BINDURA UNIVERSITY OF SCIENCE EDUCATION

FACULTY OF SCIENCE EDUCATION



Honours Degree in Computer Science

An AI Driven Framework For Personalised Career and Academic Guidance For Pre-University Students in Zimbabwe

FINAL YEAR PROJECT

BY

CLIGHTON MARUZA

B220972B

SUPERVISOR : MR H CHIKWIRIRO

JUNE 2026

Approval Form

The undersigned certify that they have supervised the student **CLIGHTON MARUZA** dissertation entitled **An AI Driven Framework For Personalised Career and Academic Guidance For Pre-University Students in Zimbabwe** submitted in partial fulfilment of the requirements for the Bachelor of Science Honours Degree in Computer Science at Bindura University of Science Education.

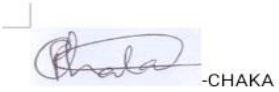
STUDENT NAME

Clighton Maruza  08-06-2026

SUPERVISORS

Mr H Chikwiriro 

CHAIRPERSON



Mr P Chaka

Dedication

I dedicate this work first and foremost to the Almighty God for his guidance, wisdom and strength which guided me through the whole research process.

I dedicate also to my parents and family for their faith, constant support, prayers, motivation and encouragement throughout my academic journey. The sacrifice they made is just invaluable and I appreciate them for that.

Acknowledgements

I extend my sincere thanks and appreciation to all individuals who made this research work a success.

I want to extend special thanks to my supervisor Mr Chikwiriro for his guidance, feedback, patience, support throughout the research work and for sharing his expertise to help me make this research better.

I would also like to convey my sincere thanks to my friends for their motivation and moral support when times were difficult during this research journey. They gave me support.

Abstract

This research paper introduces the development of an AI-Driven Framework for Personalized Career and Academic Guidance for Pre-University Students in Zimbabwe. The system seeks to solve the challenges pre-university students encounter in making appropriate decisions about their career choice due to limitations such as lack of organized guidance, personal career support and relevant career information exposure. The system utilizes Artificial Intelligence techniques (Data-Driven Recommendation methods, RAG etc.) to enable the system to generate customized career suggestions according to students interests, academic performance and abilities. The aim of the framework is to provide effective, personalized and readily available career guidance. The study employed systematic methods in the development phase: Requirements gathering, system design and implementation and testing. The system was tested based on usage and effectiveness in its recommendation ability and the study concluded that AI can assist greatly in providing appropriate career guidance to students, and further suggestions are included to enable further enhancements.

Table Of Contents

Contents

Dedication 3

Acknowledgements 3

Abstract 4

CHAPTER 1: PROBLEM IDENTIFICATION 7

1.1 Introduction 7

1.2 Background Of The Study 7

1.3 Statement Of The Problem 8

1.4 Research Objectives 8

1.5 Research Questions 9

1.6 Research Propositions/Hypothesis 9

1.7 Justification Of The Study 9

CHAPTER 2: LITERATURE REVIEW 10

2.1 Introduction 10

2.2 Relevant Theory Of The Subject Matter 11

2.3 Conclusion 12

CHAPTER 3: RESEARCH METHODOLOGY AND SYSTEM DESIGN 13

3.1 Introduction 13

3.2 Research Design 13

3.3 Data Collection Methods 14

3.3.1 Primary Data Collection 14

3.3.2 Secondary Data Collection 14

3.4 Population and Sample 15

3.5 Functional and Non-Functional Requirements 15

3.5.1 Functional Requirements 15

3.5.2 Non-Functional Requirements 17

3.6 Hardware Development Tools 18

3.7 Software Development Tools 19

3.8.1 Database Normalization (3NF) 21

3.8.2 The Hybrid AI Recommendation Framework 22

3.9 Model Data Set and Data Cleaning 23

3.10 Feature Engineering for the AI Recommendation Model 24
3.11 Model Training 26
3.12 Data Flow Diagram (DFD) 26
3.13 Entity Relationship Diagram (ERD) 28
3.14 System Architecture And Technology Stack Implementation 29
3.14.1 Presentation Layer 30
3.14.2 Application Layer 30
3.14.3 Data Layer 30
3.14.4 AI / LLM Layer 30
3.15 Chat Advisor System Prompt Design 31
3.15.1 Advice Generation Prompt. 32
3.15.2 Chat Advisor Prompt. 32
3.16 Summary of How the System Works 32
3.17 Chapter Summary 34
CHAPTER 4: SYSTEM IMPLEMENTATION AND TESTING 35
4.1 White-Box Testing 35
4.2 Black-Box Testing 36
4.3 Performance Testing 38
4.4 Confusion Matrix 39
4.5 Precision and Recall 40
4.5.1 Precision 40
4.5.2 Recall 41
4.5.3 F1 Score 41
CHAPTER 5: CONCLUSIONS AND RECOMMENDATIONS 43
5.1 Introduction 43
5.2 Major Conclusions Drawn 43
5.3 Recommendations 44
REFERENCES 45

CHAPTER 1: PROBLEM IDENTIFICATION

1.1 Introduction

The selection of a particular course of study and career path is among the critical decisions that students need to make as they are about to enter a college or university, a decision that influences their course of study and future career (Chireshe 2013). However, most adolescents struggle to choose appropriate careers because there are gaps in access to career information, competent personnel for guidance and inability to self-evaluate on strengths and weaknesses. For instance, in many underdeveloped countries like Zimbabwe, schools lack developed career guidance programs hence students rely on other people such as parents and friends in making subject choices at the tertiary level (Chireshe 2013).

According to the Zimbabwean curriculum framework, a student in Advanced Level should have appropriate career choice in the combination of subjects he or she intends to study however, many students are made to make such decisions without being properly guided and having access to adequate careers guidance systems. This leads some learners to pursue career choices which are inappropriate for them based on their interests, academic performance and other characteristics. As a consequence some students' academic motivation may be lower, which leads to poorer academic performance and career satisfaction later on. The approaches of career counselling such as sporadic individual sessions or consultation with teachers may fail to satisfy students who differ in individual needs.

The application of AI has led to the belief that the process of career guidance may be better enhanced with this kind of technology. AI can help in predicting career trajectories based on students' academic performance, individual interests, personality characteristics and behaviours. Such advanced systems have the potential to deliver a more suitable response for individual student needs compared to current methods. An AI-based career guidance system in the Zimbabwean context may help students in decision-making, providing relevant and personalized recommendations about subject combination or university courses to pursue based on the students' interests and skills.

1.2 Background Of The Study

Career guidance plays a pivotal role in the education system in making students aware of their skills, interests and various career choices. A key choice in life can be to decide what course to study at university, an issue that significantly impacts future careers. Thus it is important that learners are equipped with the appropriate guidance needed for them to relate their skills, interests and academic abilities to suitable career paths.

In Zimbabwe, the selection of Advanced Level subject combination is supposed to guide students to pursue appropriate careers later on, but a large number of students don't make such decisions with appropriate guidance and information about various career choices. In this situation some students select their preferred subjects based on peer pressure, expectations of their parents, and poor knowledge about what subject options are available, which can eventually result in the students obtaining careers which are not aligned with their talents and interests.

In the context of growing advancements in digital technology, Artificial Intelligence (AI) has emerged as an applicable tool to enhance decision-making in career guidance by processing relevant information and recommending suitable options to learners. The ability of AI to analyse data of student performance, personal interests and behaviour patterns will enable the provision of personalized career guidance recommendations to pre-university students about appropriate academic and career pathways.

