രശ്നങ്ങൾക്ക് പരിഹാരമാകുന്ന ഇത്, അഭിരുചി അടിസ്ഥാനമാക്കിയുള്ള, സമഗ്രമായ വിവരശേഖരത്തെ അടിസ്ഥാനമാക്കിയാണ് നിർമ്മിച്ചത്. ശുപാർശകളിൽ വൈവിധ്യം ഉറപ്പാക്കുന്നതിനായി വിവര ശേഖരത്തെ ശ്രേണീകൃത വർഗ്ഗീകരണത്തിന് വിധേയമാക്കുകയും, അതിലൂടെ ഓരോ വിഭാഗത്തിൽ നിന്നും ഏറ്റവും അനുയോജ്യമായ തൊഴിലുകൾ തിരഞ്ഞെടുക്കുകയും ചെയ്യുന്നു. വ്യാഖ്യാന ശേഷിയുള്ള 'ഷാപ്പ്' അൽഗോരിതത്തിന്റെ വ്യത്യസ്ത വിശകലന സാധ്യതകൾ പ്രയോജനപ്പെടുത്തി മനുഷ്യഭാഷയിൽ ശുപാർശയുടെ കാരണങ്ങൾ ലളിതമായി വിശദീകരിക്കാൻ സാധിക്കുന്നു എന്നതാണ് ഈ നവീന സംവിധാനത്തിന്റെ മറ്റൊരു പ്രധാന നേട്ടം. അമേരിക്കൻ സ്റ്റേറ്റ് തൊഴിൽ വകുപ്പിന്റെ സഹായത്തോടെ തയ്യാറാക്കിയ തൊഴിൽ വിവര ശൃംഖലാ ശേഖരത്തിൽ നിന്നും, ഇന്ത്യൻ സാഹചര്യങ്ങളിൽ പ്രാധാന്യമുള്ള 612 തൊഴിലുകൾ ഉൾക്കൊള്ളുന്ന, ഇത്തരത്തിലുള്ള ആദ്യത്തെ വിവരശേഖരം ഈ ഗവേഷണത്തിലൂടെ നിർമ്മിച്ചെടുത്തു. കേരള പൊതുവിദ്യാഭ്യാസ വകുപ്പ് രൂപപ്പെടുത്തിയെടുത്ത കെ-ഡാറ്റ് എന്ന സംവിധാനമാണ് ഈ ഗവേഷണത്തിന്റെ ഭാഗമായി അഭിരുചി നിർണയിക്കാൻ ഉപയോഗിച്ചിരിക്കുന്നത്. മുഖ്യനിർണയങ്ങളിലെല്ലാം മികച്ച കാര്യക്ഷമത തെളിയിച്ച ഈ ശുപാർശ സംവിധാനം പത്ത് തൊഴിൽ ശുപാർശകരുടെ പങ്കാളിത്തത്തോടെ നടത്തിയ വിദഗ്ധ വിലയിരുത്തലും അതിന്റെ സ്വീകാര്യത സ്ഥിരീകരിച്ചു. തൊഴിൽ മാർഗ്ഗനിർദ്ദേശ മേഖലയിലേക്ക് ഉപയുക്തമാകുന്നതോടൊപ്പം ശുപാർശ സംവിധാനത്തിന്റെ ഘടന, വ്യാഖ്യാന സംവിധാനത്തിന്റെ ഘടന എന്നിവയ്ക്ക് പുതിയ രീതികൾ ഈ ഗവേഷണ പഠനം സാധ്യമാക്കുന്നു. വൈവിധ്യം, വ്യാഖ്യാന യോഗ്യത, വിശ്വാസ്യത, വ്യാപ്തി എന്നിവ ഉൾക്കൊള്ളുന്ന ഘടകങ്ങളിലൂടെ വിഭവപരിമിത സാഹചര്യങ്ങളിലുള്ള വിദ്യാർത്ഥികളിലേക്ക് സാങ്കേതികവിദ്യ എങ്ങനെ പ്രയോജനപ്പെടുത്താമെന്ന് ഈ ഗവേഷണ ഫലങ്ങൾ വിശദീകരിക്കുന്നു.

സൂചകപദങ്ങൾ:- തൊഴിൽ മാർഗ്ഗനിർദ്ദേശം, ശുപാർശ സംവിധാനങ്ങൾ, വ്യാഖ്യാന ശേഷിയുള്ള നിർമ്മിതബുദ്ധി, ഷാപ്പ് അൽഗോരിതം, സ്വാഭാവിക ഭാഷാവിശദീകരണങ്ങൾ, വിപുലീകരണ ശേഷിയുള്ള ശുപാർശ സംവിധാനങ്ങൾ.

Noushad Rahim M.
Research Scholar

Prof. (Dr.) Mohamed Basheer K. P.
Research Guide

Prof. (Dr.) MOHAMED BASHEER K.P
PEN: 677975
Principal
Amal College of Advanced Studies
Eranhimangad P.O., Nilambur - 679329

CONTENTS

CHAPTER 1.....	1
INTRODUCTION.....	1
1.1. Background and Context	1
1.1.1. <i>Need for Personalised Career Counselling</i>	2
1.1.2. <i>Challenges in Existing Systems</i>	3
1.1.3. <i>Evolution of Recommender Systems in Career Guidance</i>	4
1.1.4. <i>The Proposed Research Solution</i>	6
1.2. Relevance of the Study	7
1.3. Scope of the Research	8
1.4. Limitations of the Research	10
1.5. Research Gap	11
1.6. Objectives of the Research	12
1.7. Methodological Overview	13
1.7.1. <i>System Architecture</i>	13
1.7.2. <i>Dataset creation</i>	14
1.7.3. <i>Performance Evaluation Metrics</i>	14
1.7.4. <i>Ethical Considerations</i>	15
1.8. Organization of the Thesis	15
1.9. Conclusion	18
CHAPTER 2.....	21
LITERATURE REVIEW.....	21
2.1. Theoretical Aspects for Career Counselling	21
2.2. Assessing Intelligence and the Evolution of Psychometric Tools in Career Counselling	25
2.2.1. <i>Limitations of Psychometric Tests and Possible Solutions</i>	32
2.3. Recommender Systems: A Comprehensive Review on Techniques, Challenges and AI-Driven Approaches	34
2.3.1. <i>Evolution of Recommender Systems</i>	34
2.3.2. <i>Applications of Recommender Systems in Diverse Domains</i>	38
2.3.3. <i>Current Approaches to Recommendation Systems in Career Guidance</i> 39	
2.4. Need for Explainable Recommendation Systems	42
2.4.1. <i>Explainability Approaches</i>	44
2.4.2. <i>Existing Implementations of Explainability in Recommendation Systems</i>	47
2.5. A Review of Relevant Datasets	50

2.6. Conclusion	53
CHAPTER 3.....	57
RESEARCH METHODOLOGY - AN OVERVIEW	57
3.1. Psychometric Foundations of the Methodology.....	57
3.2. Designing and Constructing the Dataset for Building the Career Profile.....	58
3.3. Dataset: Clustering, Sub-Clustering, and Data Augmentation	59
3.4. Generating User Profiles Through Differential Aptitude Tests	59
3.5. Structural Framework of the Career Recommendation Engine.....	60
3.6. Architectural Design of the Explainability Framework.....	60
3.7. Framework of the Dynamic Career Expansion Module	61
CHAPTER 4.....	63
THEORETICAL AND PSYCHOMETRIC FOUNDATION OF THE STUDY	63
4.1 Introduction.....	63
4.2 The Differential Aptitude Tests	64
4.2.1 Test Batteries in Differential Aptitude Tests	65
4.3 Limitations of Differential Aptitude Tests	69
4.4 Employing Differential Aptitude Tests in the Proposed Work	70
CHAPTER 5.....	71
GENERATION OF USER PROFILE USING KERALA DIFFERENTIAL APTITUDE TESTS (KDAT).....	71
5.1. Introduction to Kerala Differential Aptitude Tests.....	71
5.2. Enhancing DAT Interpretation through Automated Recommender Systems.....	73
5.3. Utilizing KDAT in the Proposed Career Recommender System	74
5.4. Normalization and Standardisation of User Profile Data.....	75
CHAPTER 6.....	79
CURATING A COMPREHENSIVE DATASET FOR THE PROPOSED CAREER RECOMMENDATION SYSTEM BASED ON DAT	79
6.1. An Overview of the Occupational Information Network Database.....	79
6.2. Ability Profile Information Related to Careers in O*NET	80
6.3. Mapping O*NET Ability Features to DAT Test Batteries	83
6.4. Converting O*NET Ability Features to DAT Test Batteries.....	83
6.5. Generating Career Clusters.....	86
6.6. Augmenting the Dataset to Address Overfitting.....	91
6.7. Evaluation of the Validity of the Augmented Dataset.....	93
CHAPTER 7.....	95
ARCHITECTURE OF THE PROPOSED CAREER RECOMMENDATION ENGINE	95

7.1. Ensemble Model Implementation Leveraging a Voting Classifier	96
7.2. The Innovative Hierarchical and Multi-Tiered Architecture of the Recommendation Engine	98
7.3. The Proposed System's Multi-Tiered Recommendation Approach	101
CHAPTER 8	103
THE HYBRID EXPLAINABILITY FRAMEWORK FOR THE PROPOSED RECOMMENDER SYSTEM USING SHAP WITH NATURAL LANGUAGE JUSTIFICATIONS	103
8.1. Using SHAP to Interpret Model Recommendations	103
8.1.1. <i>Global and Local Interpretation Using the SHAP Explainer</i>	106
8.1.2. <i>Accuracy Evaluation of the Proposed SHAP Interpretation Framework</i>	109
8.2. The Architecture of the Proposed SHAP-Based Explainability Framework with Natural Language Justifications	110
8.2.1. <i>Template-Guided Natural Language Generation of Local Shapley Explanations</i>	112
8.2.2. <i>Generating Feature-Driven Class Descriptions with Global SHAP Values</i>	114
8.2.3. <i>Generation of Natural Language Explanations via LLM by Synthesising Local and Item-Level Global Insights</i>	116
8.2.4. <i>Implementation of the Explanation Framework in the Proposed Career Recommender System</i>	117
8.3. Expert Evaluation of the Proposed System	118
8.4. Ethical Guidelines Followed in the Evaluation Process	119
8.5. The Overall Use-Case of the Proposed Hybrid Explainable Career Recommender	119
CHAPTER 9	121
DYNAMIC CAREER EXPANSION FOR FUTURE- PROOF RECOMMENDATIONS	121
9.1. Need for Dynamic Career Expansion Module	121
9.2. Architecture of the Dynamic Career Expansion Module	122
CHAPTER 10	125
RESULTS AND DISCUSSIONS	125
10.1. Introduction	125
10.2. Analysis of the Recommendation Engine's Outputs	125
10.3. Performance Evaluation of Recommendation Engine	128
10.4. SHAP Explainer Analysis: Interpreting the Recommendation Engine's Behaviour and Its Recommendations	132
10.4.1. <i>Evaluating Validity of the SHAP Explainer</i>	133
10.4.2. <i>Global Interpretation: Analyzing Model Behaviour Using SHAP Explainer</i>	134

10.4.3.	<i>Local Interpretations: Understanding Individual Career Recommendations</i>	137
10.4.4.	<i>Analysing Class-Specific SHAP Explanations</i>	141
10.5.	Analysis of the Final Natural Language Justifications Generated from the Hybrid Explainability Framework	145
10.5.1.	<i>Analysis of Textual Explanations for Individual Recommendations Using the Explainability Framework</i>	146
10.5.2.	<i>Analysis of Item-Level Textual Description Generated Via Global Class-Specific Shapley Values</i>	147
10.5.3.	<i>Analysis of Natural Language Explanation Generated Via Fusion Layer</i>	149
10.6.	Comparative Analysis: The Proposed Framework Against Existing Approaches	150
10.7.	Expert Evaluation of the Career Recommendations and Related Explanations	151
10.8.	Analysis of the Dynamic Career Expansion Module	152
CHAPTER 11		155
CONCLUSION		155
11.1	Summary of Findings	155
11.2.	Contributions of the Research	157
11.3.	Practical Implications	159
11.4.	Limitations of the Research	160
CHAPTER 12		163
RECOMMENDATIONS		163
REFERENCES		165
LIST OF PUBLICATIONS & CONFERENCE PROCEEDINGS		179
LIST OF CONFERENCES		179
Appendix A: Original Dataset (Sample entries from about 1000 careers collected from the O*NET, https://www.onetonline.org)		181
Appendix B: Generated DAT Dataset (Samples from selected 612 careers) - Published: https://doi.org/10.5281/zenodo.14192027		185
Appendix C: Generated Top-Level 8 Career Clusters		189
Appendix D: A sample from Hierarchical Clustering by sub-stratification of top-level clusters		193
Appendix E: Augmented Samples from Sub-Stratified Career Clusters		195
Appendix F: Program Code Snippets		199

LIST OF TABLES

Table 1. Evolution of Psychometric Testing Tools: Methodology, Contributions, Applicability, and Limitations	30
Table 2. Collaborative Filtering (CF), Content-Based Filtering (CBF), and Hybrid Systems	37
Table 3. Differential Aptitude Tests Batteries and Their Interpretations	69
Table 4. Test Batteries and Question Distribution in Kerala Differential Aptitude Tests	72
Table 5. Ability Parameters and their Interpretations in O*NET	80
Table 6. Mapping of DAT Test Batteries to O*NET Ability Profile Parameters.....	84
Table 7. The Process of Generating a Career Profile for the Profession of Lawyer.	85
Table 8. Samples from the Generated Dataset: DAT Parameter Values for a Few Careers	86
Table 9. Silhouette Scores for Different Clustering Algorithms.	87
Table 10. Training and Testing Accuracies of Various Models on the Original Dataset	90
Table 11. Performance of Different Classifiers on the Augmented Dataset.....	94
Table 12. Performance of Classifiers Trained on the Augmented Dataset Against the Original Dataset	94
Table 13. Accuracy Metrics for the Diverse Recommender Module and Refinement Recommendation Modules	129
Table 14. Precision Metrics for the Diverse Recommender Module and Refinement Recommendation Modules in the Proposed Model.....	130
Table 15. Recall Metrics for the Diverse Recommender Module and Refinement Recommendation Modules	131
Table 16. F1-Score Metrics for the Diverse Recommender Module and Refinement Recommendation Modules.	132
Table 17. Feature Importance Comparison Using Permutation Importance Analysis.	137
Table 18. Correlation Between Feature Values and Corresponding SHAP Values	148
Table 19. Comparison of the Proposed Explainability Framework with Existing Approaches.	151

LIST OF FIGURES

Figure 1. Internal Architecture of Deep Neural Networks	43
Figure 2. Machine Learning Models: Accuracy Versus Interpretability	43
Figure 3. The Workflow Diagram of the Proposed Study	58
Figure 4. High-Level Overview of System Development.	62
Figure 5. A Sample KDAT Test Output	73
Figure 6. Varying Interpretations of Career Counsellors on DAT Parameter Values.	74
Figure 7. Utilising KDAT in the Proposed Recommender System	75
Figure 8. Finding the Optimum Number of Clusters Using the Silhouette Method.	88
Figure 9. Finding the Optimum Number of Clusters Using Elbow Methods.....	89
Figure 10. Sample from the Generated Clusters.....	90
Figure 11. Samples from the Augmented Clusters.	93
Figure 12. Hard Voting vs. Soft Voting Classifier.	97
Figure 13. The Architecture of the Voting Classifier Adapted to the Proposed Recommender System.	98
Figure 14. The Hierarchical and Multi-tiered Architecture of the Proposed Career Recommendation Engine.....	100
Figure 15. Illustration of the Recommender System Design with its Recommendation Process.	101
Figure 16. The Hybrid Explainability Architecture of the Proposed Recommender System.....	104
Figure 17. Functional Comparison between ML Model and SHAP Explainer	105
Figure 18. A Shapley Value-Based Illustration: How Each Feature Value Shifts the Model's Base Prediction to the Final Prediction.....	106
Figure 19. Hybrid Framework for Explainable Recommendations	111
Figure 20. The Use-Case Diagram of the Proposed Hybrid Explainable Career Recommender.	120
Figure 21. The Architecture of the Dynamic Career Expansion Module	122
Figure 22. Career Recommendations from FirstPick Model	126
Figure 23. Career Recommendations from Second Pick Model	127
Figure 24. Career Recommendations from ThirdPick Model	127
Figure 25. Summary Bar Plot: Feature-Value Based Feature Importance Analysis.	135
Figure 26. Beeswarm Plot: Feature-Value Based Feature Importance Analysis....	136
Figure 27. Career Recommendations Made by the Model for the Given Candidate Profile.	139

Figure 28. Force Plot Illustrating How the Marginal Contribution of Each Feature Value in the Model's Prediction.....	139
Figure 29. The Decision Plot Illustrates How Different Feature Values Contribute to the Model's Prediction.....	140
Figure 30. Waterfall Plot Providing a Static Audit View of the Model's Reasoning	140
Figure 31. Summary Bar Plot of Career Cluster 200 Illustrating the Significant Features.....	141
Figure 32. Summary Beeswarm Plot Depicting the Feature Importance for Career Cluster Labelled 200.....	141
Figure 33. Feature Contributions to Individual Prediction: Class-Specific Analysis Using Decision Plot.	143
Figure 34. Decision Plot Explaining Why the Given Career Cluster May Not Align with the Candidate's Profile.	143
Figure 35. Recommendations Generated by the System Matching for a User Profile.	145
Figure 36. Force Plot Illustrating the Marginal Contributions of Features to the Model Prediction.	146
Figure 37. Summary Beeswarm Plot Illustrating Feature Importance for the Career Cluster Labelled 200.....	148
Figure 38. Mean Expert Ratings for Each Profile on a Five-Point Scale (1 = Not satisfactory, 2= Slightly satisfactory, 3 = Moderately satisfactory, 4 = Satisfactory,5 = Highly satisfactory).	152
Figure 39. ML Career Profile	153
Figure 40. Suggested Career Cluster for the ML Engineer Career Profile.....	153
Figure 41. Illustration of Dynamic Career Expansion.	154

LIST OF ABBREVIATIONS

AABT	:	Army Alpha and Beta Tests
AI	:	Artificial Intelligence
ANN	:	Artificial Neural Networks
APM	:	Advanced Progressive Matrices
AR	:	Abstract Reasoning
CARS	:	Context-Aware Recommendation Systems
CF	:	Collaborative Filtering
CBF	:	Content-Based Filtering
CPP	:	Core Positive Percentage
CKGE	:	Contextualized Knowledge Graph Embedding
DAT	:	Differential Aptitude Tests
DCEM	:	Dynamic Career Expansion Module
DNN	:	Deep Neural Networks
DPEBM	:	Differentially Private Explainable Boosting Machines
EBM	:	Explainable Boosting Machines
EI	:	Emotional Intelligence
ECTRS	:	Early Career Trajectory Recommender System
FP	:	False Positive
FN	:	False Negative
GAT	:	Graph Attention Network
GDPR	:	General Data Protection Regulation
GNB	:	Gaussian Naive Bayes
GPA	:	Grade Point Average
Grad-CAM	:	Gradient-weighted Class Activation Mapping
IQ	:	Intelligence Quotient
IT	:	Information Technology
KDAT	:	Kerala Differential Aptitude Tests
KG	:	Knowledge Graph
KNN	:	K-Nearest Neighbour
LDA	:	Latent Dirichlet Allocation
LIME	:	Local Interpretable Model-Agnostic Explanations

LRP	:	Layer-Wise Relevance Propagation
MBTI	:	Myers-Briggs Type Indicator
MI	:	Multiple Intelligences
ML	:	Machine Learning
MR	:	Mechanical Reasoning
MSCEIT	:	Mayer-Salovey-Caruso Emotional Intelligence Test
NA	:	Numerical Ability
NB	:	Naïve Bayes
NLP	:	Natural Language Processing
NLG	:	Natural Language Generation
NR	:	Numerical Reasoning
O*NET	:	Occupational Information Network
PSA	:	Perceptual Speed and Accuracy (PSA)
RF	:	Random Forest Classifier
RIASEC	:	Realistic, Investigative, Artistic, Social, Enterprising, Conventional
RPM	:	Raven’s Progressive Matrices
SCCT	:	Social Cognitive Career Theory
SHAP	:	Shapley Additive Explanations
SII	:	Strong Interest Inventory
SPA	:	Clerical Speed and Accuracy
SPM	:	Standard Progressive Matrices
SR	:	Space Relations
STEM	:	Science, Technology, Engineering, and Mathematics
SVM	:	Support Vector Machine
SWOT	:	Strengths, Weaknesses, Opportunities, and Threats
TF-IDF	:	Term Frequency - Inverse Document Frequency
TP	:	True Positive
TN	:	True Negative
WCSS	:	Within-Cluster Sum of Squares (WCSS)
XAI	:	eXplainable Artificial Intelligence

LIST OF SYMBOLS AND NOTATIONS

$ \phi_i^{(j)} $	: Absolute SHAP value of the feature i , for the j^{th} data point
$I(h_i(x)=c)$	: Indicator function that returns 1 if classifier i predicts class c for input x
$a(i)$	: Average distance from the i^{th} data point to other data points within the same cluster
$argmax_c f_c(x)$	: The class c that maximizes $f_c(x)$ for the input x
$b(i)$	: Smallest average distance from the i^{th} data point to data points in a different cluster
C_i	: Centroid of the cluster
DAT_{Lawyer}	: DAT score for the career <i>Lawyer</i>
$DAT_{userprofile}$	: Differential Aptitude Test Profile of a user
$f(x)$	: Model's prediction for a given instance x
$f_c(x)$	: Model's output score for class c for the given input instance x
f_S	: Model's prediction when feature i is not present along with subset S
$f_{S \cup \{i\}}$	: Model's prediction when feature i is present along with subset S
$g(x)$	: Approximation of the model prediction by the SHAP Explainer for the input x
$h_i(x)$	: Prediction of i^{th} classifier for input x
I_j	: Permutation Feature Importance value of the feature, j
I_i	: Importance score of feature, i
L_i	: Level score of the parameter, i
$M(X)$	: Model performance metric, accuracy
$P_i(c x)$	: Predicted probability of class c given input x by classifier i ,
S	: Subset of N features that does not include feature i ,
$S(k)$	: Silhouette Score for a specific number of clusters (k)
$StandardizedDAT_{userprofile}$	: Standardized DAT profile of a user profile
$WCSS_k$	: Within-Cluster Sum of Squares for a specific cluster count k
x	: Original value of the data point

x'	: The normalised value
X_{ij}	: Data point j in the cluster i
X_{perm}^j	: Dataset where feature j is randomly permuted
$x_{SU\{i\}}$	: Input values corresponding to features $SU\{i\}$
ϵ_i	: Random value drawn from a Gaussian distribution, for instance i
$\mu(X)$	: Average of scores across all test batteries, X
$\sigma(X)$	: Standard deviation of scores across all test batteries, X
ϕ_0	: Expected model output when no features are considered
Φ_i	: Shapley value for i^{th} feature
$\phi_i^{(j)}$	: SHAP value of feature i , for the j^{th} data point.
$Score_{\text{Level}}$	: Level score of a career in ONET database
$Score_{\text{Importance}}$	: Importance score of a career in ONET database
X_i	: Score in the percentile of i^{th} DAT test battery
$X_{\text{normalized}}$	: Normalized value of Importance or Level score of a career X
Z_i	: Standardized value for the i^{th} test battery
$\hat{y}$	: Predicted class with the highest vote count.

CHAPTER 1

INTRODUCTION

This section provides a comprehensive introduction to the study, addressing the background and context, highlighting the critical need for personalised career counselling and the challenges in current systems. It reviews the evolution of recommender systems within the career guidance domain in particular, identifies research gaps in current approaches, and introduces the novel solution proposed. Furthermore, it outlines the research objectives, scope, and limitations, followed by an overview of the methodology covering system architecture (encompassing the recommendation engine and the explainability framework), dataset creation, evaluation metrics, results, limitations, implementation details, deliverables, and ethical considerations, and finally concludes with the organisational structure of the thesis.

1.1. Background and Context

Investing in education is paramount for any country to build a skilled and innovative workforce. Governments allocate substantial resources to provide education to their citizens (Ramachandra, 2024); however, despite attaining even postgraduate qualifications, many students struggle to secure suitable employment. One of the primary reasons for this mismatch is that students often pursue courses misaligned with their abilities, skills, and career aspirations, leading to skill gaps and employability challenges (Wheebox et al., 2025). This underscores the critical need for personalised career counselling, ensuring that students make informed career choices based on their aptitude, competencies, and interests.

In a country like India, where the career counsellor-to-student ratio is significantly low (Balakrishnan & Karthiyayini, 2024; CG&AC State Co-ordinator, Directorate of General Education, Kerala, personal communication, November 6, 2024), delivering effective career guidance at scale remains a major challenge. A technological solution is essential to bridge this gap, enabling data-driven, personalised career

recommendations that can be scaled to reach every student across the country. This research aims to address these issues by integrating AI-driven technologies, leveraging machine learning-based recommendation systems to provide transparent, interpretable, and personalised career counselling.

1.1.1. Need for Personalised Career Counselling

It is important to note that every student should be given personalised career counselling addressing their strengths and weaknesses (Dan & Li, 2022). Students will be misguided due to parental pressure and peer group influence, and may select career paths and courses that do not align with their abilities and potential. Career selection has serious impacts on personal and professional life (Waworuntu et al., 2022), impact on society, and even the human resources of the nation, and the massive expense of spending for the next generation's education (Press Trust of India, 2024) should not be futile. Dropouts in higher educational institutions and job switches in the industry (Gallego et al., 2021; Kim & Beier, 2020; Perchinunno, Bilancia, & Vitale, 2021; Hoff et al., 2022) suggest that students are pursuing the wrong courses, not aligning their aptitude, vocational interests and personality traits. Career selection should align with the candidate's aptitude (Lang & Kell, 2020), personality traits (Mammadov, 2022) and vocational interest (Van der Linden et al., 2021; Waghmode et al., 2024; Kim & Beier, 2020).

Aptitude-based career and course selection is critical for a successful professional life. Aptitude refers to an individual's inherent ability to do certain tasks (Snow, 2021). It encompasses numerical ability, abstract reasoning skill, verbal reasoning, mechanical reasoning, spatial ability, verbal ability and many more (Bennett et al., 1947). Measuring and analysing these elements helps to identify the related career paths and courses to select in one's career life (Lohman & Nicpon, 2023). Misalignments between innate aptitude and career choices should be avoided, as a weak student in numerical ability cannot go for physics or engineering and even struggle with designing-related careers. Candidates having high logical reasoning skills when entering into the profession, and having repetitive tasks, might get bored and will eventually have to switch careers. Individuals likely perform better and achieve better

in areas that match their talents. That will provide better job satisfaction and reduce job switches (Jiang & Guo, 2022), making it good for the person, the employer and the nation.

1.1.2. Challenges in Existing Systems

The main challenge in an India-like country is the lack of enough career counsellors for the population of student candidates to provide personalised career counselling (Dheya Career Mentors, 2024). The student-to-career counsellor ratio is estimated to be 1000 to 1, compared to the recommended 250 to 1 (American School Counsellor Association, 2024). A country like India also has a huge Urban-Rural division. The main reason for this low career counsellor-to-student ratio is because of the huge student population and the lack of trained career professionals. Moreover, the Indian academic system focuses more on academics than on holistic career guidance (Economic Times, 2023). With the advancement of Machine learning, career recommendations can be automated, addressing different aspects of the candidate like aptitude, personality traits and vocational interests to provide a holistic approach to the career guidance process. An AI-based career recommendation engine can address the issue of a low career counsellor-to-student ratio in a country like India. The system can incorporate different aspects in measuring the candidate and can provide a holistic approach to career counselling by augmenting the current career counselling process or possibly replacing it.

Existing career recommendation engines have the issue of data sparsity and cold start limitations. Data sparsity refers to the missingness in the data, which will make it difficult to draw proper insights to develop mature models (Ahmadian et al., 2022). This is a common issue in data-driven applications built on machine learning models. This occurs because of limited user interactions against the large number of choices available. This will limit the recommendation engine's capacity to suggest all possible choices to a candidate. Cold-start refers to the challenge faced by the recommendation engine due to insufficient data to make accurate recommendations (Roy & Dutta, 2022). This can be because of an item-cold start due to new items interacting with the users, or a user cold start, as a new user cannot have historical data associated.

Existing career recommendation engines also suffer from a lack of transparency (Wang et al., 2024) because most of them have been built upon black boxes like Deep Neural Networks (Antamis et al., 2024). Recommendations from black boxes cannot be trusted, and career recommendations generated from a recommendation system should be transparent, as guiding a student based on the recommendation from a black box can be dangerous and should be dealt with to provide fair and unbiased recommendations. So, an explainable career recommendation system is essential to provide fair, transparent and more judicious career recommendations.

1.1.3. Evolution of Recommender Systems in Career Guidance

Recommender systems are used in various domains for information filtering for delivering personalized recommendations. They act as a tool for recommending content, items and services to users matching their tastes. The main targets of recommendation systems are personalization, decision support, information filtering, enhancing user engagements, increasing service to boost revenue, improving retention, reducing search cost, enhancing user experiences and many more (Ko et al., 2022). Recommender systems provide unique individual experiences by providing content, items and services to users' preferences (Gm et al., 2024; Wang & Hu, 2024). Recommender systems also support decision-making by suggesting the most suitable items from overwhelming choices (Zangerle & Bauer, 2022) using the information filtering technique. These systems can engage users, increasing their retention, which can be used to increase sales or views and thus increase revenue (Pajkovic, 2022). These systems support users in finding their items quickly, reducing searching costs. Users are provided with diverse items and services, increasing their experiences in respective domains. Recommender systems are extensively used in various domains like streaming platforms, social media, health care, finance, gaming, education and many more.

There are different types of recommender system frameworks applied, like content-based filtering, collaborative filtering, hybrid approaches, Ontology-based frameworks, and many more. Content-based filtering works on matching item characteristics to user preferences. It works on user-item interactions (Pazzani, 2007).

Collaborative filtering suggests items based on the interactions of similar users. Users with similar interaction patterns are grouped, and items are recommended to a user based on the preferences of others with similar behaviour (Koren et al., 2021). Ontology-based recommendation works by leveraging a structured representation of domain knowledge, where recommendations are generated based on the relationships between items, users, and associated concepts. It can provide context-aware and personalised suggestions (Arafeh et al., 2021). Hybrid recommendation systems are used to address the limitations of individual recommendation approaches. It is a blend of two or more recommendation techniques like collaborative filtering, content-based filtering, etc. (Bhaskaran & Marappan, 2023). It can provide diversity in recommendations, and scalability issues and can address cold-start problems.

Career recommendation systems have been developed using various frameworks (Kamal et al., 2024). However, most of these systems have primarily focused on limited domains, such as STEM, computer science, and information technology. Currently, few career recommendation systems comprehensively address all real-world career options. Some systems are course recommendation platforms (Xu et al., 2021) or job recommendation systems (Giabelli et al., 2021), and many face challenges such as data sparsity and cold-start issues (Panteli & Boutsinas, 2023).

Recommendation systems should not function as black boxes. Despite the high accuracy of Deep Neural Networks (DNNs), many recommendation systems are built on them, which inherently suffer from a lack of transparency due to their opaque nature (Dasari & Bhukya, 2022). To address this issue, explainable machine-learning models are introduced. These explainable frameworks are to provide transparency and fairness for their predictions. Since ML-based prediction systems are used in many critical domains like finance, health care, education, etc., it is highly important to ensure proper explanation for the predictions coming from these black boxes.

Career counselling is a critical process, especially when recommending careers to secondary school students. Recommendations from black-box systems are unacceptable, as proper explanation and evaluation are essential for ensuring trust and reliability. This highlights the need for an explainable career recommendation system

that provides transparent, fair, and justifiable recommendations. The proposed work aims to develop a career recommender system with a novel recommendation engine (recommender system) architecture and an innovative hybrid explanation framework featuring natural language justifications.

1.1.4. The Proposed Research Solution

The proposed system is a hybrid, explainable, hierarchical and multi-tiered career recommendation system based on content-based filtering. It is designed to deliver an interpretable, personalised, yet diverse set of career recommendations. Internally, the system employs a hierarchical structure to ensure diversity in the career options and provides natural language explanations for the recommendations with an innovative explainability framework.

The recommendation engine is built on a novel, hierarchical recommendation architecture that employs a suite of ML models with multi-tiered outputs to ensure diverse career recommendations aligned to the students' aptitude profiles. This approach presents candidates and career counsellors with a broad array of career choices aligned with their aptitude, helping them select careers that consider their personality and vocational interests.

To ensure transparency, fairness and trust, the system incorporates an innovative hybrid explainability framework built on Shapley Additive Explanations (SHAP) (Lundberg & Lee, 2017), which provides clear and interpretable insights into the recommendations with natural language explanations. The explainability framework integrates SHAP with NLG techniques, leveraging both local and global SHAP explanations and converting plot-based outputs into natural-language narratives that students can easily understand, thereby overcoming the technical complexity of traditional SHAP visualisations.

This work tackled the scalability challenge of the recommender system (Singh, 2020) by integrating a dynamic career expansion module, which, through its semi-supervised approach, enriches career clusters and enables real-time career recommendations in an ever-evolving career landscape.

The user profile is generated using differential aptitude tests, specifically the Kerala Differential Aptitude Test (KDAT) (Kerala Higher Secondary Education Department, 2024), developed by the Directorate of General Education, Government of Kerala. A career profile is generated using a dataset of 612 careers, the first of its kind. This dataset has been compiled by extracting ability profile information from the Occupational Network Information Database (O*NET) (National Center for O*NET Development, 2023).

The proposed system addresses the needs of students, educators, and the evolving job market. It helps bridge existing gaps in the career guidance process by offering a scalable and transparent solution. Furthermore, the research introduces novel architectures for recommender systems along with an innovative explainability framework to enhance their transparency and trustworthiness.

1.2. Relevance of the Study

This work tries to address India's severe career counsellor shortage by deploying an AI-driven, aptitude-based recommendation engine that generates standardized and objective career paths for each student. It tries to replace subjective assessments with data-driven algorithms and ensures scalable, transparent, and equitable career guidance.

This research encompasses the design of a novel hierarchical recommendation engine that generates a wide range of diverse career options tailored to candidates' aptitude. It further includes interpreting the rationale for these recommendations through an innovative SHAP-based natural language explanation framework, ensuring transparency, explainability and trust.

A first-of-its-kind career aptitude dataset comprising 612 real-world careers was curated as part of this research. This is achieved by exploiting the ability profile information from the O*NET dataset. This dataset enriches the model's ability to provide accurate and comprehensive career guidance.

This work significantly contributes to advancing educational technology and has the potential to impact the educational system in Kerala. The integration of this study's

results with the Kerala Differential Aptitude Tests (Dept. of General Education) has been formally agreed upon with an MoU, underscoring the system's relevance and applicability.

In short, this research addresses the following key challenges in the educational and professional domains:

- Tackles the critical shortage of career counsellors for secondary school students in a country like India and seeks to provide an accessible career support system for every student.
- Aims to reduce student dropout rates in higher education institutions and minimise job switches in the industry.
- Enhances the impact of government education funding by directing each student to the most suitable career path based on their skill set, fostering a creative and skilled workforce.
- Introduces an explainable career recommender system that offers diverse and transparent career options.
- Improves transparency, trust, and decision quality through natural-language justifications for recommendations.
- Develops novel architectures for both the recommendation engine and its associated explainability framework, adaptable to other domains.

1.3. Scope of the Research

- This research aims to develop an innovative, explainable and multi-tiered career recommender system targeting secondary school students in India, particularly in Kerala, to provide diverse career recommendations aligned with their skill sets, addressing the low career counsellor-to-student ratio as an attempt at an AI-powered solution to the educational environment, thereby fostering transparency and trust.

- The result of this work can be integrated into the existing aptitude testing tool, the Kerala Differential Aptitude Test (KDAT), developed by the Career Guidance and Adolescent Counselling Cell of the Government of Kerala (Kerala Higher Secondary Education Department, 2024), to automate its career counselling process. The tool KDAT is standardised, taking into account the socio-economic differences within the state to address bias (State-level official in Career Guidance and Adolescent Counselling, Directorate of General Education, official communication, November 5, 2024). Currently, KDAT is used to evaluate candidates for their aptitudes, while the counselling process is manual, which can be automated or enhanced by integrating the results of this work.
- A comprehensive career profile dataset consisting of 612 unique careers was constructed using the Occupational Information Network (O*NET) database, developed under the auspices of the U.S. Department of Labour. (National Center for O*NET Development, 2023). This enables the recommendation engine to provide career suggestions aligned with globally recognised professions, helping students navigate the global job market.
- To ensure adaptability to emerging occupations beyond those included in the training dataset, the system incorporates a semi-supervised dynamic career expansion module, enabling continuous enrichment of the career repository in line with the evolving job landscape, addressing the scalability issue.
- The methodological scope includes the development of a novel hierarchical, multi-tiered recommendation architecture, the implementation of an innovative explainability framework that generates textual explanations to ensure transparency and trust, and the creation of a first-of-its-kind dataset based on career aptitude profiling for the career recommendation engine.
- While the architectures of the proposed recommendation engine and explainability framework can be adapted to other domains, the career recommendation process is currently limited to measuring the candidate's

aptitude alone; it should be enhanced by incorporating personality traits and vocational interests.

- With respect to future research, this study provides multiple avenues for further exploration. It proposes a unique hierarchical and multi-tiered architecture capable of generating diverse yet related career recommendations, opening opportunities for advancement in recommender design. It also integrates a hybrid explainability module that combines SHAP with natural language explanations, offering a foundation for extending transparency frameworks in related domains. In addition, the study introduces a novel dataset for aptitude career profiling that can serve as a valuable resource for researchers seeking to develop innovative recommendation systems. Collectively, the study makes a significant contribution to the field and establishes multiple directions for subsequent investigations.

1.4. Limitations of the Research

Regarding the limitations and exclusions of this work, the career recommendation process would benefit from integrating candidates' vocational interests and personality traits to provide more holistic recommendations (Mammadov, 2022; van der Linden et al., 2021; Waghmode et al., 2024; Kim & Beier, 2020). The proposed framework focuses solely on aptitude-based recommendations, functioning as a hybrid man-machine recommendation system where the machine generates initial suggestions based on the candidate's aptitude alone, which are then augmented by a career counsellor based on the candidate's personality and vocational interests. Additionally, a notable limitation is the absence of real-world student feedback, which is important for evaluating and refining the system from the user's perspective. Since the proposed framework is designed for secondary school students at the Class X level, validating recommendation effectiveness through actual career outcomes is challenging because career progression develops over a long period. Therefore, large-scale longitudinal validation and incorporation of delayed career outcome feedback were beyond the scope of the present study.

Further, another limitation of the work is that while the system provides personalised and diverse recommendations for candidates, the diversity is currently achieved using the internal hierarchical structure of the recommendation engine. This diversity could be further enhanced by incorporating stream-based classifications, such as Science, Commerce, and Humanities. By integrating these classifications, the system could present candidates with recommendations from multiple stream-based career clusters that align with their profiles, thereby broadening the scope and relevance of the suggestions.

1.5. Research Gap

The traditional career counselling process faces several challenges that highlight significant gaps in existing systems. A major issue is the lack of quantitative profiling of candidates before recommending careers, which limits the precision and relevance of the suggestions provided. Considering the target group of the study, existing psychometric tools, such as the Kerala Differential Aptitude Test (KDAT), rely on manual counselling processes (State level official in Career Guidance and Adolescent Counselling, Directorate of General Education, official communication, November 6, 2024), making them insufficient to address the problem of the low career counsellor-to-student ratio prevalent in many contexts, particularly in India.

Recommendation systems face challenges in ensuring diversity in their recommendations, which is crucial in the context of career guidance, as career ambitions are not watertight compartments. This aspect is often lacking in existing implementations. (Castells, Hurley, & Vargas, 2021).

Additionally, most existing career recommendation systems operate as black-box models, lacking transparency in their recommendations. This opacity diminishes trust and prevents users from understanding the rationale behind the career suggestions. Furthermore, very few recommendation systems cater specifically to the needs of secondary school students, leaving a critical demographic underserved in career guidance. This can be addressed with explainable AI (Murdoch, Singh, Kumbier, Abbasi-Asl, & Yu, 2019)

Existing systems also face technical limitations such as data sparsity and cold-start problems (Natarajan, Vairavasundaram, Natarajan, & Gandomi, 2020), which hinder their ability to make accurate recommendations for new or underrepresented candidates.