1.3 Statement Of The Problem

Students find making career choices one of the important and complex decisions they have to make prior to pursuing a degree at university. Advanced Level subject combinations in Zimbabwe are meant to influence students' future course of study and career; nevertheless most students commencing tertiary education find themselves struggling with such decisions as they haven't been given sufficient career guidance and/or information on possible career paths. As a result some students often end up choosing courses of study which are either not aligned with their interests and strengths, or perhaps are not viable in terms of their future employability.

Despite some schools in Zimbabwe having counselling services, lack of well-trained careers counsellors, tools and materials to enable individual learners to be guided appropriately severely restricts most school counselling systems from making them adequately respond to students' needs on an individual basis with regard to their personality traits, interests and academic strengths (Chireshe 2013).

In addition, traditional means of career guidance might no longer meet the demands of today's learners and the contemporary labour market, in a world of constant evolution. Despite the potential for technological integration, particularly artificial intelligence, in decision-making processes through data analysis and personal recommendations, such approaches are underutilized in the Zimbabwean educational context (Holmes, Bialik & Fadel 2019). Therefore, an AI-driven framework needs to be designed for Zimbabwean pre-university students which will help in providing tailor-made academic and career guidance by considering their unique characteristics.

1.4 Research Objectives

- To assess the current career guidance needs for students
- To design an AI-driven framework that provides personalized career and academic recommendations.
- To evaluate the usability, accuracy, and effectiveness of the AI-driven framework in providing personalized career and academic guidance for pre-university students.

1.5 Research Questions

- What are the current career guidance needs of pre-university students?
- How can an AI-driven framework be designed to provide personalized career and academic recommendations for pre-university students?
- How usable, accurate, and effective is the AI-driven framework in providing personalized career and academic guidance for pre-university students?

1.6 Research Propositions/Hypothesis

Null Hypothesis H0:
The use of an AI-driven personalized career guidance framework does not significantly improve or influence the accuracy of career recommendations given to pre-university students.

Alternative Hypothesis H1:
The use of an AI-driven personalized career guidance framework significantly improves and influence the accuracy of career recommendations given to pre-university students.

1.7 Justification Of The Study

The study was justified because the equipment was found to be outdated and obsolete. Joseph Mtemeri (2017) has indicated that the decision on career choice among the secondary school students in Zimbabwe is often motivated by social and environmental factors as opposed to the individuals' interests and abilities. Making a choice of a career path is one of the most crucial choices that students make before entering the university, however many students of the pre-university system experience difficulties in making decisions. Many secondary schools in Zimbabwe offer limited career guidance or it isn't implemented on a consistent basis, leaving students with no support when they are seeking to discover careers which are suitable for their strengths and interests. This means that some children are placed in educational courses that are not suitable to their skills or future plans. Mtemeri (2017) and Chireshe (2013) further found that students' career choices are greatly influenced by parents' expectations, peer pressure and limited exposure to various career options.

Moreover, the world in today's era has evolved so much that students are facing increased complexities in career decision making as well due to the growth of academic programmes and career options. In many schools, career guidance is traditionally given through general guidance instead of taking into consideration students' characteristics. Regis Chireshe (2013) views career guidance as an activity that helps learners to deepen their understanding of their capabilities, interests and their possible career options, so as to make informed decisions. However, in Zimbabwean schools there is a lack of structured and technology based systems and as such many students are still left to informal means of advice when choosing a career. Hence, this study is significant due to the introduction of an AI-based framework that can help to analyse students interests, behaviours, and strengths and suggest appropriate academic and career trajectories. This system could help Zimbabwe students make better choices and assist with the alignment of education and employment needs (Mavima, 2021), as it could offer them more personalized career guidance.

CHAPTER 2: LITERATURE REVIEW

2.1 Introduction

The choice of career is one of the most important decisions that a student makes prior to entering university, which can influence career choices, school subjects and ultimately career opportunities in the future. However, many students struggle with career choice as they are unaware of career interests, strengths and options available. The use of conventional methods such as counselling is limited in most schools as personalized learning support can't be effectively given to all learners.

New methods in the area of career guidance have been developed using technology such as Artificial Intelligence. AI can be utilized in determining students' careers by collecting relevant data such as interests, behaviour , performance and skills in order to provide better career guidance recommendations tailored to the students' individual needs. J Trujillo and Suntaxi (2024) asserts that the usage of machine learning technologies is becoming more and more common in education in order to help learners with their career and academic choices. Moreover, Monika and Balaji (2025) mention that using an AI-based guidance system provides users with interactive suggestions and individualized careers guidance advice at their disposal, thereby increasing its applicability.

In this chapter, literature will be discussed on career development theories and AI-based career guidance systems. Theoretical implications underpinning decision-making in careers will be discussed, as well as the empirical research that has explored the impact of AI on career guidance. It is this literature review that will provide a foundation from which an AI-supported framework

for personalized career and academic guidance for pre-university students of Zimbabwe will be developed.

2.2 Relevant Theory Of The Subject Matter

A number of career development theories are designed to provide explanations for career choice and career development, all of which can guide the development of intelligent Career Guidance systems.

Holland's Theory of Vocational Personalities and Work Environments is a prominent career development theory. Holland (1997) suggested that jobs are typically chosen based on personality type and interests. According to Holland's theory, personalities are categorized into: realistic, investigative, artistic, social, enterprising and conventional. When choices are consistent with an individual's personality style, satisfaction is generally high and careers are likely to be stable (Gottfredson, 1999). AI-based career guidance systems may apply this by matching students with careers based on personality profiles and interests.

Another essential career theory that may be relevant is the Trait and Factor Theory introduced by Parsons. It indicates the need for understanding an individual's abilities, interests and personality traits prior to selection. The ability to make successful career decisions needs a three-step process: knowledge of self, knowledge of occupations and matching job to the person (Zunker, 2006). Career choice recommendation systems may employ an approach similar to the Trait and Factor Theory through analysing data relating to performance and interests in order to match individuals with appropriate careers.

The Social Cognitive Career Theory (SCCT) is another significant theory. It postulates that factors like self-efficacy, outcome expectation and personal goals are important for decision making in careers. Students will likely gravitate towards careers where they feel that they can excel and with high expectations of positive outcome. Using AI, the success factors can be achieved by collecting information relating to students' strengths, interests and abilities so as to help them pick their ideal career path.

Career Development Theory was also propounded by Super, asserting that career development is an ongoing process throughout a lifetime. Super (1990) classified the life span into the following developmental stages: growth, exploration, establishment, maintenance and decline. Exploration occurs in the teens and is crucial for students who are trying to determine career preferences. AI systems will be helpful for students at this stage as they can be equipped with skills and information that will assist them in exploring relevant careers.

In addition to the theories above, there have been a number of empirical studies which examine the implementation of AI in career counselling. Trujillo, Pozo and Suntaxi (2024) conducted a systematic review of machine learning techniques used in recommending aptitude careers. The studies showed that algorithms such as Random Forest, Support Vector Machines and Neural

Networks are successful in analysing student's performance and interests and in guiding them to the most fitting career.

In a similar study by Husniah, Kurniawan and Ulfa (2025), the use of AI in providing customized career and academic advice for students was investigated. Their study suggested that AI tools like chatbots and recommendation systems play an important role in improving students' awareness of careers as well as students' confidence in their decision making of their careers.

According to Saparovna (2024) AI is an essential aspect of modern career counselling. Educational data can be processed by AI systems to provide a bespoke recommendation system for careers, even exposing students to career options they might not have considered before.

Mavuso, Van Greunen and Jere (2024) developed an AI-driven career guidance framework for South African universities and found that students were very receptive to using such tools for career planning because of the tailored advice that they provided. However, challenges related to ethical use, such as data privacy and potential algorithmic bias, were highlighted.

In another study, Kariya et al. (2025) used AI for assessing students' skills and learning style in order to suggest relevant STEM careers and courses to students. Their study concludes that there are a range of approaches that can be taken to guide students toward suitable careers and learning plans.

Paddisengeng (2025) noted that predictive AI counselling systems can assist students in improving their knowledge and self-efficacy for career choice and academic progression towards employment opportunities. Their research highlighted that those who received career counselling through AI performed better career planning than those who received career guidance using traditional counselling methods.

Despite the promising implications of AI in career counselling, challenges such as inadequate technological infrastructure, issues concerning data privacy, and the potential of algorithmic bias may hinder the widespread use of AI systems, especially in developing countries.

2.3 Conclusion

The chapter covered both the theory and empirical studies on career guidance and the application of AI in education decision making. Various theories like Trait and Factor Theory, Social Cognitive Career Theory, Holland's Theory of Vocational Personalities and Super's career development theory formed the foundation for analysis of the student's choice of career based on interests, abilities and goals. The empirical research also revealed the importance of AI technologies like the machine learning algorithms, recommendation systems and AI chatbots in supporting the efficacy of career guidance systems, in delivering a targeted guidance towards the students. The algorithms use Retrieval Argument Generation and large language model to extract student's information, determine strengths and recommend a career that is in accord with the

needs of the job market. However, it was noted there were a few disadvantages including privacy and security concerns and also access to technology in certain areas. The advantages have revealed the importance of developing context-specific AI systems and thus, the research intends to contribute in this direction by proposing an AI based career and academic advisor for pre-university students in Zimbabwe.

CHAPTER 3: RESEARCH METHODOLOGY AND SYSTEM DESIGN

3.1 Introduction

This chapter outlines the technical specifications of a sample web application, CareerGuideZW which is a website and AI based system for customized guidance for the career and academic pursuits of pre-university students in Zimbabwe. The chapter covers information on research design, data gathering techniques, the targeted population, the sample, hardware and software requirements, system specifications, data flow diagram, entity-relationship diagram, system architecture, chat advisor prompt design and step-by-step illustration of system functioning. All system design decisions were guided by the research purpose, presented in Chapter 1 and the theory discussed in Chapter 2. The CareerGuideZW system was conceived to address the deficiencies found in the career guidance provided in the Zimbabwean secondary school systems where students had to choose their A-Level subject combinations and career goals without proper information or personalized, data-driven career guidance.

This need was addressed by developing a hybrid recommendation system that combined rule-based scoring with a local large language model from the Ollama platform in Zimbabwe (Lewis et al., 2020). Agile approach; specifically Scrum was used to develop the system. The development took place in sprints and testing and refinements were carried out continuously throughout the development process rather than solely at the end of the development period (Schwaber and Sutherland, 2020). Such an approach was relevant for this project as user needs varied with time and the system became refined based on users needs.

3.2 Research Design

A mixed methods research design approach was utilized, incorporating qualitative data from interviews and feedback obtained from students and various stakeholders, along with quantitative data based on structured evaluations and the assessment of system performance. Such an approach was identified as the most appropriate in order to adequately cater to both the students' unique experiences and the demonstrable efficiency, precision and ease of use of the system (Creswell and Creswell, 2018). Software development itself was conducted following a Design Science Research Methodology (DSRM), as widely accepted within Information System research for projects focused on the creation and testing of artefacts to address a particular real-

life problem, as established by Hevner et al (2004). In the case of CareerGuideZW, development of the system was considered the initial phase, with its assessment in terms of effectiveness, usability and accuracy considered the research outcome.

Agile Scrum method was adopted. The project was subdivided into development sprints of 2 weeks length. The features were continuously reviewed and changes made based on tests and feedbacks, allowing a progressive refinement rather than setting concrete requirements from the onset. This is crucial for developing AI systems as the behaviour of the AI is unpredictable until it is tested with the actual data profiles (Beck et al., 2001).

3.3 Data Collection Methods

The data for this project were sourced from two primary methods, namely the collection of primary data directly from students and school professionals, and the collection of secondary data from existing career and educational materials.

3.3.1 Primary Data Collection

Pre-university students and selected secondary school teachers from certain schools were included in the sample and asked to complete questionnaires that contained information regarding subjects they studied, marks obtained, likes and dislikes towards particular careers, and limitations in career guidance services they faced. Semi-structured interviews were conducted with career counsellors and school advisors, and their input used to guide the development of the 7-step career assessment questionnaire, a key part of the CareerGuideZW system. The questionnaire consisted of areas on student's strengths, likes, dislikes and academic information.

3.3.2 Secondary Data Collection

A range of sources was consulted in gathering secondary data; these included ministry program lists and university career guide books. The data was used to form a database on careers and programs which was used to train the recommendation system. The career dataset comprises over 358 university and college degree and diploma programs that included minimum qualification, average and maximum mark, requirements and related course combinations.

3.4 Population and Sample

The target population of this study consists of pre-university students in Zimbabwe at Form 4 and Form 6 levels with an average age range between 16 and 20 years. At this stage of education, students have to make crucial decisions such as which advanced level subjects to combine, which university or college to apply to and career path to follow. A large proportion of them do not receive professional career advice due to schools not having dedicated career counsellors, or their only counsellor being overloaded with students to offer personalized career guidance.

The target population also includes career guidance counsellors and subject teachers that have the relevant expertise to validate if the recommendations from the system are practical and contextually relevant within the Zimbabwean education system.

Since it was not practical to access all pre-university students in Zimbabwe, purposive and convenience sampling was used. Purposive sampling involved selecting individuals based on the criteria needed for the study and not necessarily random sampling (Patton, 2002). A sample of form 4 or form 6 students that had access to a web browser compatible device and were willing to undergo the career assessment was considered. Students from the schools the researcher accessed conveniently were targeted for convenience sampling as they were available and willing to participate upon request from the researcher. This method has conventionally been applied to studies within the field of educational technology research (Etikan, Musa and Alkassim, 2016).

The sample size suggested for the quantitative section of the study was between 50 and 100 students which is sufficient for meaningful statistical analysis using rating feedback, match percentage distribution and count for usage of the system. Depth over representativeness (Braun and Clarke, 2006) is also a principle applied in qualitative research so 10-15 students were targeted for qualitative questions such as open-ended feedback and interviews with counsellors.

For inclusion into the study, participants had to meet several criteria; including being at Form 4 or Form 6, access to a web browser enabled device, have given consent for their participation, and possess adequate English communication skills to interact with the system. Students from other levels were not included as it was irrelevant for them to know about career and A Level choices directly. Data from students that were obviously not valid were not included for quantitative analysis.

3.5 Functional and Non-Functional Requirements

3.5.1 Functional Requirements

The functional requirements specified the actions the CareerGuideZW system should carry out to effectively serve its purpose of helping pre-university students make appropriate academic and career choices. Through the analysis of requirements, the key functional requirements related to a specific system behaviour have been identified: (Sommerville, 2016)

Table 3.1: Functional Requirements for CareerGuide ZW

ID	Functional Requirement	System Response	Priority
FR-01	The system shall allow students to create an account and log in using a username and password	User credentials stored with bcryptjs hashing; session maintained via Express middleware	High
FR-02	The system shall provide a seven-step career assessment questionnaire capturing interests, strengths, O-Level subjects, A-Level subjects, Holland RIASEC type, hobbies, and academic performance	Multi-step form rendered via React Hook Form with Zod real-time validation at each step	High
FR-03	The system shall generate up to six personalised career recommendations ranked by match percentage	Rule-based scoring engine computes weighted scores across five dimensions; top six returned sorted descending	High
FR-04	The system shall match university and Polytechnic diploma programmes to each recommended career based on the student's qualifications	SQL query on university_programs table filtered by subject, pass count, and cut-off thresholds	High
FR-05	The system shall generate personalised AI career advice using a locally deployed LLM	Ollama API called with structured prompt; response cross-checked against programme database before display	High
FR-06	The system shall provide a conversational AI chat advisor for	Chat messages sent to LLM with system-level prompt constraints; responses	Medium

ID	Functional Requirement	System Response	Priority
	follow-up career and study guidance	stored in chat_sessions table	
FR-07	The system shall allow students to submit star-ratings and comments on career recommendations received	Feedback form linked to specific career recommendation; stored in feedback table with user ID and timestamp	Medium

3.5.2 Non-Functional Requirements

Non functional requirements are quality attributes such as performance, security, reliability, usability, scalability, data privacy etc. That the system needs to satisfy. These requirements are particularly important with regards to the deployment in a Zimbabwean context whereby the internet may be unstable and there may be scarcity in the computing resources (Pressman, 2014).