Finally, there are limited attempts to recommend unseen or emerging careers in real-world scenarios (Mashayekhi, Li, Kang, Lijffijt, & De Bie, 2024). This lack of adaptability prevents existing systems from addressing evolving job markets and new-generation careers, leaving students ill-equipped for modern career landscapes. These gaps underscore the need for an innovative, transparent, and adaptable career recommendation system designed to address these critical issues.

1.6. Objectives of the Research

The objectives of this study include multiple aspects addressing issues in the existing career guidance process, gaps in career recommendation engines, proposing a new architecture, etc. The primary objectives are listed below:

- i. To address the low career counsellor-to-students ratio by ensuring personalised career counselling for every student, thereby supporting the education system in fostering a creative and productive workforce for the country.
- ii. To provide a diverse set of career recommendations aligned with the candidate's aptitude, addressing issues such as dropouts in higher educational institutions, job switching in industries, and enabling each candidate to pursue a passionate and fulfilling career.
- iii. To develop a hierarchical, multi-tiered recommendation architecture that progressively refines candidate options across multiple levels, ensuring both intra-tier diversity and inter-tier coherence, while maintaining strict alignment with user-specific aptitude vectors and feature representations.
- iv. To develop an innovative hybrid explainable framework for recommender systems that integrates SHAP algorithms with natural language generation

(NLG) techniques, providing comprehensive insights for both model developers and end users by explaining how and why the system generates specific recommendations through interpretable justifications, thereby delivering a transparent and trustworthy career recommendation system.

- v. To develop a first-of-its-kind dataset for career profiling, enabling the creation of career recommendation engines based on the aptitude requirements for each career, and supporting next-generation researchers.
- vi. To ensure a scalable career recommendation system by enabling it to efficiently handle growing data volumes and user bases while incorporating the capability to recommend unseen careers.

1.7. Methodological Overview

1.7.1. System Architecture

The career recommendation engine is built on a novel hierarchical and multi-tiered architecture. The system targets to provide diverse career paths for a candidate, aligning with their aptitudes. The explainability module features a novel hybrid architecture incorporating both visual and natural language explanations. Additionally, the architecture includes a dynamic career expansion module.

The recommendation engine features a hierarchical architecture internally, designed to ensure diverse career recommendations. Built on a content-based filtering approach, the engine incorporates multiple machine learning models that leverage a hierarchically clustered career dataset based on aptitude profiles to provide varied and comprehensive recommendations.

The system comprises a user profile module and a career profile module. To generate the user profile, the system utilises the Kerala Differential Aptitude Tests (KDAT), a differential aptitude testing tool developed by the Directorate of General Education, Kerala. The career profile module is constructed using a first-of-its-kind dataset, created by utilising ability profile information from the Occupational Information Network (O*NET) database. The recommendation engine is built upon this dataset.

The explainability module employs a hybrid framework consisting of a SHAP-based visual explanation layer and an NLG-based textual explanation layer. The visual layer applies SHAP's local and global attribution methods to quantify and visualise feature contributions, supporting expert users and model developers in analysing the recommendation logic. The textual layer uses an NLG pipeline to convert SHAP attribution values into clear justifications for non-expert users. Together, these components explain both how recommendations are produced and why specific careers are aligned with a user profile, explaining the rationale of recommendations, ensuring transparency and trust of the recommendation engine.

The dynamic career expansion module employs a semi-supervised architecture, enabling the integration of new careers into the system. Career counsellors can input new careers along with their aptitude profiles, which are then classified by the recommendation engine. The classified career is added to the appropriate cluster, ensuring it becomes part of the system's recommendations for candidates whose profiles align with the new career.

1.7.2. Dataset creation

To develop a career recommendation engine based on aptitude parameters, a dataset containing aptitude profile information is essential. While multiple career-related datasets are available (Kaggle, 2022; Hugging Face, 2024), none provide publicly available aptitude profiles for careers comprehensively. In this research, a career dataset with aptitude profile values was created using ability profile information from the Occupational Information Network (ONET) database. From over 1,000 careers listed in ONET, 612 unique and significant careers were selected. The dataset was then clustered to form related career groups, followed by augmentation (Rebuffi et al., 2021), within each cluster to prevent overfitting (Ying, 2019). The clustering was organised hierarchically to support the development of a hierarchical recommendation engine.

1.7.3. Performance Evaluation Metrics

Multiple evaluation metrics were employed at various stages of system development to ensure robustness and effectiveness. The silhouette score was used to assess the

quality of the clusters, while the elbow method was applied to optimise the total number of clusters. To evaluate the performance of the recommendation engine, metrics such as accuracy, precision, recall, and F1-score were utilised. Additionally, local accuracy metrics were applied to validate the performance of the SHAP-based explainer module, ensuring the reliability of its explanations. Expert evaluation was conducted with a group of career specialists to assess both the quality of recommendations generated by the engine and the explanations produced by the explainability module. The system was then validated and fine-tuned based on the resulting evaluation metrics.

1.7.4. Ethical Considerations

The system employs the Kerala Differential Aptitude Test (KDAT) for user profiling, which has been standardized to address socio-economic disparities within Kerala and does not retain any private or sensitive candidate data during operation. Similarly, the career profiling dataset used in this study contains no candidate-specific information, ensuring data anonymity.

The research includes a survey of career counsellors to examine subjectivity in aptitude parameter analysis across various careers. This survey was conducted following prior informed consent protocols, ensuring ethical compliance. Counsellor anonymity and confidentiality were maintained throughout the study, with no personal information stored or disclosed in any publications. The expert evaluation was conducted only after obtaining informed consent from all participants, and their personal data were kept anonymous throughout the study.

1.8. Organization of the Thesis

Chapter 1 introduces the research area, providing the background and context for the study. It highlights the need for personalised career counselling, examines the challenges in existing systems, and traces the evolution of career recommender systems, leading to the proposed research solution.

The chapter further discusses the research gaps, the scope and relevance of the study, the objectives of the research, and a brief overview of the limitations of the proposed

system. Additionally, it offers a methodological overview, covering system architecture, dataset creation, performance evaluation metrics and ethical considerations.

Chapter 2 reviews various studies related to the proposed system, analysing theories, methodologies, architectures, advantages, and limitations to develop an improved solution. The review focuses on key areas such as career counselling, recommendation engines, their limitations, explainable artificial intelligence, explainable career recommendation systems, and available datasets. The chapter also identifies gaps and limitations in existing studies, laying the foundation for proposing a more viable solution that addresses these shortcomings effectively.

Chapter 3 provides an overview of the methodology adopted to develop the proposed system. It includes a brief discussion on dataset curation, the generation of user-profiles and career profiles, the architecture of the recommendation engine, and the design of other key modules such as the SHAP explainer, textual explainer, and dynamic career expansion module. These areas will be discussed further in Chapter 4 onwards.

Chapter 4 delves into the theoretical foundation of the proposed career recommendation system. It explores various aspects of career counselling, including aptitudes, personality traits, vocational interests, and their associated theories. The chapter also examines differential aptitude tests, their historical context, early studies, and critiques, including discussions on the eugenics-related criticisms. Additionally, the chapter focuses on the different differential aptitude parameters, their interpretations, and their relevance to career selection, providing a comprehensive understanding of how these factors influence career guidance and recommendations.

Chapter 5 discusses the generation of user profiles to evaluate students' innate abilities. This chapter provides an in-depth examination of the Kerala Differential Aptitude Tests (KDAT), a differential aptitude testing tool developed by the Department of General Education, Government of Kerala. The research utilises KDAT for user profile generation. The chapter explores KDAT in detail, including its development, test batteries, test item generation process, items and their target

measurements, standardisation and bias mitigation, current utilisation, career recommendation processes using KDAT, and the rationale for its use in user profile generation within this study.

Chapter 6 provides a detailed discussion on the curation of the dataset used for career profile creation. It covers the mapping of different ability parameters for careers from the Occupational Information Network (O*NET) database to the test batteries of Differential Aptitude Tests (DAT). The chapter outlines the methodology adopted to convert the selected ability profile parameter values into corresponding DAT parameter values for each career. Additionally, it discusses the application of clustering algorithms to generate career clusters, the augmentation of records within each cluster to generalise the patterns, and the final validation of the augmented dataset to ensure its reliability. This validation ensures the applicability and effectiveness of the career recommendation engine built upon the curated dataset.

Chapter 7 focuses on the design and architecture of the recommendation engine. It provides an in-depth discussion of the internal hierarchical and multi-tiered architecture of the recommendation engine, explaining its design and functionality. The chapter highlights how the system generates diverse career recommendations from its multi-tiered output nodes based on a given candidate profile. Additionally, it examines how various machine learning algorithms interact with the dataset and evaluates each algorithm's performance, subsequently exploring the ensemble methods employed to construct the recommendation engine.

Chapter 8 discusses the architecture and functionality of the hybrid explainability framework designed for the recommendation engine. It provides details on the implementation of the SHAP explainer, highlighting its local and global interpretability capabilities. The chapter also explains the design of the NLG-based explanation module and how SHAP's local and global explainability features can be leveraged to generate comprehensive natural language explanations for the system's recommendations.

Chapter 9 discusses the architecture and functionality of the dynamic career expansion module. It explores its semi-supervised architecture, design, integration

into the system, and human-machine hybrid workflow. The chapter also covers the module's career cluster assignment process, dynamic adaptability for recommending new careers, and how it enhances the scalability of the recommendation engine.

Chapter 10 presents the results and discussion with performance evaluation metrics of the recommendation engine, including accuracy, precision, recall, F1-score, and others. It also highlights the local accuracy metrics of the SHAP explainer, demonstrating its effectiveness in interpreting the recommendations generated by the engine. It also discusses the results obtained from the recommendation engine, the experts' evaluation of its recommendations, and the explanations generated by the system's explainability framework, which together help ensure the transparency and trustworthiness of the recommendations. Further, this chapter provides a comparison of the proposed system with the studies reviewed in the literature survey; it also evaluates its validity against domain expertise and examines the explanations generated by the hybrid explanation module. Additionally, the chapter explores the use cases of the dynamic career expansion module, addressing the scalability of the system.

Chapter 11 concludes the study by highlighting its findings, key contributions and achievements across various areas, including use cases, social impact, technological advancements, practical implications, and innovative approaches in multiple domains such as recommendation engine frameworks, explanation frameworks, and dynamic recommendation frameworks. This chapter also offers insights into potential future directions for the research community.

Chapter 12 puts forward recommendations to make future improvements within the proposed system and provides suggestions to future researchers. The chapter gives insights into potential ways of improving many areas, including holistic career recommendation models, architectural enhancements of recommendation engines, explainability models and the dynamic career expansion module.

1.9. Conclusion

In conclusion, holistic and personalised career counselling at the early stages of education plays a pivotal role in ensuring career success, personal satisfaction, and optimal utilisation of human resources. However, in a country like India, the

imbalance between a large student population and a limited number of trained career counsellors presents a significant challenge. Advances in artificial intelligence and recommendation systems offer a promising solution by enabling the development of effective and personalised career guidance tools. Equally important is the integration of responsible AI practices, ensuring transparency and accountability in machine learning predictions. This study proposes a hybrid, explainable career recommendation system that addresses these challenges by delivering personalized and diverse career recommendations while integrating dynamic prediction capabilities to adapt to the ever-evolving career landscape, thereby ensuring relevance and sustainability in its applications. Additionally, the proposed innovative recommendation engine architecture and explainability framework can be adapted to other domains, as well.

CHAPTER 2

LITERATURE REVIEW

This section examines various aspects of career recommendation systems based on existing literature. It explores the theoretical foundations of career counselling, including the assessment of intelligence, its methodologies, and evolution. Additionally, it provides a comprehensive review of recommendation systems, tracing their development and key techniques, alongside their applications across diverse domains. The discussion further extends to current approaches in career recommender systems, highlighting the need for explainability, the methodologies employed, and existing implementations of explainable career recommendation systems. Finally, an overview of relevant datasets used in career recommendation research is presented.

2.1. Theoretical Aspects for Career Counselling

There are numerous theoretical frameworks related to career counselling. Donald E. Super's (1969) Developmental Theory views career choice and development as a lifelong process. It is influenced by experiences, skills, interests, and changing life stages. Central to this theory is the evolving self-concept, which reflects how individuals gradually gain new insights through their experiences. Super identifies five distinct stages of career development. They are: Growth (0-14 years), focused on forming self-concept, attitudes, and basic skills; Exploration (15-24 years), characterized by career awareness, skill-building, and tentative choices; Establishment (25-44 years), involving career stabilization and advancement; Maintenance (45-65 years), where individuals update skills and adapt to changes in the workplace; and Decline (65+ years), which marks reduced work engagement and the transition to retirement. The theory is valued for its dynamic, lifespan approach. The theory integrates both personal growth and situational factors, although it has been critiqued for paying insufficient attention to socio-economic influences and the ever-evolving nature of modern career paths.

Krumboltz's Social Learning Theory of Career Decision Making (Krumboltz, Mitchell, & Jones, 1976) suggests that choosing a career is a lifelong, evolving

process. It is not a single, fixed event, shaped by a complex interplay of four core influences. This dynamic journey is fundamentally determined by the interaction of Genetic Endowments and Special Abilities (inherited traits like intelligence and physical characteristics); Environmental Conditions and Events (external forces such as economic, political, and cultural influences); Learning Experiences (both through direct reinforcement, or instrumental, and observation, or associative learning); and Task Approach Skills (an individual's developed problem-solving abilities and work habits). The theory's greatest strength is its adaptability and comprehensive recognition of career development as fluid. However, its primary limitation lies in this very high degree of flexibility. This approach makes the model appear unstructured and presents a challenge in practical counselling when attempting to precisely measure and apply the influence of its distinct components.

Frank Parsons, widely regarded as the father of vocational guidance, proposed the Trait-and-Factor Theory (Betz, Fitzgerald, & Hill, 1989). The theory advocates matching a person's abilities, interests, and personality traits with suitable careers. His framework consists of three fundamental stages: evaluating one's traits, gathering detailed information about various occupations, and logically correlating personal attributes with job requirements. The model's simplicity and structured methodology make it easy to apply in career counselling contexts. Despite its practical strengths, the theory has been criticized for overlooking socio-economic factors and cultural variations, restricting its effectiveness in diverse or evolving work environments.

Lent, Brown, and Hackett (1994) developed Social Cognitive Career Theory, which offers a framework wherein people become interested in career, make educational and occupational decisions, and succeed in their careers. The theory focuses on several critical constructs such as self-efficacy (belief in having the ability to do something), outcome expectations (beliefs about the consequences of doing something) and personal goals (aspirations that direct actions). It also deals with interest development, career selection, performance, persistence, and environmental and cultural influences. The main benefit of SCCT is the overall framework, wherein it combines personal, environmental, and behavioural factors. The main limitation of this approach,

however, consists in that interrelationships between the constructs suggested by it can be complex and thus difficult to apply in reality.

The RIASEC Theory of Vocational Personalities and Work Environments in Holland has been one of the foundations of career counselling, and it categorizes people and work environments into six, namely: Realistic, Investigative, Artistic, Social, Enterprising, and Conventional (Hogan, R., & Blake, R., 1999). This system represents the types on a hexagon, holding that nearest types are most compatible, and opposites least. The RIASEC model has been extensively tested and is used in career exploration as it is a fast but structured method of matching individuals to appropriate careers. Its weaknesses, however, are that it tends to simplify complex human interests into rigid categories and therefore may not necessarily capture the dynamic aspect of personal development. Moreover, the model is commonly criticized as not sufficiently considering some key success variables such as socio-economic disparities and personal aptitudes and skills required in a successful career.

In the context of getting a personality understanding, which can be applied in areas such as career development, there are two key models: Myers-Briggs Type Indicator (MBTI) (Boyle, 1995) and the Big Five (McCrae & Costa, 2003). The MBTI is based on the ideas of Jung and is immensely popular because it is a simple test; it classifies someone into one of 16 different types using four simple dichotomies, like Introversion (I) vs. Extroversion (E), Sensing (S) vs. Intuition (N), Thinking (T) vs. Feeling (F), Judging (J) vs. Perceiving (P). This makes it an excellent one to reflect on and encourage dialogues. The more academic Big Five model on the other hand is deemed to be very reliable and valid by researchers. Rather than putting you in a box, it evaluates your personality on five major dimensions: Openness, Conscientiousness, Extraversion, Agreeableness and Neuroticism. Although the Big Five has a solid evidence base to support its predictive of real-world results and remains applicable across cultures, it is not as easy or intuitive as the typical person might think. Therefore, the career counsellor himself has to decide between the ease of use offered by the MBTI and the empirical rigour presented by the Big Five.

The Circumscription and Compromise Theory by Linda Gottfredson is an excellent concept used to explain the process by which people shape and narrow down their career options based on the interplay of both self-concept and socio-economic status (Gottfredson, 2005). She suggests that career decision making is a process consisting of two correlated processes; circumscription, the process of increasingly disqualifying the options that do not conform to the self-image and compromise, the process of realigning hopes and desires with external reality, including social accessibility and opportunity constraints. Gottfredson breaks down this process into developmental stages influenced by changing perceptions, beginning with recognition of size and power in early childhood, recognition of gender role in middle childhood, recognition of social class and prestige in adolescence and finally self-evaluation of interests and competencies in later adolescence and adulthood. Her theory has provided a cohesive perspective on the development of vocations through the focus she placed on cultural, social and psychological aspects of vocational development as adaptive and situational. However, critics note that some of its elements are not empirically accurate and overemphasize societal norms at the expense of personal agency.

These conceptual frameworks present the career counsellor with an understanding of how to conduct career counselling as well as being aware of the strengths and weaknesses of each concept. The Trait-and-Factor Theory developed by Frank Parsons is a fairly broadly-based foundational strategy, simple to apply but does not answer socio-economic issues and considerations when prescribing a career. The Developmental Theory of Super, which emphasizes career development as a dynamic and lifelong process and includes it as one of its major strengths, nevertheless cannot respond to socio-economic influences and cannot offer the means to adapt to the trends of current career development patterns. The RIASEC Theory developed by Holland is popular, tested and simple to apply but it is constrained because it is inflexible in its classification of the types of personalities and fails to take into account changing interests. The Social Learning Theory by Krumboltz integrates both individual experiences and environmental factors in career-building but is difficult to use because of its excessive malleability. In the same spirit, the Circumscription and Compromise Theory by Gottfredson can be used in effectively incorporating socio-

cultural aspects and developmental levels, although it relies so heavily on the influence of social factors that it cannot be practically applied. Finally, the Social Cognitive Career Theory of Lent is more holistic in that it takes into account the aspects of personality, environment, and behaviour, but it is still a challenge to implement Lent. The combination of these theories helps establish a strong basis of career counselling, emphasizing the necessity of combined efforts that focus on individual and contextual aspects of career counselling in dynamic and inclusive modes.

Based on theories of career development, it is evident that career assessment and career advice to candidates needs a holistic psychometric framework which goes beyond any simplistic matching. This theoretical overview highlights the complexity of career paths as multi-factorial, implying that no single technological solution can be effective. Rather, an effective strategy should critically assess aptitude, vocational interests, and personality and understand that these characteristics are dynamic and change over time. Although it is quite difficult to completely incorporate all the ideas presented in theory into a basic technological solution, any progress in the field of machine learning can hold great potential for developing a complex psychometric instrument to be able to fully incorporate these evolving aptitudes, personality characteristics, and interests to provide more accurate career advice.

2.2. Assessing Intelligence and the Evolution of Psychometric Tools in Career Counselling

Measuring intelligence can be traced back to 1869, particularly the Hereditary Genius by Francis Galton (Galton, 1891). Subsequently, Alfred Binet and Théodore Simon created the original intelligence test, the Binet-Simon Scale (Schwegler, 1914; Houcan and Kan, 2023), which was aimed at identifying students in France in need of special education. The Intelligence Quotient (IQ), or the ratio of mental age to chronological age, was introduced by William Stern (Stern, 1914; Cherry, 2022). Later, the Binet-Simon Scale was modified by Lewis Terman as the Stanford-Binet Intelligence Scale, which was an important improvement to the intelligence testing instrument (Terman, 1916; Pomplun & Custer, 2005). The Army Alpha and the Army

Beta Tests were used in the United States during World War I to test military recruits (Terman, 1918). The Differential Aptitude Tests (DAT), later by George K. Bennett, Harold G. Seashore, and Alexander G. Wesman (Bennett, Seashore, and Wesman, 1947) were designed to assess the different aptitudes of the candidates. Howard Gardner introduced a more comprehensive view in his theory of Multiple Intelligences (Gardner, 1987), which includes linguistic, logical-mathematical and spatial intelligence among others. In addition to intelligence testing, other tests appeared to test personality and vocational interests including Myers-Briggs Type Indicator (MBTI) (Briggs, 1976; Boyle, 1995) to test personality types and Strong Interest Inventory (SII) (Hansen and Bourne, 1988; Carter and Swanson, 1990) to test vocational interests, improved by John Holland as RIASEC model (Holland, 1959; De Fruyt and Mervielde, 1997). All of these tools demonstrate how psychometric methods have evolved to understand human potential and how to direct career development.

Hereditary Genius by Francis Galton (Galton, 1891) is an early study of human intelligence. He claimed that intelligence is genealogical and inherited. He suggested a radical notion of selective breeding towards a superior human race. He employed statistical tools such as correlation and regression as a measure of hereditary traits. He assessed the Blacks as being less intelligent than the Whites because of the difference in genetic makeup. His suggestion was attacked on the grounds of its eugenic view. His methodological approach of assessing intelligence based on genes and heredity alone, without other influences such as socioeconomic and cultural differences, was not acceptable (Constable, 1905). In spite of these valid criticisms, his work is regarded as having formed the basis of the study of intelligence.

In 1905 Alfred Binet and Théodore Simon presented the Binet-Simon Scale, which became commonly known as the first formal intelligence testing instrument (Binet and Simon, 1961). The scale was developed to recognize French school children who require special education, and it contains consecutive assignments that evaluate cognition, problem-solving, memory and attention. Items were placed in a progression of difficulty which permits categorization of intelligence levels using the idea of

mental age. Another significant advantage of the scale was the focus on practical thinking skills, as opposed to memorisation, and the variety of question types, addressing various fields. However, this test had numerous limitations: it was first validated on a small sample of French, the standardisation methods were not developed, cultural bias was observed, and only mental skills were evaluated, disregarding the wider skill range. All these aspects highlight the innovative importance as well as the natural limitations of the Binet-Simon Scale in the early years of intelligence testing.

Intelligence Quotient (IQ) originated with William Stern, modifying the concept of Mental Age introduced by Binet-Simon (Stern, 1914). He standardised intelligence in terms of the formula:
$$IQ = (\text{Mental Age} / \text{Chronological Age}) \times 100$$
in this way establishing a standard measure of intelligence. It offered a standardised number score as a measure of intelligence. Some criticism involves the test not working on at least adults at other ages, since after some age, intelligence does not change. It also gauged intelligence through one lens as opposed to the wider notion of multiple intelligences (Davis, Christodoulou, Seider, & Gardner, 2011).

With the Stanford-Binet Intelligence Scale (1916), a complex development of the Binet-Simon scale, Lewis Terman designed an elegant leap forward in psychometrics by introducing the notion of the Intelligence Quotient (IQ), independently developed by William Stern (Terman, 1916). Under a strict standardization process, Terman created the norms which, to the first time, would not be restricted only to children but also to adult populations. Even then, the Stanford-Binet, despite its innovativeness, received substantial criticism because of its strong cultural bias as its normative sample was composed of mostly white, native-born Americans which invalidated the use of the scale with other ethnocultural groups (Jensen, 1980). Moreover, its unidimensional emphasis on an individual IQ score was denounced as a reductionist measure that did not reflect the multidimensional and multifaceted quality of human intelligence.

During world war I, Alfahari Yerkes and a team of psychologists designed the Army Alpha and Beta Tests (AABT), which had capabilities of mass assessment; this

examination was mainly employed in army recruiting (Yoakum & Yerkes, 1920). The Alpha tests, which tested literate recruits, tested verbal and numeric reasoning, general knowledge, practical judgment, etc. The Beta tests, in their turn, were designed to test and filter out the non-literate candidates, based on their ability to recognize patterns, analyzing cubes, mazes, and the like. The tests were standardised and could be applied to groups of people at the same time. AABT, similar to most psychometric testing instruments, was not without its fair share of critics who accused it of being socio-economically biased, which, by proxy, disfavored black American populations.

The Wechsler-Bellevue Intelligence Scale is regarded as the first intelligence assessment test that targets adults to determine both the verbal and non-verbal cognitive skills (Wechsler, 1940). It was a breakthrough of the child-centred tests of Stanford-Binet as it offered an inter-dimensional approach to intelligence testing. Whereas the verbal tests evaluated vocabulary, comprehension, arithmetic, and similarities, the non-verbal tests tested block design, object assembly, picture completion, among others. The intelligence scale also proposed a deviation IQ, unlike the ratio IQ in the Stanford-Binet, which can be standardised by age. The shortcomings were cultural bias, fixed measurement, and a time-consuming management procedure. Although the Wechsler-Bellevue Intelligence Scale borrowed several components of the Binet-Simon Tests and the AABT, its greatest input was the deviation IQ (Boake, 2002).

Developed by George K. Bennett, Harold G. Seashore and Alexander G. Wesman, the Differential Aptitude Tests is a set of test batteries that measure numerous aptitudes, including numerical ability, abstract reasoning, 15, mechanical reasoning, space relations, language use (spelling and grammar), and clerical speed and accuracy. The tests develop an in-depth aptitude identities of a candidate and help him in the choice of suitable courses and careers. Cultural bias, static fixed measurement and time-consuming administration are the limitations (Wang, 1993).

Theory of Multiple Intelligences (MI) was presented by Howard Gardner in his famous book *Frames of Mind* in 1983. Gardner outlined eight types of intelligence: Linguistic, Logical-Mathematical, Musical, Bodily-Kinesthetic, Spatial,

Interpersonal, Intrapersonal and Naturalistic (Gardner, 2011). He has also developed another intelligence field known as existential intelligence. Gardner employed a qualitative over a quantitative cognitive measurement. Its benefits are: since it is a holistic method, it can be adapted to teaching and learning activities and not to career counselling, but it has some relevance in counselling process. His methodology is interdisciplinary and comprehensive. Weaknesses include the unavailability of empirical evidence, testing difficulties and consequently inapplicability in real-life situations (Waterhouse, 2006).

James R. Flynn has discovered a phenomenon known as the Flynn Effect, describing the steady increase in IQ test scores between generations (Flynn, 1984). The finding suggests that IQ scores rose about 3 points per decade on different tests. This is because of longitudinal studies of various IQ tests. There are more factors involved in increasing the IQ scores such as better education, nutrition, and changes in society instead of genetic evolution. The dynamic IQ scores contest the traditional approach of intelligence testing as a fixed level. All psychometric tests must be revised with new norms and standardisation procedures to be correct and relevant to stay valid.

Emotional Intelligence (EI) was proposed by John Mayer and Peter Salovey as the capacity to be aware of, control, and shape emotions to achieve better individual and interpersonal performance (Salovey, Mayer, 1990). Their EI model has four components: Perception of Emotions, which is the ability to identify emotions correctly; Facilitation of Thinking, which is the ability to use emotions to enhance thinking and problem-solving; Understanding of Emotions, which is the ability to perform a correct identification of the cause and progression of emotions; and Management of Emotions, which is the ability to manage emotions adequately in order to achieve goals. Their stance is based on emotion as a cognitive process, and heavily influenced by the study of cognition and psychology. They also invented the Mayer-Salovey-Caruso Emotional Intelligence Test (MSCEIT), an instrument to objectively assess EI as a set of skills. EI can be applied to various age groups, both in observation and building the background of basic emotional skills in children and with regard to academic achievement in teenagers and workplace relationships in

adults. EI in career counselling can be used to assist job seekers in their pursuit of emotionally flexible jobs. There are, however, constraints such as the inaccuracy of the measure of EI despite the use of tools such as the MSCEIT. EI has been accused of not having a standardised and universally accepted method of measurement (Murphy, 2014).

C. Raven, developed and revised his Progressive Matrices (RPM) to assess non-verbal reasoning, including abstract reasoning as a measure of general intelligence (Raven, 2000). It has a Standard Progressive Matrices (SPM) that is used by the general population and an Advanced Progressive Matrices (APM) that is used in academic and research environments. Its non-verbal pattern was meant to deal with cultural bias. However, its evaluation solely on abstract reasoning gives it a limited scope in measuring intelligence. A summary of the evolution of different psychometric testing tools is provided in Table 1.

Table 1. Evolution of Psychometric Testing Tools: Methodology, Contributions, Applicability, and Limitations

Psychometric Test/Concept	Year	Methodology	Contribution	Applicability as a Testing Tool	Limitations
Francis Galton's Hereditary Genius	1869	Statistical analysis of genealogical data to assess hereditary traits.	Introduced statistical tools like correlation and regression; first systematic study of intelligence.	Foundation for studying intelligence and its heritability.	Focused solely on genetics, ignoring socio-economic and cultural factors; eugenic bias.
Binet-Simon Scale	1905	Age-graded tasks assessing reasoning, memory, and problem-solving skills.	First practical intelligence test; introduced concept of <i>mental age</i> .	Identifying children needing special education.	Limited scope to children; lacked standardisation; cultural bias.
Intelligence Quotient (IQ)	1912	Ratio of mental age to chronological age, multiplied by 100 (William Stern).	Provided a standardised numerical measure for intelligence.	Widely used in educational and psychological assessments.	Static ratio ineffective for adults; oversimplifies intelligence.

Stanford-Binet Intelligence Scale	1916	Revised Binet-Simon Scale; introduced standardisation and deviation IQ.	Adapted IQ testing for the American population; applicable for adults.	Educational placement, psychological research.	Cultural bias; focused heavily on academic intelligence; time-consuming.
Army Alpha and Beta Tests (AABT)	1917	Group-administered written (Alpha) and non-verbal (Beta) tests for military recruits.	First large-scale group intelligence test; addressed literacy challenges with the Beta test.	Military personnel selection and classification.	Cultural bias; limited generalisability; oversimplified intelligence measurement.
Wechsler-Bellevue Intelligence Scale	1939	Combined verbal and non-verbal subtests; introduced deviation IQ.	First test tailored for adults; provided a multidimensional view of intelligence.	Educational and clinical assessments, workplace evaluations.	Cultural bias; static measurement; time-intensive.
Differential Aptitude Tests (DAT)	1947	Battery of tests measuring specific aptitudes (e.g., numerical, verbal, spatial, mechanical).	Focused on vocational guidance and career suitability.	Career counselling, educational planning.	Narrow focus; static assessment; cultural bias in test content.
Multiple Intelligences (MI)	1983	Framework identifying eight (later nine) distinct intelligences.	Expanded understanding of intelligence to include non-academic abilities (e.g., musical, kinesthetic).	Education, personalised teaching strategies.	Lack of empirical validation; no standardised measurement tool; theoretical overlap.
Emotional Intelligence (EI)	1990	Ability-based model measuring perception, understanding, and management of emotions.	Highlighted role of emotions in personal and social functioning; MSCEIT developed for measurement.	Leadership development, workplace training, career counselling.	Measurement challenges; overlaps with personality traits; cultural and contextual variability.
Flynn Effect	1984	Analysis of IQ trends over decades, showing generational increases in IQ scores.	Challenged static view of intelligence; emphasised environmental impact on	Highlighting the need for updating psychometric test norms over time.	Does not explain uneven gains across cognitive domains; plateau/reversal in recent trends.

			cognitive abilities.		
Raven's Progressive Matrices	Original: 1936 (Revised later)	Non-verbal test using matrix patterns to measure abstract reasoning.	Minimised cultural bias; focused on general intelligence (g-factor).	Widely used in cross-cultural settings, research, and educational contexts.	Solely measures abstract reasoning; static snapshot of intelligence; and ceiling effects for high performers.

2.2.1. Limitations of Psychometric Tests and Possible Solutions

Almost all psychometric tests exhibit cultural bias, making them ineffective when applied to culturally diverse populations from those on which they were standardised. This issue can be addressed by adapting the tests, including modifying test items within test batteries, removing items that function differently across cultural groups, or generating cross-cultural testing frameworks (Fernández & Abe, 2018). Another limitation of psychometric tests is their static measurement of a predefined set of abilities. This can be addressed either by updating the ability categories over time or by developing dynamic testing methodologies. Test validity and reliability are another concern (DeVon et al., 2007), necessitating rigorous validation studies to ensure accuracy and consistency. Furthermore, the limited scope of individual psychometric tests highlights the need for a more holistic approach, considering multiple dimensions such as aptitudes, personality traits, vocational interests, and emotional intelligence.

Personality traits can be measured using psychometric tests like the Myers-Briggs Type Indicator (MBTI) and Big-Five personality tests (Furnham, 1996). MBTI is based on Carl Jung's psychological types (Introversion (I) vs. Extroversion (E), Sensing (S) vs. Intuition (N), Thinking (T) vs. Feeling (F), Judging (J) vs. Perceiving (P)). Based on these, it defines sixteen different personality types. MBTI is criticised for its lack of empirical support. The Five-Factor model defines five broad traits: Openness to Experience (O), Conscientiousness (C), Extraversion (E), Agreeableness (A) and Neuroticism (N). Both MBTI and Five-Factor use self-report questionnaires

with traits. The five-factor method has been widely validated and supported by extensive psychological research (Stein & Swan, 2019).

Vocational interests can be tested using tools like the Strong Interest Inventory (SII), based on Holland's RIASEC model, which measures vocational interests across six categories: Realistic, Investigative, Artistic, Social, Enterprising, and Conventional (Carter & Swanson, 1990). The SII employs self-reported questionnaires with a five-point scale. The advantage of SII is its applicability in career counselling, as it aligns an individual's interests with occupational profiles.

Thus, a holistic framework for career counselling based on psychometric tools can be derived by using the Differential Aptitude Tests (DAT) for aptitude measurement, complemented by a Five-Factor Model (Big Five Personality Test) for assessing personality traits, and a tool like the Strong Interest Inventory (SII) (The Myers-Briggs Company, n.d.) for measuring vocational interests. Career recommendations generated from a DAT-based tool can be enhanced by incorporating insights from the Five-Factor Model and SII, ensuring a more comprehensive and personalised approach to career counselling.

In this study, the career recommendation system will be built on the Differential Aptitude Tests (DAT). DAT is specifically selected as it measures both verbal and non-verbal reasoning skills of candidates. DAT includes a comprehensive array of eight distinctive test batteries that assess various areas of aptitude, making it more robust compared to other psychometric intelligence tests like Multiple Intelligences, Emotional Intelligence, and Raven's Progressive Matrices. While Multiple Intelligences provides a holistic approach to intelligence analysis, it lacks a standardised measurement framework. Similarly, Emotional Intelligence lacks a standardised measurement approach and makes it challenging to implement in an Artificial Intelligence system. Raven's Progressive Matrices, on the other hand, has limited utility due to its exclusive focus on abstract reasoning alone.

On the other hand, Differential Aptitude Tests are well-defined, offering a comprehensive aptitude measurement approach. Its measurable characteristics are relevant to career paths and aligned with artificial intelligence systems. Limitations of

Differential Aptitude Tests, such as bias, static measurement, and limited scope, will be addressed in this study. To address bias, we utilise the Kerala Differential Aptitude Tests (KDAT), a DAT tool developed by the Directorate of General Education, Kerala. KDAT has been standardised across various socio-economic groups within the X standard population in Kerala. To overcome the issue of static measurement, the study advocates using DAT alongside other career path selection criteria, such as vocational interests and personality traits. Essentially, the career recommendations by our proposed system should be augmented by these additional factors, with inputs from the candidates and the career counsellor. In the future, these factors could also be automated and integrated into the recommendation system.

2.3. Recommender Systems: A Comprehensive Review on Techniques, Challenges and AI-Driven Approaches

Recommendation systems are algorithms designed to provide personalised suggestions for relevant items by analysing user preferences and behaviour (Resnick & Varian, 1997; Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994). These systems play a crucial role in decision-making (Chen et al., 2013) by offering diverse choices and enhancing operational efficiency (Ge et al., 2010; Konstan & Riedl, 2012) through the automatic delivery of recommendations and support. Additionally, they can be employed to improve user retention and engagement (Tarnowska et al., 2020), among other applications. They utilise machine learning techniques to filter and predict items and are applied across various domains, including video recommendations, healthcare, medicine, and career counselling. Examples include Amazon.com (Amazon, n.d.) for books and other products, Netflix (Netflix, n.d.) for streaming video, MovieLens (MovieLens, n.d.) for movies, Last.fm (Last.fm, n.d.) for music, Google News (GoogleNews, n.d.) for news recommendations, Google Search (Google, n.d.) for advertisements, YouTube (YouTube, n.d.) for online videos, and TripAdvisor for travel products.

2.3.1. Evolution of Recommender Systems

The initial development in recommendation approaches was the Collaborative Filtering (CF) technique, introduced by Paul Resnick and first implemented in the

GroupLens project (Resnick et al., 1994) for news recommendations. The methodology used in Collaborative Filtering is to recommend items to users based on the item interactions of similar users. The key feature of Collaborative Filtering is its user-centric approach, which builds on the idea that an item liked by a user can be suggested to similar users. For example, if users X and Y like item I_1 , and if user X also likes item I_k , then user Y may also like item I_k . CF does not require domain knowledge of items for recommendations; it relies solely on user-item interactions and user similarity.

Collaborative Filtering is categorised into user-based and item-based Collaborative Filtering. User-based CF finds similarities between users using metrics like cosine similarity (Liu et al., 2014) or Pearson correlation coefficients (Koochi & Kiani, 2016). However, user-based Collaborative Filtering faces scalability issues (Kumar & Sharma, 2013). Item-based CF suggests items that are frequently rated similarly by users. This approach reduces computational costs compared to the user-based Collaborative Filtering method (Sarwar et al., 2001). Despite its advantages, Collaborative Filtering suffers from data sparsity, the cold start problem, along with scalability issues (Guo, 2012).

To address some of the limitations of Collaborative Filtering, the Content-Based Filtering (CBF) technique was introduced in recommender systems. One of the earlier influential works in this area is attributed to Michael Pazzani (Pazzani & Billsus, 1997; Pazzani, 2007). In Content-Based Filtering, items are presented using a set of features called attributes. The user profile is created by analysing user-item interactions and the features of interacted items are aggregated to create the user profile. Users are recommended items based on user-profile item-profile similarities which is calculated using cosine similarity or Euclidean distance. Items are recommended based on the ranking of similarity metrics. Content-based filtering provides better-personalised recommendations compared to CF due to its inherent methodology. The user cold start issue is resolved using some initial user-item interaction techniques. Item cold start issue persists to recommend new items to the

user. CBF may advocate the same types of items due to its overspecialization problem (Thorat et al., 2015).

To address the limitations of individual recommendation approaches, hybrid recommender systems have been designed by combining two or more techniques, such as Content-Based Filtering and Collaborative Filtering. Integrating multiple techniques can improve accuracy, scalability, and robustness. Robin Burke's initial work on hybrid recommendation techniques (Burke, 2002) discussed different methods to implement hybrid recommender systems. In weighted hybridization, the weighted sum of scores generated from each recommender system is taken to provide the final recommendation. In the switching approach, based on a specific condition, recommendations are generated from one of the recommender systems. In the feature combination approach, features from multiple models are combined to create a uniform model. In the cascade hybridization model, recommendations generated from one model are fed to the next to refine the recommendations. In the feature augmentation approach, the output of one model is used as input to the next model as an additional feature. In meta-hybridization, one method is used to train the model for another, such as using CBF to develop a user profile for CF. While hybrid recommendation systems address critical challenges like cold start, data sparsity, and over-specialisation, their implementation is challenged by high complexity and significant resource requirements. Table 2 provides a comparison of the three CF, CBF and hybrid approaches discussed.

Recent advances in recommender systems have brought other directions, such as Context-Aware Recommendation Systems (CARS) (Adomavicius and Tuzhilin, 2010). Context-aware recommendation systems use contextual data like time, place, and social context in the recommendation process to make it more relevant and personalized. Several approaches have been adopted to include contextual information, which include pre-filtering, post-filtering and contextual modelling. In pre-filtering, the dataset is filtered using context, and then model is built around it. In post-filtering approach, once the model has been constructed, context is used to filter the recommendations created by the recommendation engine. Contextual modelling is a method involving the use of contextual information as an additional feature in model building. Context-Aware recommendation methods can offer better

personalization and higher user satisfaction, but are complex and require that the contextual information is of high quality.

Table 2. Collaborative Filtering (CF), Content-Based Filtering (CBF), and Hybrid Systems

Aspect	Collaborative Filtering (CF)	Content-Based Filtering (CBF)	Hybrid Systems
Personalisation	Leverages user-item interactions, providing personalised recommendations based on collective user behaviour.	Builds a profile unique to the user by analysing item attributes they interacted with.	Combines multiple techniques for greater personalisation, balancing strengths of CF and CBF.
Serendipity	Recommends novel items by leveraging preferences of similar users.	Offers accurate but narrow recommendations closely aligned with the user's known preferences.	Ensures both serendipity and relevance by integrating CF's novelty and CBF's accuracy.
Cold Start (Items)	Struggles with recommending new items due to lack of interaction data.	Capable of recommending new items using item metadata or features.	Partially mitigates item cold start by using CBF for new items and CF for established ones.
Cold Start (Users)	Performs poorly for new users with no interaction history.	Can handle new users better through initial profiles, such as explicit ratings or preferences.	Mitigates new-user cold start by combining demographic, content-based, and collaborative data.
Data sparsity	Struggles with sparse interaction matrices, limiting recommendation quality.	Less affected as it depends on item attributes rather than interaction data.	Reduces the impact of sparsity by leveraging multiple data sources and methodologies.

There are deep learning implementations also in recommender systems. He et al. (He et al., 2017) suggested a Neural Collaborative Filtering model based on deep neural networks. Generalized matrix factorization was included in the model to learn linear relationships with a multi-layer perceptron to learn non-linear interactions and presented a framework based on deep neural networks that can enhance personalized recommendations. The costs to implement this approach are hyperparameter tuning and the cost of computation. A similar (session-based) recommendation method based on recurrent neural networks was suggested by Hidasi (Hidasi, 2015).

Another implementation of a recommender system is by employing knowledge graphs (Wang et al., 2019). Knowledge Graph-Based Recommendation Systems represent objects (things, users, categories) as nodes and their connection to each other as edges. As an example, entities can be movies, directors and genres and relationships using “directed by”, “belongs to” and so on. This characterization helps KG-based recommender systems to give recommendations which are more semantically meaningful. Recommendations are generated based on path-based reasoning, exploring multi-path hops in the knowledge graph. Although it can solve the data sparsity and cold-start problems, it has limitations such as scalability and computational cost.

2.3.2. Applications of Recommender Systems in Diverse Domains

Linden, Smith, and York (2003) proposed a novel collaborative filtering technique called item-to-item collaborative filtering. Instead of user-user or item-user collaborative filtering, they used item-item collaborative filtering by computing similarities between items rather than users. This helped Amazon provide highly personalized and contextually relevant recommendations in real-time, even for their vast database. This is achieved by precomputing item similarities against user baskets, so that recommendations can be provided in real time. The limitation of the model is its expectation of consistent user behaviour.

Philip, Shola, and Ovyte (2014) proposed a research paper recommendation system based on a content-based filtering approach for personalised suggestions for researchers. The dataset for the study included research papers from the Federal University Kashere and open-access sources on the internet. User profiles were created based on users' transactional behaviour, such as downloading, liking, and opening research papers. Research papers were represented as vectors using TF-IDF weighting techniques to identify significant keywords related to the papers. Item profiles were constructed using research paper attributes such as titles, keywords, research areas, and author information. The system utilised TF-IDF and cosine similarity to recommend papers to users by matching user profiles with item profiles. While the approach provided recommendations with high precision, it suffered from cold-start problems for new users and limited diversity in recommendations.

Walek and Fojtik (2020) introduced a hybrid recommendation system to suggest movies by integrating content-based filtering, collaborative filtering, and an expert system. The rule-based expert system augments recommendations generated by the collaborative filtering and content-based filtering modules. The collaborative filtering module creates a user-item matrix containing user movie ratings, rankings, and item (movie) ratings, which is used to rank the most suitable movies based on the user's ratings and the ratings of other users. The content-based filtering module calculates movie similarity by processing user-rated movie descriptions using a bag-of-words model. TF-IDF is then applied to extract features and compute cosine similarity between the features of unrated movies and those of movies rated by the user. Final recommendations are filtered and ranked by the expert system, which considers factors such as the average movie rating, the total number of ratings, similarity to already-rated movies, and outputs from the content-based recommender module. The expert system applies these criteria using if-then-else rules. The key contribution of this hybrid system is its ability to provide high-precision yet diverse recommendations, while limitations include the complexity of implementation and scalability challenges.

2.3.3. Current Approaches to Recommendation Systems in Career Guidance

Zameer Gulzar et al. (2018) published a hybrid course recommendation system that combines information retrieval methods and ontology to represent knowledge related to courses. The system uses N-gram classification as well as WordNet, query expansion. Ontology representation is used to construct the knowledge model. The query is then expanded with the help of WordNet, and the input query words are pre-processed and N-grams generated. Subsequently, courses related to these queries are extracted from the database based on ontology relations and recommended to the user. This method was much superior to the previous methods of search by keyword. But one of the most important weaknesses of the system is its knowledge base which is limited to the courses that are related to computer science. This specialization introduces the problem of data sparsity and the possible challenge of cold start problems. Furthermore, the system was built to the level of research students and

cannot be used to recommend courses at the secondary level of school education, since it does not consider the aptitude of students, which is an essential factor there.

Mohamed Ezz et. al (2020) proposed in their work an adaptive recommendation system to estimate appropriate engineering education paths of students once they have finished their preparatory course. It achieves this by using data mining techniques to extract relevant features to each engineering discipline. A variety of machine learning models including KNN, SVM, and NB are constructed on each discipline given these selected features, and the most successful model identified on each discipline. The training data is composed of the historical data of final-year students in each discipline. After the models have been constructed, the academic achievement of preparatory-year students is entered into the adaptive recommendation engine, which is then used to advise students on the engineering disciplines in which they are likely to achieve success. This paper introduces a promising recommendation engine architecture, and shows good results, especially in areas such as urban planning and mining and petroleum engineering, where the models showed the highest F-measures. Nonetheless, there are still difficulties in selecting features per department and imbalanced data.

In the study by Atienza John Robert D. et al. (2022), a Deep Neural Network (DNN) framework was elucidated to predict the academic trajectory of students, combining academic performance of students, in each subject, age, sex, salary levels of parents, and number of siblings. The study took into account a large group of 1500 junior high school students in the Philippines. The web-based system as a useful tool supports career counsellors in assisting students to make sound decisions when it comes to their career paths. Although the paper illustrates the creation of the candidate profile through academic track records and demographic data fed into the DNN model, it does not explain how the candidate profile is mapped to a suitable career path. The system also fails to factor in whether the student is apt for the career being projected.

The research of Saeed, Takreem, et al. (2021) presented a novel job recommender system targeting engineering students in six different disciplines. They surveyed engineering students of several universities to create the user dataset, prioritizing four engineering areas. The data feature includes qualifications, field of engineering, key subjects, GPA, programming languages, frameworks, and platforms. An equivalent

job profile database, compiled by practitioners in related sectors, offered job titles and elaborate descriptions in those engineering jobs. Based on these datasets, the authors have created a Deep NLP Career Recommender System. The system divides the data into training and testing data, tokenizes and lemmatizes job descriptions, and finds features by using an embedding layer. These characteristics are run through a convolutional layer and fully connected neural network, and job suggestions are formed as a result of the learned weights of the features. Although the system can help engineering students in getting job recommendations, it cannot be used to give career recommendations to the secondary students on basis of their abilities. Moreover, it fails to address problems such as data sparsity or cold start problems.

Rashid, Adib Hakimi Abdul, et al. (2022) developed a content-based filtering approach to advising computer science students on career choices. Their primary objective was to create an automatic mechanism through which people in the computer science profession are recommended individual careers. The system applied web scraping technology to collect associated career information on the JobStreet site. The content-based filtering was developed into a recommendation mechanism in which user preferences served as the point of comparison. The candidates were introduced to different fields of computer science and related careers first, which allowed them to determine their interests. The system then matched the areas chosen by the candidate with various computer science career paths derived out of JobStreet. This comparison was conducted using cosine similarity on TF-IDF vectors of the interests of the candidate and the career path data available within the system. Depending on the similarity measure, the system suggested to the candidate an appropriate career. However, while the system effectively suggests a set of computer science-related career paths, it faces a significant challenge with data sparsity due to its reliance on web scraping, which can result in incomplete or missing data.

Ngiam, Avery Tzeng Chiet, et al. (2022) presented an Early Career Trajectory Recommender System (ECTRS) designed specifically with youth and high school students in mind, yet aimed at Information Technology (IT) careers. The system used the occupational information network (O*NET) database to develop user profiles. The profiles were created when the Holland Code Test, created by the Open-Source Psychometrics Project, was administered to learners in order to determine their interests. Also, the skillset and knowledge base of each student were established with

the help of a questionnaire. Career profiles were then built around the same attributes as the user profiles, but limited to the IT domain. The work used a Bag-of-Words model to identify features, and the cosine similarity measure was used to match the user with the appropriate IT careers. Although this work can be of great help to students interested in IT careers, there are two limitations to it. One reason is that the system does not objectively determine the skills of the user, but rather uses self-reported skills as provided by the questionnaire. Second, the career suggestions are limited to the IT industry, and not the wider range of career choices out there in the real world.

Obeid, Charbel, et al. (2022) developed a hybrid recommendation system that combines collaborative filtering and knowledge-based methods to help high school students choose suitable universities, majors, and career paths. The collaborative filtering component uses graduates' career ratings to find cases similar to a student's profile, generating an initial set of recommendations. These are then refined through a knowledge-based system built on case-based reasoning and an ontology that links information about institutions, students, and careers. The final output blends both approaches to produce a personalised recommendation. Data for the model were collected through an online graduate survey that included demographic details, education history, interests, and career evaluations. While the system offers a structured and personalised approach, it faces certain limitations, such as cold-start issues, data bias, quality inconsistencies, and scalability challenges.

2.4. Need for Explainable Recommendation Systems

Recommender systems have found applications in various and critical fields such as education, healthcare, finance, e-commerce, and entertainment, among others. The most essential underlying algorithms are based on Deep Neural Networks (DNNs) to provide more accurate results. Nevertheless, these machine-learning models can be regarded as black boxes (Ribeiro, Singh, and Guestrin, 2016; Amarasinghe and Manic, 2019), as illustrated in Figure 1, and both model developers and end users struggle with the reasoning behind the suggestions.




Figure 1. Internal Architecture of Deep Neural Networks

There is a trade-off between accuracy and interpretability in machine learning models, where the more interpretable the model, the less accurate it is (Freitas, 2019), as illustrated in Figure 2. In the algorithmic society of today, where algorithms affect nearly all facets of life, it is important that the predictions and suggestions made by such black boxes be interpretable and transparent. Despite the predictive accuracy of deep neural networks, the networks also contain bias and can fail to make responsible predictions (Nguyen, Yosinski, and Clune, 2015). This shows the necessity of clear and responsible algorithms. The global community has pointed out that artificial systems not only must be precise but also just and reliable (Council, 2017; Viswan et al., 2024). To mitigate these issues, it is necessary to incorporate key elements of transparency, including quality of data, combating algorithmic bias, ethical norms, and human control (Hassija et al., 2024). These are important steps towards ensuring that machine learning models are not discriminatory against individuals or groups of people.




Figure 2. Machine Learning Models: Accuracy Versus Interpretability

There are several reasons why it is crucial to create transparent and equitable machine learning systems such as recommender systems. Due to the fact that these systems operate within critical decision-support systems, user acceptance and trust are of the utmost importance (Rechkemmer and Yin, 2022). The system needs to be able to explain the decisions proposed by the system. The lack of transparency among the users of these systems may result in distrust and therefore their rejection. Reduction of bias and fairness is significant (Hort et al., 2024) since models are trained using training data, and when the dataset is biased, the predictions made using the dataset will also be biased. Recommendations based on gender, race, socioeconomic and geographical location can be unfair hence leading to issues of fairness. It requires transparency to meet ethical and legal standards such as the General Data Protection Regulation (GDPR) (Privacy Regulation, n.d.). Personal data must be gathered with consent, utilized to a narrow intention, and must not be abused. The data must be handled in a legal manner so that the information is correct and that adequate security has been taken to shield against unauthorized entry. Because machine learning models are frequently seen as black boxes, the aforementioned measures can be broken without notice unless explainability is incorporated. The refinement of the systems will also be supported by interpretability and transparency (Sotgiu et al., 2022), allowing the designers to debug and refine them. Furthermore, it is only open systems that can keep up with the evolving norms and expectations in society.

A machine learning model must address explainability from multiple dimensions, including trustworthiness in data preparation, robust and explainable data representations, fair, transparent, and responsible predictions, and trustworthy evaluation (Wang et al., 2024). Data preparation should ensure data quality, fairness, and privacy by addressing biases and ensuring compliance with ethical standards. Data representations should be resistant to manipulations and provide interpretable insights. The generated recommendations should satisfy the user while upholding fairness, ethical principles, and accountability.

2.4.1. Explainability Approaches

There are different types of explainability methods for machine learning models: model-specific versus model-agnostic, global versus local explainable frameworks, intrinsic versus post-hoc explainability, and many more.

The model-specific explainability approaches are designed to be model-specific and use the internal structures of the model to explain. A model-based implementation is Layer-Wise Relevance Propagation (LRP) (Bach et al., 2015). Nonetheless, model-agnostic applications can be run on any machine learning model, regardless of its internal construction. Examples of model-agnostic explainable frameworks include Local Interpretable Model-Agnostic Explanations (LIME) (Ribeiro, Singh, and Guestrin, 2016; Amarasinghe and Manic, 2019) and Shapley Additive Explanations (SHAP) (Lundberg and Lee, 2017).

Global explainability refers to the ability of the model to predict its overall behaviour, as feature importance lists the most important features of the model in its overall predictions. Local interpretability focuses on the model's ability to interpret a particular instance of prediction for a given input, like instance-specific feature importance of a single prediction (Clough et al., 2019). Some algorithms, like SHAP, can provide both global interpretability and local interpretability.

Post-hoc explainability refers to the process of explaining model behaviour after the model is built. Visualization, feature attribution, or rule-based feature extraction techniques are used in this case. Intrinsic explainability refers to the inherent ability of the model's structure to provide interpretability for possible predictions coming from the system (Ai & Narayanan, 2021).

Some machine learning models are intrinsically explainable, such as Linear Regression, Decision Trees, and Rule-Based algorithms (Molnar, 2020). These are often called white-box models. Decision trees have a human-readable internal structure, and each node precisely defines a rule in the decision-making process. The importance of a feature can be identified by its position in the tree. In linear regression, the coefficient represents the impact of a feature on the model's prediction.

The interpretation of black-box models can be achieved using rule-based explanation models. This can be achieved using if-then-else structures to define rules by analyzing the dataset. Additionally, leveraging domain knowledge expertise or extracting rules using a white-box model like a decision tree trained on the dataset (Yang, 2022) can also be explored. For example, training a decision tree to mimic the behaviour of a

DNN and extracting rules from the trained decision tree can help break down the black-box nature of DNNs (Chou, 2019).

To enhance transparency and accountability in black-box machine learning models like DNNs, various eXplainable Artificial Intelligence (XAI) algorithms have been developed. The algorithms include SHapley Additive exPlanations (SHAP), Local Interpretable Model-agnostic Explanations (LIME), Gradient-weighted Class Activation Mapping (Grad-CAM), and rule-based systems (Vimbi, Shaffi, & Mahmud, 2024; Jahmunah, Ng, Tan, Oh, & Acharya, 2022; Macha, Kozielski, Wróbel, & Sikora, 2022).

Local Interpretable Model-Agnostic Explanations (LIME) (Ribeiro, Singh, and Guestrin, 2016) is based on the principle of locally approximating a complex model using an interpretable but simpler model, like a decision tree or linear regression model. LIME generates an interpretable model for each input instance to provide local explanations. The process begins with perturbing the input instance to create a set of samples. These samples are then fed to the original black-box model to obtain predictions. Based on these predictions, an interpretable surrogate model (such as a linear model or decision tree) is trained on the generated dataset. By analysing the feature importance, such as the coefficient of this surrogate model, LIME identifies the key factors influencing the original model's prediction. LIME has the limitation of instability. Further, LIME may generate inconsistent explanations (Zhao et al., 2021).

SHAP (SHapley Additive exPlanations) is a Model-Agnostic XAI framework based on cooperative game theory. It assigns each feature a value called the SHAP value for a given prediction, calculating the marginal contribution of each feature to explain how it influences the model's prediction (Lundberg & Lee, 2017). SHAP's Additive feature attribution property ensures that the sum of all feature contributions equals the model's predicted value. This makes its explanation consistent with the model's prediction. Even though SHAP is model-agnostic, it has different specialized implementations like Tree SHAP for tree-based models. Other implementations are: Kernel SHAP for model-agnostic implementations, and Deep SHAP for neural network-based models. SHAP provides both local and global explanations. SHAP

uses visualization plots, such as bee swarm plots, force plots, and decision plots, to explain both local and global explanations. Limitations of the SHAP explainer include its computational cost and the lack of end-user-friendly explanations.

Gradient-weighted Class Activation Mapping (Grad-CAM) utilizes the gradients of target class scores flowing into the final convolutional layer of a Convolutional Neural Network to generate a coarse localization map (Selvaraju et al., 2020). This map visualizes the important regions of the input image, explaining the model's prediction. GRAD-CAM focuses on a specific class for which an explanation is needed. It computes the gradient of the selected class score with respect to the feature map of the last convolutional layer. GRAD-CAM provides human-readable visual interpretations. It can be applied in image or object detection applications for explainability. However, when deploying it for critical applications, its main limitations, such as dependence on convolutional layers and coarse resolution, must be considered.

2.4.2. Existing Implementations of Explainability in Recommendation Systems

Explainable AI-based recommendation systems have been tried in various fields. A proposal of a Grad-CAM-based explainable deep transfer learning algorithm was presented by Panwar et al. to detect COVID-19 early, using chest X-ray and CT scan images' datasets (Panwar et al., 2020). They visualised the infected areas of the lung using colour to make their predictions more interpretable by plotting class activation maps according to the Grad-CAM technique. Grad-CAM was applied to the last convoluted layer of their CNN, since it can determine high-level spatial information about image features, and gives better visual explanations of the model prediction.

Tahir, Alayed, Hassan, and Haider (2024) introduced a hybrid explainable AI (XAI) model that integrates LIME and SHAP to achieve real-time interpretability in autonomous vehicles. This approach exploited the computational efficiency of LIME while utilizing SHAP for comprehensive feature importance analysis. Specifically, the computationally expensive SHAP is applied to the top n features identified by LIME. This approach allowed the model to maintain high prediction accuracy while providing interpretability. This method enhanced both computational and memory

efficiency of the system. The approach helped to make it suitable for onboard implementation in real-time applications like autonomous vehicles.

Shams, M. Y., Gamel, S. A., and Talaat, F. M. (2024) proposed an explainable crop recommendation algorithm called LIME-based XAI-CROP. This algorithm takes into account various parameters such as soil characteristics, weather conditions, and historical crop performance. Selecting a sample from the validation dataset, applying perturbations to the data, and training a linear model on the perturbed data, the LIME algorithm was employed to generate feature importance explanations for the prediction of the selected sample.

In their prostate cancer treatment recommender system, Tang et al. (2024) used the SHAP algorithm to discover principal parameters like tumour stage (T stage), cancer stage, age, core positive percentage (CPP), prostate-specific antigen, and Gleason score that had a significant impact on the prediction ability of the model on surgical and non-surgical treatment of prostate cancer. This analysis was conducted using SHAP's global explanation ability employing beeswarm plots. They applied the SHAP algorithm to explain predictions across three different models they developed. The system identified the best model to use in its prediction procedure by conducting feature importance analysis using a SHAP beeswarm plot. This decision was based on the degree of coherence between the feature importance explanations provided by SHAP with domain expertise in the field.

There is a strong need to include explainability in the career recommendation system, particularly among students in the early career decision-making, where the end user of the system can understand the logic behind career recommendations. Although black-box based recommender systems may be very accurate, it does not give confidence to the career counsellor and the end-user unless appropriate interpretations of the reasons why the system made some recommendations. Interpretability is beneficial to both end-users and model developers. The system offers developers the opportunity to debug and optimise the system to eliminate bias, inaccuracies, and inconsistencies (Xu et al., 2019; Dwivedi et al., 2023). The students can be provided with information on why a particular career path is recommended and hence make

sound decisions. The key aspects of the long-term success and widespread use of the system are trust and confidence.

Research on explainable career recommendation systems is still emerging. Schellingerhout, Barile, and Tintarev (2023) developed an XAI framework employed on job recommendations. Their work targeted candidates, recruiters, and company managers and built on a Graph Attention Network (GAT). Their system incorporated the candidate's features and vacancy features to produce three explanation types: i) visual knowledge graph paths, ii) textual summaries using large language models, and iii) attribution charts. Company managers preferred visual explanations for clarity. For actionable insights, candidates and recruiters favoured textual explanations. In the job recommendation domain, the study provides foundational guidelines for stakeholder-specific explainability.

Tran (2023) provided an empirical approach to an explainable job recommendation system. Their approach applied both local and global explainability techniques. For local explainability, the study employed LIME and SHAP algorithms. With inherently self-explanatory models like Explainable Boosting Machines (EBM) and Differentially Private Explainable Boosting Machines (DPEBM), global explainability was achieved. Text data from job descriptions was used to calculate TF-IDF and Latent Dirichlet Allocation (LDA). For building the recommendation engine, ensemble algorithms like XGBoost, factorization machines, etc., were employed. For the global explanation evaluation of the system, the model fidelity rate was used, and a feature importance fidelity rate was used for local explanation evaluation. Even though fidelity-based measures were suggested in order to measure the soundness of explanations, the system primarily converted feature weights to keywords or topical labels, instead of producing prose length natural language explanations to make explanations more comprehensible to the final users.

Yang et al. (2023) came up with an explainable trainee course recommender system that is based on a Contextualized Knowledge Graph Embedding (CKGE) model. Their work employed a Knowledge Graph, which comprised of entities, including talents, skills, and courses. The meta-graph of any talent-course contains a meta-path

between the two entities and the neighbours of each entity. The model used a KG-based Transformer and customized to serialise entities and paths on the meta-graph into a structured representation, which represented a local and global interdependence. Explainability is achieved through the exploration of the meta-graph that demonstrates how user talents are related to the suggested courses. Although very useful in demonstrating the line of reasoning, the framework did not involve processes that generate explanations in natural language form.

The literature review shows that there has been a great development in ensuring explainability in the recommender systems in various areas of application. Even so, there still exists a critical gap in creating consistent accounts in natural language that can be interpreted by non-specialist users. Currently used methods such as those supported by SHAP, LIME, and knowledge graphs mostly generate visual or structural justifications, which in many cases require technical knowledge to interpret properly.

2.5. A Review of Relevant Datasets

To come up with efficient career recommendation systems, datasets are essential. The quality of the datasets is necessary to make unbiased and fair decision-making when using machine-learning models based on data. A recommendation engine can be enhanced with a rich dataset having relevant features, which can considerably improve the performance of such an engine. There is a high level of importance in considering the relevance of a dataset to a certain topic, quality (minimal missing values and outliers), diversity (to overcome biases), and ethical implications (privacy, fairness) when choosing a dataset.

Regarding the career counselling domain, datasets should contain diverse career-related information to build a holistic career recommendation system. But none of the existing career datasets fully address this requirement. Some datasets focus on ability profiles, and some others prioritize vocational interests. Most career datasets are limited in scope and incomplete. Another concern is with datasets built using web scraping techniques, which suffer from data sparsity issues.

The CareerMap dataset by Pimplikar et al. (2023) is particular to the computer science field alone. It consists of 9,180 records having 28 features and covers up to 16 careers within the computer science field. Besides the characteristics that are relevant to technical professions, the dataset contains the features concerning the Big Five personality traits (openness, conscientiousness, extraversion, agreeableness, and neuroticism), and therefore has both technical and personality aspects to the recommendation systems constructed based on this dataset. The key issues that can be identified with the dataset are its limitations, since it has been narrowed to 16 computer science-related careers only, which restricts its general applicability. It also lacks aptitude profile information that is usually essential in making holistic career recommendations.

Computer Science Students Career Prediction Dataset (Patil, n.d.) is a small dataset with the data on 180 computer science students. It contains the fields of student ID, name, gender, age, GPA, major, interested domain, projects, future career goals, and Python, SQL, and Java skills with the future career goals being the target field. Although the dataset can be utilized in terms of career recommendations in the computer science field, it does not have a wide coverage and is limited. The synthetic character of the dataset can decrease its usefulness in the context of the career recommendations in a real-world scenario. Moreover, the dataset does not consider such important factors like aptitude, professional interests, and personality traits, limiting its scope and overall utility.

The Career Recommendation Dataset, developed by Breejesh Dhar (Dhar, n.d.), contains records of 1,316 professionals with diverse attributes. Attributes include demographic information (gender), academic details (undergraduate course, specialisation, CGPA, and certifications), professional data (employment status, first job title), and competencies (skills and interests). The dataset takes a holistic approach, considering skill sets, vocational interests, achievements, and specialisations. But it is limited to a narrow set of careers, and leads to issues such as data sparsity and cold start limitations. The dataset, even though it contains the skill

set features, does not adequately address the innate aptitude of candidates, that essential for providing more efficient and personalised recommendations.

The Career Prediction Dataset (Shrivastav, n.d.) combines the Big-Five personality factors (openness, conscientiousness, extraversion, agreeableness, and neuroticism) with the types of differences in aptitudes (numerical, spatial, perceptual, abstract reasoning and verbal reasoning) in 105 different careers. It is a unique asset that can be used to develop career recommender systems because it is a combination with both personality and aptitude features associated with each career. The dataset, however, has significant weaknesses. It lacks mechanical reasoning and verbal skills, language-based tests, which have to be considered as the core part of differential aptitudes to select a career in full. Its size is another weakness as it comprises of 105 distinct careers, which do not cover the scope of career opportunities in the real world. Irrespective of these weaknesses, this dataset can still be a helpful resource for researchers and analysts developing career recommender systems.

The analysis of existing career-related datasets shows that career-related datasets remain at their early phases. This leaves the research society without a mature and comprehensive dataset. The majority of the available career recommender systems are constructed based on scraped data on job portals (Rashid, Mohamad, Masrom, and Selamat, 2022) or social network websites such as LinkedIn (Kumalasari and Susanto, 2020). The simplest solution to sealing this gap involves making a new dataset based on the data provided in the Occupational Information Network (ONET). The ONET database consists of in-depth information about 1,000 professions. It enumerates skills, interests, working styles, education, skills, training requirements, work experience, daily work, job situation, titles, and technology skills. Such in-depth information may be utilized to construct a high quality and trustworthy career recommendation system. In our proposed system, we utilised the ability profile information of all careers available in O*NET to create a unique, first-of-its-kind dataset containing 612 records. This approach effectively addressed the issue of data sparsity, significantly enhancing the system's utility. The weakness of the dataset is that it did not consider features like personality traits and vocational interests of the

candidate. Another issue to address, considering datasets, is the evolving and ever-changing careers in the real world. There should be a dynamic career expansion approach to enrich existing datasets, and the recommender systems should incorporate such changes in the career domain.

2.6. Conclusion

The multidimensionality of career counselling process is seen in literature review. It places much importance on considering various theoretical factors including aptitude, occupational interests, personality traits, multiple intelligences, and emotional intelligence. Given the intricacy of these dimensions, the review examines the application of psychometric instruments using different methods of measuring these dimensions. And lastly, the review has considered how machine learning technologies have been incorporated into systems that can develop recommendation engines. In order to deal with the multivaried nature of career counselling, psychometric instruments like Differentiated Aptitude Tests, Five-Factor Model and Strong Interest Inventory were considered to gain a systematic approach to assessment. This research is the foundation of the creation of career recommendation systems providing personalised advice. The development of machine learning algorithms has also made it possible to develop automated solutions when it comes to career development. Recommendation engines based on AI provide personalised and scalable career advice. Their capacity to process large datasets efficiently makes them indispensable for modern solutions.

The main limitation identified in the existing recommender systems is their restricted capacity to generate diverse recommendations that are meaningfully aligned with the user's profile and preferences. This gap is significant and requires innovative methods of architecture that can literally incorporate diversity without diminishing the relevance of the recommendations. Here, the proposed approach suggests a new hierarchical recommendation structure that will categorize the recommendations into multi-tiered outputs. The system combines a collection of internal models organized on the basis of the natural hierarchy of career clusters in the dataset. This guarantees breadth as well as contextual coherence at the time of suggesting career pathways.

Such a hierarchical process of recommendation and multi-tiered outputs allows the recommendation engine to trade in personalization with diversity, and eventually provides the users with more detailed and differing career paths.

The factors of bias, fairness, and transparency are very significant in career recommender systems. The recommendations provided by AI will directly affect the formative decisions of students and these decisions need to be explained in natural language to ensure that the candidate can make informed and meaningful decisions. Transparent explanations not only favour informed decisions but also build trust and ensure adherence to ethical guidelines, such as GDPR. Gender, race, and socioeconomic biases should be considered to enhance the concept of inclusiveness and equity further. Although explainable algorithms like SHAP and LIME enhance the interpretability, since they require technical conciseness in their explanations, they are not suitable for providing a rationale that is comprehensible to the layperson. To address it, the current work proposes a hybrid explainability framework that integrates SHAP-guided insights with structured descriptions of natural language explanation. Therefore, the proposed work, addressing the responsible AI concepts, attempts to increase transparency and make it more accessible to non-specialist users.

Regarding career datasets, the review found several shortcomings in the currently available ones. Most of them are too narrow and fail to represent the wide range of real-world careers, making them unsuitable for complete career recommendation engines. The O*NET database offers a better foundation, as it includes detailed information on abilities, interests, and other career characteristics that can help build a more comprehensive dataset. Studying the information in O*NET guided the development of the dataset used in this work. The review also showed the need for a flexible approach that can keep adding new and changing careers, so the recommendation system stays useful and up to date.

The proposed system spans multiple dimensions of inquiry, integrating theoretical frameworks, psychometric evaluations, advanced recommendation methodologies and explainable AI Algorithms. It addresses critical challenges to ensuring diversity in recommendations, designing explainable architectures, and enabling dynamic

career pathway expansion. Central to this effort is the curation of a first-of-its-kind dataset tailored for comprehensive career guidance. Building on this foundation, the work proposes a suite of novel architectural components-including the core recommendation module, an explainable framework layer, and a dynamic expansion module-that collectively establish an innovative and holistic framework for career recommendation with direct real-world applicability.

CHAPTER 3

RESEARCH METHODOLOGY - AN OVERVIEW

This section provides an overview of the methodology adopted for the proposed work. The discussion starts with detailing the workflow from theoretical foundations to system implementation. It begins with the theoretical aspects that form the basis of the study, followed by the creation of a comprehensive dataset serving as the backbone for generating career profiles, and then covers the process of user profile generation, explaining how individual aptitude data is incorporated to align recommendations with the user's abilities. The novel architecture of the proposed recommendation engine is briefly introduced, focusing on its hierarchical, multi-tiered design, discussing how it enables the generation of personalised and diverse career options. The innovative explanation framework, with its hybrid architecture, is presented to highlight the system's transparency and interpretability, enabling users to understand the rationale behind the recommendations. Additionally, to address the ever-evolving nature of the career domain, the scalable architecture of the dynamic career expansion module (DCEM) is described, showcasing its ability to incorporate new and emerging career options to ensure the system's adaptability.

This section provides only a concise understanding of the design and functionality of the system, as this chapter is targeted at laying the groundwork. Each component of the methodology, like dataset creation, career and user profile generation, recommendation engine architecture, explanation framework, and career expansion module, is examined comprehensively later in subsequent chapters. Those chapters will provide a deeper insight into their development and implementation, ensuring a structured and logical progression throughout the study. The workflow diagram presented in Figure 3 depicts the overall process of the research work.

3.1. Psychometric Foundations of the Methodology

A comprehensive study on the psychometric foundation of the work was carried out, and the proposed Differential Aptitude Test (DAT) was selected as the psychometric

tool to identify candidates' aptitudes. A thorough exploration of DAT was conducted, examining all its test batteries and analysing how each identifies different innate abilities of the candidates and how these are related to various real-world careers. The analysis also included DAT's weaknesses, criticisms against it, issues of bias and fairness, and the methodology used to address these concerns. The discussion further acknowledges the limitations of DAT. It provides recommendations to enhance the career counselling process by integrating other aspects, such as vocational interests and personality traits, to create a holistic career recommendation system. The details of this analysis will be presented in the respective chapter that follows.




Figure 3. The Workflow Diagram of the Proposed Study.

3.2. Designing and Constructing the Dataset for Building the Career Profile

As there are few publicly available datasets based on differential aptitude tests encompassing a large number of real-world careers, it was necessary to build a DAT-based dataset from scratch. This led to the curation of a comprehensive dataset using DAT parameters. To create the dataset, the ability profile information from the

Occupational Information Network (O*NET) database was utilised. By analysing the ability profile information from O*NET, key features were identified and mapped to the corresponding DAT test batteries. From over 1,000 careers, 612 significant careers were selected. The parameter values were converted into percentile values. This forms the basis DAT dataset for the proposed work. The systematic methodology followed in the creation of this dataset will be discussed in detail in the respective chapter ahead.

3.3.Dataset: Clustering, Sub-Clustering, and Data Augmentation

The careers in the dataset were grouped into related career clusters using various clustering techniques. The resultant clusters were evaluated for effectiveness. Each cluster's records were labelled to enable the application of classification algorithms. The initial clusters generated were further stratified to create second-level clusters, forming the basis for a hierarchical recommendation engine capable of offering a diverse array of career options aligned with a candidate's innate aptitude. The records in both the top-level and secondary-level clusters were augmented to generate additional related records, aiming to learn more generalised patterns from the dataset. This approach helps mitigate overfitting and ensures the system provides more realistic and robust career recommendations. The augmented clusters were validated to ensure the models did not learn patterns from noise. The entire process will be thoroughly explored in the corresponding chapter ahead.

3.4.Generating User Profiles Through Differential Aptitude Tests

Generating user profiles is a crucial task for a career recommendation engine based on DAT. Psychometric tests must be systematically developed and standardised for the target population, considering the socio-economic and cultural diversity within that population. The test items should be specifically designed to address these factors. A psychometric test known as the Kerala Differential Aptitude Test (KDAT), developed by the Directorate of General Education, Kerala, has already been standardised for students in Standard X in Kerala. The proposed work utilises KDAT for user profile generation based on DAT, with the department's approval as per the Memorandum of Understanding (MoU). The details of KDAT will be thoroughly discussed in the chapter dedicated to user profile generation.

3.5. Structural Framework of the Career Recommendation Engine

A Content-Based Filtering (CBF) approach is employed to design the proposed recommendation engine. The recommendation engine is built upon machine learning algorithms, including deep learning. The study investigates methods to deliver diverse yet personalised career recommendations aligned with a candidate's user profile. A novel hierarchical and multi-tiered architecture is proposed and designed for the recommendation engine. The design applies hierarchical clustering to the dataset. The dataset is initially clustered into an optimal number of groups, and these initial clusters undergo further stratification to generate second-level clusters. This internal hierarchical structure, along with the multi-tiered output endpoints of the recommender, ensures diversity in the career recommendation process aligned to the given user profile. Additionally, an innovative explainability module and a dynamic career expansion module are integrated into the recommendation engine to enhance transparency, trust, fairness and adaptability in real-world scenarios. The chapter on the recommendation engine architecture will provide a detailed discussion of its implementation.

3.6. Architectural Design of the Explainability Framework

Since an explainable machine learning-based system is highly advocated and critical in a sensitive area like career counselling, a robust explainability module has been integrated into the recommendation engine. Machine learning-based recommender engines built on DNN-like algorithms often function as black boxes. They do not offer insights into their decision-making process. To address this, we employed a hybrid explainability framework that integrates SHAP (Shapley Additive Explanations) with a textual explanation module. SHAP was chosen for its ability to provide both global and local interpretability. Additionally, SHAP is model-agnostic in functionality. However, while SHAP offers valuable insights, its plot-based visual interpretations may not be easily understood by general users. This work introduces an innovative explainability framework employing SHAP's local and global explainability features and NLG techniques to produce natural language explanations for the recommendations made by the system. This hybrid approach ensures that both expert

users and general users can understand the rationale behind the career recommendations generated by the engine. Thus, the framework fosters transparency, fairness and trust in the recommendation engine.

3.7. Framework of the Dynamic Career Expansion Module

A career recommendation system should be adaptable to predict new careers in the ever-evolving real-world career domain. Systems built on static datasets are typically limited to predicting careers within the dataset. This can lead to the obsolescence of the recommendation engine and reduce its long-term usability. To address this limitation, the proposed career recommendation engine integrates a dynamic career expansion module that incorporates new careers into its career clusters, enabling the system to recommend up-to-date, real-world career options. The dynamic career expansion module is developed using a semi-supervised approach. A career counsellor can profile new careers for DAT parameters, and the recommendation engine classifies these careers, updating the related career clusters accordingly. This process ensures the adaptability and relevance of the system over time. The detailed process is illustrated in the chapter on the dynamic career expansion module.

The overall system development process, depicted in Figure 4, encompasses every stage from dataset creation, career clustering, and data augmentation to model building, the explainability framework, and finally, DCEM integration.




Figure 4. High-Level Overview of System Development.

CHAPTER 4

THEORETICAL AND PSYCHOMETRIC FOUNDATION OF THE STUDY

4.1 Introduction

Providing a holistic career recommendation to a candidate, especially at the early stages of their education, while considering multiple factors such as aptitudes (Savickas, 2019), personality traits (Mammadov, 2022), and vocational interests (Kim & Beier, 2020), is very important.

Aptitude measurement is pivotal in career recommendations, as different careers require different types of intelligence, such as numerical ability, abstract reasoning, space relations, and mechanical reasoning. By evaluating a candidate's aptitude, recommending related careers and course majors can guide students toward appropriate career paths. This approach enhances both course and job satisfaction, reducing dropouts in educational institutions and job switches in the industry. These arguments should be augmented by insights into the need for unbiased and fair aptitude tests, as aptitude assessments are often criticised for socio-economic and cultural biases (Reynolds et al., 2021; Washington et al., 2016).

The analysis further clarifies that career recommendations should also take into account the candidate's personality traits and vocational interests. Personality traits such as conscientiousness, openness, neuroticism, extraversion, and agreeableness (McCrae and Costa, 2003) are significant predictors of career success in various domains. For instance, conscientiousness fosters self-discipline and aligns with careers such as accountant, doctor, or systems analyst. Openness to experience encourages creativity and suits careers like writer, designer, research scientist, or journalist. Candidates with high neuroticism may excel as therapists, counsellors, or software testers, where work pressure is typically lower. Individuals high in extraversion thrive in roles such as sales manager or brand ambassador, while those with high agreeableness, known for their empathy, are well-suited for careers such as customer support specialist, mediator, personal manager, or fundraising specialist.

These arguments are critically evaluated in various studies (Zakaria & Yusof, 2021; Angelini, 2023).

Vocational interest also plays an important role in one's career success. Based on John Holland's RIASEC theory (Holland, 1959), a realist personality may excel in fields such as engineering, while an investigative individual might pursue careers in mathematics or scientific research. An artistic personality is well-suited for jobs like designing, whereas a social personality may thrive as a doctor. Enterprising individuals often choose careers in business, and those with a conventional personality may find success in roles such as accounting (Phan & Rounds, 2018; Xu, 2021).

These evaluations highlight the need for a comprehensive approach to career counselling that integrates various aspects such as aptitude, personality, and vocational interests. Recognising the importance of these factors, the proposed study narrows its focus to aptitude-based career recommendations alone and needs to expand the work to include other factors in the future, to address this limitation.

4.2 The Differential Aptitude Tests

The career recommendation system proposed in this study is based on Differential Aptitude Tests (DAT) (Bennett, Seashore, & Wesman, 1956) conceptualized by George K. Bennet and his collaborators and subsequently published by the Psychological Corporation (Bennett, Seashore, & Wesman, 1991).

An aptitude test is a tool used to measure an individual's ability to acquire, through future training, a specific set of skills. The Differential Aptitude Test (DAT) is an instrument designed to assess various abilities, including cognitive processing, reading comprehension, language proficiency, attention and working memory, executive functioning, processing speed, and complex reasoning and problem-solving skills. The DAT is conducted under strict conditions and is not an achievement test. It is specifically designed to evaluate natural talents rather than knowledge or skills gained through training. Additionally, no specific preparation or prior training is required to undertake the DAT.