Table 3.2: Non-Functional Requirements for CareerGuide ZW

ID	Category	Requirement	Implementation Strategy
NFR-01	Performance	The system shall return career recommendations within 5 seconds under normal load	Rule-based engine is O(n) where n=12 careers; SQLite indexed queries; LLM timeout set to 30s with graceful fallback
NFR-02	Security	Passwords shall never be stored in plaintext; user sessions shall expire after inactivity	bcryptjs with salt rounds=10; Express-session with secure cookie flags and session expiry configured
NFR-03	Reliability	The system shall continue operating in degraded mode if the LLM component is unavailable	Try-catch wraps all Ollama calls; rule-based results returned with a notice if LLM fails

ID	Category	Requirement	Implementation Strategy
NFR-04	Data Privacy	Student profile data shall not be transmitted to external third-party AI services	Ollama runs locally on the deployment server; no external API calls for LLM inference
NFR-05	Usability	The interface shall be navigable by users with basic computer literacy without formal training	Radix UI accessible components; step-by-step wizard form; descriptive error messages
NFR-06	Portability	The system shall be deployable on any server with Node.js 20+ without a separate database server	SQLite is file-based and serverless; pnpm workspace monorepo; no cloud dependencies required

3.6 Hardware Development Tools

One of the most important aspects of the hardware requirements of CareerGuideZW is the use of a locally deployed large language model through the Ollama platform. Unlike a typical web app, running a lightweight LLM like Gemma 3:1b or Llama 3.2 on the server requires a lot more memory. The minimum and recommended hardware is given below for a server and student client device.

Table 3.3: Hardware Requirements for CareerGuide ZW

Component	Minimum Specification	Recommended Specification
Processor (Server)	Dual-core 64-bit CPU, 2.0 GHz	Quad-core 64-bit CPU, 3.0 GHz (for Ollama LLM inference)
RAM (Server)	4 GB RAM	8 GB RAM minimum for running Gemma 3:1b or Llama 3.2 locally
Storage (Server)	20 GB SSD	50 GB SSD (LLM model files are 1–4 GB each)

Component	Minimum Specification	Recommended Specification
Network (Server)	10 Mbps upload bandwidth	100 Mbps, low latency for concurrent student users
Client Device	Any device with a modern web browser (Chrome 90+, Firefox 88+)	Smartphone or laptop with stable internet connection (2G minimum)
Client RAM	1 GB RAM on client device	2 GB RAM for smooth React rendering and real-time form validation

The most crucial server hardware resource is the amount of RAM it needs to have. In terms of hardware, the selected language model is fully loaded into RAM by Ollama when inference needs to be done, hence for the Gemma 3:1b model this amounts to around 900MB RAM, for the Llama 3.2 model this ranges from 2 to 4GB depending on how it was quantized. Providing 8 GB RAM for the server ensures that there is enough for the Operating System, the Node.js process, the SQLite engine and the LLM, all at the same time. Client side demands are kept as low as possible so as to allow students who only have entry-level mobile phones or computer laboratory facilities at school.

3.7 Software Development Tools

Three requirements were taken into consideration for the selection of the development tools to be used for the CareerGuide ZW; it needed to be open source so as not to have licensing fees associated with it, it needed to work well with TypeScript for end-to-end type-safety, and it needed to be suitable for deployment to low-resource environments. The core tools and technologies used are displayed in Table 3.4

Table 3.4: Software Development Tools and Technologies

Tool / Technology	Category	Justification and Role in CareerGuide ZW
React 18	Frontend Framework	Component-based UI architecture; efficient virtual DOM rendering for the multi-step assessment form

Tool / Technology	Category	Justification and Role in CareerGuide ZW
TypeScript 5.9	Programming Language	Static typing catches errors at compile time; improves maintainability across the full-stack monorepo
Node.js 24 + Express 5	Backend Runtime / Framework	Asynchronous I/O enables concurrent request handling; lightweight REST API development
SQLite 3 + Drizzle ORM	Database / ORM	Serverless, zero-config database; Drizzle provides type-safe schema and queries without a separate database server
Ollama	LLM Inference Platform	Local LLM deployment; avoids external API dependency; protects student data privacy (NFR-04)
Gemma 3:1b / Llama 3.2	AI Language Model	Compact models suitable for resource-constrained deployment; sufficient for career advisory text generation
Zod	Validation Library	Schema-based validation shared between frontend and backend; prevents invalid data from entering the system
Tailwind CSS + Radix UI	Styling / Components	Utility-first styling with accessible headless component primitives for WCAG compliance
bcryptjs	Security	Password hashing with salt; prevents plaintext credential storage; OWASP recommended approach
pnpm Workspaces	Monorepo Manager	Efficient shared dependency management; faster installation than npm; single repository for frontend and backend
Git / GitHub	Version Control	Source code management; branching strategy for feature development; backup and collaboration

Tool / Technology	Category	Justification and Role in CareerGuide ZW
Visual Studio Code	IDE	TypeScript IntelliSense; ESLint integration; debugging tools; primary development environment

All of the tools that have been listed have been classified as open-source, which indicates that there is no associated software licensing fee which in line with the development of a system that will be maintained and advanced by local developers without any licensing fees. A system which works with sensitive student data needs a closed system in the front-end and back-end to work with, in which any errors in terms of data are discovered at compile time rather than at runtime and this therefore decreases the possibility of errors in handling data.

3.8.1 Database Normalization (3NF)

Database normalization is a technique by which to design and structure the database such that data redundancy is minimized and data consistency is maximized. The database used in the CareerGuideZW system has been structured to comply with Third Normal Form (3NF). This is done to store information efficiently and not repeat information needlessly.
A table is said to comply with Third Normal Form if:

- It is already in Second Normal Form (2NF).
- All non-key attributes depend only on the primary key.
- Transitive dependencies between non-key attributes do not exist.

3NF ensures that only one entity is stored per table and all attributes of a table are related to its primary key.

Table 3.5 below shows examples of tables in the CareerGuide ZW database that satisfy Third Normal Form.

Table 3.8: Database Table in 3NF

Table Name	Primary Key	Non-Key Attributes	3NF Condition
users	user_id	username, email, password_hash	All attributes depend directly on user_id only

Table Name	Primary Key	Non-Key Attributes	3NF Condition
student_profiles	profile_id	interests, strengths, subjects, holland_type	Profile details depend only on profile_id
feedback	feedback_id	rating, comment, created_at, user_id	Feedback information depends only on feedback_id
university_programs	program_id	program_name, institution, required_subjects, cutoff_points	Programme details depend only on program_id

The users table is in 3NF as all non key attributes username, email, pwdhash depend on the only key attribute user_id. None of the non key attributes depend upon each other.

The student_profiles table is also in 3NF since student interests, strengths, Holland personality type, subjects are all describe solely with the profile_id. The user login is not duplicated in this table, so this reduces redundancy.

The feedback table is in 3NF as feedback_id functionally determines all attributes in the table. In particular comment and rating depend on feedback_id but not other non-key attributes.

The university_programs table is in 3NF because all programme details like institution name, taking subjects and cut off points are fully functionally dependent on program_ID and student data is not consolidated with programmes data.

The CareerGuide ZW is enhanced due to implementation of database in 3NF. This database model widely helps the CareerGuideZW to enhance the data consistency, minimize the redundant data, get easier management and also improve the database performance to some extent.

3.8.2 The Hybrid AI Recommendation Framework

The CareerGuideZW recommendation model is a hybrid model, whereby two distinct fundamental computation models, the deterministic rule based scoring system and the probabilistic LLM based system are integrated. The two models individually cannot achieve the required aim of the system. A pure rule based system has consistent and explainable results but

cannot produce an understandable natural language recommendation to an individual student. An unmodified LLM can generate text which appears human-like but is prone to hallucination and cannot guarantee the existence of mentioned programmes. The hybrid model however harnesses benefits of both systems whilst overcoming drawbacks of the individual models.

The rule based scoring engine determines a weighted score for every career contained within the data set for each of the five categories: mapping of Interest to category, matching of subjects, Holland RIASEC affinity, qualification category and affinity of strengths. For each of these parameters points may be added to the score to represent match levels which fall within the pre-defined boundaries. The result is a rank of n careers, with n being 6, which can be said to represent a 100 percent match to the profile. This score is deterministic and is crucial in an educational tool to maintain consistency. Students of the same profile must obtain a similar career ranking (Hevner et al., 2004).

An LLM model (Gemma 3:1b or Llama 3.2) is trained using student data and output from the rule-based engine in order to produce an enriched context, for which the LLM is trained to develop a four-to-five-sentence personalized recommendation text. The risk of a program name being hallucinated by the LLM is mitigated by performing a cross reference with all university programs contained within the database following the generation of the text.

3.9 Model Data Set and Data Cleaning

The data set upon which the CareerGuide ZW recommendation engine is founded is embedded and comprised of a combination of two datasets developed specifically for the educational context of Zimbabwe. The first dataset represents the careers which encompass a total of 358 careers distributed across careers sectors: technology, business and finance, engineering, education, health sciences, law, agriculture, public service, sports, tourism, media and communication and social sciences. Each of the career profiles contain data on the descriptions of the career, skill requirements, A level subjects which have most relevance, sample University programmes available in Zimbabwe, the career salary expectation, an evaluation of job outlook, and keywords for matching.

The second is a Zimbabwe university program dataset that contains all degree and diploma programs offered by the various universities in Zimbabwe, with each program containing information on the prerequisite A and O level subjects required, a pass threshold for each level of education and minimum cut-off score, a program type (degree or diploma) classification and a note on the award of the HEXCO diploma to indicate alternative options available to those without an A level.

This data needed to be cleaned extensively before it can be used by the recommendation engine. The following data cleaning steps were taken:

- All missing data for all career and program records was either filled in with second sources or excluded from the data set to avoid having incomplete data to bias the scoring engine.
- Duplicate removal: Multiple data sources generated duplicate entries, which were identified and removed, leaving the most complete and accurate version of each record.
- Subject alias normalization: There are multiple alternate names for the same subject and these were converted to a normalized form. For instance ‘Maths', ‘Mathematics' and ‘Pure Mathematics' are all standardised to ‘Mathematics'. Likewise, the following terms are mapped to ‘Accounting’: ‘Accounts’ and ‘Accounting’. If a student had not taken this step, then a student who had entered their profile and written ‘Maths' would not have matched with a program that was looking for ‘Mathematics' although the qualification might be the same.
- All free-text fields were normalized in terms of standardizing text to lowercase, removing special characters and punctuation, and making it consistent across all records.

3.10 Feature Engineering for the AI Recommendation Model

Feature engineering aims to convert raw input data into numerical or categorical data in a structured way that can be fed into the recommendation algorithm (Zheng and Casari, 2018). The rule-based recommendation engine is the main model in CareerGuideZW, and the feature engineering step is the one that transforms the students' freeform answers to their assessments into dimensions that can be scored over by the rule-based recommendation engine. Table 3.5 records each of the engineered features, the fields where they are found, the transformation they undergo, and their uses in the model.

Table 3.5: Feature Engineering for the CareerGuideZW Recommendation Engine

Feature	Source Field(s)	Transformation Applied	Use in Model
Interest-Category Score	interests (multi-select)	Map interest strings to career category weights; sum per career	Primary weighting (0–30 pts)
Subject Match Score	o_subjects, a_subjects	Alias normalization then binary match per required subject; aggregate	Eligibility scoring

Feature	Source Field(s)	Transformation Applied	Use in Model
Holland RIASEC Score	holland_type	One-hot encode 6 types; map to career category affinity matrix; dot product	Personality fit bonus
Academic Tier Flag	o_passes, a_passes	Rule: A_passes ≥2 AND o_passes ≥5 → Degree tier; else Diploma tier	Pathway branching logic
Demand Level Tag	job_outlook (career dataset)	Categorical encode: Excellent/High → High; Moderate → Medium; else Low	Display label and re-ranking
Strength Affinity Score	strengths (multi-select)	Map strength strings to career skill requirements; weighted sum	Supplementary scoring
Cutoff Eligibility Score	cutoff_points (ZIMSEC 0–15)	Compare student cutoff against programme min cut off threshold.	Program eligibility filter
LLM Context Vector	All profile fields	Concatenate structured profile as natural language prompt template	Prompt engineering for LLM

Interest-Category Score is the most complicated feature engineering. This process translates the student's selected natural-language interests (chosen from a pre-defined list of interests) into a weighted scoring matrix of career categories, where each of the student's interests can contribute to a number of different career categories with a given weight representing the degree of contribution. LLM Context Vector is the final engineering process, which is composed of serializing the student's entire structured profile into the LLM-formatted natural language prompt template for the LLM to get the context to generate coherent and personalized advise.

3.11 Model Training

Unlike training the language model from scratch, CareerGuideZW takes advantage of a pre-trained language model already deployed via Ollama. The 'training' in this system is conducted in the process of knowledge base construction, rule engine calibration, and retrieval optimization, involving these stages:

- Career and University Programs Datasets Cleaning and Preprocessing, as described in Section 3.9, the data used by the scoring engine is accurate, consistent, and representative of the Zimbabwean context.
- Interest-to-Category Weighting Matrix calibration, taking consultation from career advisors and empirical literature in Holland's RIASEC theory as guide (Holland, 1997) in assigning weights based on realism of the interests' correlation with the career categories.
- Subject alias normalization table compilation, by iterating through all combinations of ZIMSEC O-Level and A-Level subjects, and matching each variant subject name to its normalized form.
- Academic tier branching logic design and test, this logic is verified by running a set of students profiles with different academic standings and identifying that the Degree tier and Diploma tier were activated at correct boundaries.
- The entire recommendation pipeline is test with a set of 30 sample student profiles and scoring weights, prompt template are adjusted iteratively until it produces consistent and accurate career recommendations.