Through a thorough analysis of theoretical constructs, empirical evidence, and practical applications, eight distinct test batteries were systematically designed for DAT (Bennett, Seashore, & Wesman, 1956). These batteries are designed to measure a spectrum of aptitudes, ranging from Verbal Reasoning and Language Proficiency to Abstract Reasoning (AR), Mechanical Reasoning (MR), Numerical Ability (NA), Space Relations (SR), and Perceptual Speed and Accuracy (PSA). The Differential Aptitude Tests, therefore, serve as a sophisticated tool for comprehensive aptitude assessment in educational contexts. Each test battery, used alone or in combination, can provide candidates with clear, measurable insights into career paths that align with their differential cognitive abilities. These insights help candidates choose appropriate courses and pursue suitable careers, thereby supporting sustained personal and professional success.

4.2.1 Test Batteries in Differential Aptitude Tests

There are eight distinct test batteries in Differential Aptitude Tests (Bennett, Seashore, & Wesman, 1956). Each test battery is targeted to evaluate a certain aptitude or ability of the candidates. The Test batteries are:

1. Abstract Reasoning (AR)
2. Numerical Ability (NA) or Numerical Reasoning (NR)
3. Verbal Reasoning (VR)
4. Mechanical Reasoning (MR)
5. Spatial Ability
6. Language Usage – Spelling (LS)
7. Language Usage – Grammar (LU)
8. Clerical Speed and Accuracy (SPA)

Abstract Reasoning (AR): Abstract-reasoning tests are strong indicators of fluid intelligence and of an individual's capacity to grasp new concepts quickly. As non-

verbal assessments, they measure the ability to detect relations among abstract figures and to reason inductively. Items present a series of patterns that vary in size, shape, position or quantity. The examinee must infer the rule governing the sequence and select the figure that follows logically. Performance on these tasks therefore indexes logical, analytical and conceptual problem-solving skills.

This skill is essential in courses and careers that require recognising relationships between objects by size, shape, position and other attributes, and analysing how those relations change over time. Fields that demand advanced abstract reasoning include mathematics, computer programming and software development, the sciences, architecture, engineering, law and medicine. It is also valuable in economics, where predicting outcomes depends on analysing causes and effects, for example when forecasting market trends. In management, especially at senior levels, the ability helps leaders see how different parts of a complex organisation fit together.

Numerical Ability (NA) or Numerical Reasoning (NR): Numerical reasoning assesses an individual's ability to understand numerical relationships, apply quantitative concepts, and use numerical data to solve problems and make decisions. It measures basic arithmetic proficiency, estimation skills, speed and accuracy of calculation, sustained concentration, and critical thinking. The assessment also captures creative problem-solving, since candidates must analyse numerical problems efficiently and draw logical conclusions. People who show strength in numerical reasoning typically enjoy working with numbers and statistics and often find novel ways to interpret quantitative information.

Critical thinking and numerical reasoning are essential across almost all professions in today's data-driven economy. They underpin work in mathematics, physics, chemistry, accounting, actuarial science, economics and engineering. They are also important factors affecting skilled trades such as banking, insurance, computing and surveying, and for clerical and administrative roles in local government and accounts management.

Verbal Reasoning (VR): The items in this subtest are presented as analogies and measure the ability to identify relationships between words. The candidate must infer

the connection in the first word pair and apply the same relation to the second pair. Verbal reasoning is essential for conveying ideas and interpreting written information. It is central to analytical professions such as law, medicine, teaching, journalism and to technical roles like engineering and computer programming. It is also important in business areas such as management, sales, marketing and advertising, and in roles that depend on strong communication, such as training and administrative posts.

Mechanical Reasoning (MR): Mechanical aptitude assesses an individual's capacity to grasp and apply elementary mechanical and physical principles. It measures the ability to reason about motion, force, energy transfer and simple mechanisms, and to understand how everyday devices operate, for example levers, pulleys, gears, springs, wheels and axles, inclined planes, simple electrical circuits and screws. This aptitude helps predict suitability for technical trades, engineering roles and maintenance tasks that demand practical problem solving.

High mechanical reasoning indicates strong practical aptitude and a readiness to learn how machines and equipment operate. People with this skill grasp procedures and tools more quickly, visualise how parts interact, and troubleshoot faults in a logical way. The aptitude matters in engineering and skilled trades such as carpentry, vehicle maintenance and plant operation. Specific roles that rely on mechanical reasoning include aircraft technician, auto mechanic, electrician, plumber, machine operator, plant operator and firefighter. In these jobs, the capacity to read technical diagrams, anticipate mechanical failure and apply hands-on problem-solving reduces training time and improves safety and equipment reliability.

Spatial Ability (SA): Spatial ability refers to the capacity to visualise and mentally manipulate objects in two and three dimensions. It covers skills such as mental rotation, recognising space relationships, judging relative size and position, and picturing how parts fit together. These abilities are tested by tasks that require interpreting diagrams, reading blueprints, imagining assemblies, or predicting the appearance of an object after rotation. High spatial ability supports accurate interpretation of technical drawings and effective use of CAD tools, and it aids fine motor tasks that depend on precise mental imagery. Professions that rely heavily on

spatial skill include architecture, engineering, surgery, graphic and interior design, navigation, and trades such as carpentry and mechanics, where anticipating spatial arrangements is a routine requirement.

Language Usage – Spelling (LS): This test battery contains items that assess a candidate's ability to recognise correct spelling and appropriate word use, including homophones, word forms, and common orthographic rules. It measures accuracy in written communication and the capacity to select words that fit context and tone. Strong performance on these tasks indicates readiness for roles that demand precise writing, such as editing, proofreading, journalism, copywriting, advertising, public relations, technical writing and legal or administrative drafting.

Language Usage – Grammar (LU): This test battery contains items that assess a candidate's command of grammar, including punctuation, sentence construction, syntax, register and word choice. It evaluates the ability to write clearly, edit text for accuracy, and select phrasing appropriate to the purpose and audience. Strong performance on these tasks indicates suitability for roles that demand precise written communication, such as teaching, the civil service, public relations, advertising, editing, scriptwriting, proofreading, legal drafting and technical writing.

Perceptual Speed and Accuracy (SPA): This test battery assesses candidates' ability to combine speed with accuracy when identifying and marking simple letter and number patterns. It evaluates rapid visual scanning, sustained attention, hand-eye coordination and a low error rate under time pressure. The measure reflects how well an individual can perform routine, detail-focused tasks quickly and reliably. Such skills are important in clerical and secretarial work, data-entry and indexing, quality-control inspection, proofreading, and technical data analysis, where both pace and precision matter.

Table 3 provides a detailed breakdown of the test battery components. Studies are there to ensure the predictive validity of Differential Aptitude Tests (Setiawati, 2020).

Differential Aptitude Test results are reported as percentile scores rather than raw totals. Percentiles place an individual's performance in relation to a standardised reference group, allowing a norm-referenced interpretation. This process accounts for

question difficulty through normalisation, removes bias inherent in absolute scores, and gives career counsellors a clearer basis for comparing candidates. This supports making well-grounded career recommendations.

Table 3. Differential Aptitude Tests Batteries and Their Interpretations

Ability	Description	Related Area
Numerical Ability (NA)	Proficiency in comprehending numerical relationships and executing arithmetic calculations.	Mathematics, physics, engineering, Market Research, Designing, Actuary, Statistician, etc.
Abstract Reasoning (AR)	Aptitude for recognizing concepts conveyed not through numbers or words but via sequences of figures. It measures the ability to see patterns and relationships in objects.	Architecture, Computer Programming, Law, Economics, Engineering, Mathematics, Mechanics, etc.
Verbal Reasoning (VR)	Ability to conceive ideas framed in words and to reason with words.	Law, Business, Journalism, Science, Social Science, Public Relations, Advertising, Civil Service, Authors, Management, etc.
Mechanical Reasoning (MR)	Proficiency in comprehending fundamental mechanical principles and utilizing tools.	Engineer, Mechanic, Electrician, Carpenter, Quality Controller. Marine Engineer, Assembly Workers, etc.
Space Relations (SR)	Ability to visualize or think in spatial terms. Ability to visualize 3D from 2D	Architecture, Designing, Dentistry, Civil Engineering, Animators, Photographers, etc.
Perceptual Speed and Accuracy(PSA)	Speed and accuracy of response in simple perceptual tasks	Clerical Works, Filing, Coding, Banking, Legal Executives, etc.
Language Spelling (SP)	Ability to use spelling	Proofreading, Law, Secretarial Work, Journalism, Public Relations, Authors, Civil Service, Advertising, etc.
Language Usage (LU)	Ability to use language Grammar	Proofreader, Lawyer, Secretarial Work, Journalism, Script Writer, etc.

4.3 Limitations of Differential Aptitude Tests

Psychometric tests, including Differential Aptitude Tests, have faced substantial criticism due to their historical roots in the eugenics movement (Lowe, 1998). The eugenics movement propagated the false belief that certain racial groups were intellectually superior. Additionally, a test like the DAT, which is standardised on one

population, can yield inaccurate and biased results when applied to other groups. Research indicates that tools designed for one cultural group cannot be equitably used with others. Additionally, various factors such as differing definitions of intelligence, cultural variations, access to quality education, and influences during test administration can significantly affect a candidate's performance (Washington, K., Malone, C., Briggs, C., & Reed, G. K., 2016). These concerns underline the importance of ensuring that tests are specifically standardised for the populations they are intended to assess. Test items must be culturally appropriate. This will reduce bias and promote fair evaluation across diverse groups (Mellenbergh, 1989). The limitations of the DAT should be addressed in terms of its potential bias and limited scope. This can be achieved by standardising the test across the cultural and socio-economic groups to which it is applied. Finally, the insights provided by DAT, being it is a single dimension of the career selection, should be augmented with additional dimensions such as vocational interest and personality of the candidate, to ensure more comprehensive career recommendations.

4.4 Employing Differential Aptitude Tests in the Proposed Work

The proposed work utilises the Kerala Differential Aptitude Tests (KDAT), a psychometric testing tool based on the Differential Aptitude Tests. KDAT was developed by the Department of General Education, Government of Kerala. KDAT is used for user-profile generation to evaluate candidates to recommend careers tailored to their aptitude based on the test result. KDAT has been standardised for the Class X population of Kerala, taking into account socio-economic and cultural differences. The test items are specifically designed for students in Kerala to ensure fairness and mitigate potential bias. The details of KDAT will be discussed in the relevant chapter that follows.

CHAPTER 5

GENERATION OF USER PROFILE USING KERALA DIFFERENTIAL APTITUDE TESTS (KDAT)

In our content-based career recommendation system user profile is generated by employing Differential Aptitude Tests (DAT). DAT is used to evaluate candidates' inherent abilities, and based on that, the recommendation engine will recommend suitable careers to a candidate aligned to their abilities. We followed this assessment while acknowledging and addressing the known criticisms and concerns related to this approach. Specifically, we utilised the Kerala Differential Aptitude Tests (KDAT), an aptitude testing tool developed by the Department of General Education, Govt. of Kerala (Kerala Higher Secondary Education Department, 2024). KDAT is specifically selected as it has been standardised across socio-economic and cultural differences across Kerala Standard X students.

5.1. Introduction to Kerala Differential Aptitude Tests

Kerala Differential Aptitude Tests is a psychometric Testing Tool based on Differential aptitude Tests. KDAT has been designed and standardised specifically for 10th-grade students in Kerala, considering the region's unique socio-economic factors. To achieve this, KDAT has been developed under the guidance of experts in the relevant fields and test items are customised specifically for Kerala students. This is targeted at minimizing potential biases in the testing process (Personal Communication, Career Guidance and Adolescent Counselling, Directorate of General Education, official communication, November 5, 2024).

While DAT is made up of a set of eight distinct test batteries, containing Abstract Reasoning (AR), Numerical Ability (NA), Verbal Reasoning (VR), Mechanical Reasoning (MR), Spatial Relations (SR), Perceptual Speed and Accuracy (PSA), Language Spelling (SP), and Language Usage (LU), KDAT simplifies this structure. It blends the Language Spelling (SP) and Language Usage (LU) batteries into a single category termed Verbal Ability. Further, KDAT omits the Perceptual Speed and

Accuracy (PSA) battery and consolidates the test into a set of six test batteries. The KDAT test batteries are: Abstract Reasoning (AR), Numerical Ability (NA), Verbal Reasoning (VR), Mechanical Reasoning (MR), Space Relations (SR), and Verbal Ability.

KDAT comprises 180 unique questions strategically distributed across the six aptitude test batteries, as detailed in Table 4 (Kerala Higher Secondary Education Department, 2024). These questions play a crucial role in evaluating the skill sets of candidates in their respective domains. The raw scores obtained by candidates are converted into percentile scores for each category. It employed significant item-analysis metrics such as the difficulty index, discrimination factor, and guess factor associated with each question. Percentile scores are used instead of percentage scores to normalise the results and determine a candidate's rank within the sample population. This enables a comparison of the candidate's skill set against a standardised population. Thus, KDAT helps to identify the key aptitude aspects of candidates, helping career counsellors to recommend careers aligned with their outstanding abilities.

Table 4. Test Batteries and Question Distribution in Kerala Differential Aptitude Tests

SI No	Test Category	Number of Questions
1	Numerical Ability	28
2	Abstract Reasoning	28
3	Verbal Reasoning	30
4	Mechanical Reasoning	35
5	Space Relations	24
6	Verbal Ability	35
Total Questions		180

KDAT is conducted in an online mode under strict monitoring. Candidates are provided with pre-counselling to explain why KDAT is not an achievement test and to emphasise the importance of answering each question honestly. Although each question has a time limit, candidates are allowed to take breaks between different categories of questions. This enables them to complete the test without undue time pressure. Upon completion, a graph, as illustrated in Figure 5 (Kerala Higher Secondary Education Department, 2024), is generated, displaying the percentile

scores for each category of questions. A career counsellor skilled in Differential Aptitude Tests can analyse this graph to interpret the candidate's ability profile and recommend suitable careers. While making career recommendations based on the KDAT ability profile graph, it is strongly advised to incorporate the candidate's other aspects, like personality traits and vocational interests, to ensure a comprehensive and personalised career recommendation.




Figure 5. A Sample KDAT Test Output

5.2.Enhancing DAT Interpretation through Automated Recommender Systems

Manual interpretation of DAT results and career recommendations by a career counsellor has both advantages and limitations. The primary advantage lies in the counsellor's ability to enhance the recommendations by considering DAT results. They can augment career decisions, considering other candidates' factors like personality traits and vocational interests. However, a significant limitation is the counsellor's subjectivity involved in interpreting the DAT results. Different career counsellors may provide varying interpretations depending on their expertise. Gaps in their knowledge can lead to misinterpretation of test results. As part of this study, a survey-based analysis of career counsellors' potential subjectivity and related

interpretation bias in the interpretation of DAT results was conducted among career counsellors. The study employed KDAT test results in conjunction with career counsellors trained in the use of the KDAT. (Kerala Higher Secondary Education Department, 2024). The analysis revealed significant variation in the interpretation of KDAT results among different career counsellors. A sample result of a single career and a counsellor's aptitude profiling for the career is illustrated in Figure 6. The figure shows that for a career in Agricultural Engineering, there are both comparable and conflicting parameter value assignments by different counsellors, highlighting the subjective nature of parameter interpretation. Another important aspect is that a counsellor's alertness and cognitive load can affect the quality of advice. An AI-based decision-support tool that generates evidence-based, aptitude-aligned career options provides consistent, reviewable support to a career counsellor with data-driven direction to their counselling process. An AI-based tool can assist, but not replace a career counsellor. Additionally, an AI-based system can address the poor career counsellor-to-student ratio, making the accessibility of psychometric tests and career recommendations available to every student.




Figure 6. Varying Interpretations of Career Counsellors on DAT Parameter Values.
As Presented in the Author's Published Work (Rahim & Basheer, 2025a) Under a CC BY License.

5.3.Utilizing KDAT in the Proposed Career Recommender System

The proposed career recommender system utilises the Kerala Differential Aptitude Tests (KDAT) to generate user profiles for candidates. The user profile produced by

KDAT, i.e., the test output, is fed into the proposed system to automatically recommend careers that align with the candidate's profile. Traditionally, KDAT results are interpreted by a career counsellor, relying on their expertise in Differential Aptitude Tests. Our research, however, focuses on integrating AI algorithms with novel recommendation architectures and an innovative explainability framework to automate the career recommendation process. This can address a significant gap in psychometric testing tools like KDAT, which lack automated career recommendation capabilities.

The proposed career recommendation system aims to address the limitations discussed above by automating the career recommendation process. It ensures a standardised interpretation of aptitude test results, as generated by a Differential Aptitude Testing tool like KDAT, and recommends careers aligned with the candidate's aptitude profile. Figure 7 illustrates the integration of KDAT within our proposed career recommendation system.




Figure 7. Utilising KDAT in the Proposed Recommender System

5.4. Normalization and Standardisation of User Profile Data

Normalisation is a technique used to rescale feature values to a specific range, addressing the problem of varying feature importance caused by disproportionate value scales. This process enhances the efficiency of machine learning models such as KNN, SVM, and DNN (Han et al., 2022) by ensuring that features with larger value

ranges do not dominate the learning process, thereby reducing model bias towards such features. The equation for normalisation, which rescales data to a specified range, is given in Equation 1.

$$x' = \frac{x - \min(x)}{\max(x) - \min(x)} \quad (1)$$

Where:

- x : The original value of the data point,
- $\min(x)$: The minimum value in the dataset,
- $\max(x)$: The maximum value in the dataset,
- x' : The normalised value.

Standardisation is a technique used to transform data points in a dataset to have a mean of 0 and a standard deviation of 1. It assumes a standard distribution and is applied in machine learning models, particularly in statistical models, to enhance efficiency as these models are highly sensitive to feature variance. Standardisation is also beneficial for gradient descent optimisation. Since each test battery score of DAT is on a percentile scale ranging from 0 to 100, normalisation is not applied to the user profile data. However, standardisation is employed on the DAT scores using the formula shown in Equation 2 (Ali, Faraj, & Koya, 2014).

$$Z_i = \frac{(X_i - \mu(X))}{\sigma(X)} \quad (2)$$

Where,

- X_i is the percentile score obtained by the candidate for each test battery,
- $\mu(X)$ is the average of scores across all test batteries,
- $\sigma(X)$ is the standard deviation of scores across all test batteries, and
- Z_i is the standardized value for the i^{th} test battery.

Following the administration of the Differential Aptitude Test (DAT), the user profile is acquired as represented below:

$$DAT_{userprofile} = [X1, X2, X3, X4, X5, X6]$$

Where each X_i represents DAT score for each test battery.

After the standardisation, the final user profile record can be represented as:

$$StandardizedDAT_{userprofile} = [Z1, Z2, Z3, Z4, Z5, Z6]$$

This standardised user profile will be input into the recommendation engine to generate relevant career recommendations for the candidate.

In conclusion, KDAT is utilised for generating candidate profiles based on Differential Aptitude Tests. It is specifically chosen because it has been standardised for Class X students in Kerala, taking into account the socio-economic and cultural differences of the population. Once the KDAT results are obtained, they undergo standardisation and are subsequently fed into the recommendation engine.

CHAPTER 6

CURATING A COMPREHENSIVE DATASET FOR THE PROPOSED CAREER RECOMMENDATION SYSTEM BASED ON DAT

To develop a career recommender system based on a candidate's aptitude profile using the Content-Based Filtering (CBF) approach, a career profile dataset aligned with aptitudes similar to those of the user profile is required. As discussed in the previous section, in the review of existing datasets, although multiple datasets (Pimplikar, Patil, Dhar, Shrivastav, Kaggle, Hugging Face) on careers are available, there is a lack of a comprehensive dataset based on aptitude profiles, encompassing all probable real-world careers. To address this challenge, this study utilised the ability profile information from the Occupational Information Network Database (O*NET) to create a comprehensive aptitude-based career dataset.

6.1. An Overview of the Occupational Information Network Database

The Occupational Information Network database (O*NET) was created with the support of the US Department of Labor, Employment and Training Administration (National Center for ONET Development, 2024). The O*NET database encompasses numerous standardised and job-specific descriptors for nearly 1,000 careers. This publicly accessible resource, offered free of charge, is regularly refreshed through a multifaceted data collection process. Data sources include jobholders, industry experts, occupational evaluators, employer job advertisements, and feedback from customers and professional organisations.

Data content areas in O*NET include:

- Worker Characteristics (e.g., Abilities, Interests, Work Styles)
- Worker Requirements (e.g., Education, Knowledge, Skills)
- Experience Requirements (e.g., On-the-Job Training, Work Experience)

- Occupational Requirements (e.g., Detailed Work Activities, Work Context)
- Occupation-Specific Information (e.g., Job Titles, Tasks, Technology Skills)

The O*NET database is continuously updated with information on various job zones, emerging and in-demand technologies, essential skills, and more. It is available under a Creative Commons License. It provides valuable support to researchers and developers in the career domain. In summary, the Occupational Information Network Database offers a dynamic and comprehensive resource on occupations, featuring extensive, up-to-date information about a wide range of careers. It aids job seekers, employers, educators, and researchers in making well-informed decisions related to career development, educational planning, and economic forecasting.

6.2.Ability Profile Information Related to Careers in O*NET

The ability profile of the Occupational Information Network database (O*NET) contains 52 different ability features for each career in the database. All the ability features and their interpretation are given in Table 5 (ONET Resource Center, 2024)).

Table 5.Ability Parameters and their Interpretations in O*NET

S.No	Ability	Interpretation
1	Oral Comprehension	The ability to listen to and understand information and ideas presented through spoken words and sentences.
2	Written Comprehension	The ability to read and understand information and ideas presented in writing.
3	Oral Expression	The ability to communicate information and ideas in speaking so others will understand.
4	Written Expression	The ability to communicate information and ideas in writing so others will understand.
5	Fluency of Ideas	The ability to come up with a number of ideas about a topic (the number of ideas is important, not their quality, correctness, or creativity).
6	Originality	The ability to come up with unusual or clever ideas about a given topic or situation, or to develop creative ways to solve a problem.
7	Problem Sensitivity	The ability to tell when something is wrong or is likely to go wrong. It does not involve solving the problem, only recognizing that there is a problem.
8	Deductive Reasoning	The ability to apply general rules to specific problems to produce answers that make sense.

9	Inductive Reasoning	The ability to combine pieces of information to form general rules or conclusions (includes finding a relationship among seemingly unrelated events).
10	Information Ordering	The ability to arrange things or actions in a certain order or pattern according to a specific rule or set of rules (e.g., patterns of numbers, letters, words, pictures, mathematical operations).
11	Category Flexibility	The ability to generate or use different sets of rules for combining or grouping things in different ways.
12	Mathematical Reasoning	The ability to choose the right mathematical methods or formulas to solve a problem.
13	Number Facility	The ability to add, subtract, multiply, or divide quickly and correctly.
14	Memorization	The ability to remember information such as words, numbers, pictures, and procedures.
15	Speed of Closure	The ability to quickly make sense of, combine, and organize information into meaningful patterns.
16	Flexibility of Closure	The ability to identify or detect a known pattern (a figure, object, word, or sound) that is hidden in other distracting material.
17	Perceptual Speed	The ability to quickly and accurately compare similarities and differences among sets of letters, numbers, objects, pictures, or patterns. This includes comparing a presented object with a remembered object.
18	Spatial Orientation	The ability to know your location in relation to the environment or to know where other objects are in relation to you.
19	Visualization	The ability to imagine how something will look after it is moved around or when its parts are moved or rearranged.
20	Selective Attention	The ability to concentrate on a task over a period of time without being distracted.
21	Time Sharing	The ability to shift back and forth between two or more activities or sources of information (such as speech, sounds, touch, or other sources).
22	Arm-Hand Steadiness	The ability to keep your hand and arm steady while moving your arm or while holding your arm and hand in one position.
23	Manual Dexterity	The ability to quickly move your hand, your hand together with your arm, or your two hands to grasp, manipulate, or assemble objects.
24	Finger Dexterity	The ability to make precisely coordinated movements of the fingers of one or both hands to grasp, manipulate, or assemble very small objects.
25	Control Precision	The ability to quickly and repeatedly adjust the controls of a machine or a vehicle to exact positions.
26	Multilimb Coordination	The ability to coordinate two or more limbs (for example, two arms, two legs, or one leg and one arm) while sitting, standing, or lying down. It does not involve performing the activities while the whole body is in motion.
27	Response Orientation	The ability to choose quickly between two or more movements in response to two or more different signals (lights, sounds, pictures). It includes the speed with which the correct response is started with the hand, foot, or other body part.

28	Rate Control	The ability to time your movements or the movement of a piece of equipment in anticipation of changes in the speed and/or direction of a moving object or scene.
29	Reaction Time	The ability to quickly respond (with the hand, finger, or foot) to a signal (sound, light, picture) when it appears.
30	Wrist-Finger Speed	The ability to make fast, simple, repeated movements of the fingers, hands, and wrists.
31	Speed of Limb Movement	The ability to quickly move the arms and legs.
32	Static Strength	The ability to exert maximum muscle force to lift, push, pull, or carry objects.
33	Explosive Strength	The ability to use short bursts of muscle force to propel oneself (as in jumping or sprinting), or to throw an object.
34	Dynamic Strength	The ability to exert muscle force repeatedly or continuously over time. This involves muscular endurance and resistance to muscle fatigue.
35	Trunk Strength	The ability to use your abdominal and lower back muscles to support part of the body repeatedly or continuously over time without "giving out" or fatiguing.
36	Stamina	The ability to exert yourself physically over long periods of time without getting winded or out of breath.
37	Extent Flexibility	The ability to bend, stretch, twist, or reach with your body, arms, and/or legs.
38	Dynamic Flexibility	The ability to quickly and repeatedly bend, stretch, twist, or reach out with your body, arms, and/or legs.
39	Gross Body Coordination	The ability to coordinate the movement of your arms, legs, and torso together when the whole body is in motion.
40	Gross Body Equilibrium	The ability to keep or regain your body balance or stay upright when in an unstable position.
41	Near Vision	The ability to see details at close range (within a few feet of the observer).
42	Far Vision	The ability to see details at a distance.
43	Visual Color Discrimination	The ability to match or detect differences between colors, including shades of color and brightness.
44	Night Vision	The ability to see under low light conditions.
45	Peripheral Vision	The ability to see objects or movement out of the corner of your eye.
46	Depth Perception	The ability to judge which of several objects is closer or farther away from you, or to judge the distance between you and an object.
47	Glare Sensitivity	The ability to see objects in the presence of glare or bright lighting.
48	Hearing Sensitivity	The ability to detect or tell the differences between sounds that vary in pitch and loudness.
49	Auditory Attention	The ability to focus on a single source of sound in the presence of other distracting sounds.
50	Sound Localization	The ability to tell the direction from which a sound originated.
51	Speech Recognition	The ability to identify and understand the speech of another person.
52	Speech Clarity	The ability to speak clearly so others can understand you.

From the analysis of the ability features in the Occupational Information Network (O*NET) database, it is evident that certain features closely align with the Differential Aptitude Test batteries. However, not all the 52 features of the O*NET ability profile are relevant to the parameters of the Differential Aptitude Test. Consequently, a detailed analysis and comparative mapping of O*NET ability features to the Differential Aptitude Test batteries were conducted in this study. This mapping facilitated the creation of career profiles for 612 relevant, unique careers, selected from the more than 1,000 careers available in the O*NET database. For each ability feature in the O*NET database, two related numerical values are provided: *importance* and *level*. The *importance score* represents how critical a particular ability is for performing a specific occupation, while the *level score* indicates the degree of proficiency or expertise required for that ability in the context of the occupation.

6.3.Mapping O*NET Ability Features to DAT Test batteries

The analysis of 52 features of the ability profile information of the Occupational Information Network (O*NET) database helped to identify Differential Aptitude-related features available in the database. This analysis facilitated the creation of a DAT-based career dataset and supported the building of the proposed career recommender engine. For each career entry in O*NET, comprising 52 unique features (National Center for ONET Development, 2024), we separated and mapped the most consequential ones to our specific DAT test batteries. The mapping of the O*NET ability profile to DAT test batteries is given in Table 6. The process of mapping is guided by expert domain knowledge of the Differential Aptitude Test batteries.

6.4.Converting O*NET Ability Features to DAT Test batteries

Once the mapping has been done, the next task is to convert the ability profile value for each of the selected features of all careers in the O*NET to their DAT parameter values. The study considered 612 unique significant careers out of more than 1000 from the O*NET database, excluding some insignificant ones. For each career in the O*NET database, its ability profile provides two most important parameter values: *Importance* and *Level*. The Importance score refers to “*how necessary the ability is*”, and the Level score refers to “*How much the ability is needed*”.

Table 6. Mapping of DAT Test Batteries to O*NET Ability Profile Parameters

DAT Parameter	Ability Profile of Occupational Information Network Database
Numerical Ability	Mathematical Reasoning, Number Facility
Abstract Reasoning	Inductive Reasoning, Deductive Reasoning, Fluency of Ideas
Verbal Reasoning	Oral Comprehension, Written Comprehension
Mechanical Reasoning	Problem Sensitivity, Manual Dexterity, Finger Dexterity, Control Precision
Space Relations	Visualisation
Verbal Ability	Written Expression
Clerical Speed and Accuracy	Perceptual Speed

To convert the ability profile parameter values into their corresponding DAT test battery values, we utilized the *Importance* and *Level* scale values from the O*NET ability profile parameters. Normalization was necessary due to the differing scales: *Importance* values ranged from 1 to 5, and *Level* values ranged from 0 to 7 (National Center for ONET Development, 2024). The normalization process is done, employing the standard normalization equation, as shown in Equation 3 (Ali, Faraj, & Koya, 2014).

$$X_{normalized} = \frac{X - \min(X)}{\max(X) - \min(X)} \quad (3)$$

Where, X represents the original Importance score or Level score.

Given that *Importance* values fall between 1 to 5 and *Level* values span from 0 to 7, we computed normalized scores using Equations 4 and 5.

$$Score_{Importance} = \frac{I_i - 1.0}{5.0 - 1.0} \quad (4)$$

$$Score_{Level} = \frac{L_i - 0.0}{7.0 - 0.0} \quad (5)$$

Where I_i and L_i represent the Importance score Level score for parameter i respectively.

After normalizing the values for each mapped ability profile parameter, and factoring in the percentile score of the respective DAT parameter values, we employ Equation 6 to calculate the corresponding DAT parameter values.

$$X_i = \frac{Score_{Importance} + Score_{Level}}{2.0} * 100 \quad (6)$$

Where X_i is the score in the percentile of i^{th} DAT test battery.

Table 7 outlines the process of generating a career profile for the profession ‘Lawyer’ using the procedure above.

Hence, the skill set for the career Lawyer in terms of DAT parameters is determined as:

$$DAT_{Lawyer} (NA, AR, VR, MR, SR, VA) = [37, 67, 80, 30, 37, 75]$$

This approach was used to calculate the DAT parameter values for 612 distinct careers. This process established the foundational career profile dataset for our content-based recommender system. An illustrative sample is presented in Table 8.

Table 7. The Process of Generating a Career Profile for the Profession of Lawyer

DAT Test Battery	Mapped O*NET Ability Parameters	Importance Score	Level Score	Conversion Process	Computed Value (Rounded Up)	Finding Average Value	Final DAT Score
Numerical Ability	Mathematical Reasoning	2.62	2.38	$((2.62-1)/(5-1))*100+((2.38-0)/(7-0))*100)/2$	38	$(38+36)/2$	37
	Numerical Facility	2.38	2.5	$((2.38-1)/(5-1))*100+((2.5-0)/(7-0))*100)/2$	36		
Abstract Reasoning	Inductive Reasoning	4	4.25	$((4.0-1)/(5-1))*100+((4.25-0)/(7-0))*100)/2$	68	$(68+69+63)/3$	67
	Deductive Reasoning	4	4.38	$((4.0-1)/(5-1))*100+((4.38-0)/(7-0))*100)/2$	69		
	Fluency of Ideas	3.75	3.88	$((3.75-1)/(5-1))*100+((3.88-0)/(7-0))*100)/2$	63		
Verbal Reasoning	Oral Comprehension	4.5	5	$((4.5-1)/(5-1))*100+((5.0-0)/(7-0))*100)/2$	80	$(80+80)/2$	80
	Written Comprehension	4.5	5	$((4.5-1)/(5-1))*100+((5.0-0)/(7-0))*100)/2$	80		

				$0/(7-0)*100/2$			
Mechanical Reasoning	Problem Sensitivity	3.88	4	$((((3.88-1)/(5-1))*100+((4.0-0)/(7-0))*100)/2$	65	$(65+8+30+15)/4$	30
	Manual Dexterity	1.38	0.38	$((((1.38-1)/(5-1))*100+((0.38-0)/(7-0))*100)/2$	8		
	Finger Dexterity	2.25	2	$((((2.25-1)/(5-1))*100+((2.0-0)/(7-0))*100)/2$	30		
	Control Precision	1.75	0.75	$((((1.75-1)/(5-1))*100+((0.75-0)/(7-0))*100)/2$	15		
Space Relations	Visualization	2.5	2.5	$((((2.5-1)/(5-1))*100+((2.5-0)/(7-0))*100)/2$	37	37	37
Verbal Ability	Written Expression	4.12	5	$((((4.12-1)/(5-1))*100+((5.0-0)/(7-0))*100)/2$	75	75	75

Table 8. Samples from the Generated Dataset: DAT Parameter Values for a Few Careers

SI No	Career	NA	AR	VR	MR	SR	VA
1	Lawyer	37	67	80	30	37	75
2	Human Resources Managers	50	64	71	26	36	67
3	Robotics Engineers	65	70	74	58	69	69
4	Mechanical Engineers	67	65	73	62	59	58
5	Statisticians	81	67	74	30	42	73
6	Physicists	85	78	81	32	64	78
7	Economists	65	68	74	13	34	74
8	Journalists	25	60	70	30	34	67
9	Business Intelligence Analysts	55	66	72	16	30	70
10	Poets Lyricists and Creative Writers	16	62	66	43	45	81

6.5. Generating Career Clusters

After generating ability profiles for a comprehensive dataset comprising 612 distinct occupational roles, a clustering methodology is applied to identify career clusters. The process of clustering is characterized by its shared ability profiles. This is required since there can be a set of careers, and not a single career, whose aptitude profiles align with the candidate's aptitude profile, so that a candidate can be recommended a similar set of careers. Various clustering algorithms, including Agglomerative Clustering, Gaussian Mixture, Birch, and K-Means, were employed. This was attempted with the insight that no single clustering algorithm is universally effective

across all datasets (Xu & Wunsch, 2005). The ultimate choice of the K-Means clustering approach is determined based on an evaluation involving Silhouette Scores (Shahapure & Nicholas, 2020). The finding was supplemented by a visual assessment of the resultant cluster structures. Table 9 provides silhouette Scores for different clustering algorithms. The relatively higher, yet still low, Silhouette score (K-Means: 0.23) is acceptable, as career clusters are not expected to exhibit sharply defined boundaries. In practice, the attributes of different careers often overlap, rather than forming discrete, mutually exclusive categories.

After the clustering algorithm was chosen, the next step involved determining the optimal number of clusters for the dataset. This was accomplished by applying both the Silhouette Score method and the Elbow Method.

Table 9. Silhouette Scores for Different Clustering Algorithms.

SI No	Clustering Algorithm	Silhouette Score
1	K-Means	0.23
2	Agglomerative	0.18
3	Birch	0.18
4	Gaussian Mixture	0.11

The Silhouette Score, utilized to determine the optimal number of clusters, is computed via the following mathematical formula given in Equation 7.

$$S(k) = \frac{1}{N} \sum_{i=1}^N \frac{b(i) - a(i)}{\max\{a(i), b(i)\}} \quad (7)$$

Where,

$S(k)$ is the Silhouette Score for a specific number of clusters (k),

N represents the total number of data points,

$a(i)$ denotes the average distance from the i^{th} data point to other data points within the same cluster,

$b(i)$ signifies the smallest average distance from the i^{th} data point to data points in a different cluster, minimized over clusters.

The elbow method involves drawing a graph that shows the Within-Cluster Sum of Squares (WCSS) against the number of clusters. The goal is to find the point on the graph where the curve bends, called the “*elbow point*”. This point can indicate the optimal number of clusters to be determined while clustering the dataset. The mathematical expression given in Equation 8 is used for calculating $WCSS$.

$$WCSS_k = \sum_{i=1}^k \sum_{j=1}^n (X_{ij} - C_i)^2 \quad (8)$$

where,

$WCSS_k$ is the $WCSS$ for a specific cluster count k , X_{ij} represents data point j in cluster i , C_i is the centroid of the cluster, n signifies the number of data points.

This resulted in the selection of 8 as the optimal number of clusters. This selection procedure of 8 optimal cluster employing silhouette score and the elbow method is illustrated in Figures 8 and 9. As a consequence, the K-Means clustering algorithm, selected from multiple clustering algorithm based on silhouette score, is employed to generate 8 distinct clusters from the dataset.




Figure 8. Finding the Optimum Number of Clusters Using the Silhouette Method.

Upon conducting a visual analysis of the generated clusters, it was ascertained that certain clusters exhibit significant variances in attribute values among the included

careers, suggesting the need for further stratification. Sub-clustering was subsequently employed on these clusters, leveraging the elbow method to systematically partition them into additional sub-clusters. As a result, a total of 49 clusters were generated.




Figure 9. Finding the Optimum Number of Clusters Using Elbow Methods

Figures 10a and 10b illustrate a subset of the generated clusters, presenting them in conjunction with their respective DAT parameter values.

With the clusters partitioned into sub-clusters, the process adopts a semi-supervised approach, initially clustering the unlabelled dataset, and manually labelling the records in the clusters later. Systematically labelling each record within these sub-clusters facilitated the construction of machine learning models. Multiple models were built on the dataset using different classification algorithms, including Artificial Neural Networks (ANN), Random Forests Classifier (RF), Support Vector Machine (SVM), K-Nearest Neighbour (KNN), Decision Tree Classifier, and Gaussian Naive Bayes (GNB). Table 10 provides a comprehensive overview of the training and testing accuracy of these models when applied to the dataset.




Figure 10a. Sample from the Generated Clusters




Figure 10b. Another Sample from the Generated Clusters.

Table 10. Training and Testing Accuracies of Various Models on the Original Dataset

Sl No	Classifier	Training Accuracy	Testing Accuracy
1	ANN	1.00	0.72
2	Random Forests Classifier	1.00	0.75
3	Decision Tree	1.00	0.63
4	GNB	0.95	0.87
5	KNN	0.86	0.70
6	SVM	0.95	0.87

The data presented in Table 10 highlights a notable risk of over-fitting (Ying, 2019), and this causes the substantial disparity between training and testing accuracies. This discrepancy can be attributed to the dataset's limited size, containing just 612 records. To resolve the challenges posed by overfitting, we took measures to augment the records (Rebuffi et al., 2021) within each cluster of the dataset, resulting in the generation of approximately fifty thousand additional records. The subsequent section provides a comprehensive explanation of the augmentation process.

6.6. Augmenting the Dataset to Address Overfitting

The primary objective of enhancing the dataset's records arises from its limited size. The dataset comprises roughly 612 records distributed across 44 distinct clusters. Each cluster contains only a small number of records, which hinders the classification algorithm's ability to generalize patterns within these clusters. This disparity between training and testing accuracy is a direct result of this issue. Thus, it becomes necessary to expand the dataset by generating additional records that are similar to the existing ones, thereby creating a larger dataset for building classifiers. Moreover, it's important to recognize that each record in the original dataset corresponds to a specific career. So, it would be overly simplistic to base a career profile solely on a single record from the dataset. In reality, there are likely to be some variations in the DAT career profiles for a particular profession. Hence, augmenting each record related to a career within the dataset is essential to capture the nuanced patterns associated with those careers. In essence, data augmentation addressing these factors is to be applied to enhance the model's generalization prowess by systematically accounting for the real-world data variations.

We utilize Gaussian Noise data augmentation (He et al., 2019) to generate extra related records. Gaussian Noise is deliberately selected as it can introduce random fluctuations into the dataset. The aim is to bolster the model's robustness and enhance its adaptability to real-world scenarios. This approach fosters resilience to noise, promotes the learning of more generalized features. It also injects diversity into the

dataset through random disturbances and thus serves as a regularization technique. This enables the cultivation of the model's invariance to minor input variations.

We apply Gaussian Noise data augmentation within our dataset using the following mathematical expression:

$$Y_i = X_i + \epsilon_i$$

Where:

Y_i represents i^{th} element of the augmented data point,

X_i represents i^{th} element of the original data point,

ϵ_i is the random value drawn from a Gaussian distribution with a mean of 0 and a variance of σ^2 .