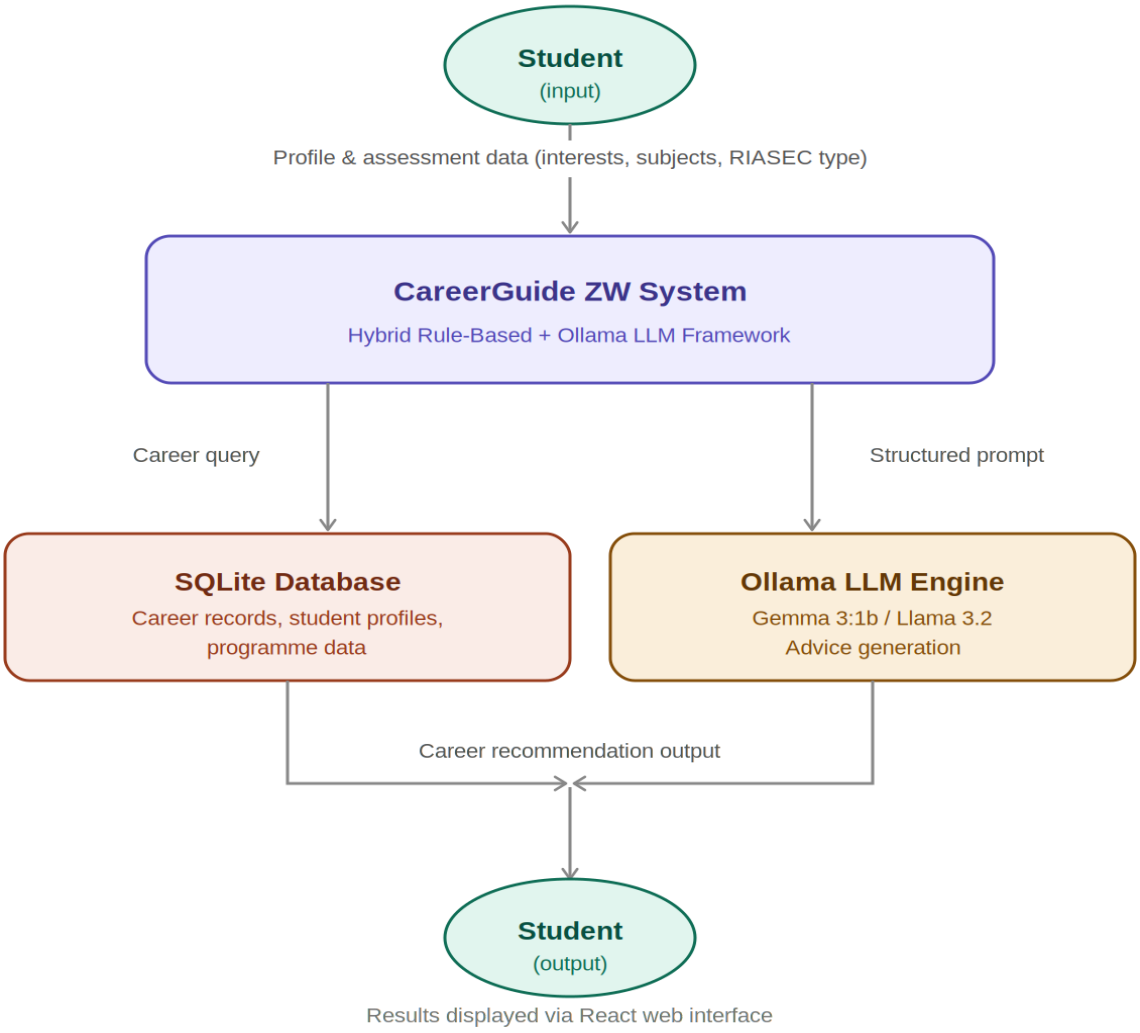
3.12 Data Flow Diagram (DFD)

Data Flow Diagram (DFD) is a diagram that illustrates the flow of information within a system. The overall context of an application can be presented by its level which treats the application as a single process and shows its external interfaces with which data are exchanged (Yourdon, 1989).

In CareerGuideZW, the primary external interface is the Student, which sends in student's profile and assessment information, while interacting with the SQLite Database, where it stores career records, student profiles and program information and Ollama LLM Engine to generate the recommendation text. The recommended career and suggestion are then returned to the Student through the web interface, as shown in Figure 3.1 below.

Figure 3.1: Level 0 Data Flow Diagram (Context Diagram) — CareerGuide ZW

Figure 3.1 — Level 0 Data Flow Diagram (Context Diagram)



The flow is triggered when the student inputs their answers into the assessment page in the React web interface. These inputs pass through the Node.js/Express application layer, through the rule-based scoring engine and out to the SQLite database for the matching career and program entries. The student profile and the top career outputs are packaged into a string to act as a prompt and sent to the Ollama LLM engine where the advisory text is generated. It then goes through the cross reference check before being returned to the student as the final recommendation output.

3.13 Entity Relationship Diagram (ERD)

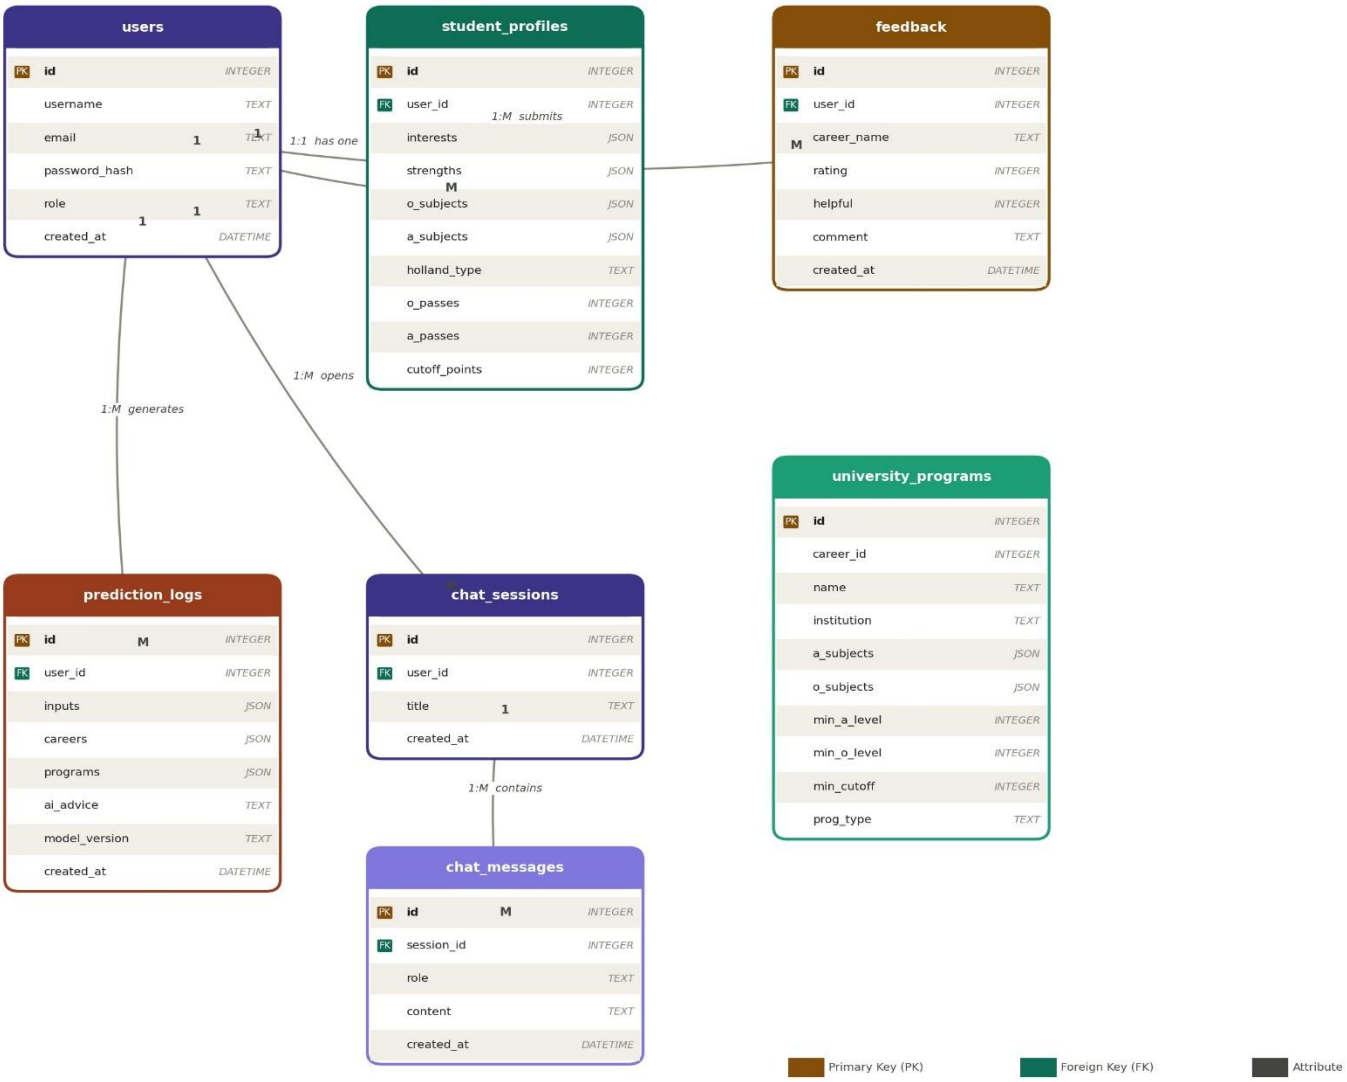
The Entity Relationship Diagram (ERD) shows the logical structure of the CareerGuide ZW database and how the main data entities relate to one another (Chen, 1976). Six main tables are used: users, studentprofiles, prediction_logs, university_programs, feedback, and chatsessions/chat_messages.