This expression illustrates that each element of the augmented data point, denoted as Y_i , is created by incorporating a random value ϵ_i , drawn from a Gaussian distribution, into the corresponding element of the original data point X_i . The degree of noise introduced is determined by the standard deviation, denoted as σ . Algorithm 1 outlines the entire process:

Algorithm 1: Data_Augmentation_Using_Guassain_Noise

Inputs : Input data point X_i , Augmentation Factor F , Noise standard deviation σ ,

Output: Augmented data list ADL

Repeat Augmentation Factor F times

Generate Gaussian noise for each data point i:

Noise, ϵ_i = randomly sample from a Gaussian distribution with a mean of -2.0 and a standard deviation σ of 5.0

Generate Augmented Record:

$$Y_i = X_i + \epsilon_i$$

Append the record Y_i to Augmented data list ADL.

Return Augmented data list ADL

For our specific augmentation procedure, we drew random values from a Gaussian distribution with a mean(μ) of -2.0 and a noise standard deviation (σ) of 5.0. This process was iterated for each element of the data point with an augmentation factor (F) of 100, resulting in a dataset enriched with controlled noise. These values were finalized after repeated evaluation of the performance of the models built upon the resultant dataset.

By applying the algorithm described above, we augmented each record within every cluster of the original dataset, consisting of 612 records. As a result, our dataset has been expanded to include over 57,000 records. We systematically removed any potential duplicate records from the dataset. A selection of sample clusters is illustrated in Figure 11.




Figure 11. Samples from the Augmented Clusters.

6.7.Evaluation of the Validity of the Augmented Dataset

Augmenting and expanding a dataset from 612 records to over 57,000 records can introduce potential issues. The augmented dataset may fail to capture the inherent patterns present in the original dataset. As a result, a model trained on the augmented dataset may learn from the noise-introduced patterns present in the augmented data. It is important that the augmented dataset accurately reflects the patterns found in the original data. This underlines the need for a thorough evaluation of the augmented dataset's validity. This section will analyse the performance and validity of the augmented dataset.

To assess the performance of the augmented dataset, we constructed machine learning models using the ANN, SVM, and Random Forest Classifier algorithms on the dataset. Table 11 showcases the performance of these models on the train-test split from the augmented dataset.

Table 11. Performance of Different Classifiers on the Augmented Dataset

SI No	Classifier	Training Accuracy	Testing Accuracy
1	ANN	0.97	0.94
2	Random Forests Classifier	1.00	0.99
3	KNN	0.98	0.97

The method we employed to assess the validity of the augmented dataset consisted of constructing several machine-learning models using the augmented data alone, keeping the original dataset as test data. We then compared the testing accuracies of these models on the original dataset. If the testing accuracies of the models on the original dataset surpass those on the train-test split, it confirms that the augmented dataset clearly captures the patterns present in the original data. As depicted in Table 12, the models demonstrated exceptionally high testing accuracies (100%) on the original dataset.

Table 12. Performance of Classifiers Trained on the Augmented Dataset Against the Original Dataset

SI No	Classifier	Testing Accuracy
1	ANN	1.00
2	Random Forests Classifier	1.00
3	KNN	1.00

This observation confirms that the augmented dataset clearly captures the patterns present in the original dataset. With this confirmation, we can confidently proceed with the augmented dataset to build the career recommendation system.

CHAPTER 7

ARCHITECTURE OF THE PROPOSED CAREER RECOMMENDATION ENGINE

The architecture of our career recommender system follows the Content-Based Filtering (CBF) approach. Even if there exist other approaches like collaborative filtering, knowledge-based systems, etc. (Khanal, Prasad, Alsadoon, & Maag, 2020), CBF is selected as this is particularly suitable considering the intrinsic attributes of the user profile and the career profile in our problem domain. Collaborative filtering recommends items by identifying users with similar profiles and leveraging their shared interactions to suggest relevant content (Jalili, Ahmadian, Izadi, Moradi, & Salehi, 2018). In contrast, content-based filtering focuses on individual user preferences and item characteristics, recommending new items with attributes similar to those the user has engaged with in the past (Thorat, Goudar, & Barve, 2015). The proposed system employs a content-based filtering approach, using the Kerala Differential Aptitude Test (KDAT) to generate the user profile. The career profile is constructed based on ability profile information extracted from the Occupational Information Network (O*NET) database, as explained in the previous section.

While designing the proposed career recommendation system, the primary objective was to generate diverse career recommendations aligned with the candidate's profile. To achieve this, multiple strategies were integrated into the architecture of the proposed recommendation engine. The system is built on machine learning models trained on the developed dataset. The architecture follows a novel hierarchical and multi-tiered architecture to enhance diversity in recommendations. Additionally, a voting classifier mechanism is incorporated to further improve this aspect. Both strategies contribute to increasing diversity in recommendations, enhancing robustness to variance and bias. This also facilitates improving flexibility in decision-making and addressing model-specific weaknesses.

7.1. Ensemble Model Implementation Leveraging a Voting Classifier

The career recommendation engine is constructed using an ensemble of machine learning models, employing a voting classifier (Manconi et al., 2022). The proposed recommendation architecture integrates predictions from multiple algorithms, including Artificial Neural Networks (ANN), K-Nearest Neighbours (KNN), Support Vector Machines (SVM), Random Forest, Decision Tree, and Gaussian Naïve Bayes.

The main objective of building the model based on the ensemble method is to achieve diversity in recommendations, which can also provide flexibility in decision-making. There are two main approaches to building a voting classifier: Hard voting and soft voting. Both strategies are illustrated in Figure 12. In hard voting, each base classifier will come up with a discrete label prediction, and the final prediction will be decided by the majority vote. The process can be explained as shown in the mathematical Equation 9.

$$\hat{y} = \arg \max_{c \in C} \sum_{i=1}^M 1(h_i(x) = c) \quad (9)$$

Where:

C is the set of all possible classes,

M is the number of base classifiers,

$h_i(x)$ is the prediction of i^{th} classifier for input x ,

$1(h_i(x)=c)$ is an indicator function that returns 1 if classifier i predicts class c , otherwise 0,

The final predicted class $\hat{y}$ is the one with the highest vote count.




Figure 12. Hard Voting vs. Soft Voting Classifier.

In soft voting, the probability distributions of the classifiers over different classes are considered. The final prediction is made based on the sum of these probabilities. Equation 10 explains the concept mathematically.

$$\hat{y} = \arg \max_{c \in C} \sum_{i=1}^M w_i P_i(c|x) \quad (10)$$

Where:

$P_i(c|x)$ is the predicted probability of class c given input x by classifier i ,

w_i is the weight assigned to classifier i ,

The class $\hat{y}$ with the highest weighted probability sum selected.

We selected soft voting for our model prediction due to its ability to capture uncertainty and nuanced patterns. Additionally, the approach can balance contributions from diverse models, handle overlapping career profiles, and generate top-N recommendations instead of a single prediction.

The overall voting classifier architecture of our model is illustrated in Figure 13.




Figure 13. The Architecture of the Voting Classifier Adapted to the Proposed Recommender System.

Considering the methodology adopted for model training, the dataset was divided into training and testing subsets using an 80:20 split. To improve evaluation reliability and reduce sampling bias, k-fold cross-validation was additionally performed during model evaluation. Data augmentation was applied exclusively to the training set. Model testing was conducted using the original non-augmented data points.

7.2. The Innovative Hierarchical and Multi-Tiered Architecture of the Recommendation Engine

The proposed system employs a novel hierarchical and multi-tiered recommendation approach. The target is to offer diverse career recommendations tailored to a candidate's DAT results. As previously discussed, the original dataset underwent an initial division into eight distinct clusters based on insights derived from the silhouette and elbow methods. Subsequently, these clusters were further subdivided into 49 clusters, labelled, and augmented to yield the final dataset. Thus, we have a higher-level set of career clusters and a fine-tuned set of lower-level career clusters. This will help us to generate a hierarchical career recommender system.

The recommendation system consists of a *Diverse Recommender Module* followed by a *Refined Recommendation Module*. The *Diverse Recommendation Module* tries to provide diversity in career recommendations at a higher level. The *Refined*

Recommendation Module fine-tunes each set of recommendations delivered by the *Diverse Recommendation Module*.

The *Diverse Recommendation Module* is constructed using the initially generated dataset with top top-level eight clusters. Simultaneously, the *Refined Recommendation Module*, a set of eight second-level classifiers, are developed on datasets generated by further sub-clustering the original top-level eight clusters. These second-level refinement classifier models, named from *RefMod1* to *RefMod8*, act as a *Refinement Suite* to fine-tune the recommendations made by the *Diverse Recommendation Module*. While the *Diverse Recommendation Module* provides diversity in recommendations, these recommendations are fine-tuned using the *Refinement Suite*. Both the *Diverse Recommendation Module* and *Refined Recommendation Module* are modelled employing voting classifiers - explained in the previous section- incorporating machine learning algorithms such as ANN, KNN, SVM, Random Forest Classifier, Decision Tree, and Gaussian Naive Bayes Classifier.

The candidate's DAT results serve as input for the *Diverse Recommendation Module*. This classifier module identifies the best three diverse career classes from the eight available in the higher-level model. These three classes are fed to the *Refined Recommendation Module*. The module then loads the corresponding best three second-level classifiers from the *Refinement Suite*, containing eight different classifiers (named *RefMod1* to *RefMod8*), to fine-tune the recommendations made by the *Diverse Recommendation Module*. Each of these identified three sub-models, called *FirstPick*, *SecondPick* and *ThirdPick* models, in turn, provides the best three sets of career clusters to a candidate aligned to their aptitude profile. The multiple output ends, called *First Tier*, *Second Tier* and *Third Tier*, thus provide career recommendations in the order of their significance for the candidate, resulting in a total of nine different sets of career clusters. This comprehensive array of options provides the candidate with diverse career paths aligned with their DAT results. Additionally, career counsellors can derive valuable insights from the recommender system to enhance career counselling sessions. The entire process is illustrated in Figure 14, and the recommendation process is further detailed in Pseudocode 1, as

presented in author's published work (Rahim & Basheer, 2025a) under a CC BY license.




Figure 14. The Hierarchical and Multi-tiered Architecture of the Proposed Career Recommendation Engine. As Presented in the Author's Published Work (Rahim & Basheer, 2025a) Under a CC BY License

The overall recommendation process starts with the candidate being tested for abilities in different aptitude dimensions using the Kerala Differential Aptitude Test (KDAT). The results are then fed into the proposed hierarchical and multi-tiered career recommendation engine, which provides diverse career recommendations aligned with the candidate's aptitude profile. The entire recommendation process is depicted in Figure 15.

Pseudocode 1. The Multi-Tiered Recommendation Process for Providing Diverse Career Recommendations.

1. *user_profile* = GetUserProfile() // DAT result for the candidate
2. *top_classes* = GetBest3Classes(DiverseRecommenderModule, *user_profile*)
3. For each class in *top_classes*:

```

refinedModel = LoadRefinedModel(RefinementSuite, class)

// Obtain predictions from refinedModel

tier1_predictions = GetTier1Predictions(refinedModel, user_profile)
tier2_predictions = GetTier2Predictions(refinedModel, user_profile)
tier3_predictions = GetTier3Predictions(refinedModel, user_profile)

```




Figure 15. Illustration of the Recommender System Design with its Recommendation Process. As Presented in the Author's Published Work (Rahim & Basheer, 2025a) Under a CC BY License

7.3.The Proposed System's Multi-Tiered Recommendation Approach

The model is designed to generate diverse career recommendations for a candidate by processing their DAT results within the system. The system, via its *Diverse Recommendation Module*, initially identifies three diverse sets of best-suited career clusters for the candidate's profile. Since initial recommendations (career clusters) contain nearly eighty careers on average, the *Refined Recommendation Module* further

refines these suggestions through specialised sub-models. Each sub-model finally provides multi-level recommendations to candidates, from its *Tier 1*, *Tier 2*, and *Tier 3* output ends, that are aligned with their ability profile and ranked according to relevance.

Thus, the system offers a diverse range of career options aligned with the candidate's aptitude profile, allowing for flexible yet well-suited career choices based on the recommendations provided by the recommendation engine. A career cannot be treated as a fixed or mutually exclusive category like cat and rat classification. Careers often overlap in the skills and abilities they demand. The proposed novel hierarchical and multi-tiered architecture, supported by ensemble modelling, is designed to reflect this diversity. It can present a candidate with multiple, diverse career paths that are nevertheless consistent with the individual's underlying ability profile, offering a balanced view of suitability and variety.

CHAPTER 8

THE HYBRID EXPLAINABILITY FRAMEWORK FOR THE PROPOSED RECOMMENDER SYSTEM USING SHAP WITH NATURAL LANGUAGE JUSTIFICATIONS

Since the proposed recommendation engine is based on an ensemble approach comprising multiple machine-learning models, including Deep Neural Networks (DNNs), it inherently functions as a black-box system. It is difficult to directly interpret the rationale behind recommendations made by such black boxes. As highlighted in the introductory section, ensuring transparency in a career recommendation system is crucial to fostering user trust, enabling informed decision-making, and enhancing the reliability of the recommendations.

This section provides a detailed discussion of the innovative methodology employed to integrate explainability into the proposed career recommender system. This approach aims to enhance transparency and accountability (Akhtar, Kumar, & Nayyar, 2024), enabling users to make more informed career decisions. We outline the step-by-step integration process of a hybrid explainer, which combines a SHAP explainer (Lundberg & Lee, 2017) with natural language explanations. Thus, the explanations generated by the framework in natural language for the recommendations made by the recommendation engine can be easily understood by common users, such as students, addressing the challenge of requiring technical expertise to interpret visual explanations, typically generated by SHAP. Diagrams, plots, procedures, and tables are used to clarify each stage, with natural language explanations for the recommendations. The overall methodology architecture is illustrated in Figure 16.

8.1. Using SHAP to Interpret Model Recommendations

SHAP (Shapley Additive explanations) is based on Shapley values derived from cooperative game theory (Lundberg & Lee, 2017). It is a powerful method for providing interpretability for machine learning predictions based on feature importance analysis. Shapley values distribute the total gain (model prediction) across

all the players (features) involved in the model prediction process. This allows us to interpret the prediction based on the marginal feature importance for a given input. The SHAP algorithm offers better local accuracy, addresses missing data through imputation methods and ensures fairness. These are achieved by evenly distributing the feature contributions against the total gain (Lundberg & Lee, 2017; Li, Sun, Huang, & Chen, 2024).




Figure 16. The Hybrid Explainability Architecture of the Proposed Recommender System.

While the machine learning model provides predictions based on the input and the model parameter values without offering insights into how the decision is made, a SHAP explainer can explain the model's predictions. This helps the end user make more informed decisions, as illustrated in Figure 17. The figure shows how different parameters in our problem domain, like Numerical Ability (NA), Abstract Reasoning (AR), Verbal Reasoning (VR), Mechanical Reasoning (MR), Space Relations (SR), and Verbal Ability (VA), contribute to the recommendations generated by the black-box recommendation engine using the SHAP explainer.

The main advantages of SHAP are its model-agnostic nature and its ability to provide both local and global explanations, regardless of the underlying ML algorithm of the

recommendation model (Chromik, 2021). While local explanations help us to understand individual predictions, global explanations provide insights into feature importance across the entire dataset. These explanations not only help the model developer debug their model (Xu et al., 2019; Dwivedi et al., 2023) but also enable end users to understand why the model provided certain predictions.




Figure 17. Functional Comparison between ML Model and SHAP Explainer

SHAP provides various plots, including the summary plot, dependence plot, and interaction plot, for enhanced global interpretability. Plots like force plot, decision plot, and waterfall plot offer local interpretations for the model's predictions.

As mentioned earlier, SHAP Explainer employs the concept of Shapley values derived from cooperative game theory to explain predictions made by machine-learning models. Shapley value is a unified measure of feature importance, calculated as the average of a feature's marginal contributions across all possible subsets of features to the model's prediction. The Shapley value ϕ is calculated using Equation 11.

$$\phi_i = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|! (|N| - |S| - 1)!}{|N|!} * (f_{S \cup \{i\}}(x_{S \cup \{i\}}) - f_S(x_S)) \quad (11)$$

Where:

ϕ_i is the Shapley value for i^{th} feature,

N is the set of all features (i.e, NA , AR , VR , MR , SR , and VA as of our dataset)

S is any subset of N that does not include feature i ,

$f_{S \cup \{i\}}$ is trained with that feature i present and f_S is trained with the feature withheld.

Figure 18. visualizes how the Shapley values of each feature - in our case, Numerical Ability, Abstract Reasoning, Verbal Reasoning, Mechanical Reasoning, Space Relations, and Verbal Ability- of the input x contribute to the changes in the expected model prediction $E[f(z)]$ to reach the predicted output $f(x)$, of the model. The expected model prediction, $E[f(z)]$, also known as the base value (Lundberg & Lee, 2017), represents the average model prediction when no feature information is provided to the model.




Figure 18. A Shapley Value-Based Illustration: How Each Feature Value Shifts the Model's Base Prediction to the Final Prediction.

8.1.1. Global and Local Interpretation Using the SHAP Explainer

The SHAP explainer provides both global and local explanations (Salih et al., 2024). It offers a global view of the model's behaviour, providing insights into feature importance across the dataset, and this can also be utilized at the class level (item

level). The primary goal of the global interpretation is to find out the most influential features in the dataset guiding its recommendations. This helps the model builder to identify overall model behaviour. It helps to even fine-tune the model and verify if the model follows the domain knowledge and expected behaviour. The local explainer module offers insights into individual predictions made by the system for specific input instances. This helps the users to understand the rationale behind the recommendation made by the system (Vimbi, Shaffi, & Mahmud, 2024).

This study employed SHAP summary plots, including bar plots and beeswarm plots, to analyse feature importance and achieve global interpretability of the model. For local interpretability, visualisations such as force plots, decision plots, and waterfall plots were used to illustrate how different feature values influence individual predictions - in our case, the career recommendations generated by the recommendation engine for a given user's profile. The results and discussion section includes illustrative examples.

The SHAP bar plot uses the mean absolute SHAP values for each feature to illustrate the feature's importance (global feature importance) across all the data points in the dataset. Mathematically, the mean absolute SHAP value for a feature can be computed using equation 12.

$$\text{Global Importance of Feature, } i = \frac{1}{N} \sum_{j=1}^N |\phi_i^{(j)}| \quad (12)$$

Where:

$|\phi_i^{(j)}|$ is the SHAP value of the feature i , for the j^{th} data point,

N is the total number of samples in the dataset.

In the SHAP bar plot, each feature is represented by a bar. The longer the bar, the higher the effect of the feature on the model's prediction. However, it does not tell whether the feature increases or decreases the likelihood of the prediction.

There comes the importance of the SHAP Beeswarm plot. It provides direction of the impact of the feature (positive or negative). In the beeswarm plot, each dot

corresponds to the SHAP value for a particular feature. It visualizes the distribution of SHAP values for each feature for a data point from its base value. Mathematically, this distribution can be calculated using Equation 13.

$$\text{Variance of SHAP values for feature } i = \frac{1}{N} \sum_{j=1}^N \left(\phi_i^{(j)} - \mu_i \right)^2 \quad (13)$$

Where:

μ_i represents the mean SHAP value for the feature i ,

N is the number of samples,

$\phi_i^{(j)}$ is the SHAP value of feature i , for the j^{th} data point.

For providing local interpretability, we utilised SHAP visualisation plots such as force plots, decision plots, and waterfall plots. All of these help in understanding how the model makes individual predictions. While the force plot illustrates how each feature, in terms of its Shapley values, contributes to the model's prediction, the decision plot presents the cumulative SHAP values in an ordered sequence, and the waterfall plot breaks down the incremental contribution of each feature using stacked SHAP values. All these plots depend on the Shapley value ϕ_i , and follow the mathematical formula given in Equation 14.

$$f(x) = \phi_0 + \sum_{i=1}^M \phi_i \quad (14)$$

Where:

$f(x)$ is the model's prediction for a given instance.

ϕ_0 is the base value (expected model output when no features are considered).

ϕ_i represents the SHAP value of feature i , showing its contribution to the prediction.

M is the total number of features.

8.1.2. Accuracy Evaluation of the Proposed SHAP Interpretation Framework

We performed a local accuracy analysis on the individual predictions generated by the recommendation engine to validate the interpretive accuracy of the SHAP explainer (Lundberg & Lee, 2017). Equation 15 was used to calculate local accuracy, ensuring that the Shapley values generated for the user profile align with the model's prediction for that specific profile.

$$\text{Local Accuracy} = f(x) = g(x) = \phi_0 + \sum_{i=0}^M \phi_i \quad (15)$$

where,

$f(x)$: The predicted value of the recommendation engine,

$g(x)$: The approximation of the model prediction by the SHAP Explainer,

ϕ_0 : The expected output or the base value of the model across the dataset,

ϕ_i : The Shapley value of feature i , to the model's prediction,

M : Total number of features present in the model.

The local accuracy evaluation is done to verify if the SHAP explainer faithfully explains model behaviour, such as its decision-making process. If the sum of SHAP values, computed using Equation 15, significantly deviates from the model's output, it suggests that the explainer has potential issues and cannot be used to explain the model's prediction. If the values are consistent, the explainer is considered valid and can be reliably employed to interpret the model's behaviour.

In summary, this study employed the SHAP algorithm to generate both global and local explanations for the model's predictions. Various visualization techniques were used to illustrate the model's overall behaviour and to provide insights into individual predictions. The performance of the SHAP explainer was evaluated using local accuracy metrics to assess the reliability of its attributions. The results and discussion section will offer a clearer understanding of the methodology and its effectiveness.

8.2. The Architecture of the Proposed SHAP-Based Explainability Framework with Natural Language Justifications

The visual explanations generated by the SHAP explanation framework using different plotting mechanisms may not be suitable for common users. This is due to the need for higher technical expertise to interpret these patterns. The SHAP explainer is more suitable for model builders and advanced users. To bridge this gap and to provide meaningful and easily understandable explanations to users like students, a textual explanation framework is proposed for the career recommendation engine. A hybrid approach is used to integrate the textual explainer, which takes inputs (both local and item-level global Shapley value-based explanations) from the SHAP-based framework and translates them into natural language explanations.

The proposed system employs an innovative approach that utilises SHAP's local and global interpretability features expressed via their numerical Shapley values to generate a comprehensive and human-readable textual interpretation for the recommendations. Local explanations are produced using SHAP's local Shapley values to quantify each feature's contribution for individual user instances. This can offer precise, instance-specific justifications for the generated recommendations. On the other hand, global explanations are generated by aggregating Shapley values across all instances at the class level (item-level). These aggregated item-level Shapley values are exploited to produce feature-driven summaries characterising the recommended class (item). A central innovation of this approach is the integration of local and class-level global insights to produce unified textual explanations that merge personalised reasoning with broader class-level behavioural patterns. By combining user-specific justifications (why a particular class is recommended to an individual) with class-level evidence (why the class is generally recommended based on its item-level global Shapley value-driven profile), the architecture delivers a comprehensive explanation that strengthens interpretability and user trust. The innovative architecture of the proposed system is illustrated in Figure 19.

As shown in the system architecture (Figure 19), the SHAP explainer operates atop the recommendation engine and its corresponding training data. Initialization of the

SHAP module is performed by selecting a model-agnostic explainer, such as Kernel Explainer. The choice of explainer is adaptable, depending on the underlying classification model employed for the recommendation engine. For example, Deep Explainer is preferable for neural network frameworks like TensorFlow or Keras. And, Kernel Explainer can be employed for any model when internal access is restricted. The explainer uses SHAP's cooperative game theory principles to distribute each feature's contribution fairly among the model's predictions. These contributions, called Shapely values, are computed through appropriate approximation techniques, such as weighted linear regression, in the case of the Kernel Explainer (SHAP, n.d.).




Figure 19. Hybrid Framework for Explainable Recommendations. As Presented in the Author's Published Work (Rahim & Basheer, 2025b) Under a CC BY License.

8.2.1. Template-Guided Natural Language Generation of Local Shapley Explanations

This module forms the first layer of the proposed system’s textual explainability framework. Its purpose is to provide local interpretability by clarifying the reasoning behind the model’s recommendation for a particular user instance. In this layer, the explanation is generated through an instance-specific textual module that exploits the SHAP explainer’s local interpretability capability.

For multi-class models, SHAP provides instance-specific interpretability by calculating a vector of Shapley values. Each value in this vector represents the marginal contribution of a feature to a specific prediction. For interpretability, the explanation is produced for the class, $c = \operatorname{argmax}_c(f_c(x))$. SHAP returns a vector of numerical values known as Shapley values, represented as $\phi^{(c)} = [\phi_1^{(c)}, \phi_2^{(c)}, \dots, \phi_M^{(c)}]$, for M features. SHAP typically visualizes numerical Shapley values for local interpretation using plots such as force, decision, and waterfall plots. However, the technical complexity of these visuals poses a significant barrier to interpretation for non-expert users.

In the proposed system, Shapley values are applied to generate explanatory text that clarifies the reasoning behind a specific recommendation for a given user profile. The explanation module uses a template-driven natural language generation (NLG) approach (Pino, Parra, Besa, & Lagos, 2021), which interprets the instance-specific SHAP values associated with the predicted class. It determines the proportional contribution of each feature by normalising the absolute SHAP values. This enables the identification and ranking of the most influential features in terms of Shapley values. Using the sign and magnitude of each feature’s standardised value along with its SHAP impact, the module populates predefined templates to create clear, human-readable explanations describing how individual features increased or decreased the likelihood of the recommended class. The full procedure is outlined in Algorithm 1.

Algorithm 1: *Generate Textual Explanation for Local SHAP Values.* (as presented in the author's published work (Rahim & Basheer, 2025b), under a CC BY license)

Input:

feature_names: List of M feature names
instance_values: Feature values of the instance (length M)
shap_values: SHAP values for a single instance (array of size M)
predicted_class_idx: Index of the predicted class

Output:

recommendation_explanation: A natural language explanation detailing proportional feature contributions

Steps:

1. Extract *feature_impact* $\leftarrow$ *shap_values* corresponding to the predicted class.
 2. Compute *abs_shap_values* $\leftarrow$ absolute values of *feature_impact* feature-wise.
 3. Compute *total_impact* $\leftarrow$ sum of all *abs_shap_values*.
 4. If *total_impact* $\neq 0$:
 - a. Calculate *percentages* $\leftarrow$ (*abs_shap_values* / *total_impact*) $\times 100$
 - Else:
 - a. Set *percentages* $\leftarrow$ vector of zeros.
 5. Sort features by descending order of percentages, storing their indices in *sorted_indices*.
 6. Initialise an empty list *explanation_list*.
 7. For each index i in *sorted_indices*:
 - a. Retrieve *feature* $\leftarrow$ *feature_names*[i]
 - b. Retrieve *value* $\leftarrow$ *instance_values*[i]
 - c. Retrieve *impact* $\leftarrow$ *feature_impact*[i]
 - d. Retrieve *percentage* $\leftarrow$ *percentages*[i]
 - e. If *impact* > 0 and *value* > 0 :

Append:

“The high feature (value: *value*, impact: *impact*) contributed *percentage*%, increasing the likelihood of recommending this career cluster.”
 - f. Else if *impact* > 0 and *value* < 0 :

Append:

“The low feature (value: *value*, impact: *impact*) contributed *percentage*%, increasing the likelihood of recommending this career cluster.”
 - g. Else:

Append:

“The feature (value: *value*) with low (impact: *impact*) contributed *percentage*%, decreasing the likelihood of recommending this career cluster.”
 8. Optionally sort *explanation_list* alphabetically for consistency.
 9. Concatenate all entries in *explanation_list* to form *recommendation_explanation*.
 10. Return *recommendation_explanation*.
-

8.2.2. *Generating Feature-Driven Class Descriptions with Global SHAP Values*

The framework's second module provides a global, data-driven explanation for why a class (item) is generally recommended. It leverages aggregated, class-level Shapley values to identify the features most characteristic of that class. To achieve this, Algorithm 2 identifies key features that have a directly proportional influence on the class recommendation. The algorithm adapts the Pearson correlation coefficient, a proven method for feature selection (Gong, Li, Zhang, Zhang, & Wang, 2024), to measure the relationship between a feature's values and its corresponding Shapley values across instances. Only features with a positive correlation are selected, ensuring they consistently contribute to the class prediction. These features are then ranked by correlation strength to build the final class-level (item-level) description. This item-level global explanation is then combined with the instance-specific text from the first layer (Algorithm 1) to construct a comprehensive explanation for the recommendation made by the system.

After identifying the key features of the class using global Shapley values (Algorithm 2), this information is transformed into a natural language summary (description) of the class(item), through a rule-based NLG method (Algorithm 3). Feature correlation values are mapped to qualitative significance levels using predefined thresholds (e.g., ≥ 0.5 for “very strong significance,” ≥ 0.3 for “high importance,” etc.). These thresholds are adapted from Cohen’s rule-of-thumb for interpreting Pearson correlation coefficients (Cohen, 2013). The resulting feature-based narrative explains why the class is generally recommended by highlighting its most influential attributes.

Algorithm 2: *Identifying Class-Level Important Features Using Global Shapley Values.* (as presented in the author’s published work (Rahim & Basheer, 2025b), under a CC BY license)

Input:

Class_level_shap_values: SHAP values for all samples belonging to a given class

Feature_values: Feature values for the samples of the class

Feature_names: List of feature names

Output:

Top_features: Top features with the positive correlation between feature values and SHAP values

Steps:

1. Initialise an empty list: *Correlations*.
 2. For each feature index $i=1$ to M :
 - a. Extract the SHAP values and feature values for the feature.
 - b. Compute the Pearson correlation coefficient between these two vectors.
 - c. If the correlation is positive, append it to *Correlations*.
 3. Create a map with each feature name and corresponding correlation value.
 4. Sort the map in descending order of correlation.
 5. Return the map containing the features and their significance.
-

Algorithm 3: Generating Class-Level Textual Description Based on Class-Level Global Shapley Values (as presented in the author's published work (Rahim & Basheer, 2025b), under a CC BY license)

Input:

Feature_importance_map: Important features and their degree of importance

Output:

Class_level_description: Textual summary describing significant features of the class

Steps:

1. Initialize *Calss_level_description*
 2. For each record from *Feature_importance_map*
 - a. $Feature = record[feature]$
 - b. $feature_impact = record[feature_importance]$
 - c. If $feature_impact \geq 0.5$

$desc = "Feature \{feature\} \text{ has very strong significance as far as the class is concerned.}"$
 - Else if $feature_impact \geq 0.3$

$desc = "Feature \{feature\} \text{ has high importance as far as the class is concerned.}"$
 - Else if $feature_impact \geq 0.0$

$desc = "Feature \{feature\} \text{ has moderately high significance as far as the class is concerned.}"$
 - $Calss_level_description.append(desc)$
 3. Return *Class_level_description*.
-

By automatically generating data-driven class descriptions, this methodology eliminates the need for domain-specific expert input, ensuring its adaptability across various recommendation systems. This automated, feature-importance-based class profile is then fused with the local textual explanation from Algorithm 1. Together, they form the robust foundation for the final, comprehensive explanation of any given recommendation.

8.2.3. *Generation of Natural Language Explanations via LLM by Synthesising Local and Item-Level Global Insights*

The framework's final layer serves as a synthesis engine, utilising a Large Language Model, LLaMA 2, implemented as a sentence transformer, to generate a comprehensive explanation (Said, 2025). The LLM processes two key inputs: the instance-specific text (from Algorithm 1) and the feature-driven class descriptions (from Algorithms 2 and 3). By integrating these local and global perspectives, the LLM produces a final, meaningful explanation for the system's recommendation.

The fusion layer uses a carefully structured, systematic prompt to steer the model's reasoning, drawing on insights from both instance-specific explanations and class-level feature descriptions generated through local and global SHAP analyses. By grounding its process in these SHAP-based insights, the approach allows a controlled language model to produce clear and meaningful explanations.

The Proposed Systematic Prompt for the Explanation Generation via LLM,
as Presented in the Author's Published work (Rahim & Basheer, 2025b) Under a CC BY License:

Context:

Two types of feature-based textual explanations are provided:

Instance-Specific Justification: A textual explanation for the recommendation derived from the feature contribution scores corresponding to an individual user profile.

Class-Level Feature Summary:

A description of characteristic feature patterns associated with the recommended class.

Inputs:

Instance_Specific_Justification: {instance_specific_recommendation_explanation}

Class_Level_Feature_Summary: {class_level_description}

Objective:

Generate a final explanation with an analytical evaluation of how the instance-specific justification aligns with or contradicts the class-level feature summary.

Constraint:

Constrain the explanation strictly to the provided inputs. Do not add external information or assumptions.

8.2.4. Implementation of the Explanation Framework in the Proposed Career Recommender System

The SHAP explainer, employing a Deep Explainer for its model-agnostic capability, was built on the proposed career recommender engine and its underlying dataset (O*NET-based DAT dataset). To obtain a local explanation, the candidate profile generated from KDAT is fed to the recommendation engine. The user profile and the predicted career cluster are then provided to the SHAP explainer. The SHAP explainer computes the Shapley values representing the marginal contribution of each feature (NR, AR, VR, MR, SA, and VA). These Shapley values are then translated into a human-readable explanation using Algorithm 1.

After producing the local explanation in natural language, the class-level (item-level) global description module - the second layer of the explainability framework - examines the recommended career class to determine which features most strongly shape the model's decision. It does this by checking the global SHAP values for that class and identifying features - Numerical Ability (NA), Abstract Reasoning (AR), etc. - whose values consistently show a positive correlation with the class prediction. A stronger positive correlation means that higher values for that feature reliably increase the likelihood of recommending the class. These influential features are then ranked by their correlation strength and translated into an easy-to-read summary using a rule-based natural language generation method. Algorithms 3 and 4 carry out this process, turning the feature patterns into a clear explanation of why the career class tends to be recommended.

To generate the final natural language explanation, the fusion layer integrating local and global explanations was implemented using LLaMA 2, a Large Language Model (LLM). The LLM module's task is to generate the final textual explanation by comparing and contrasting the explanations generated from the local explanation module and the global explanation module to provide a rational justification for the recommendations made by the recommendation engine. The prompt was customised for the career recommender system, with the constraints kept as they are, but with domain-specific elements added to it. This implementation helped to generate a

rational explanation within the constraints, eliminating hallucinations in the explanation generation process while preserving domain information.

This study used Python 3.10 for model building, along with TensorFlow 2.x and Keras libraries. We utilised the SHAP explainer library, specifically the Deep Explainer, to build the explanation framework. To construct the fusion layer, we employed the LLaMA 2 large language model through the Hugging Face interface as a sentence transformer to fine-tune the explanation in natural language. Data handling was done using the Pandas and NumPy libraries, and plots were created using Matplotlib. The work was carried out on Ubuntu OS, integrating a GPU through Google Colab.

8.3. Expert Evaluation of the Proposed System

As part of the validation of the work, an expert evaluation was conducted both for the recommendations generated by the system and the explanations generated for the recommendations by the explanation framework. This is done in addition to the traditional performance metrics like accuracy, precision, recall, f1-score, etc., for the recommendation engine, and local accuracy metrics related to the explanation framework. The explanation evaluation process used 16 distinct user profiles spanning across 33% of the career clusters in the dataset. Ten career experts, who are well-versed in career counselling and KDAT from the Department of General Education, Kerala, were selected for the expert evaluation process. Each of the career counsellor were provided 16 user profiles and the career clusters recommended by the recommendation engine, along with the natural language explanations for these recommendations generated by the explanation framework. They were instructed to evaluate the recommendations and the related explanations using a five-point Likert scale (Joshi, Kale, Chandel, & Pal, 2015):

- 1 – Not satisfactory
- 2 – Slightly satisfactory
- 3 – Moderately satisfactory
- 4 – Satisfactory
- 5 – Highly satisfactory

A Google Form was provided to the career experts for this evaluation, containing anonymous candidate profiles along with their corresponding recommendations generated by the recommender system and related explanations generated by the explanation framework.

8.4. Ethical Guidelines Followed in the Evaluation Process

All the experts considered in the evaluation process were approved career counsellors and KDAT experts of the Department of General Education, Government of Kerala. All experts were provided with detailed information about the study in the introductory part of the Google form used to evaluate the study. No personal, health or sensitive data were part of this evaluation process. Pre-approved consent was obtained for all the career experts. No personal data of candidates was shared with the career experts. Only the DAT results of anonymised candidates were supplied for evaluation. The Ethics Committee of Sullamussalam Science College, Areekode, reviewed and approved this consent procedure.

8.5. The Overall Use-Case of the Proposed Hybrid Explainable Career Recommender

Once a candidate is evaluated using the Differential Aptitude Test, their aptitude profile is fed into the recommendation engine. The recommender system generates career clusters aligned with the candidate's aptitude profile. This recommendation, along with the user profile, is then fed into the SHAP-based explanation framework, which provides visual and textual interpretations for the recommendations. The use-case diagram of the system given in Figure 20 clearly illustrates how the system helps the candidate to make well-informed career decisions.




Figure 20. The Use-Case Diagram of the Proposed Hybrid Explainable Career Recommender.

CHAPTER 9

DYNAMIC CAREER EXPANSION FOR FUTURE-PROOF RECOMMENDATIONS

This module aims to enable the proposed recommendation engine to suggest careers that are not present in the existing career clusters derived from the dataset on which the engine is built. Without this dynamic career expansion module, the system may not be able to recommend emerging careers. This will lead to the system becoming obsolete in the long run.

9.1. Need for Dynamic Career Expansion Module

Career landscapes are constantly evolving. New roles are emerging due to advancements in technology, shifts in industry demands, and changes in skill requirements. Traditional career recommendation systems, which rely on static datasets and predefined career clusters, struggle to accommodate these rapid developments. Once a recommendation engine is built on a fixed dataset, it becomes limited to suggesting only the careers already present in its training data. As a result, such systems fail to address the needs of students and professionals seeking guidance in newly emerging career paths.

To overcome this limitation, a dynamic career expansion module is essential. This module enables the recommendation engine to integrate new careers that were previously unseen within the existing career clusters. Instead of treating career recommendations as a static process, it introduces a mechanism for continuous learning and adaptation. The approach incorporates a semi-supervised human-machine hybrid workflow. The system allows career experts to define new career profiles based on aptitude parameters such as Numerical Ability (NA), Abstract Reasoning (AR), Verbal Reasoning (VR), Mechanical Reasoning (MR), Space Relations (SR), and Verbal Ability (VA). These manually curated career profiles are then processed by the recommendation engine to determine which existing cluster it can be mapped to, thereby integrating the career profile into that cluster.

9.2. Architecture of the Dynamic Career Expansion Module

The architecture of the dynamic career expansion module follows a semi-supervised approach. It has basically three elements:

- a. Career Profile Interface
- b. Career Cluster Mapping
- c. Career Cluster Integration

The Career Profiler Interface allows a career expert to define new career profiles based on aptitude parameters. A career counsellor can profile a new career using aptitude parameter values such as Numerical Ability (NA), Abstract Reasoning (AR), Verbal Reasoning (VR), Mechanical Reasoning (MR), Space Relations (SR), and Verbal Ability (VA). This process should be conducted by an expert in career counselling with knowledge of Differential Aptitude Tests, as it is crucial for ensuring that the predictions are valid and meaningful.




Figure 21. The Architecture of the Dynamic Career Expansion Module

Career Cluster Mapping is performed by the proposed career recommendation engine. Once the career expert profiles a new career based on its required aptitude parameter

values, it is fed into the recommendation engine, disguised as a user profile. This approach is valid since both user profiles and career profiles rely on the same set of parameters. When the career profile is processed by the recommendation engine, it recommends the most relevant career cluster aligned with the given profile. This recommendation serves as the mapping of the new career to the existing career clusters.

The Career Cluster Integration module integrates newly defined careers into the mapped career cluster along with their aptitude parameters. This enables the career recommendation engine to suggest these new careers to candidates whose profiles match the corresponding career cluster. The overall architecture of the dynamic career expansion module is illustrated in Figure 21.

This approach ensures that the recommendation engine remains relevant over time, guiding users toward not only traditional careers but also emerging opportunities across various industries. By enabling real-time career expansion, the system becomes more robust and scalable.

CHAPTER 10

RESULTS AND DISCUSSIONS

10.1. Introduction

The following section offers a detailed discussion of the results produced by the recommendation engine and the explainability framework. Its primary considerations are the validity, relevance and performance of the recommendation engine; and the extent of its explainability. It initially reviews the proposed system's career recommendations, comparing them with established standards and domain expectations. These recommendations are investigated in terms of accuracy and consistency to guarantee their applicability to the real-world career guidance situations. This part also explores the soundness of the explanations generated by the explainability paradigm, which encompasses SHAP's local and global explanations, and the final fusion layer's natural language explanations. These explanations are evaluated based on domain knowledge to make sure they are meaningful and consistent with the principles of professional career counselling. Also, the accuracy measures are examined to identify the credibility and usefulness of the explanations. This analysis makes the process of recommendation more transparent and understandable. This section also presents the expert's evaluation of the recommendation engine and its explanations to showcase its acceptability; and a comparison of the work with current implementations. Finally, the performance of the dynamic career expansion module is also presented here.