The users entity is the centre point; it has a one-to-one relationship with student_profiles (a user has one profile) and a one-to-many relationship with prediction_logs, feedback and chat sessions (a student can produce multiple recommendations, several bits of feedback and hold multiple chat sessions). Each chat session has one-to-many relationship with chat_messages. The university_programs entity does not directly connect with users; it is the reference table that the recommendation engine queries against.

The six entities are illustrated below in Figure 3.2. The university_programs table is designed as a stand-alone reference table as that will make updates to programs simpler without impacting student data. Prediction_logs records the complete JSON response for each recommendation session including the student's inputs, career output and model version. This allows analysis of recommendation quality over time.

Figure 3.2: Entity Relationship Diagram — CareerGuideZW

Figure 3.3: CareerGuide ZW – Entity Relationship Diagram



3.14 System Architecture And Technology Stack Implementation

The CareerGuide ZW application is built as a four-layered web application; using a client-server architecture. Each layer has its own clearly defined responsibility, is called "Presentation Layer, Application Layer, Data Layer, and AI / LLM Layer", and communicates with its neighboring layers through defined interfaces (Fowler, 2002).

3.14.1 Presentation Layer

The presentation layer is developed using the React 18 framework, utilizing the TypeScript language to enable type safety and provide components, and is supported by Zustand which stores the student profile across the 7 steps of the assessment page so no server round trips are needed between each step. The presentation layer also contains the registration and login pages as well as the chat advisor.

3.14.2 Application Layer

The application layer is written in TypeScript and uses the Node.js 24 runtime and the Express 5 framework. All browser HTTP requests are routed through this layer, it stores the student's session using the Express-session middleware, validates the form inputs using Zod schemas and orchestrates the entire recommendation pipeline. The application layer exposes an API specified according to the OpenAPI 3.1 specification; this is then parsed by Orval in order to auto-generate typed code that synchronizes automatically with the backend without any need for maintenance of two different definitions, thereby eliminating one source of errors in the system.

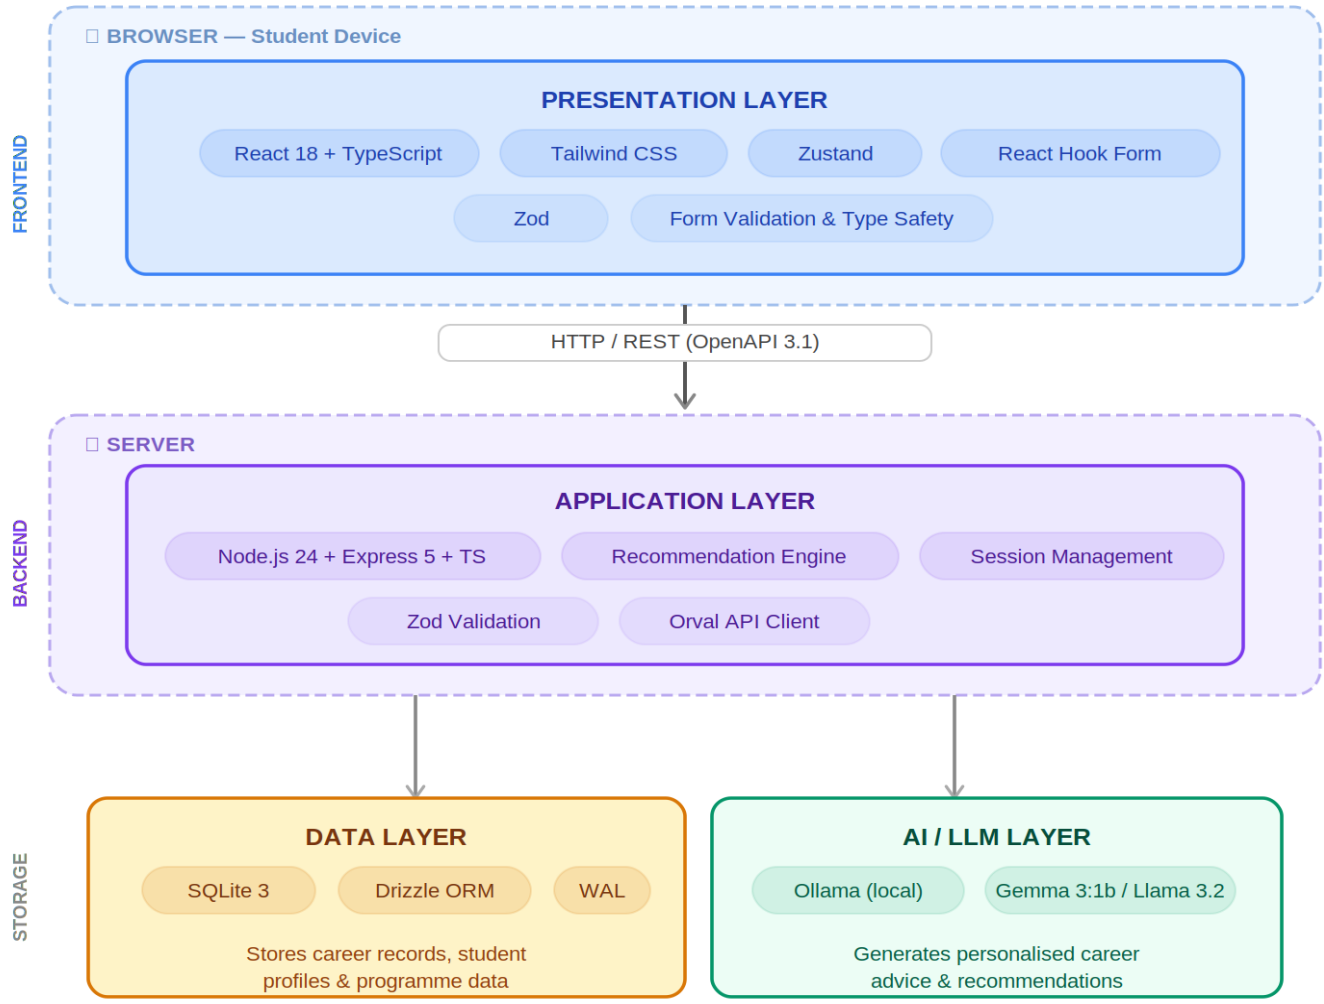
3.14.3 Data Layer

The data layer is an SQLite 3 database managed by the Drizzle ORM. SQLite was chosen for its zero-configuration deployment, as there is no need for a database server process that must run alongside the rest of the application. Drizzle supports type safe schema creation and type safe querying, meaning that compile time errors are raised if any attempt is made to query a non-existent column within the schema. Write-Ahead Logging (WAL) has been enabled on the database to allow reads to occur in parallel with writes, reducing potential blocking.