10.2. Analysis of the Recommendation Engine's Outputs

The recommendation engine generates diverse career recommendations for a given user profile. These diverse recommendations are fine-tuned using the refinement suite present in the system. Thus, the user is provided a diverse but fine-tuned set of recommendations from its multi-tiered output ends. The recommendations are ranked by their significance.

Consider the user profile:

$$DAT_{user_profile}(NA, AR, VR, MR, SR, VA) = [50, 80, 77, 50, 50, 70]$$

The profile was supplied to the recommendation engine, which produced First-Tier, Second-Tier and Third-Tier suggestions as shown in Figures 22, 23 and 24, respectively. The system produces these recommendations using its First Pick, Second Pick, and Third Pick fine-tuned models according to the user profile.

As demonstrated by the outputs (Figures 22, 23, and 24), the system provides a vast array of career recommendations. A total of nine distinct career clusters aligned with the candidate's aptitude profile are presented to the candidate. The first career cluster, shown in Figure 22, represents the most suitable recommendation for the candidate based on their career profile. The other two clusters are the next best recommendations, ranked accordingly. As shown in Figures 23 and 24, the multi-tiered system further provides similar yet diverse recommendations from its second and third tiers.




Figure 22. Career Recommendations from FirstPick Model

This is a multi-level suggestion strategy that offers flexibility in career choice. It provides various options aligned with the aptitude profile of the candidate. The system is superior to current recommendation systems since it offers a wider and more organised array of career choices. Therefore, it can be a useful resource for career counsellors, as it allows them to suggest and identify a variety of career paths for a

candidate according to their ability profile. This is quite difficult in traditional counselling environments.




Figure 23. Career Recommendations from Second Pick Model




Figure 24. Career Recommendations from Third Pick Model

Moreover, the visualization of the candidate profile next to the recommended career clusters enables direct comparison of the aptitude parameters associated with the candidate and the profile associated with the career cluster, thereby increasing transparency. This visual clarity offers a good indication that the proposed cluster of careers certainly fits the given candidate profile, since an increase in Abstract Reasoning and Verbal Reasoning scores clearly justifies the recommended careers. This affirms the validity and effectiveness of the recommendation engine.

It is important to note that the recommendations generated by the proposed framework are primarily based on the candidate's cognitive and aptitude profile derived from DAT measures. Although aptitude is an important indicator of occupational suitability, career decision-making is inherently multidimensional and may also depend on personality traits, vocational interests, and other factors. Consequently, aptitude-only recommendations may occasionally produce mismatches with candidate interests or bias toward careers aligned strongly with measurable cognitive abilities. Therefore, the current framework is intended to function as a decision-support tool for career counsellors rather than a fully autonomous recommendation system. Counsellors are expected to further interpret the recommendations by incorporating personality characteristics and vocational interests to provide more holistic career guidance. Integration of personality and vocational interest frameworks such as Big Five and Holland/RIASEC models remains an important direction for future work.

10.3. Performance Evaluation of Recommendation Engine

The accuracy, precision, recall and F1-score were some of the performance metrics used to measure the performance of the recommendation engine. Accuracy measures the general correctness of the system's predictions. Precision examines the dependability of positive forecasts in terms of the cost of false positives. Recall, in turn, measures the system capacity to recognize all the relevant positive cases and can emphasize the cost of false negatives. F1-score is an attempt to strike a balance between precision and recall as the harmonic mean of the two. It guarantees a

thorough assessment of the model's performance under the situations where false positives and false negatives are important factors.

Accuracy gives a first gauge of the validity of a system. It is the ratio of correct predictions to the total number of predictions made. To evaluate the general reliability of our recommendation engine, we relied on the accuracy metric. The standard formula expressed in Equation 16 is employed to calculate the accuracy of the recommendation engine.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (16)$$

Where:

TP = True Positives

TN = True Negatives

FP = False Positives

FN = False Negatives

Table 13 showcases the overall accuracy metrics of the proposed career recommendation engine, considering its Diverse Recommender Module and Refinement Recommendation Modules.

Table 13. Accuracy Metrics for the Diverse Recommender Module and Refinement Recommendation Modules

Models	Accuracy
Diverse Recommender	0.99220
RefMod1	0.99468
RefMod2	0.99729
RefMod3	0.99615
RefMod4	0.99511
RefMod5	0.99703
RefMod6	0.99292
RefMod7	0.99488
RefMod8	0.99434

Precision is an important metric in a multi-class career recommender system. It is important that candidates are not guided towards career suggestions that do not align with their abilities. This is significant as the system provides several career cluster recommendations ranked by their level of relevance. A high precision indicates that the recommended clusters are highly relevant and that the system produces few false positives. This helps candidates trust the career suggestions generated by the recommendation engine. The precision of the recommendation engine is calculated using Equation 17.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (17)$$

Table 14 shows that the precision of all modules in the recommendation engine exceeds 99%. This is a strong indication that candidates can reliably depend on the career recommendations provided by the system.

In our study, the metric recall is also significant. It guarantees all relevant career clusters are recommended to the candidate. A low recall indicates that the system fails to include pertinent careers in its recommendation outputs. This will potentially limit the candidate's options. Recall is computed using the standard formula, given in Equation 18.

Table 14. Precision Metrics for the Diverse Recommender Module and Refinement Recommendation Modules in the Proposed Model.

Models	Precision
Diverse Recommender	0.99223
RefMod1	0.99472
RefMod2	0.99729
RefMod3	0.99618
RefMod4	0.99512
RefMod5	0.99704
RefMod6	0.99297
RefMod7	0.99491
RefMod8	0.99436

$$\text{Recall} = \frac{TP}{TP + FN} \quad (18)$$

Table 15 demonstrates the system's effectiveness, as indicated by its exceptionally high recall measure, in suggesting all relevant career clusters in its recommendations.

Table 15. Recall Metrics for the Diverse Recommender Module and Refinement Recommendation Modules

Models	Recall
Diverse Recommender	0.99221
RefMod1	0.99468
RefMod2	0.99729
RefMod3	0.99615
RefMod4	0.99511
RefMod5	0.99703
RefMod6	0.99292
RefMod7	0.99488
RefMod8	0.99434

The recommendation engine should provide quality career recommendations, which can be measured in terms of precision. It is also important that the system does not fail to consider any relevant careers thus providing sufficient recall. The F1-score is a balanced score that represents both precision and recall at the same time. Therefore, the F1-score evaluation is a robust technique of gauging general performance of the recommendation engine. F1-score of our system is estimated as harmonic mean of precision and recall, as defined in Equation 19.

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (19)$$

Table 16 reveals that the recommendation engine has maintained a strong balance between precision and recall, thus showing that it is able to suggest all the relevant career paths at the same time maintaining high precision during the recommendation process.

In short, the findings show the system demonstrated an excellent performance whereby testing accuracy was over 99% in the top-level Diverse Recommender Module and the sub-level Refined Recommendation Module. Moreover, the key performance metrics such as precision, recall and F1 score were all above 99, indicating that there was a high degree of harmony between the measurements. This high consistency supports the strength and stability of the predictive functions of the model in a wide performance dimension.

Table 16. F1-Score Metrics for the Diverse Recommender Module and Refinement Recommendation Modules.

Models	F1 Score
Diverse Recommender	0.99220
RefMod1	0.99468
RefMod2	0.99729
RefMod3	0.99615
RefMod4	0.99510
RefMod5	0.99703
RefMod6	0.99292
RefMod7	0.99487
RefMod8	0.99434

The proposed system showed very strong predictive performance - precision, accuracy, and recall were all close to 0.99. We used an 80:20 train - test split; without augmentation the model's accuracy was 0.76. To improve generalisation and reduce overfitting, Gaussian noise was added to the training data and duplicates from augmentation were removed to avoid memorisation. Testing on the original non-augmented samples produced consistent results, suggesting the model preserved underlying data patterns. 5-fold cross-validation was performed on individual classifiers. Experts reviewed the recommendations; however, additional evaluation is recommended to ensure results are not overly optimistic and to check for any unintended data leakage.

10.4. SHAP Explainer Analysis: Interpreting the Recommendation Engine's Behaviour and Its Recommendations

This section examines different features of the SHAP explainer module, including the validity of the explainer in terms of local accuracy measures, global explanations of the overall behaviour of the recommendation engine, and local explanations of

individual recommendations made by the SHAP explainer. We use different SHAP visualisation techniques to offer insight into the decision-making process of the model. This includes summary plots and beeswarm plots to examine overall model behaviour, as well as force plots, decision plots, and waterfall plots to illustrate individual predictions made by the system, providing transparency in its decision-making process.

10.4.1. Evaluating Validity of the SHAP Explainer

To use the SHAP explainer that has been built on the recommendation engine and the underlying dataset, it is essential to validate the explainer to verify if it faithfully explains the recommendations generated by the recommendation engine. It should ensure that the sum of SHAP values correctly reconstructs the model's prediction. If the local accuracy does not hold, the SHAP explainer will be inconsistent with the model's behaviour. This section does a local accuracy analysis of the generated SHAP explainer.

We considered the user profile:

$$\text{USER_PROFILE (NA, AR, VR, MR, SR, VA)} = [65, 60, 66, 25, 24, 59]$$

Once this particular record was fed to our recommendation engine, the system predicted a career cluster labelled 200.

To evaluate local accuracy, we used the standard equation 20:

$$\text{Local Accuracy} \quad = f(x) = g(x) = \phi_0 + \sum_{i=0}^M \phi_i \quad (20)$$

Where,

$f(x)$: The predicted value of the recommendation engine,

$g(x)$: The approximation of the model prediction by the SHAP Explainer,

ϕ_0 : The expected output or the base value of the model across the dataset,

ϕ_i : The Shapley value of feature i , to the model's prediction,

M : Total number of features present in the model.

Here we have:

$f(x) = 200$, predicted label.

SHAP explainer's base value for the model,

$\phi_0 = 557.1110703812317$

Shapley values of the features (NA , AR , VR , MR , SR , VA) generated for the predicted class

$= [-52.00565738, -45.4643695, -9.70420332, -68.92581867, -165.69422043, -15.31680108]$

Applying Equation 20,

$$\begin{aligned} g(x) &= 557.1110703812317 + (-52.00565738 + -45.4643695 + - \\ &\quad 9.70420332 + -68.92581867 + -165.69422043 + -15.31680108) \\ &= 200 \\ &= f(x) \end{aligned}$$

$f(x)=g(x)=200$ ensures that the generated SHAP Explainer provides an accurate and faithful approximation of the model and that the explanation is locally faithful. Furthermore, it suggests the model is transparent, as its decision-making process is interpretable for this specific input using our SHAP explainer object. Thus, the explainer can be used to explain the recommendations generated by the recommendation engine. Based on this evaluation and the established validity of the explainer, the following section will present the global and local interpretations of the recommendation engine using the explainer.

10.4.2. Global Interpretation: Analyzing Model Behaviour Using SHAP Explainer

Using SHAP summary plots, including bar plots and beeswarm plots, the most influential features among Numerical Ability (NA), Abstract Reasoning (AR), Verbal

Reasoning (VR), Mechanical Reasoning (MR), Space Relations (SR), and Verbal Ability (VA) across the entire dataset were identified. As illustrated in Figure 25, MR, SR, and NA emerge as the most important decision-making features in career prediction across the entire dataset. This indicates that features, Mechanical Reasoning (MR), Space Relations (SR) and Numerical Ability (NA) are the most discriminative features that help the recommendation engine to differentiate between career clusters when recommending careers in general. This also suggests that variations in features, MR, SR, and NA will produce larger shifts in the model's prediction.




Figure 25. Summary Bar Plot: Feature-Value Based Feature Importance Analysis.

The beeswarm plot, Figure 26, reveals that higher values of most of the features—especially that of MR and SR—have a positive impact on class discrimination, specifically in career cluster selection from the dataset. While features VA, AR, and VR also have a positive impact on class selection, NA exhibits a unique behaviour. In some cases, its higher value guides the recommendation engine towards a specific cluster of careers, while in some other cases, its lower value favours a different set of careers, indicating its varied influences across different career cluster recommendations.




Figure 26. Beeswarm Plot: Feature-Value Based Feature Importance Analysis

To further confirm feature importance across the entire dataset, we employed Permutation Importance analysis (Altmann et al., 2010), while acknowledging that the SHAP summary plot and beeswarm plot can exhibit better insights than Permutation Importance feature analysis.

Permutation Feature Importance evaluates how shuffling the values of a feature affects the model's performance. The importance score of a feature j is calculated using Equation 21.

$$I_j = M(X) - M(X_{perm}^j) \quad (21)$$

Where:

X is the dataset,

X_{perm}^j is the the dataset where feature j is randomly selected,

$M(X)$ is the model performance metric (accuracy).

This analysis, as illustrated in Table 17, confirms that Mechanical Reasoning (MR), Space Relations (SR) and Numerical Ability (NA) are the most influential characteristics in career cluster selection. The above 3 features (MR, SR, and NA) had scores more than 0.3 meaning that they are very essential in the prediction process of the model. The values lie between 0.1420 (VR) and 0.3838 (MR), which suggests that

every feature has a moderate level of significance in the model's output. The analysis validates the observations provided by the SHAP summary plot and beeswarm plot.

Therefore, the Permutation Importance analysis supports the fact that MR, SR, and NA play a central role in career cluster determination as well. To sum up, both the Summary plot analysis and the Permutation Importance analysis point to the fact that Mechanical Reasoning (MR), Space Relations (SR), and Numerical Ability (NA) are the most powerful in the model recommendation. Other features like the Abstract Reasoning (AR) and Verbal Reasoning (VR) also affect career selection, but in a smaller degree and more evenly distributed through all the career clusters.

Table 17. Feature Importance Comparison Using Permutation Importance Analysis.

Feature Name	Permutation Importance Score
Mechanical Reasoning (MR)	0.38
Numerical Ability (NA)	0.36
Space Relations (SR)	0.33
Verbal Ability (VA)	0.26
Abstract Reasoning (AR)	0.21
Verbal Reasoning (VR)	0.14

10.4.3. Local Interpretations: Understanding Individual Career Recommendations

This section attempts to justify why the model makes particular predictions of the class labels for individual instances. In order to do so, we instantiated a SHAP explainer object through the predict method of the classifier and used the SHAP plots including the force plot, decision plot, and waterfall plot. These plots explain how the feature-specific SHAP values of the input instance cause it to shift the model's prediction towards the desired predicted label from its base value, giving a better insight into how the model makes a decision (Lundberg and Lee, 2017).

We considered a sample candidate test result:

User profile = {

Numerical Ability: 65,

Abstract Reasoning: 60,
Verbal Reasoning: 66,
Mechanical Reasoning: 25,
Space Relations: 24,
Verbal Ability: 59
}
i.e., USER_PROFILE (NA, AR, VR, MR, SR, VA) = [65, 60, 66, 25, 24, 59]

Once this particular record was fed to our recommendation engine, the system predicted the career cluster labelled *200* as depicted in Figure 27.

Next, we applied SHAP Explainer's visualisation tools to explain the recommendations, shown in Figure 27, made by our recommendation engine for this particular user profile. Figure 28 demonstrates how the force plot explains the model's prediction for the given user profile. It specifically shows how higher values of NA and AR, along with lower values of MR and SR, significantly contribute to shifting the model's prediction from its base value (557.1) to the predicted value (200). The explainer clearly highlights that features such as Numerical Ability (NA) and Abstract Reasoning (AR) are key aptitudes for the recommended career cluster, as depicted in Figure 28. Almost all careers in this cluster, according to established domain knowledge, require mathematical and problem-solving skills, while Mechanical Reasoning (MR) and Space Relations (SR) are less relevant. If MR and SR had higher positive values, the recommendation engine might have suggested a different career cluster for the candidate. Thus, the explainer provides a mature explanation of the recommendation generated by the recommendation engine.




Figure 27. Career Recommendations Made by the Model for the Given Candidate Profile.




Figure 28. Force Plot Illustrating How the Marginal Contribution of Each Feature Value in the Model's Prediction

Figure 29 presents the same idea using a decision plot by extending the explanation. The plot illustrates how feature contributions accumulate along the prediction path. Furthermore, Figure 30 reiterates this information in a static audit view of the reasoning using a waterfall plot. All these plots (force plot and decision plot, and waterfall plot) illustrate the prediction process in terms of the marginal contributions of feature values through their Shapley values.




Figure 29. The Decision Plot Illustrates How Different Feature Values Contribute to the Model's Prediction.




Figure 30. Waterfall Plot Providing a Static Audit View of the Model's Reasoning

To verify the validity of this explanation, we did a feature importance analysis of the predicted career cluster labelled 200 using class-level(item-level) SHAP summary plots. The summary bar plot (Figure 31) indicates that, for the selected career cluster, labelled 200, the features such as Numerical Ability (NA) and Space Relations (SR) are the most influential. However, based on well-established domain knowledge in differential aptitude tests, we know that SR is not an important aptitude needed for this recommended career cluster labelled 200. This issue was further investigated using a beeswarm plot, shown in Figure 32, which clarifies that a higher value of NA positively contributes to the career cluster selection, whereas a lower feature value of SR generates a higher Shapley value contributing positively to this selection. This indicates that if SR were high, the recommendation engine would not have recommended this career cluster for the given user profile.




Figure 31. Summary Bar Plot of Career Cluster 200 Illustrating the Significant Features




Figure 32. Summary Beeswarm Plot Depicting the Feature Importance for Career Cluster Labelled 200.

Additionally, the beeswarm plot shows that as the value of MR decreases, its SHAP value increases for this career cluster, and this observation is binding for this career cluster as per DAT. Therefore, the system's prediction for the input instance and the explanation provided by the explainer are largely valid according to the feature importance analysis of the selected career cluster. This ensures the reliability of the SHAP explainer and, thus, it can be used for providing explanations for the recommendations generated by the recommendation engine.

10.4.4. Analysing Class-Specific SHAP Explanations

In this section, with the SHAP explainer, we aim to explain why a particular class is predicted and why another is not, thus providing insights to the user about why they are recommended for one career cluster over another.

In the previous section, we utilized an explainer built on the *predict* method of the underlying classifier. This approach helped explain why the final class label was predicted by showing how feature values influenced the overall prediction. However, it did not support class-specific analysis using the SHAP explainer, as it only showed how feature values shifted the model's prediction from the base value to the final predicted class label without breaking down the contributions for each possible class. Since the model is a multiclass classifier, it is more effective to build the SHAP explainer object using the *predict_proba* method of the classifier for class-specific analysis based on the expected values and SHAP values for each class. This approach helps us understand the model's probabilistic reasoning by showing how features shift the model's prediction from class-specific base values to the predicted value. This benefits the end user to make more informed decisions when selecting their career path. We used force plots, decision plots, and waterfall plots to illustrate this.

For class-specific analysis, we utilized the same standardised user profile which we used in the previous analysis:

USER_PROFILE (NA, AR, VR, MR, SR, VA) =
[1.65079783, 0.1751207, -0.06732215, -0.89514629, -1.77507484, -
0.13278312]

When input into the model, it predicted a probability of 1.00 for the career cluster labelled 200. We did a class-specific analysis for this predicted career cluster using our SHAP Explainer.

Employing SHAP visualisation techniques like force plot, decision plot and waterfall plot, we demonstrated the impact of each feature on the output probability. The decision plot, shown in Figure 33, visualises how these features contribute to shift the prediction from the base value (0.02) of the career cluster to the predicted probability value of 1.00.

Observing each feature's spread, we noticed that NA, AR, and VR have relatively higher(positive) values, while MR and SR have lower (negative) values. And the resultant Shapley value-based plots show that all the feature values marginally

contributed to the predicted probability. Notably, lower values of MR and SR increased the likelihood of selecting this career cluster.




Figure 33. Feature Contributions to Individual Prediction: Class-Specific Analysis Using Decision Plot.

We further explored whether a candidate with the same user profile, analyzed in the previous section, might be suitable for a quite different career, like an *Architect*. This investigation involved analyzing the candidate profile against the Shapely values for the target career cluster containing the career *Architect*, labelled 900. The SHAP plot, shown in Figure 34, clearly illustrates why the candidate's profile may not align well with the target career. The likelihood of selecting the career cluster labelled 900 for the given candidate profile is approximately 0.00.




Figure 34. Decision Plot Explaining Why the Given Career Cluster May Not Align with the Candidate's Profile.

The plot reveals that low scores in the Space Relations (SR), along with the poor scores of Abstract Reasoning (AR) and Mechanical Reasoning (MR), do not justify

the candidate to choose a career in Architecture. This observation is consistent with the domain knowledge. This fact indicates that people with low spatial reasoning and abstract reasoning abilities are probably not fit to work in Architecture.

The analysis of the results revealed that the SHAP explainer is a useful tool that can be used to understand the recommendations that our proposed system produces. The explainer module provided acceptable explanations for each prediction using local explanations. The explainer also gives insights into the overall model behaviour through global interpretations. These explanations are consistent with established domain knowledge. This suggests that the explainer model can be successfully applied to provide insightful information to the model builder and, to a degree, to an expert end user, like a career counsellor. This supports the usefulness of SHAP in bringing better transparency and interpretability to recommender systems.

Considering the proposed recommendation engine, the SHAP explainer provides valid and meaningful global and local interpretations for its recommendations. The global interpretation reveals that Mechanical Reasoning (MR), Space Relations (SR), and Numerical Ability (NA) are the most crucial features in its overall career recommendation process. Abstract Reasoning (AR) and Verbal Reasoning (VR) also contribute to career suggestions, but their impact is comparatively lower than MR, SR, and NA. This is because almost all careers inherently require AR and VR to a significant extent. But changes in MR, SR, and NA will make a significant impact on the model's recommendation process. This indicates that candidates with strong MR, SR, and NA have a diverse set of career paths. The local explanations via different SHAP plots also showed how different features influence the recommendations for a given user profile. These influences were consistent with the insights provided by domain experts. This is also supported by the local accuracy evaluation. These facts highlight the explainer's ability to accurately interpret recommendations. Thus, the explainer enhances transparency in the recommendation engine's decision-making process.

Finally, the study identifies different levels of explanation for different stakeholders. Model developers can use SHAP visualisations such as bar plots and beeswarm plots

to understand global feature importance and fine-tune the model. Force plots, decision plots, and waterfall plots can help both model developers and technically trained career experts analyse individual predictions, whereas SHAP-guided natural-language explanations can help career counsellors, students, and parents understand the rationale behind the system recommendations.

10.5. Analysis of the Final Natural Language Justifications Generated from the Hybrid Explainability Framework

While the SHAP explainer provides explanations for the recommendations through visualisation plots, a typical user, such as a student, may find these difficult to interpret. This section explains how the textual explainer enhances interpretability by providing clear, natural language explanations for the reasoning behind the recommendations.

Considering the sample candidate test result, i.e, the user profile: USER_PROFILE (NR, AR, VR, MR, SR, VA) = [65, 60, 66, 25, 24, 59].

The system generated the recommendations as shown in Figure 35 for this user profile.




Figure 35. Recommendations Generated by the System Matching for a User Profile.

The following sections discuss the results of local and global interpretations based on Shapley values, their resultant textual explanations via Algorithm 1, Algorithm 2 and Algorithm 3, and final comprehensive natural language justification generated by the hybrid explainability framework for the above recommendation.

10.5.1. Analysis of Textual Explanations for Individual Recommendations Using the Explainability Framework

For the USER_PROFILE (NA, AR, VR, MR, SR, VA) = [65, 60, 66, 25, 24, 59], and the corresponding career recommendations (given in Figure 35), this section will analyze the local explanation generated by the system in natural language.

After applying standardization, the transformed feature values are:

[1.69090851, 0.19389258, -0.05227239, -0.90342015, -1.79073424, -0.11755611]

The Shapley values of the features for the given user profile for the predicted class are:

[0.44310025, 0.07924637, 0.0609078, 0.09494073, 0.24430919, 0.05990119].

And, the expected SHAP value (base value) for the recommended class is: 0.01759431

The force plot (Figure 36) illustrates the rationale of the systems prediction for the user profile:




Figure 36. Force Plot Illustrating the Marginal Contributions of Features to the Model Prediction.

The instance-specific local explanation module employing Algorithm 1 generated the following explanation:

- *The high Numerical Ability skill (value: 1.6909, impact: 0.4431) contributed 45.1%, increasing the likelihood of recommending this career cluster.*
- *The high Abstract Reasoning skill (value: 0.1939, impact: 0.0792) contributed 8.1%, increasing the likelihood of recommending this career cluster.*

- *The low Mechanical Reasoning skill (value: -0.9034, impact: 0.0949) contributed 9.7%, increasing the likelihood of recommending this career cluster.*
- *The low Spatial Relation skill (value: -1.7907, impact: 0.2443) contributed 24.9%, increasing the likelihood of recommending this career cluster.*
- *The low Verbal Ability skill (value: -0.1175, impact: 0.0599) contributed 6.1%, increasing the likelihood of recommending this career cluster.*
- *The low Verbal Reasoning skill (value: -0.0523, impact: 0.06091) contributed 6.2%, increasing the likelihood of recommending this career cluster.*

The explanation aligns with the visual interpretation illustrated in the force plot (Figure 35). The textual explanation correctly informs the candidate that their Numerical Ability (NA), and Abstract Reasoning skill (SR) are the primary factors for recommending this particular career cluster by the system. Mechanical Reasoning (MR) and Space Relations (SR) are less influential in recommending this career. This explanation provides a clearer and more meaningful interpretation for ordinary users, such as students, who may not have the technical background to understand graphical outputs.

10.5.2. Analysis of Item-Level Textual Description Generated Via Global Class-Specific Shapley Values

For the career cluster recommended in Figure 35 by the recommendation engine, corresponding to the user profile USER_PROFILE (NA, AR, VR, MR, SR, VA) = [65, 60, 66, 25, 24, 59], a feature importance analysis was performed using the Shapley values of the cluster's features. The results are illustrated with a SHAP beeswarm plot in Figure 37.




Figure 37. Summary Beeswarm Plot Illustrating Feature Importance for the Career Cluster Labelled 200.

For this career cluster (Figure 35), using Algorithm 2, a feature importance analysis was carried out leveraging class-level global Shapley values. Table 18 presents the Pearson correlation between each feature and its corresponding Shapley value, using positive correlations only, to assess feature importance.

Table 18.Correlation Between Feature Values and Corresponding SHAP Values

Feature	Correlations (SHAP vs Feature values)
Numerical Ability	0.55
Verbal Ability	0.21
Abstract Reasoning	0.10

Algorithm 3 translated these features' importance details, presented in Table 18, into the natural language description of the career cluster. As shown below:

Feature Numerical Ability has very strong significance as far as this cluster is concerned.

Feature Verbal Ability has moderately high significance as far as this cluster is concerned.

Feature Abstract Reasoning has a moderately high influence as far as this cluster is concerned.

This natural language, feature-driven description of the career cluster is largely accurate from a domain knowledge perspective, as it identifies the key features.

However, the description of Abstract Reasoning (AR) does not accurately reflect its significance. This can be improved by augmenting the description with additional domain knowledge.

10.5.3. Analysis of Natural Language Explanation Generated Via Fusion Layer

For the user profile, USER_PROFILE (NA, AR, VR, MR, SR, VA) = [65, 60, 66, 25, 24, 59] and the generated recommended career cluster (Figure 39), the instance-specific natural language explanation (local explanation presented in section 10.5.1) and item-level textual description (class-level global explanation presented in section 10.5.2), were fed to the LLM-assisted fusion layer with its systematically created prompt. The following final natural language explanation was generated for the recommendation:

“The recommendation is well justified by the close alignment between the student’s aptitude profile and the characteristics of the target career cluster. The student demonstrates a particularly high score in Numerical Ability (NA), which serves as the key determinant for this cluster and therefore holds the greatest influence on the outcome. The Abstract Reasoning (AR) score is moderately above the mean, satisfactorily matching the cluster’s requirement. Although Verbal Ability (VA) falls slightly below average, it remains within an acceptable range for a feature of moderate relevance. The remaining attributes, such as Mechanical Reasoning (MR), Space Relations (SR), and Verbal Reasoning (VR), are relatively lower. But their reduced values exercise a favourable effect on the overall recommendation. This indicates that these factors are either of lesser importance or inversely related to the features defining the cluster. Hence, the recommendation is supported by the strong correspondence between the student’s key strengths and the critical attributes of the recommended career cluster.”

This explanation is largely consistent with established domain knowledge. The framework identifies NA and AR as important factors for the recommended career cluster, which aligns with domain knowledge. While VR is considered important for the cluster, the system identifies the feature as having a limited scope, a slight

contradiction of domain knowledge. This inconsistency is due to the data-driven explanation mechanism used in its explainability architecture, operating without domain-specific inputs, which is a limitation of the framework.

The proposed SHAP-NLG explainability framework primarily relies on data-driven feature attribution without incorporating explicit domain constraints or expert-defined reasoning rules. The explanations generated by the framework are based on SHAP values computed from domain-specific aptitude features and therefore capture statistically significant feature contribution patterns relevant to the recommendation context. Consequently, the framework is capable of addressing domain-related feature interactions to a considerable extent. However, certain explanations may occasionally differ from established domain expectations regarding the relative importance of specific aptitude factors for particular recommendation clusters, since SHAP explanations reflect statistical learning patterns rather than explicit expert psychological reasoning processes. To address this issue, the proposed work incorporated expert evaluation of both the recommendations and the generated explanations. Therefore, although the framework improves transparency and interpretability, further integration of domain expert knowledge and expert-guided validation mechanisms may enhance the semantic alignment and contextual reliability of the generated explanations.

10.6.Comparative Analysis: The Proposed Framework Against Existing Approaches

The interpretability framework that has been proposed employs data-driven (feature-driven) procedures to generate natural language explanations to provide a clear justification of how and why the system makes the recommendations it does. This offers significant benefits compared to the existing explainability systems. The differences between the proposed framework and the older methods are shown in Table 19.

Table 19. Comparison of the Proposed Explainability Framework with Existing Approaches.

Study	Explanation method	Output format	Limitations	The proposed work
(Yang et al., 2023)	Contextualized Knowledge Graph + Transformer over meta-paths	Graph reasoning (meta-paths)	Context-aware but not user-friendly explanation	✓ Provides context-aware and user-friendly explanations that are easy to interpret. ✓ Combines visual and textual outputs to suit both expert and non-expert users. ✓ Integrates SHAP's local and global interpretability features for detailed insight. ✓ Highlights key feature contributions through clear, natural-language explanations. ✓ Enhances transparency and strengthens user confidence in the recommendations.
(Tran, 2023)	LIME, SHAP (local), EBM/DPEBM (global), TF-IDF, LDA	Feature Importance + Topic keywords	Designed for technical interpretability rather than end-user accessibility	
(Shams et al., 2024)	LIME-based XAI with contextual features	Numeric feature importances	Limited to expert users; lacks a textual narrative	
(Tang et al., 2024)	SHAP for feature attribution	Visual explanations (SHAP plots)	Not accessible to non-specialists	

The comparison shows that existing explainable recommendation systems mainly focus on technical interpretability through feature importance scores, SHAP plots, or graph-based reasoning, which are often difficult for non-expert users to understand. The proposed work improves accessibility by combining visual explanations with natural-language interpretations that clearly describe feature contributions. In addition, the integration of both local and global SHAP-based interpretability enhances transparency, usability, and user confidence in the recommendation process.

10.7. Expert Evaluation of the Career Recommendations and Related Explanations

An expert evaluation was carried out on the recommendations and related explanations. Ten career experts, well-versed in the counselling process and

Differential Aptitude Tests, were provided with 16 unique candidate profiles. The candidate profiles are systematically selected to include diverse career paths, covering approximately 33% of career clusters in the dataset. The expert's evaluation provided a generally positive response with the average rating values on a five-point Likert scale ranging from 3.4 to 4.6 for the systems' recommendations and from 3.7 to 4.4 for the related explanations across the 16 profiles. The distribution of mean expert satisfaction scores on a five-point Likert scale for the recommendations and the related explanations across the 16 user profiles is illustrated in Figure 38.




Figure 38. Mean Expert Ratings for Each Profile on a Five-Point Scale (1 = Not satisfactory, 2 = Slightly satisfactory, 3 = Moderately satisfactory, 4 = Satisfactory, 5 = Highly satisfactory). As Presented in Author's Published Work (Rahim & Basheer, 2025b) Under a CC BY License

There are slight variations in the system-generated explanations against experts' knowledge. Additional mechanisms such as expert-reviewed explanation templates, reliability thresholds on SHAP values, and rule-based validation to detect explanations that conflict with domain knowledge can be included as possible future enhancements to strengthen explanation reliability.

10.8. Analysis of the Dynamic Career Expansion Module

The Dynamic Career Expansion Module is introduced to make the system scalable. This will make the system capable of recommending ever-evolving careers that are not present in the training dataset. This is done through a semi-supervised approach. When a new career is introduced, its ability profile is determined by a career expert and fed to the recommendation engine in disguise as a user aptitude profile. The

recommendation engine then suggests the most suitable existing career cluster for this career. The career is automatically added to the suggested cluster, making it available for recommendation to a suitable user profile in the future.

To evaluate the validity of the process, a new career, *Machine Learning Engineer*, which is not in the training dataset, is selected, and its aptitude parameters are determined by a career expert as shown below:

ML Engineer (NA, AR, VR, MR, SR, VA) = $[81, 67, 74, 30, 42, 73]$, which is depicted in Figure 39.




Figure 39. ML Career Profile

When this particular profile was fed to the system, the recommendation engine provided target career cluster label 903 as shown in Figure 40.




Figure 40. Suggested Career Cluster for the ML Engineer Career Profile.

The Dynamic Integration Module present in the DCEM automatically added this new career to the suggested career cluster 903. Then, when a new user profile aligned with the newly added career was fed to the system, it recommended the career cluster labelled 903 containing the newly added career, ML Engineer, as illustrated in Figure 41.




Figure 41. Illustration of Dynamic Career Expansion.

Thus, the DCEM module effectively addresses the scalability issue in the ever-evolving career landscape. The limitation of the system is the need for a career expert proficient in the differential aptitude test to profile new careers for their aptitude profile. The system is semi-supervised, where a man-machine combination is needed for proper functioning, which needs to be addressed further.

CHAPTER 11

CONCLUSION

This chapter summarises the key findings of this research, highlights its methodological and practical contributions, acknowledges its limitations, and proposes directions for future work.

The study presents an explainable, hierarchical, and multi-tiered career recommender system. It provides a new architecture for recommendation engines, which can be adapted to other domains as well. Its novel hierarchical and multi-tiered recommendation framework facilitates data-driven diversity in the recommendation process. The innovative explainability framework introduced in this work can provide data-driven (feature-driven) natural language explanations for recommendations generated by recommendation engine. This framework can also be adopted in other domains. The study also introduces a first-of-its-kind career dataset based on aptitude profiles to address the classical problems of data sparsity and cold start limitations in career recommender systems. To address the ever-evolving nature of the career landscape, the study proposes a semi-supervised solution, thus making the recommender system scalable and adaptable for future use. While maintaining high accuracy, precision, and F1-score, the system showcased its usability through experts' evaluation.

With the improved access to career advice, this work has the potential to deal with important issues such as high dropouts in higher educational institutions and job switches in industries. This system can fill the career counselling gap in a country like India, where the career counsellor-to-student ratio is extremely low and thus, all students can be given a personalized and data-driven career advice, helping them to make informed choices and an employable and skilled workforce.

11.1 Summary of Findings

This study aimed to develop an explainable, hierarchical and multi-tiered career recommender system with dynamic career expansion to provide personalised

recommendations to secondary students aligned with their aptitude profile. The main findings are:

1. ***Enhanced Career Diversity Through a Novel Hierarchical and Multi-Tiered Recommendation Engine Architecture:*** The recommendation engine, which is developed based on a novel hierarchical and multi-tiered architecture, outperforms the conventional recommendation models because the architecture enables providing wider and more diverse sets of career recommendations. Whereas traditional recommender systems offer a small pool of choices, the new multi-tiered framework can guarantee that candidates are presented with a broader range of various career opportunities that are relevant to their aptitude profiles. This systematic method increases career exploration with various options instead of choosing from a limited set of suggestions.
2. ***Improved Interpretability and Transparency Through an Innovative Hybrid Explainability Framework:*** The innovative explainability framework made the proposed recommendation engine highly interpretable. Through the use of SHAP-based natural language explanations, the framework makes the recommendation process transparent to the users, including those who do not have technical knowledge. Moreover, the framework can be customized to a wide range of recommendation systems in various fields. And thus, the system complies with the tenets of responsible artificial intelligence, promoting equity, trust, and user confidence in its recommendation process.
3. ***High-Accuracy Career Recommendations Using Content-Based Filtering:*** The proposed recommendation engine with its content-based filtering approach showcased its efficiency via different performance metrics like accuracy, precision, recall, and f1-score, all exceeding 99%. This ensures that the recommendations are reliable with low false positives, and the system can retrieve all relevant career choices for the candidate, proving its comprehensive coverage in career recommendations.
4. ***Comprehensive Career-Aptitude Dataset Encompassing Real-World Career Options:*** The study developed a comprehensive career dataset containing 612

unique careers based on differential aptitude tests. The dataset includes almost all real-world careers and is built by utilising the O*NET ability profiles. Thus, the recommendation engine built upon this dataset can recommend almost all real-world careers to a candidate based on their aptitude patterns. This ensures the system's ability to provide diverse career choices to a candidate, allowing him or her to further fine-tune these choices based on vocational interests and personality traits.

5. ***Scalable Career Recommendations Through a Semi-Supervised Dynamic Career Expansion Module:*** The dynamic career expansion module, which is made with the help of a semi-supervised architecture, enables the incorporation of continually evolving career paths into its spectrum. This aspect makes the system relevant in providing career guidance over time. This results in the system being resistant to long-term obsolescence. Thus, it can be trusted by the candidates and career counsellors as a long-term career guidance tool. Although additional improvements may be needed, the system will continue to be aligned to the ever-evolving career landscape.
6. ***Strategies for Addressing Data Sparsity and Cold Start Problems:*** With a rich dataset containing 612 unique real-world careers, a dynamic career expansion module to keep pace with the ever-evolving career landscape, and the effective use of the Kerala Differential Aptitude Test tool, which is standardised across the class X population in Kerala considering its socio-economic diversities, the system can address data sparsity and cold-start limitations making it able to recommend nearly all real-world careers aligned with the candidate profile.

11.2. Contributions of the Research

1. ***Advancement in Recommendation Frameworks:*** This study contributed a novel architecture for a recommendation system targeting diversity in its recommendation process. The novelty is achieved by adopting a hierarchical recommendation process that internally uses its Diverse Recommendation Module to diversify the recommendations, the Refinement Recommendation Module to refine these recommendations, and further provides a multi-tiered

output-ends with a vast array of career choices for a candidate. This architecture can be adapted to other domains where diversifying the recommendations is essential.

2. ***Enhancing Explainability in Career Recommendations:*** With an innovative hybrid explainability framework, integrating SHAP's feature-driven explanations with NLG techniques, the proposed system enhances transparency in its recommendation process, targeting non-expert users and expert users alike. SHAP provides feature attribution-based insights via Shapley values, which are mathematically grounded, and these Shapley values are translated into natural language justifications that even a non-technical user can understand. This hybrid explainability approach ensures transparency and trust in the recommendation process.
3. ***Creation of a Comprehensive Career-Aptitude Dataset:*** The research contributed a first-of-its-kind dataset containing 612 unique real-world careers. The careers have been profiled with their aptitude parameters. The dataset is adapted from the Occupational Information Network (O*NET) database. From more than 1000 careers, a significant 612 careers were selected. Their ability profile features are mapped and adapted to corresponding differential aptitude parameters to create the dataset. This dataset has been published with open access permission and can be a valuable resource for future researchers in the career guidance domain.
4. ***Aptitude-Driven Career Recommendation Addressing Data Sparsity and Cold-Start Limitations:*** The suggested recommendation engine, which is based on the content-based filtering model, provides recommendations based on the aptitude profiles of the candidate. The system produces satisfactory recommendations in accordance with the candidate profile, and the recommendations are enriched by diversity. The recommendation engine addresses the data sparsity and cold-start problems by training it on a high-quality aptitude-profiled dataset of 612 unique real-world careers. This is also

reinforced through the integration of the Kerala Differentiated Aptitude Tests, which is a standardised testing tool used for generating user profiles.

5. ***Improving System Scalability and Adaptability:*** Another contribution is a semi-supervised scalable architecture for the recommendation engine. The dynamic career expansion module is a man-machine system where new careers are profiled by a career expert for their aptitude profile. The career profile is then fed to the recommendation engine in disguise as a user profile. The recommendation engine recommends a matching career cluster for the new career, which is then automatically added to the recommended cluster. This approach enables the career recommendation engine to recommend ever-evolving careers, which are not present in the training dataset. This makes the system relevant over time, and candidates and career counsellors can depend on it as the recommendations are not obsolete in the long run.

11.3. Practical Implications

1. ***Real-World Implementation and Adoption:*** The Department of General Education, Government of Kerala, has agreed to integrate the outcomes of this research work with KDAT, and an MoU has been signed in this regard. The proposed system can support approximately 1,680 career counsellors in improving the counselling process for more than seven lakh students across Kerala. However, the current implementation is intended to function as an additional support tool and may require further enhancement through the integration of additional counselling dimensions.
2. ***Enhancing Career Guidance Accessibility:*** The work will increase the accessibility of career guidance to every student in a country like India, where the career counsellor-to-student ratio is very poor. Once the proposed recommendation engine is integrated into KDAT, every standard X student in Kerala can be provided with career guidance based on their aptitude.
3. ***Technology-Driven Solution:*** This work showcases how an AI-driven solution can be integrated into the education sector. This will facilitate personalized

education for every student. A recommender system with natural language interpretable solutions can tackle challenges in the education sector, where large-scale implementation without technology is challenging.

4. ***Provides Direction for Future Research:*** The proposed work paves the way for researchers in this domain, with the integration of theoretical foundations in career counselling, AI-driven solutions, and explainability frameworks. All of these elements can be further enhanced and can come up with a more holistic career recommender system. Along with the aptitude profile of the candidate, vocational interests and personality traits of the candidate can also be considered in future research.

11.4. Limitations of the Research

1. ***Dependence on Standardised Aptitude Test Alone:*** The main limitation of this work is that the career recommendation process considered only the aptitude of the candidate and bypassed other aspects such as vocational interests and personality traits. A holistic career recommender system should consider these important factors as well.
2. ***Limitations in Candidate Engagement and Response Authenticity:*** Honesty and commitment of the test taker are very critical in an assessment test. Random or careless answers in an objective-type test will jeopardize the entire process. Ensuring student engagement and honesty while taking the test is essential in an AI-based automated career recommender system.
3. ***Limitations in SHAP-Based Explanations and Interpretability:*** The proposed innovative explainability framework is purely based on feature-driven (data-driven) explanations. The Shapley value-based explanation may contradict well-established domain knowledge. Thus, if the framework is employed in any other domain, as the framework is adaptable to other domains also, it is essential to evaluate the explanation by experts in the respective domain. Bias in the dataset can also badly affect the explanation generated. Additionally, the sentence

transformer-based fusion layer can add hallucination in the textual explanations, which has been currently addressed using a systematic prompt.

4. ***Computational Cost and Development Challenges of SHAP Explainer:*** Building a SHAP explainer is highly computationally intensive. This creates obstacles for the model builder as it will affect debugging and fine-tuning the models due to the time delay between each process. This affects the iterative improvement of the system as far as a developer is concerned. Thus, the SHAP-based system creates challenges for optimizing the explainability framework.
5. ***Challenges in Dynamic Career Expansion and Expert Oversight:*** The dynamic career expansion module has limitations. Even if the module is effective in delivering scalability to the recommendation engine, an expert's intervention is required to profile new careers before it is integrated into the system. This semi-supervised approach restricts the system's potential to recommend evolving careers.
6. The proposed system is conceptually adaptable in terms of its recommendation engine architecture and explainability framework to similar domains. Within the career recommendation domain, adaptability is supported by the use of occupational profiles derived from the O*NET database. However, effective deployment in other educational or demographic contexts other than Kerala would require customised aptitude assessment tools similar to KDAT to address regional, cultural, and socioeconomic diversity in user profiling.

CHAPTER 12

RECOMMENDATIONS

1. ***Need for a Holistic Career Recommendation Approach:*** A holistic career recommender system should evaluate multiple factors like aptitude, personality traits, and vocational interests of a candidate, before recommending careers to a candidate. A system based on aptitude alone will have limitations in recommending personalized careers to a candidate. Future work can incorporate these additional dimensions to provide holistic career recommendations.
2. ***Incorporating a Multimodal Career Recommendation Framework:*** In the proposed system, diversity in career recommendations is achieved using a data-driven approach. By employing a hierarchical clustering approach and generating multiple models, the system provides a diverse array of careers through its multi-tiered endpoints. This approach can be refined by clustering the careers into different streams like science, commerce, and humanities. Based on these divisions of careers, along with the existing data-driven clustering, more meaningful recommendations can be provided across different domains.
3. ***Enhancing the Dynamic Career Expansion Model for Scalability:*** There is scope to improve the proposed dynamic career expansion module. The current implementation is based on a semi-supervised approach. An expert's intervention is required in the current implementation. Analyzing labour market data and career-related websites, new careers can be automatically identified and added to the system to keep it up-to-date dynamically.

REFERENCES

- Adomavicius, G., & Tuzhilin, A. (2010). *Context-aware recommender systems*. In Recommender systems handbook (pp. 217-253). Boston, MA: Springer US.
- Ahmadian, S., Joorabloo, N., Jalili, M., & Ahmadian, M. (2022). *Alleviating data sparsity problem in time-aware recommender systems using a reliable rating profile enrichment approach*. Expert Systems with Applications, 187, 115849.
- Ai, Q., & Narayanan, R. L. (2021, October). *Model-agnostic vs. model-intrinsic interpretability for explainable product search*. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management (pp. 5-15).
- Akhtar, M. A. K., Kumar, M., & Nayyar, A. (2024). *Transparency and Accountability in Explainable AI: Best Practices*. In Towards Ethical and Socially Responsible Explainable AI: Challenges and Opportunities (pp. 127-164). Cham: Springer Nature Switzerland.
- Ali, P. J. M., Faraj, R. H., Koya, E., Ali, P. J. M., & Faraj, R. H. (2014). *Data normalization and standardization: a technical report*. Mach Learn Tech Rep, 1(1), 1-6.
- Altmann, A., Tološi, L., Sander, O., & Lengauer, T. (2010). *Permutation importance: a corrected feature importance measure*. Bioinformatics, 26(10), 1340-1347.
- Amarasinghe, K., & Manic, M. (2019, June). *Explaining what a neural network has learned: Toward transparent classification*. In 2019 IEEE International Conference on Fuzzy Systems (FUZZ-IEEE) (pp. 1-6). IEEE.
- Amazon. (n.d.). Amazon.in. Retrieved December 4, 2024, from <https://www.amazon.in/>
- American School Counselor Association. (2024). *Student-to-school-counselor ratio 2022–2023*. <https://www.schoolcounselor.org/getmedia/a988972b-1faa-4b5f-8b9e-a73b5ac44476/ratios-22-23-alpha.pdf>
- Angelini, G. (2023). *Big five model personality traits and job burnout: a systematic literature review*. BMC psychology, 11(1), 49.
- Antamis, T., Drosou, A., Vafeiadis, T., Nizamis, A., Ioannidis, D., & Tzovaras, D. (2024). *Interpretability of deep neural networks: A review of methods, classification and hardware*. Neurocomputing, 128204.
- Apoorva Ramachandra. (2024, February 27). *Indian Budget Series: Budget for Education. Centre for Development, Policy & Practice (CDPP)*. Retrieved from <https://www.cdpp.co.in/articles/indian-budget-series-budget-for-education>.
- Arafeh, M., Ceravolo, P., Mourad, A., Damiani, E., & Bellini, E. (2021). *Ontology based recommender system using social network data*. Future generation computer systems, 115, 769-779.
- Atienza, J. R. D., Hernandez, R. M., Castillo, R. L., De Jesus, N. M., & Buenas, L. J. E. (2022). *A Deep Neural Network in a Web-based Career Track Recommender System*

-
- for Lower Secondary Education. 2022 2nd Asian Conference on Innovation in Technology (ASIANCON), 1–6. IEEE.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K. R., & Samek, W. (2015). *On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation*. PloS one, 10(7), e0130140.
- Balakrishnan, P., & Karthiyayini, S. (2024). *Challenges of career counseling in schools: A quantitative investigation of counselors opinions*. Scholarly Research Journal for Interdisciplinary Studies, 13(83), 172–179. <https://doi.org/10.5281/zenodo.12799290>.
- Bennett, G. K., Seashore, H. G., & Wesman, A. G. (1947). *Differential aptitude tests*.
- Bennett, G. K., Seashore, H. G., & Wesman, A. G. (1956). *The Differential Aptitude Tests: an overview*. Personnel & Guidance Journal.
- Bennett, G. K., Seashore, H. G., & Wesman, A. G. (1991). *Differential Aptitude Tests for personnel and career assessment: Technical manual*. San Antonio, TX: Psychological Corporation/Harcourt Brace.
- Betz, N. E., Fitzgerald, L. F., & Hill, R. E. (1989). *Trait-factor theories: Traditional cornerstone of career theory*. Handbook of career theory, 26-40.
- Bhaskaran, S., & Marappan, R. (2023). *Design and analysis of an efficient machine learning based hybrid recommendation system with enhanced density-based spatial clustering for digital e-learning applications*. Complex & Intelligent Systems, 9(4), 3517-3533.
- Binet, A., & Simon, T. (1961). *The development of intelligence in children*.
- Boake, C. (2002). *From the Binet–Simon to the Wechsler–Bellevue: Tracing the history of intelligence testing*. Journal of clinical and experimental neuropsychology, 24(3), 383-405.
- Boyle, G. J. (1995). *Myers-Briggs type indicator (MBTI): some psychometric limitations*. Australian Psychologist, 30(1), 71-74.
- Briggs, K. C. (1976). *Myers-Briggs type indicator*. Palo Alto, CA: Consulting Psychologists Press.
- Burke, R. (2002). *Hybrid recommender systems: Survey and experiments*. User modeling and user-adapted interaction, 12, 331-370.
- Carter, R. T., & Swanson, J. L. (1990). *The validity of the Strong Interest Inventory with Black Americans: A review of the literature*. Journal of Vocational Behavior, 36(2), 195-209.
- Castells, P., Hurley, N., & Vargas, S. (2021). *Novelty and diversity in recommender systems*. In *Recommender systems handbook* (pp. 603-646). New York, NY: Springer US.
- Chen, L., De Gemmis, M., Felfernig, A., Lops, P., Ricci, F., & Semeraro, G. (2013). *Human decision making and recommender systems*. ACM Transactions on Interactive Intelligent Systems (TiiS), 3(3), 1-7.
- Cherry, K. (2022). *Theories of intelligence in psychology*. Verywell mind. Available at: <https://www.verywellmind.com/theories-of-intelligence-2795035>. Accessed: May, 7, 2023.
-

-
- Chou, T. N. (2019, July). *An explainable hybrid model for bankruptcy prediction based on the decision tree and deep neural network*. In 2019 IEEE 2nd international conference on knowledge innovation and invention (ICKII) (pp. 122-125). IEEE.
- Chromik, M. (2021). *Making SHAP Rap: Bridging local and global insights through interaction and narratives*. In Human-Computer Interaction–INTERACT 2021: 18th IFIP TC 13 International Conference, Bari, Italy, August 30–September 3, 2021, Proceedings, Part II 18 (pp. 641-651). Springer International Publishing.
- Clough, J. R., Oksuz, I., Puyol-Antón, E., Ruijsink, B., King, A. P., & Schnabel, J. A. (2019, October). *Global and local interpretability for cardiac MRI classification*. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 656-664). Cham: Springer International Publishing.
- Cohen, J. (2013). *“Statistical power analysis for the behavioral sciences”*. routledge.
- Constable, F. C. (1905). *Poverty and Hereditary Genius: A Criticism of Mr. Francis Galton's Theory of Hereditary Genius*. AC Fifield.
- Council, A. U. P. P. (2017). *Statement on algorithmic transparency and accountability*. Commun. ACM.
- Dan, Y., & Li, M. (2022, December). *SWOT analysis and strategies of high-quality career guidance education in universities in the context of big data*. In 2022 19th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP) (pp. 1-4). IEEE.
- Dasari, C. M., & Bhukya, R. (2022). *Explainable deep neural networks for novel viral genome prediction*. Applied Intelligence, 52(3), 3002-3017.
- Davis, K., Christodoulou, J., Seider, S., & Gardner, H. E. (2011). *The theory of multiple intelligences*. Davis, K., Christodoulou, J., Seider, S., & Gardner, H. (2011). The theory of multiple intelligences. In RJ Sternberg & SB Kaufman (Eds.), Cambridge Handbook of Intelligence, 485-503.
- De Fruyt, F., & Mervielde, I. (1997). *The five-factor model of personality and Holland's RIASEC interest types*. Personality and individual differences, 23(1), 87-103.
- DeVon, H. A., Block, M. E., Moyle-Wright, P., Ernst, D. M., Hayden, S. J., Lazzara, D. J., ... & Kostas-Polston, E. (2007). *A psychometric toolbox for testing validity and reliability*. Journal of Nursing scholarship, 39(2), 155-164.
- Dhar, B. (n.d.). *Career Recommendation Dataset* [Data set]. Kaggle. Retrieved from <https://www.kaggle.com/datasets/breejeshdhar/career-recommendation-dataset>
- Dheya Career Mentors. (2024, September 26). *Scope & growth of career counseling in India*. Dheya. <https://www.dheya.com/growth-of-career-counseling-in-india/>
- Dwivedi, R., Dave, D., Naik, H., Singhal, S., Omer, R., Patel, P., ... & Ranjan, R. (2023). *Explainable AI (XAI): Core ideas, techniques, and solutions*. ACM Computing Surveys, 55(9), 1-33.
- Economic Times. (2023, September 7). *India's education system vs job market reality: A call for reform*. The Economic Times.
-

- Ezz, M., & Elshenawy, A. (2020). *Adaptive recommendation system using machine learning algorithms for predicting student's best academic program*. Education and Information Technologies, 25(4), 2733-2746.
- Fernández, A. L., & Abe, J. (2018). *Bias in cross-cultural neuropsychological testing: problems and possible solutions*. Culture and Brain, 6, 1-35.
- Flynn, J. R. (1984). *The Mean IQ of Americans: Massive Gains 1932 to 1978*. Psychological Bulletin, 95(1), 29-51. [<https://doi.org/10.1037/0033-2909.95.1.29>]
- Foundation for Responsive Governance. (2025). *Trends in finances for Ministry of Education: Budget insights (Vol. 2, Issue 11)*. Retrieved from https://educationforallinindia.com/wp-content/uploads/2025/02/MoE-Budget-Insights-2025_foundation_by_responsive-governance.pdf
- Freitas, A. A. (2019, August). *Automated machine learning for studying the trade-off between predictive accuracy and interpretability*. In International cross-domain conference for machine learning and knowledge extraction (pp. 48-66). Cham: Springer International Publishing.
- Furnham, A. (1996). *The big five versus the big four: the relationship between the Myers-Briggs Type Indicator (MBTI) and NEO-PI five factor model of personality*. Personality and individual differences, 21(2), 303-307.
- Gallego, M. G., Perez de los Cobos, A. P., & Gallego, J. C. G. (2021). *Identifying students at risk to academic dropout in higher education*. Education Sciences, 11(8), 427.
- Galton, F. (1891). *Hereditary genius*. D. Appleton.
- Gardner, H. (1987). *The theory of multiple intelligences*. Annals of dyslexia, 19-35.
- Gardner, H. E. (2011). *Frames of mind: The theory of multiple intelligences*. Basic books.
- Ge, Y., Xiong, H., Tuzhilin, A., Xiao, K., Gruteser, M., & Pazzani, M. (2010, July). *An energy-efficient mobile recommender system*. In Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 899-908).
- Giabelli, A., Malandri, L., Mercorio, F., Mezzanzanica, M., & Seveso, A. (2021). *Skills2Job: A recommender system that encodes job offer embeddings on graph databases*. Applied Soft Computing, 101, 107049.
- Gm, D., Goudar, R. H., Kulkarni, A. A., Rathod, V. N., & Hukkeri, G. S. (2024). *A digital recommendation system for personalized learning to enhance online education: A review*. IEEE Access, 12, 34019-34041.
- Gong, H., Li, Y., Zhang, J., Zhang, B., & Wang, X. (2024). *"A new filter feature selection algorithm for classification task by ensembling pearson correlation coefficient and mutual information"*. Engineering Applications of Artificial Intelligence, 131, 107865.
- Google. (n.d.). Google Search. Retrieved December 4, 2024, from <https://www.google.co.in/>
- GoogleNews. (n.d.). Google News. Retrieved December 4, 2024, from <https://news.google.com/>

- Gottfredson, L. S. (2005). Applying Gottfredson's theory of circumscription and compromise in career guidance and counseling. *Career development and counseling: Putting theory and research to work*, 1, 71-100. John Wiley & Sons, Inc.
- Gulzar, Z., Leema, A. A., & Deepak, G. (2018). *Pcrs: Personalized course recommender system based on hybrid approach*. *Procedia Computer Science*, 125, 518–524.
- Guo, G. (2012). *Resolving data sparsity and cold start in recommender systems*. In User Modeling, Adaptation, and Personalization: 20th International Conference, UMAP 2012, Montreal, Canada, July 16-20, 2012. *Proceedings 20* (pp. 361-364). Springer Berlin Heidelberg.
- Han, J., Pei, J., & Tong, H. (2022). *Data mining: concepts and techniques*. Morgan kaufmann.
- HANSEN, J. I. C., & Bourne, B. (1988). *Edward Kellog Strong, Jr.: First author of the Strong interest inventory*. *Journal of Counseling & Development*, 67(1), 9-16.
- Hassija, V., Chamola, V., Mahapatra, A., Singal, A., Goel, D., Huang, K., ... & Hussain, A. (2024). *Interpreting black-box models: a review on explainable artificial intelligence*. *Cognitive Computation*, 16(1), 45-74.
- 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).
- He, Z., Rakin, A. S., & Fan, D. (2019). *Parametric noise injection: Trainable randomness to improve deep neural network robustness against adversarial attack*. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 588-597).
- Hidasi, B. (2015). *Session-based Recommendations with Recurrent Neural Networks*. arXiv preprint arXiv:1511.06939.
- Hoff, K. A., Chu, C., Einarsdóttir, S., Briley, D. A., Hanna, A., & Rounds, J. (2022). *Adolescent vocational interests predict early career success: Two 12-year longitudinal studies*. *Applied Psychology*, 71(1), 49-75.
- Hogan, R., & Blake, R. (1999). *John Holland's vocational typology and personality theory*. *Journal of Vocational Behavior*, 55(1), 41-56.
- Holland, J. L. (1959). *A theory of vocational choice*. *Journal of counseling psychology*, 6(1), 35.
- Hort, M., Chen, Z., Zhang, J. M., Harman, M., & Sarro, F. (2024). *Bias mitigation for machine learning classifiers: A comprehensive survey*. *ACM Journal on Responsible Computing*, 1(2), 1-52.
- Houcan, Z., & Kan, Z. (2023). *Binet-Simon Intelligence Scale*. In *The ECPH Encyclopedia of Psychology* (pp. 1-2). Singapore: Springer Nature Singapore.
- Hugging Face. (2024). *Job recommendation system* [Data set]. Retrieved from <https://huggingface.co/datasets/job-recommendation-system>

- Jahmunah, V., Ng, E. Y., Tan, R. S., Oh, S. L., & Acharya, U. R. (2022). *Explainable detection of myocardial infarction using deep learning models with Grad-CAM technique on ECG signals*. Computers in Biology and Medicine, 146, 105550.
- Jalili, M., Ahmadian, S., Izadi, M., Moradi, P., & Salehi, M. (2018). *Evaluating collaborative filtering recommender algorithms: a survey*. IEEE access, 6, 74003-74024.
- Jensen, A. R. (1980). *Bias in Mental Testing*.
- Jiang, S., & Guo, Y. (2022). *Reasons for college major-job mismatch and labor market outcomes: evidence from China*. China Economic Review, 74, 101822.
- Joshi, A., Kale, S., Chandel, S., & Pal, D. K. (2015). "Likert scale: Explored and explained". British journal of applied science & technology, 7(4), 396.
- Kaggle. (2022). *Career recommendation dataset*[Dataset]. Retrieved from <https://www.kaggle.com/datasets/breejeshdhar/career-recommendation-dataset>
- Kamal, N., Sarker, F., Rahman, A., Hossain, S., & Mamun, K. A. (2024). *Recommender system in academic choices of higher education: a systematic review*. IEEE Access, 12, 35475-35501.
- Kerala Higher Secondary Education Department. (2024, September 28). *Kerala Differential Aptitude Tests*. Retrieved September 28, 2024, from <https://www.careerguidance.dhse.kerala.gov.in>
- Khanal, S. S., Prasad, P. W. C., Alsadoon, A., & Maag, A. (2020). *A systematic review: machine learning based recommendation systems for e-learning*. Education and Information Technologies, 25, 2635–2664.
- Kim, M. H., & Beier, M. E. (2020). *The college-to-career transition in STEM: An eleven-year longitudinal study of perceived and objective vocational interest fit*. Journal of Vocational Behavior, 123, 103506.
- Kim, M. H., & Beier, M. E. (2020). *The college-to-career transition in STEM: An eleven year longitudinal study of perceived and objective vocational interest fit*. Journal of Vocational Behavior, 123, 103506.
- Kim, M. H., & Beier, M. E. (2020). *The college-to-career transition in STEM: An eleven-year longitudinal study of perceived and objective vocational interest fit*. Journal of Vocational Behavior, 123, 103506.
- Ko, H., Lee, S., Park, Y., & Choi, A. (2022). *A survey of recommendation systems: recommendation models, techniques, and application fields*. Electronics, 11(1), 141.
- Konstan, J. A., & Riedl, J. (2012). *Recommender systems: from algorithms to user experience*. User modeling and user-adapted interaction, 22, 101-123.
- Koohi, H., & Kiani, K. (2016). *User based collaborative filtering using fuzzy C-means*. Measurement, 91, 134-139.
- Koren, Y., Rendle, S., & Bell, R. (2021). *Advances in collaborative filtering*. Recommender systems handbook, 91-142.

- Krumboltz, J. D., Mitchell, A. M., & Jones, G. B. (1976). *A social learning theory of career selection*. The counseling psychologist, 6(1), 71-81.
- Kumalasari, L. D., & Susanto, A. (2020). *Recommendation system of information technology jobs using collaborative filtering method based on LinkedIn skills endorsement*. Sisforma, 6(2), 63-72.
- Kumar, A., & Sharma, A. (2013). *Alleviating sparsity and scalability issues in collaborative filtering based recommender systems*. In Proceedings of the International Conference on Frontiers of Intelligent Computing: Theory and Applications (FICTA) (pp. 103-112). Springer Berlin Heidelberg.
- Lang, J. W., & Kell, H. J. (2020). *General mental ability and specific abilities: Their relative importance for extrinsic career success*. Journal of Applied Psychology, 105(9), 1047.
- Last.fm. (n.d.). Last.fm. Retrieved December 4, 2024, from <https://www.last.fm/>
- Lent, R. W., Brown, S. D., & Hackett, G. (1994). *Toward a unifying social cognitive theory of career and academic interest, choice, and performance*. Journal of vocational behavior, 45(1), 79-122.
- Li, M., Sun, H., Huang, Y., & Chen, H. (2024). *Shapley value: from cooperative game to explainable artificial intelligence*. Autonomous Intelligent Systems, 4(1), 2.
- Linden, G., Smith, B., & York, J. (2003). *Amazon. com recommendations: Item-to-item collaborative filtering*. IEEE Internet computing, 7(1), 76-80.
- Liu, H., Hu, Z., Mian, A., Tian, H., & Zhu, X. (2014). *A new user similarity model to improve the accuracy of collaborative filtering*. Knowledge-based systems, 56, 156-166.
- Lohman, D. F., & Nicpon, M. F. (2023). *Ability testing & talent identification*. In Identification (pp. 287-335). Routledge.
- Lowe, R. (1998). *The educational impact of the eugenics movement*. International Journal of Educational Research, 27(8), 647-660.
- Lundberg, S. (2017). *A unified approach to interpreting model predictions*. arXiv preprint arXiv:1705.07874.
- Macha, D., Kozielski, M., Wróbel, Ł., & Sikora, M. (2022). *RuleXAI—A package for rule-based explanations of machine learning model*. SoftwareX, 20, 101209.
- Mammadov, S. (2022). *Big Five personality traits and academic performance: A meta analysis*. Journal of personality, 90(2), 222-255.
- Manconi, A., Armano, G., Gnocchi, M., & Milanese, L. (2022). *A soft-voting ensemble classifier for detecting patients affected by COVID-19*. Applied Sciences, 12(15), 7554.
- Mashayekhi, Y., Li, N., Kang, B., Lijffijt, J., & De Bie, T. (2024). *A challenge-based survey of e-recruitment recommendation systems*. ACM Computing Surveys, 56(10), 1-33.
- McCrae, R. R., & Costa, P. T. (2003). *Personality in adulthood: A five-factor theory perspective*. Guilford press.

- Mellenbergh, G. J. (1989). *Item bias and item response theory*. *International journal of educational research*, 13(2), 127-143.
- Molnar, C. (2020). *Interpretable machine learning*. Lulu. Com.
- MovieLens. (n.d.). MovieLens. Retrieved December 4, 2024, from <https://movielens.org/>
- Murdoch, W. J., Singh, C., Kumbier, K., Abbasi-Asl, R., & Yu, B. (2019). *Interpretable machine learning: definitions, methods, and applications*. arXiv preprint arXiv:1901.04592.
- Murphy, K. R. (Ed.). (2014). *A critique of emotional intelligence: What are the problems and how can they be fixed?* Psychology Press.
- Natarajan, S., Vairavasundaram, S., Natarajan, S., & Gandomi, A. H. (2020). *Resolving data sparsity and cold start problem in collaborative filtering recommender system using linked open data*. *Expert Systems with Applications*, 149, 113248.
- National Center for O*NET Development. *Browse by Abilities*. O*NET OnLine. Sponsored by the U.S. Department of Labor, Employment and Training Administration (USDOL/ETA). Retrieved 02-12-2023, from <https://www.onetonline.org/find/descriptor/browse/1.A>
- National Center for ONET Development. (n.d.). *Browse by ONET data descriptor: Abilities (1.A)*. O*NET Online. Retrieved December 14, 2024, from <https://www.onetonline.org/find/descriptor/browse/1.A>
- National Center for ONET Development. (n.d.). *O*NET OnLine Help Scales, Ratings, and Standardized Scores*. O*NET Online. Retrieved December 14, 2024, from <https://www.onetonline.org/help/online/scales>
- Netflix. (n.d.). Netflix India. Retrieved December 4, 2024, from <https://www.netflix.com/in/>
- Ngiam, A. T. C., Islam, M. R., & Xu, G. (2022). ECTRS: *A Personalised Early Career Trajectory Recommender System for Youths*. 2022 9th International Conference on Behavioural and Social Computing (BESC), 1–7. IEEE.
- Nguyen, A., Yosinski, J., & Clune, J. (2015). *Deep neural networks are easily fooled: High confidence predictions for unrecognizable images*. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 427-436).
- Obeid, C., Lahoud, C., El, K. H., & Champin, P.-A. (2022). *A novel hybrid recommender system approach for student academic advising named COHRS, supported by case-based reasoning and ontology*. *Computer Science and Information Systems*, 19(2), 979–1005.
- ONET Resource Center. (n.d.). *Abilities*. ONET OnLine. Retrieved December 14, 2024, from <https://www.onetonline.org/find/descriptor/browse/1.A>
- Pajkovic, N. (2022). *Algorithms and taste-making: Exposing the Netflix Recommender System's operational logics*. *Convergence*, 28(1), 214-235.
- Panteli, A., & Boutsinas, B. (2023). *Addressing the cold-start problem in recommender systems based on frequent patterns*. *Algorithms*, 16(4), 182.
-

- Panwar, H., Gupta, P. K., Siddiqui, M. K., Morales-Menendez, R., Bhardwaj, P., & Singh, V. (2020). *A deep learning and grad-CAM based color visualization approach for fast detection of COVID-19 cases using chest X-ray and CT-Scan images*. *Chaos, Solitons & Fractals*, 140, 110190.
- Patil, R. (n.d.). *Computer Science Students Career Prediction Dataset* [Data set]. Kaggle. Retrieved from <https://www.kaggle.com/datasets/devildyno/computer-science-students-career-prediction>
- Pazzani, M. J. (2007). *Content-Based Recommendation Systems*.
- Pazzani, M., & Billsus, D. (1997). *Learning and revising user profiles: The identification of interesting web sites*. *Machine learning*, 27, 313-331.
- Perchinunno, P., Bilancia, M., & Vitale, D. (2021). *A statistical analysis of factors affecting higher education dropouts*. *Social Indicators Research*, 156, 341-362.
- Phan, W. M. J., & Rounds, J. (2018). *Examining the duality of Holland's RIASEC types: Implications for measurement and congruence*. *Journal of vocational behavior*, 106, 22-36.
- Philip, S., Shola, P., & Ovyne, A. (2014). *Application of content-based approach in research paper recommendation system for a digital library*. *International Journal of Advanced Computer Science and Applications*, 5(10).
- Pimplikar, A., Pati, R., Singh, A., & Shah, A. (2023). *CareerMap- Mapping Tech Roles With Personality & Skills (Version 1)* [Dataset]. Mendeley Data. <https://doi.org/10.17632/5z68cvxssn.1>
- Pino, P., Parra, D., Besa, C., & Lagos, C. (2021, September). *"Clinically correct report generation from chest x-rays using templates"*. In *International Workshop on Machine Learning in Medical Imaging* (pp. 654-663). Cham: Springer International Publishing.
- Pomplun, M., & Custer, M. (2005). *The construct validity of the Stanford-Binet 5 measures of working memory*. *Assessment*, 12(3), 338-346.
- Press Trust of India. (2024, February 1). *Interim Budget 2024-25: School education gets record allocation of over Rs 73,000 crore*. *The Economic Times*. <https://economictimes.indiatimes.com/news/india/interim-budget-2024-25-school-education-gets-record-allocation-of-over-rs-73000-crore/articleshow/107331620.cms>
- Privacy Regulation. (n.d.). Retrieved from <https://gdpr.eu/what-is-gdpr/>
- Rahim, M. N., & Basheer, K. P. M. (2025a). A Hierarchical and Multi-Tiered Personalized Career Recommender System Tailored to Individual Aptitudes. *Indian Journal of Science and Technology*, 18(3), 231-244.
- Rahim, M. N., & Basheer, K. P. M. (2025b). A hybrid explainability framework for recommender systems using SHAP with natural language interpretations. *International Journal of Innovative Research and Scientific Studies*, 8(6), 687-703. <https://doi.org/10.53894/ijirss.v8i6.9669>
- Rashid, A. H. A., Mohamad, M., Masrom, S., & Selamat, A. (2022). *Student Career Recommendation System Using Content-Based Filtering Method*. 2022 3rd

- International Conference on Artificial Intelligence and Data Sciences (AiDAS), 60–65. IEEE.
- Raven, J. (2000). *The Raven's progressive matrices: change and stability over culture and time*. Cognitive psychology, 41(1), 1-48.
- Rebuffi, S. A., Goyal, S., Calian, D. A., Stimberg, F., Wiles, O., & Mann, T. A. (2021). *Data augmentation can improve robustness*. Advances in neural information processing systems, 34, 29935-29948.
- Rechtemmer, A., & Yin, M. (2022, April). *When confidence meets accuracy: Exploring the effects of multiple performance indicators on trust in machine learning models*. In Proceedings of the 2022 chi conference on human factors in computing systems (pp. 1-14).
- Resnick, P., & Varian, H. R. (1997). *Recommender systems*. Communications of the ACM, 40(3), 56-58.
- Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994, October). *Grouplens: An open architecture for collaborative filtering of netnews*. In Proceedings of the 1994 ACM conference on Computer supported cooperative work (pp. 175-186).
- Reynolds, C. R., Altmann, R. A., & Allen, D. N. (2021). *The problem of bias in psychological assessment*. In Mastering modern psychological testing: Theory and methods (pp. 573-613). Cham: Springer International Publishing.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August). *"Why should i trust you?" Explaining the predictions of any classifier*. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining (pp. 1135-1144).
- Roy, D., & Dutta, M. (2022). *A systematic review and research perspective on recommender systems*. Journal of Big Data, 9(1), 59.
- Saeed, T., Sufian, M., Ali, M., & Rehman, A. U. (2021). *Convolutional Neural Network Based Career Recommender System for Pakistani Engineering Students*. 2021 International Conference on Innovative Computing (ICIC), 1–10. IEEE.
- Said, A. (2025). *On explaining recommendations with Large Language Models: a review*. Frontiers in Big Data, 7, 1505284.
- Salih, A. M., Raisi-Estabragh, Z., Galazzo, I. B., Radeva, P., Petersen, S. E., Lekadir, K., & Menegaz, G. (2024). *A perspective on explainable artificial intelligence methods: SHAP and LIME*. Advanced Intelligent Systems, 2400304.
- Salovey, P., & Mayer, J. D. (1990). *Emotional intelligence. Imagination, cognition and personality*, 9(3), 185-211.
- Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001, April). *Item-based collaborative filtering recommendation algorithms*. In Proceedings of the 10th international conference on World Wide Web (pp. 285-295).
- Savickas, M. (2019). *Career counseling* (pp. xvi-194). Washington, DC: American Psychological Association.

- Schellingerhout, R., Barile, F., & Tintarev, N. (2023, July). *A Co-design Study for Multi-stakeholder Job Recommender System Explanations*. In World Conference on Explainable Artificial Intelligence (pp. 597-620). Cham: Springer Nature Switzerland.
- Schwegler, R. A. (1914). *The Binet-Simon Scale of Intelligence*. Topeka, Kansas: State Printing Office, 56.
- Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., & Batra, D. (2020). *Grad-CAM: visual explanations from deep networks via gradient-based localization*. International journal of computer vision, 128, 336-359.
- Setiawati, F. A. (2020). *Aptitude Test's Predictive Ability for Academic Success in Psychology Student*. Psychological research and intervention, 3(1), 1-12.
- Shams, M. Y., Gamel, S. A., & Talaat, F. M. (2024). *Enhancing crop recommendation systems with explainable artificial intelligence: a study on agricultural decision-making*. Neural Computing and Applications, 36(11), 5695-5714.
- SHAP. (n.d.). shap.KernelExplainer. Retrieved November 14, 2024, from <https://shap.readthedocs.io/en/latest/generated/shap.KernelExplainer.html>
- Shrivastav, U. (n.d.). *Career Prediction Dataset* [Data set]. Kaggle. Retrieved from <https://www.kaggle.com/datasets/utkarshshrivastav07/career-prediction-dataset>
- Singh, M. (2020). *Scalability and sparsity issues in recommender datasets: a survey*. Knowledge and Information Systems, 62(1), 1-43.
- Snow, R. E. (2021). *Aptitude processes*. In *Aptitude, learning, and instruction* (pp. 27-64). Routledge.
- Sotgiu, A., Pintor, M., & Biggio, B. (2022, August). *Explainability-based debugging of machine learning for vulnerability discovery*. In Proceedings of the 17th International Conference on Availability, Reliability and Security (pp. 1-8).
- Stein, R., & Swan, A. B. (2019). *Evaluating the validity of Myers-Briggs Type Indicator theory: A teaching tool and window into intuitive psychology*. Social and Personality Psychology Compass, 13(2), e12434.
- Stern, W. (1914). *The psychological methods of testing intelligence* (No. 13). Warwick & York.
- Super, D. E. (1957). *The psychology of careers; an introduction to vocational development*.
- Super, D. E. (1969). Vocational development theory: Persons, positions, and processes. *The counseling psychologist*, 1(1), 2-9.
- Tahir, H. A., Alayed, W., Hassan, W. U., & Haider, A. (2024). *A Novel Hybrid XAI Solution for Autonomous Vehicles: Real-Time Interpretability Through LIME-SHAP Integration*. Sensors, 24(21), 6776.
- Tang, S., Zhang, H., Liang, J., Tang, S., Li, L., Li, Y., ... & Zhou, Y. (2024). *Prostate cancer treatment recommendation study based on machine learning and SHAP interpreter*. Cancer Science.

- Tarnowska, K., Ras, Z. W., & Daniel, L. (2020). *Recommender system for improving customer loyalty* (Vol. 1). Cham: Springer International Publishing.
- Terman, L. M. (1916). *The measurement of intelligence* (Vol. 191). Boston: Houghton Mifflin Company.
- Terman, L. M. (1918). *The use of intelligence tests in the army*. Psychological bulletin, 15(6), 177.
- The Myers-Briggs Company. (n.d.). *Strong Interest Inventory*®. Retrieved December 1, 2024, from <https://www.themyersbriggs.com/en-US/Products-and-Services/Strong>
- Thorat, P. B., Goudar, R. M., & Barve, S. (2015). *Survey on collaborative filtering, content-based filtering and hybrid recommendation system*. International Journal of Computer Applications, 110(4), 31-36.
- Tran, T. H. A. (2023). *Explainable Artificial Intelligence in Job Recommendation Systems* (Master's thesis, University of Twente).
- van der Linden, D., Dunkel, C. S., De Zeeuw, E. J., Wu, P., & Pelt, D. H. (2021). *The general factor of personality (GFP) and vocational interests: a test of social effectiveness at the behavioral and genetic level*. Journal of Business and Psychology, 1-22.
- Vimbi, V., Shaffi, N., & Mahmud, M. (2024). *Interpreting artificial intelligence models: a systematic review on the application of LIME and SHAP in Alzheimer's disease detection*. Brain Informatics, 11(1), 10.
- Viswan, V., Shaffi, N., Mahmud, M., Subramanian, K., & Hajamohideen, F. (2024). *Explainable artificial intelligence in Alzheimer's disease classification: A systematic review*. Cognitive Computation, 16(1), 1-44.
- Waghmode, M., Awati, J., & Shukla, M. (2024). *Synergizing Aptitude, Interest, and Intelligence: A Systematic Approach to Career Selection Decision-making in Higher Education*. Journal of Engineering Education Transformations, 37(Special Issue 2).
- Walek, B., & Fojtik, V. (2020). *A hybrid recommender system for recommending relevant movies using an expert system*. Expert Systems with Applications, 158, 113452.
- Wang, L. (1993). *The Differential Aptitude Test: A Review and Critique*.
- Wang, S., Zhang, X., Wang, Y., & Ricci, F. (2024). *Trustworthy recommender systems*. ACM Transactions on Intelligent Systems and Technology, 15(4), 1-20.
- Wang, X., & Hu, B. (2024). *Machine Learning Algorithms for Improved Product Design User Experience*. IEEE Access.
- Wang, X., He, X., Cao, Y., Liu, M., & Chua, T. S. (2019, July). *Kgat: Knowledge graph attention network for recommendation*. In Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining (pp. 950-958).
- Washington, K., Malone, C., Briggs, C., & Reed, G. K. (2016). *Testing and African Americans: Testing monograph from the association of black psychologists*. Testing and assessment with persons of and communities of color, 3-11.

- Waterhouse, L. (2006). *Multiple intelligences, the Mozart effect, and emotional intelligence: A critical review*. Educational psychologist, 41(4), 207-225.
- Waworuntu, E. C., Kainde, S. J., & Mandagi, D. W. (2022). *Work-life balance, job satisfaction and performance among millennial and Gen Z employees: a systematic review*. Society, 10(2), 384-398.
- Wechsler, D. (1940). *The measurement of adult intelligence*. The Journal of Nervous and Mental Disease, 91(4), 548.
- Wheebox, AICTE, CII, AIU, & UNDP. (2025). *India Skills Report 2025*. Wheebox. Retrieved from https://wheebox.com/assets/pdf/ISR_Report_2025.pdf.
- Xu, F., Uszkoreit, H., Du, Y., Fan, W., Zhao, D., & Zhu, J. (2019). *Explainable AI: A brief survey on history, research areas, approaches and challenges*. In Natural language processing and Chinese computing: 8th CCF international conference, NLPCC 2019, dunhuang, China, October 9–14, 2019, proceedings, part II 8 (pp. 563-574). Springer International Publishing.
- Xu, G., Jia, G., Shi, L., & Zhang, Z. (2021). *Personalized course recommendation system fusing with knowledge graph and collaborative filtering*. Computational Intelligence and Neuroscience, 2021(1), 9590502.
- Xu, H. (2021). *Career decision-making from a dual-process perspective: Looking back, looking forward*. Journal of Vocational Behavior, 126, 103556.
- Yang, L. H., Liu, J., Ye, F. F., Wang, Y. M., Nugent, C., Wang, H., & Martínez, L. (2022). *Highly explainable cumulative belief rule-based system with effective rule-base modeling and inference scheme*. Knowledge-Based Systems, 240, 107805.
- Yang, Y., Zhang, C., Song, X., Dong, Z., Zhu, H., & Li, W. (2023). *Contextualized knowledge graph embedding for explainable talent training course recommendation*. ACM Transactions on Information Systems, 42(2), 1-27.
- Ying, X. (2019, February). *An overview of overfitting and its solutions*. In Journal of Physics: Conference series (Vol. 1168, p. 022022). IOP Publishing.
- Yoakum, C. S., & Yerkes, R. M. (Eds.). (1920). *Army mental tests* (p. 303). New York, NY: H. Holt.
- YouTube. (n.d.). YouTube. Retrieved December 4, 2024, from <https://www.youtube.com/>
- Zakaria, N., & Yusof, H. M. (2021). *Relationship between the 'Big Five' personality and career success*. Journal of Advanced Research in Business and Management Studies, 24(1), 69-76.
- Zangerle, E., & Bauer, C. (2022). *Evaluating recommender systems: survey and framework*. ACM computing surveys, 55(8), 1-38.
- Zhao, X., Huang, W., Huang, X., Robu, V., & Flynn, D. (2021, December). *Baylime: Bayesian local interpretable model-agnostic explanations*. In Uncertainty in artificial intelligence (pp. 887-896). PMLR.

LIST OF PUBLICATIONS & CONFERENCE PROCEEDINGS

1. Rahim MN, Basheer KPM. (2025) A Hierarchical and Multi-Tiered Personalized Career Recommender System Tailored to Individual Aptitudes. *Indian Journal of Science and Technology*. 18(3):231-244. <https://doi.org/10.17485/IJST/v18i3.3979> (WoSS).
2. Rahim, M. N., & Basheer, K. P. M. (2025). A hybrid explainability framework for recommender systems using SHAP with natural language interpretations. *International Journal of Innovative Research and Scientific Studies*, 8(6), 687-703. <https://doi.org/10.53894/ijirss.v8i6.9669> (Scopus, Scimago).
3. Rahim, N. (2025). Principles and methods for transparency, interpretability, trust, accountability, and responsibility in AI-based predictive modelling. *Computer Science, Engineering and Technology*, 3(3), REST Publisher. ISSN: 2583-9179. <https://doi.org/10.46632/cset/3/3/10>.
4. Noushad Rahim M., & Ashika Mumthas P. (2022). An exploration of the possibility of explainable artificial intelligence systems for achieving the right to explanation in AI-based decision support systems. In *Proceedings of the International Conference on Right to Privacy in the Digital Era – Challenges and Solutions* (pp. 85-90). Kerala Law Academy. ISBN 978-93-5777-096-5.

LIST OF CONFERENCES

1. Presented a paper titled A Systematic Approach to Build Career Recommender Systems Based on Computational Intelligence to Solve Data Sparsity and Cold Start Limitations in the *Two-Day International Conference on Artificial Intelligence and Law*, organized by Rajiv Gandhi National Law Cyber Centre, National Law Institute University, Bhopal, 17-18 November 2022.
2. Presented a paper titled Exploring an Explainable Career Recommender System: A Rule based Approach in the *Third National Conference on Computational Intelligence and Data Analytics (CIDA'24)* organized by the

Department of Computer Science, University of Calicut, held on 2024 March 13,14.

3. Presented a paper titled A Hybrid Interpretable Career Recommender System: Ensuring Transparency and Fairness with a SHAP and Rule-Based Explanation Framework in the *Three-Day AMALIT International Conference*, organized by the Department of Computer Science and IQAC, Amal College of Advanced Studies in collaboration with Kerala State Higher Education Council, Govt. of Kerala, on 30th, 31st January and 01st February 2025.

Appendix A: Original Dataset (Sample entries from about 1000 careers collected from the O*NET, <https://www.onetonline.org>)

Title	Element Name	Scale Name	Data Value	N	Standard Error	Lower CI Bound	Upper CI Bound
Chief Executives	Oral Comprehension	Importance	4.5	8	0.19	4.13	4.87
Chief Executives	Oral Comprehension	Level	4.88	8	0.13	4.63	5.12
Chief Executives	Written Comprehension	Importance	4.25	8	0.16	3.93	4.57
Chief Executives	Written Comprehension	Level	4.62	8	0.18	4.27	4.98
Chief Executives	Oral Expression	Importance	4.38	8	0.18	4.02	4.73
Chief Executives	Oral Expression	Level	5	8	0.5		
Chief Executives	Written Expression	Importance	4.12	8	0.13	3.88	4.37
Chief Executives	Written Expression	Level	4.62	8	0.18	4.27	4.98
Chief Executives	Fluency of Ideas	Importance	3.88	8	0.13	3.63	4.12
Chief Executives	Fluency of Ideas	Level	4.62	8	0.18	4.27	4.98
Chief Executives	Originality	Importance	3.88	8	0.13	3.63	4.12
Chief Executives	Originality	Level	4.25	8	0.16	3.93	4.57
Chief Executives	Problem Sensitivity	Importance	4	8	0.4		
Chief Executives	Problem Sensitivity	Level	5	8	0.5		
Chief Executives	Deductive Reasoning	Importance	4.12	8	0.13	3.88	4.37
Chief Executives	Deductive Reasoning	Level	5	8	0.5		
Chief Executives	Inductive Reasoning	Importance	4	8	0.4		
Chief Executives	Inductive Reasoning	Level	5	8	0.5		
Chief Executives	Information Ordering	Importance	3.62	8	0.18	3.27	3.98
Chief Executives	Information Ordering	Level	4	8	0.4		
Chief Executives	Category Flexibility	Importance	3.5	8	0.19	3.13	3.87
Chief Executives	Category Flexibility	Level	4.12	8	0.13	3.88	4.37
Chief Executives	Mathematical Reasoning	Importance	3.38	8	0.18	3.02	3.73
Chief Executives	Mathematical Reasoning	Level	3.88	8	0.23	3.43	4.32
Chief Executives	Number Facility	Importance	3.25	8	0.16	2.93	3.57
Chief Executives	Number Facility	Level	4.12	8	0.13	3.88	4.37
Chief Executives	Memorization	Importance	3	8	0.3		
Chief Executives	Memorization	Level	3.12	8	0.13	2.88	3.37
Chief Executives	Speed of Closure	Importance	3.12	8	0.13	2.88	3.37
Chief Executives	Speed of Closure	Level	3.38	8	0.18	3.02	3.73
Chief Executives	Flexibility of Closure	Importance	3.38	8	0.18	3.02	3.73
Chief Executives	Flexibility of Closure	Level	3.5	8	0.19	3.13	3.87
Chief Executives	Perceptual Speed	Importance	3.12	8	0.13	2.88	3.37
Chief Executives	Perceptual Speed	Level	2.88	8	0.23	2.43	3.32
Chief Executives	Spatial Orientation	Importance	1.12	8	0.13	1	1.37
Chief Executives	Spatial Orientation	Level	0.12	8	0.13	0	0.37
Chief Executives	Visualization	Importance	3.12	8	0.13	2.88	3.37
Chief Executives	Visualization	Level	3.88	8	0.35	3.19	4.56
Chief Executives	Selective Attention	Importance	3	8	0.3		
Chief Executives	Selective Attention	Level	3.12	8	0.13	2.88	3.37
Chief Executives	Time Sharing	Importance	3	8	0.19	2.63	3.37
Chief Executives	Time Sharing	Level	2.88	8	0.23	2.43	3.32
Chief Executives	Arm-Hand Steadiness	Importance	1	8	0.1		
Chief Executives	Arm-Hand Steadiness	Level	0	8	0		
Chief Executives	Manual Dexterity	Importance	1	8	0.1		
Chief Executives	Manual Dexterity	Level	0	8	0		
Chief Executives	Finger Dexterity	Importance	2.25	8	0.16	1.93	2.57
Chief Executives	Finger Dexterity	Level	2	8	0.19	1.63	2.37
Chief Executives	Control Precision	Importance	1.75	8	0.16	1.43	2.07

Video Game Designers,Multilimb Coordination,Importance,1,8,0,1,1
Video Game Designers,Multilimb Coordination,Level,0,8,0,0,0
Video Game Designers,Response Orientation,Importance,1,8,0,1,1
Video Game Designers,Response Orientation,Level,0,8,0,0,0
Video Game Designers,Rate Control,Importance,1,8,0,1,1
Video Game Designers,Rate Control,Level,0,8,0,0,0
Video Game Designers,Reaction Time,Importance,1,8,0,1,1
Video Game Designers,Reaction Time,Level,0,8,0,0,0
Video Game Designers,Wrist-Finger Speed,Importance,1.5,8,0.19,1.13,1.87
Video Game Designers,Wrist-Finger Speed,Level,0.62,8,0.26,0.11,1.14
Video Game Designers,Speed of Limb Movement,Importance,1,8,0,1,1
Video Game Designers,Speed of Limb Movement,Level,0,8,0,0,0
Video Game Designers,Static Strength,Importance,1,8,0,1,1
Video Game Designers,Static Strength,Level,0,8,0,0,0
Video Game Designers,Explosive Strength,Importance,1,8,0,1,1
Video Game Designers,Explosive Strength,Level,0,8,0,0,0
Video Game Designers,Dynamic Strength,Importance,1.12,8,0.13,1,1.37
Video Game Designers,Dynamic Strength,Level,0.12,8,0.13,0,0.37
Video Game Designers,Trunk Strength,Importance,1.62,8,0.18,1.27,1.98
Video Game Designers,Trunk Strength,Level,0.88,8,0.3,0.3,1.45
Video Game Designers,Stamina,Importance,1,8,0,1,1
Video Game Designers,Stamina,Level,0,8,0,0,0
Video Game Designers,Extent Flexibility,Importance,1,8,0,1,1
Video Game Designers,Extent Flexibility,Level,0,8,0,0,0
Video Game Designers,Dynamic Flexibility,Importance,1,8,0,1,1
Video Game Designers,Dynamic Flexibility,Level,0,8,0,0,0
Video Game Designers,Gross Body Coordination,Importance,1,8,0,1,1
Video Game Designers,Gross Body Coordination,Level,0,8,0,0,0
Video Game Designers,Gross Body Equilibrium,Importance,1,8,0,1,1
Video Game Designers,Gross Body Equilibrium,Level,0,8,0,0,0
Video Game Designers,Near Vision,Importance,4,8,0,4,4
Video Game Designers,Near Vision,Level,4,8,0,4,4
Video Game Designers,Far Vision,Importance,2.88,8,0.13,2.63,3.12
Video Game Designers,Far Vision,Level,3,8,0,3,3
Video Game Designers,Visual Color Discrimination,Importance,2.88,8,0.13,2.63,3.12
Video Game Designers,Visual Color Discrimination,Level,3.12,8,0.13,2.88,3.37
Video Game Designers,Night Vision,Importance,1,8,0,1,1
Video Game Designers,Night Vision,Level,0,8,0,0,0
Video Game Designers,Peripheral Vision,Importance,1,8,0,1,1
Video Game Designers,Peripheral Vision,Level,0,8,0,0,0
Video Game Designers,Depth Perception,Importance,1.62,8,0.18,1.27,1.98
Video Game Designers,Depth Perception,Level,0.62,8,0.18,0.27,0.98
Video Game Designers,Glare Sensitivity,Importance,1,8,0,1,1
Video Game Designers,Glare Sensitivity,Level,0,8,0,0,0
Video Game Designers,Hearing Sensitivity,Importance,2.12,8,0.13,1.88,2.37
Video Game Designers,Hearing Sensitivity,Level,2.12,8,0.13,1.88,2.37
Video Game Designers,Auditory Attention,Importance,2,8,0,2,2
Video Game Designers,Auditory Attention,Level,1.12,8,0.13,0.88,1.37
Video Game Designers,Sound Localization,Importance,1,8,0,1,1
Video Game Designers,Sound Localization,Level,0,8,0,0,0
Video Game Designers,Speech Recognition,Importance,3.25,8,0.16,2.93,3.57

Mechanical Engineers, Response Orientation, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Response Orientation, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Rate Control, Importance, 1.38, 8, 0.18, 1.02, 1.73
 Mechanical Engineers, Rate Control, Level, 0.38, 8, 0.18, 0.02, 0.73
 Mechanical Engineers, Reaction Time, Importance, 2.12, 8, 0.13, 1.88, 2.37
 Mechanical Engineers, Reaction Time, Level, 2.12, 8, 0.13, 1.88, 2.37
 Mechanical Engineers, Wrist-Finger Speed, Importance, 1.5, 8, 0.19, 1.13, 1.87
 Mechanical Engineers, Wrist-Finger Speed, Level, 0.62, 8, 0.26, 0.11, 1.14
 Mechanical Engineers, Speed of Limb Movement, Importance, 1.12, 8, 0.13, 1, 1.37
 Mechanical Engineers, Speed of Limb Movement, Level, 0.12, 8, 0.13, 0, 0.37
 Mechanical Engineers, Static Strength, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Static Strength, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Explosive Strength, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Explosive Strength, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Dynamic Strength, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Dynamic Strength, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Trunk Strength, Importance, 1.5, 8, 0.19, 1.13, 1.87
 Mechanical Engineers, Trunk Strength, Level, 1, 8, 0.42, 0.17, 1.83
 Mechanical Engineers, Stamina, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Stamina, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Extent Flexibility, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Extent Flexibility, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Dynamic Flexibility, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Dynamic Flexibility, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Gross Body Coordination, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Gross Body Coordination, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Gross Body Equilibrium, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Gross Body Equilibrium, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Near Vision, Importance, 3.75, 8, 0.16, 3.43, 4.07
 Mechanical Engineers, Near Vision, Level, 4, 8, 0.19, 3.63, 4.37
 Mechanical Engineers, Far Vision, Importance, 2.88, 8, 0.13, 2.63, 3.12
 Mechanical Engineers, Far Vision, Level, 3.25, 8, 0.25, 2.76, 3.74
 Mechanical Engineers, Visual Color Discrimination, Importance, 2.75, 8, 0.16, 2.43, 3.07
 Mechanical Engineers, Visual Color Discrimination, Level, 2.88, 8, 0.23, 2.43, 3.32
 Mechanical Engineers, Night Vision, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Night Vision, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Peripheral Vision, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Peripheral Vision, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Depth Perception, Importance, 2, 8, 0, 2, 2
 Mechanical Engineers, Depth Perception, Level, 2.38, 8, 0.18, 2.02, 2.73
 Mechanical Engineers, Glare Sensitivity, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Glare Sensitivity, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Hearing Sensitivity, Importance, 2.5, 8, 0.19, 2.13, 2.87
 Mechanical Engineers, Hearing Sensitivity, Level, 2.75, 8, 0.25, 2.26, 3.24
 Mechanical Engineers, Auditory Attention, Importance, 2.5, 8, 0.19, 2.13, 2.87
 Mechanical Engineers, Auditory Attention, Level, 2.62, 8, 0.18, 2.27, 2.98
 Mechanical Engineers, Sound Localization, Importance, 1, 8, 0, 1, 1
 Mechanical Engineers, Sound Localization, Level, 0, 8, 0, 0, 0
 Mechanical Engineers, Speech Recognition, Importance, 3.25, 8, 0.16, 2.93, 3.57
 Mechanical Engineers, Speech Recognition, Level, 3.75, 8, 0.16, 3.43, 4.07
 Mechanical Engineers, Speech Clarity, Importance, 3.25, 8, 0.16, 2.93, 3.57

Appendix B: Generated DAT Dataset (Samples from selected 612 careers) -


Published: <https://doi.org/10.5281/zenodo.14192027>


Career	NA	AR	VR	MR	SR	VA
Chief Executives	58	73	77	30	55	72
Chief Sustainability Officers	45	68	70	20	43	71
General and Operations Managers	44	54	67	20	40	67
Advertising and Promotions Managers	47	64	68	16	51	65
Marketing Managers	53	66	68	19	47	66
Sales Managers	54	65	69	19	36	65
Administrative Services Managers	41	56	67	34	44	65
Computer and Information Systems Managers	48	64	70	29	47	64
Financial Managers	61	64	72	18	30	66
Treasurers and Controllers	67	66	69	29	36	67
Investment Fund Managers	66	65	69	30	42	66
Industrial Production Managers	52	65	68	37	50	60
Quality Control Systems Managers	50	64	68	33	44	67
Geothermal Production Managers	54	60	67	44	47	60
Biofuels Production Managers	49	58	70	34	44	62
Biomass Power Plant Managers	49	61	68	39	45	63
Hydroelectric Production Managers	53	60	71	38	53	63
Purchasing Managers	53	63	67	14	36	64
Transportation Storage and Distribution Managers	48	59	67	26	38	65
Supply Chain Managers	50	60	68	21	41	67
Compensation and Benefits Managers	49	52	67	19	25	67
Human Resources Managers	50	64	71	26	36	67
Training and Development Managers	45	67	68	27	43	71
Farmers Ranchers and Other Agricultural Managers	44	58	64	47	45	60
Construction Managers	53	61	66	20	66	64
Education and Childcare Administrators Preschool and Daycare	44	63	67	25	36	70
Education Administrators Kindergarten through Secondary	48	68	76	20	42	74
Education Administrators Postsecondary	46	65	73	16	37	73
Architectural and Engineering Managers	63	66	74	31	65	69
Biofuels/Biodiesel Technology and Product Development Managers	57	64	68	45	47	66
Food Service Managers	46	51	64	46	36	56
Medical and Health Services Managers	48	63	74	10	30	67


Clinical Data Managers	53	62	68	23	31	66
Bioinformatics Technicians	57	60	68	57	48	63
Architects	58	70	71	33	73	67
Landscape Architects	46	64	67	31	67	63
Cartographers and Photogrammetrists	47	56	67	45	53	59
Surveyors	67	58	67	48	54	63
Geodetic Surveyors	70	61	73	42	54	67
Aerospace Engineers	64	67	74	30	53	70
Agricultural Engineers	59	65	72	35	59	68
Bioengineers and Biomedical Engineers	63	69	74	40	60	69
Chemical Engineers	66	69	75	26	55	58
Civil Engineers	67	68	74	26	64	67
Transportation Engineers	67	69	72	30	67	67
Water/Wastewater Engineers	66	69	72	17	63	67
Computer Hardware Engineers	51	64	71	39	56	64
Electrical Engineers	58	62	71	36	50	66
Electronics Engineers Except Computer	59	65	69	49	60	66
Radio Frequency Identification Device Specialists	47	62	66	35	46	57
Environmental Engineers	58	67	72	16	51	72
Health and Safety Engineers Except Mining Safety Engineers and Ins	51	67	72	25	54	68
Fire-Prevention and Protection Engineers	47	63	66	28	54	65
Industrial Engineers	52	63	69	27	55	68
Human Factors Engineers and Ergonomists	57	69	74	19	56	71
Validation Engineers	54	63	70	28	44	69
Manufacturing Engineers	68	69	72	37	67	65
Marine Engineers and Naval Architects	65	67	73	35	67	67
Materials Engineers	61	67	73	38	56	71
Mechanical Engineers	67	65	73	62	59	58
Fuel Cell Engineers	61	66	73	45	56	70
Automotive Engineers	67	70	72	27	64	65
Mining and Geological Engineers Including Mining Safety Engineers	63	69	74	20	59	71
Nuclear Engineers	61	65	75	20	48	73
Petroleum Engineers	53	66	69	25	52	71


Astronomers	71	70	75	26	49	74
Physicists	85	78	81	32	64	78
Atmospheric and Space Scientists	52	63	77	22	40	72
Chemists	64	65	74	53	51	67
Materials Scientists	55	66	75	31	55	73
Environmental Scientists and Specialists Including Health	56	65	75	35	42	70
Climate Change Policy Analysts	47	61	70	13	39	67
Environmental Restoration Planners	54	65	72	31	51	69
Industrial Ecologists	57	68	70	0	43	73
Geoscientists Except Hydrologists and Geographers	60	67	75	27	49	69
Hydrologists	63	67	74	32	51	71
Remote Sensing Scientists and Technologists	64	67	73	39	49	69
Economists	65	68	74	13	34	74
Environmental Economists	70	63	75	14	44	72
Survey Researchers	67	63	69	24	30	70
Industrial-Organizational Psychologists	57	68	76	27	44	76
Clinical and Counseling Psychologists	35	66	77	3	29	71
School Psychologists	45	63	76	24	35	72
Sociologists	40	63	75	10	27	74
Urban and Regional Planners	50	65	70	24	54	69
Anthropologists and Archeologists	44	71	78	39	54	81
Geographers	50	64	73	16	45	71
Historians	21	58	77	19	26	76
Political Scientists	40	62	78	10	29	73
Transportation Planners	52	65	67	22	60	65
Agricultural Technicians	49	52	63	64	41	59
Precision Agriculture Technicians	51	61	67	63	53	64
Food Science Technicians	48	54	63	61	41	62
Biological Technicians	49	59	66	64	48	61
Chemical Technicians	49	58	66	67	49	65
Environmental Science and Protection Technicians Including Health	52	57	67	64	45	63
Geological Technicians Except Hydrologic Technicians	45	51	66	61	45	56
Nuclear Technicians	50	59	63	66	48	48

Appendix C: Generated Top-Level 8 Career Clusters













Appendix D: A sample from Hierarchical Clustering by sub-stratification of top-level clusters




Appendix E: Augmented Samples from Sub-Stratified Career Clusters









Appendix F: Program Code Snippets

Python code to generate the DAT dataset from ONET ability profiles for feature Numerical Reasoning (same procedure applied to other features)

```
def standardizeImpScore(originalValue):
    #S = ( (O - L) / (H - L) ) * 100
    stdImpScore = (( originalValue-1)/(5-1))*100
    return stdImpScore

def standardizeLevelScore(originalValue):
    #S = ( (O - L) / (H - L) ) * 100
    stdLevelScore = (( originalValue-0)/(7-0))*100
    return stdLevelScore

def getImpScore(career, elementName):
    eleImpResultset = dfabilities[(dfabilities['Title']==career ) & \
    (dfabilities["Element Name"]==elementName) & \
    (dfabilities["Scale Name"]=="Importance")]

    if eleImpResultset.empty:
        return
    print("Here .....",eleImpResultset)
    print(eleImpResultset.iloc[0,2], eleImpResultset.iloc[0,3])
    eleImpValue=(eleImpResultset.iloc[0,3])
    print(elementName, eleImpValue)

    stdEleImpScore=standardizeImpScore(eleImpValue)
    return stdEleImpScore

def getLevelScore(career, elementName):
    eleLevelResultset = dfabilities[(dfabilities['Title']==career ) & \
    (dfabilities["Element Name"]==elementName) & \
    (dfabilities["Scale Name"]=="Level")]

    if eleLevelResultset.empty:
        return
    print(eleLevelResultset.iloc[0,2], eleLevelResultset.iloc[0,3])
    eleLevelValue=(eleLevelResultset.iloc[0,3])
    print(elementName, eleLevelValue)

    stdEleLevelScore=standardizeLevelScore(eleLevelValue)
    return stdEleLevelScore

def getNR(career):
    mrImpValue=getImpScore(career,"Mathematical Reasoning")
    mrLevelValue=getLevelScore(career,"Mathematical Reasoning")

    if mrImpValue is None:
        return
    if mrLevelValue is None:
        return

    avgMrScore= math.ceil((mrImpValue+ mrLevelValue)/2)
    print("Average Math Reasoning: ", avgMrScore)

    nfImpValue=getImpScore(career, "Number Facility")
    nfLevelValue=getLevelScore(career, "Number Facility");
    avgNfScore=math.ceil((nfImpValue+nfLevelValue)/2)

    print("Average Numerical Facility: ", avgNfScore)

    datNRScore= math.ceil((avgMrScore+avgNfScore)/2)
    print("DAT NA Score:", datNRScore)

    return datNRScore
```

Code: Evaluation of Clustering Algorithms via Silhouette score

```
from sklearn.preprocessing import StandardScaler as SS
from sklearn.cluster import KMeans
from sklearn.mixture import GaussianMixture
from sklearn.cluster import Birch
from sklearn.cluster import AgglomerativeClustering
from sklearn.cluster import DBSCAN

scaler = SS()
DNP_dfdat_standardized = scaler.fit_transform(odfdat)
df_dfdat_standardized = pd.DataFrame(DNP_dfdat_standardized, columns=["NR", "AR", "VR", "MR", "SR", "VA"])
df_dfdat_standardized = df_dfdat_standardized.set_index(odfdat.index)
print(".....odfdat.....", odfdat.head())
print(df_dfdat_standardized.head())

algorithms = {
    'AgglomerativeClustering': AgglomerativeClustering(n_clusters=8, metric='euclidean',
    linkage='ward'),
    'GMM': GaussianMixture(n_components=8, random_state=0),
    'BIRCH': Birch(threshold=0.03, n_clusters=8),
    'KMEANS': KMeans(n_clusters=8, random_state=0)
}

for algo_name, algo in algorithms.items():
    if algo_name == 'GMM':
        labels = algo.fit_predict(df_dfdat_standardized)
    else:
        labels = algo.fit_predict(df_dfdat_standardized)
    silhouette_avg = silhouette_score(df_dfdat_standardized, labels)
    print(f'Silhouette Score for {algo_name}: {silhouette_avg:.4f}')
```

Python code to generate clusters (Code for KMeans clustering only provided here)

```

from sklearn.preprocessing import StandardScaler as SS
from sklearn.cluster import KMeans

odf=pd.read_csv('DAT_Career_dataset.csv')
career_values = odf['Career'].values
odfdat=odf.drop(columns=["Career"])

row_dict=odf.iloc[:,0:7]

#.....

scaler = SS()
DNP_dfdat_standardized = scaler.fit_transform(odfdat)
df_dfdat_standardized = pd.DataFrame(DNP_dfdat_standardized, columns=["NR", "AR", "VR", "MR", "SR", "VA"])
df_dfdat_standardized = df_dfdat_standardized.set_index(odfdat.index)


algorithm="KMeans"
kmeans_model = KMeans(n_clusters=8, random_state=0)
kmeans_model.fit(df_dfdat_standardized)
labels = kmeans_model.fit_predict(df_dfdat_standardized)
#.....
clusters = {}
n = 0
for item in labels:
    if item in clusters:
        clusters[item].append(row_dict.iloc[n].tolist())
    else:
        clusters[item] = [row_dict.iloc[n].tolist()]
    n +=1

#=====
for item in clusters:
    clusterFile=open("GeneratedClusters/"+algorithm+"/"+"careerClusters"+str(item+1)+".csv", "w")
    #clusterFile.write("Career,NR,AR,VR,MR,SR,VA,cluster")
    clusterFile.write("Career,NR,AR,VR,MR,SR,VA")
    for i in clusters[item]: #printing each element of the cluster from the dictionary
        clusterFile.write("\n")
        clusterFile.write(str(i) [1:-1]) # [1:-1] to remove square bracket from list values
        #clusterFile.write(str(i)+" "+str(item))
    clusterFile.close();

```

Soft-Voting Classifier Implementation

```
def buildVotingClassifierModel():
    # Instantiate your classifiers (SVM, KNN, Random Forest)
    clfDNN = MLPClassifier(solver='lbfgs', activation='relu',
                           alpha=1e-5,
                           hidden_layer_sizes=(100,60,44),
                           max_iter=1000,
                           random_state=1)

    clfKNN = KNeighborsClassifier(n_neighbors=5, leaf_size=36,p=1, weights='uniform', metric='manhattan')
    clfSVM = svm.SVC(kernel='rbf', probability=True)
    clfRF = RandomForestClassifier(random_state=1)
    clfDT=DecisionTreeClassifier(random_state=10)
    # Create a VotingClassifier with the individual classifiers
    voting_classifier = VotingClassifier(estimators=[
        ('DNNclf', clfDNN),
        ('KNNclf', clfKNN),
        ('RFclf', clfRF),
        ('SVMclf', clfSVM),
        ('DTclf', clfDT),
        #('GNBclf', clfGNB)
    ], voting='soft') # Use 'soft' for probability-based voting
```

Code snippets: Hierarchical and Multi-Tiered Recommendation Engine

```

#.....MAIN PROGRAM - Get Diverse Recommendations from Top-Level Career Clusters.

loadedBaseVC=pickle.load(open('TopLevelModels/myBaseVotingClassifierModel.sav','rb'))
loadedBaseVCScaler=pickle.load(open('TopLevelModels/scalerBaseVotingClassifierScaler.sav','rb'))

record=[[61, 64, 72, 18, 30, 66]] #candidate profile
originalRecord=record

inputRec = loadedBaseVCScaler.transform(record)
vc_proba = loadedBaseVC.predict_proba(inputRec)

record=record[0]
recordddf=pd.DataFrame(pd.Series(record))

firstKey, secondKey, thirdKey=getDiverseRecommendations("VC", vc_proba, recordddf)

finetuneRecommendations(firstKey,secondKey,thirdKey, originalRecord)

def finetuneRecommendations(firstKey,secondKey,thirdKey,record):
    #.....load sub level models

    sub1=pickle.load(open('SubModels/Model1/myVotingClassifierModel1.sav','rb'))
    scaler1=pickle.load(open('SubModels/Model1/scalerVotingClassifierScaler1.sav','rb'))
    sub1Df=pd.read_csv('SubModels/Model1/careerClusters1combined_clusters.csv')

    sub2=pickle.load(open('SubModels/Model2/myVotingClassifierModel2.sav','rb'))
    scaler2=pickle.load(open('SubModels/Model2/scalerVotingClassifierScaler2.sav','rb'))
    sub2Df=pd.read_csv('SubModels/Model2/careerClusters2combined_clusters.csv')

    sub3=pickle.load(open('SubModels/Model3/myVotingClassifierModel3.sav','rb'))
    scaler3=pickle.load(open('SubModels/Model3/scalerVotingClassifierScaler3.sav','rb'))
    sub3Df=pd.read_csv('SubModels/Model3/careerClusters3combined_clusters.csv')

    sub4=pickle.load(open('SubModels/Model4/myVotingClassifierModel4.sav','rb'))
    scaler4=pickle.load(open('SubModels/Model4/scalerVotingClassifierScaler4.sav','rb'))
    sub4Df=pd.read_csv('SubModels/Model4/careerClusters4combined_clusters.csv')

    sub5=pickle.load(open('SubModels/Model5/myVotingClassifierModel5.sav','rb'))
    scaler5=pickle.load(open('SubModels/Model5/scalerVotingClassifierScaler5.sav','rb'))
    sub5Df=pd.read_csv('SubModels/Model5/careerClusters5combined_clusters.csv')

    sub6=pickle.load(open('SubModels/Model6/myVotingClassifierModel6.sav','rb'))
    scaler6=pickle.load(open('SubModels/Model6/scalerVotingClassifierScaler6.sav','rb'))
    sub6Df=pd.read_csv('SubModels/Model6/careerClusters6combined_clusters.csv')

    sub7=pickle.load(open('SubModels/Model7/myVotingClassifierModel7.sav','rb'))
    scaler7=pickle.load(open('SubModels/Model7/scalerVotingClassifierScaler7.sav','rb'))
    sub7Df=pd.read_csv('SubModels/Model7/careerClusters7combined_clusters.csv')

    sub8=pickle.load(open('SubModels/Model8/myVotingClassifierModel8.sav','rb'))
    scaler8=pickle.load(open('SubModels/Model8/scalerVotingClassifierScaler8.sav','rb'))
    sub8Df=pd.read_csv('SubModels/Model8/careerClusters8combined_clusters.csv')

```

```
#select first best fined tune model
if firstKey== 1:
    firstModel=sub1
    firstScaler=scaler1
    firstDf=sub1Df
elif firstKey== 2:
    firstModel=sub2
    firstScaler=scaler2
    firstDf=sub2Df
elif firstKey== 3:
    firstModel=sub3
    firstScaler=scaler3
    firstDf=sub3Df
elif firstKey== 4:
    firstModel=sub4
    firstScaler=scaler4
    firstDf=sub4Df
elif firstKey== 5:
    firstModel=sub5
    firstScaler=scaler5
    firstDf=sub5Df
elif firstKey== 6:
    firstModel=sub6
    firstScaler=scaler6
    firstDf=sub6Df
elif firstKey== 7:
    firstModel=sub7
    firstScaler=scaler7
    firstDf=sub7Df
elif firstKey== 8:
    firstModel=sub8
    firstScaler=scaler8
    firstDf=sub8Df

#... similarly select secondKey and thirdKey as second & third best fine tuned Models

#.....1. get best 3 predictions from first model
inputRec = firstScaler.transform(record)
vc_proba = firstModel.predict_proba(inputRec)

recordCopy=record[0]
recorddf=pd.DataFrame(pd.Series(recordCopy))

predictBestThreeCareerClusters("", vc_proba, recorddf, firstDf)

#.....2. get best 3 predictions from second model
inputRec = secondScaler.transform(record)
vc_proba = secondModel.predict_proba(inputRec)

recordCopy=record[0]
recorddf=pd.DataFrame(pd.Series(recordCopy))

predictBestThreeCareerClusters("", vc_proba, recorddf, secondDf)

#.....3. get best 3 predictions from third model
inputRec = thirdScaler.transform(record)
vc_proba = thirdModel.predict_proba(inputRec)

recordCopy=record[0]
recorddf=pd.DataFrame(pd.Series(recordCopy))

predictBestThreeCareerClusters("", vc_proba, recorddf, thirdDf)
#.....
```

Code snippets: SHAP-Guided hybrid explanation framework

```

import numpy as np
import pandas as pd
import shap
import joblib
from sklearn.preprocessing import StandardScaler
from sklearn.pipeline import Pipeline
from scipy.stats import pearsonr

# Creating the SHAP Explainer, model is the recommendation engine
# =====
explainer = shap.Explainer(model.predict_proba, xtrain,
                           feature_names=feature_names)
shap_values = explainer(xtrain)

#class specific summary plots
# =====
shap.summary_plot(shap_values[:, :, class_index], features=xtrain, feature_names=feature_names)
shap.summary_plot(shap_values[:, :, class_index], features=xtrain, feature_names=feature_names, plot_type='bar')

#Instance level explanation via SHAP plots
# =====
shap.initjs()
f_names=['NR', 'AR', 'VR', 'MR', 'SR', 'VA']

predicted_proba = clf.predict_proba(input)

predicted_class_index = np.argmax(predicted_proba, axis=1)[0]
max_proba_values = np.max(predicted_proba, axis=1)
class_index=predicted_class_index

shap_values_input = explainer(input)
class_specific_base_value = explainer.expected_value[class_index]
selected_class_shap_values = specific_class_shap_values.values

fp=shap.force_plot(class_specific_base_value,selected_class_shap_values, features=input, feature_names=f_names)
display(fp)
dp=shap.decision_plot(class_specific_base_value,selected_class_shap_values, features=input, feature_names=f_names)
display(dp)
shap_explanation = shap.Explanation(
    values=specific_class_shap_values.values,
    base_values=class_specific_base_value,
    data=input,
    feature_names=f_names
)
print("Shap_explanation",shap_explanation)
wf=shap.plots.waterfall(shap_explanation[0])
display(wf)

```

```

#class-level(item-level) global shapley value based feature importance analysis.
#=====
def get_cluster_feature_importance_directly_proportional(
    shap_values_cluster, xtrain_cluster, feature_names, top_n=10
):
    assert shap_values_cluster.shape == xtrain_cluster.shape, \
        "Mismatch between SHAP and input values"

    correlations = []
    for i in range(shap_values_cluster.shape[1]):
        shap_vals = shap_values_cluster[:, i]
        feature_vals = xtrain_cluster[:, i]

        if np.std(feature_vals) == 0 or np.std(shap_vals) == 0:
            corr = 0
        else:
            corr, _ = pearsonr(feature_vals, shap_vals)

        correlations.append(corr if corr > 0 else 0)

    importance_df = pd.DataFrame({
        'Feature': feature_names,
        'Correlation(SHAP vs Value)': correlations
    })

    return importance_df.sort_values(
        by='Correlation(SHAP vs Value)', ascending=False
    ).head(top_n).reset_index(drop=True)

xtrain_np = xtrain # Convert DataFrame to NumPy array if needed
print("Xtrain shape", xtrain_np.shape)
print("shap shape", shap_values[:, :, class_index].values.shape)
feature_names = ['NR', 'AR', 'VR', 'MR', 'SR', 'VA']
feature_importance_df = get_cluster_feature_importance_directly_proportional(
    shap_values_cluster=shap_values[:, :, groupIndex].values,
    xtrain_cluster=xtrain,
    feature_names=feature_names,
    top_n=6
)
print(feature_importance_df)

# Item-level global explanation:item-level feature description
#=====
def describe_cluster_from_feature_importance(importance_df, cluster_name="this cluster"):
    descriptions = []
    for idx, row in importance_df.iterrows():
        feature = row['Feature']
        corr = row['Correlation(SHAP vs Value)']
        if corr >= 0.5:
            desc = f"Feature {feature} has very strong significance as far as {cluster_name} is concerned."
        elif corr >= 0.3:
            desc = f"Feature {feature} has high importance as far as {cluster_name} is concerned."
        elif corr > 0:
            desc = f"Feature {feature} has a moderately high importance as far as {cluster_name} is concerned."
        else:
            continue # Skip zero correlations (already filtered by your function)
        descriptions.append(desc)

    if not descriptions:
        return f"No features show a significant directly proportional relationship in {cluster_name}."

    header = f"### Cluster Profile Description for {cluster_name}\n\n"
    body = "\n".join(descriptions)
    return header + body

description = describe_cluster_from_feature_importance(feature_importance_df)
print(description)

```

```
#Local Explanation:prediction explanation for a given instance
#=====
def generate_proportional_explanation(shap_values, feature_names, instance_values, predicted_class_idx):
    feature_impact = shap_values[0][:, predicted_class_idx] # instance specific SHAP values for the prediction

    # Compute absolute SHAP values and their percentage contributions
    abs_shap_values = np.abs(feature_impact)
    total_impact = np.sum(abs_shap_values)
    percentages = (abs_shap_values / total_impact) * 100 if total_impact != 0 else np.zeros_like(abs_shap_values)

    # Sort features by contribution percentage
    sorted_indices = np.argsort(-percentages)

    print(f"The recommendations are based on the following factor:")
    explanation = []
    for i in sorted_indices:
        feature = feature_names[i]
        value = instance_values[i]
        impact = feature_impact[i]
        percentage = percentages[i]

        #print("ALL", feature,value,impact,percentage)
        if impact > 0 and value > 0 : # if impact (shapely value) increases, as feature value increases
            explanation.append(f"\n The high {feature} (value: {value}) (impact:{impact}) contributed \
            {percentage:.1f}%, increasing the likelihood of recommending this career cluster.")
        elif impact > 0 and value < 0 :
            explanation.append(f"\n The low {feature} (value: {value}) (impact: {impact}) contributed \
            {percentage:.1f}%, increasing the likelihood of recommending this career cluster.")
        else :
            explanation.append(f"\n The {feature} (value: {value}) (has low impact: {impact}) and contributed \
            {percentage:.1f}%, decreasing the likelihood of recommending this career cluster.")
    explanationSorted=sorted(explanation)
    return " ".join(explanationSorted)

#Fusion Layer: Final Explanation Generation
#=====

from transformers import AutoModelForCausalLM, AutoTokenizer
import torch
#recommendation_explanation: explanation for instance specific recommendation
#career_cluster_description: Recommended Item-level description
def finalExplanation(recommendation_explanation, career_cluster_description):

    model_name = "meta-llama/Llama-2-7b-chat-hf" # # using LLaMA 2 model and tokenizer
    tokenizer = AutoTokenizer.from_pretrained(model_name)
    model = AutoModelForCausalLM.from_pretrained(model_name)

    prompt = f"""
    Analyze the following student's profile and compare it with the career description.
    Identify the matching skills and their strengths.

    Why the student recommended this career (student profile): {recommendation_explanation}
    What the career expects as per its requirements (career description): {career_cluster_description}

    Note that NR stands for Numerical Reasoning, AR stands for Abstract Reasoning, VR stands for Verbal Reasoning,
    MR stands for Mechanical Reasoning, SR stands for Spatial Relations, and VA stands for Verbal Ability.
    Explain why this career is suitable (or not) for the student, specifying matching skills and strengths. Do not hallucinate.
    """

    inputs = tokenizer(prompt, return_tensors="pt")
    outputs = model.generate(inputs, max_new_tokens=300)
    explanation = tokenizer.decode(outputs[0], skip_special_tokens=True)

    return explanation
```