3.14.4 AI / LLM Layer

The AI layer is contained within the Ollama process, which serves either the Gemma 3:1b or Llama 3.2 language model through its local HTTP API. Since Ollama runs locally within the deployment environment, no student information is sent outside of the network, which ensures NFR-04 is met. The CareerGuide ZW backend makes calls to the Ollama API on the local network and then receives the resulting advisory text to check against the hallucination verification system before returning it to the student.

Figure 3.3 — CareerGuide ZW System Architecture



In Figure 3.3 we show the four layer architecture and data flows between layers. The browser only ever interacts with the application layer via the OpenAPI-described REST API. The application layer interacts with the data layer and AI layer within the locally running server, such that only internal data is processed.

3.15 Chat Advisor System Prompt Design.

One of the core features of the CareerGuideZW application is the conversational chat advisor where students may pose career-related questions in natural language and receive appropriate responses. The responses generated by the LLM depend on how well the system prompt has been designed. A system prompt is the hidden instruction given to the model at the start of each conversation; it informs the model about what is expected of it, the scope of its responses, its

persona, and behavioural guidelines (Brown et al., 2020). Two unique prompt templates are used in the CareerGuideZW system, both configured with specific inference parameters optimized for their respective use-case.

3.15.1 Advice Generation Prompt.

For each student the advice generation prompt is constructed dynamically to include the student's full profile (interests, strengths, subjects, Holland type, top career matches, top program matches), into the system prompt. It is instructed to respond as a Zimbabwean career advisor who is knowledgeable of the local education system and to provide 4 to 5 sentences of practical, actional advice directly to the student that references university programs by name. There is a post processing step which checks any program names that are output against the university_programs database.

3.15.2 Chat Advisor Prompt.

The chat advisor prompt is used as the system instruction for the chat session. It specifies that the LLM can provide answers to questions about career advice, A/O Level subject requirements, Zimbabwean universities and colleges, and general pathway planning but if the student provides out of scope questions then it should gently steer the student back toward the available topics. The student profile is provided as context for the system prompt at the start of each chat, so that the LLM can refer to the student's selected subjects, interests, and top careers when appropriate in responses. The overall design aligns with common prompt engineering principles for LLMs regarding specifying the role, grounding in context, and constraining output.

3.16 Summary of How the System Works

The following steps explain how the system responds to student requests:

Table 3.6: Step-by-Step Summary of the CareerGuideZW Process

Step	Stage	Description
1	Registration and Login	Student creates account with details: username, email, password. Password hashed with bcryptjs and saved. Student logs in and session is established with Express-session middleware.

Step	Stage	Description
2	Seven-Step Assessment	Student complete the 7-step career questionnaire which requires them to provide details about their interests, personal strengths, O/A Level subjects, Holland RIASEC personal type, hobbies, and performance. Zod provides validation on all steps in real time.
3	Feature Engineering	Student's input from the questionnaire is sent to the application layer where the feature engineering pipeline is responsible for normalizing subject aliases, computing scores for all 5 dimensions, calculating the academic tier, and constructing the context vector for the LLM.
4	Rule-Based Scoring	The scoring engine calculates a weighted sum across all 12 possible careers. It then takes the top 6 of these careers by score and uses this to populate the recommendation package (career name, score, reason string in plain english).
5	Programme Matching	Each career result found by the scoring engine is used to query the university_programs table for courses that match student's choice of subjects, number of passes and minimum grades (cut-off points). These matching courses are then added to the student's results and displayed under each career title.
6	LLM Advice Generation	Student's profile and top career results are serialized into the advice generation prompt and sent to the locally running Ollama LLM. 4 to 5 sentences of personalized advisory are returned and a cross check is performed against the program database for output consistency before presentation.
7	Results Display	Recommendation package is sent back to the React frontend which displays the recommended careers and advice to the student on the results page.
8	Chat Follow-Up	Student can then send further questions which will be directed at the Ollama LLM alongside the entire conversation

Step	Stage	Description
		history and student's profile. LLM will respond within its allowed context.
9	Feedback Submission	Student can rate the suggestion or provide further comments for the recommendation they have received. This information is stored in the feedback table associated with the user's profile and prediction session
10	Session Logging	complete record of the user session is saved in the prediction_logs table. This will provide data for the evaluation of the system as described in the next chapter

This end-to-end process should run within 8 seconds for all 3 steps of scoring, programme matching and LLM advice generation combined. What distinguishes CareerGuideZW from an 'unintelligent' keyword search or static FAQ are two distinct aspects; firstly, a deterministic, rule-based engine to provide precise and explainable career matching, and secondly, a locally-deployed LLM for a tailored natural-language advice component. CareerGuideZW can offer a response that is both based on actual data, and delivered in the kind of 'empathetic' language a human career advisor would provide.

3.17 Chapter Summary

This chapter provided a detailed account of the research methodology, system design and implementation approach that went into the construction of CareerGuideZW. Starting with research design and data collection methods, the chapter explained how the system is grounded in student requirements and actual Zimbabwean career data. The section on population and sample specified who the system is for and how the participants were sampled. Functional and non-functional requirements laid down what the system must do and how well. Hardware and software requirements outlined the running environment and entire technology stack (React, Node.js, SQLite, Ollama LLM platform). The section on model design outlined the hybrid rule-based/LLM architecture that gave CareerGuideZW both the accuracy and predictability of a deterministic engine, along with the responsiveness of a language model. Dataset and data cleaning discussed how career and programme data was gathered and prepared. Feature engineering section focused on how student inputs are mapped onto score dimensions. The DFD and ERD presented visual representations of data flow and database structure respectively, while the system architecture diagram showed communication across the four layers. Chat advisor prompt design elaborated on the LLM prompt engineering efforts. Finally, the step-by-step

system flow process, displayed in Table 3.6, summarized the entirety of how a student interacts with CareerGuideZW from login to feedback submission.

CHAPTER 4: SYSTEM IMPLEMENTATION AND TESTING

This chapter details the testing, evaluation, and implementation outcomes for the CareerGuideZW system. Testing was undertaken to validate the system's correct behavior under various conditions and measure the accuracy of career recommendations provided. Four types of testing were performed, namely white-box, black-box, performance testing, and statistical assessment using a confusion matrix, precision and recall values.

4.1 White-Box Testing

White-box testing, or structural testing, is concerned with the internal logic and execution paths of the system (Myers et al., 2011). It involves analysis of the code, allowing testers full visibility of its operation as opposed to black-box testing. For CareerGuideZW, white-box testing was employed to confirm that critical functions like database retrieval, input processing, and AI recommendation logic ran without error and effectively handled unusual or unexpected data inputs.

The following five tests were designed to confirm functionality on each of the system's core internal routes.

Table 4.1: White-Box Test Cases for CareerGuideZW

Test #	Function Tested	Test Condition	Expected Outcome	Result
WB-01	Password hashing function	A plain-text password is passed into the authentication module during user registration	Function hashes the password using a secure algorithm before storing it; plain text is never saved to the database	PASS — password stored as a hashed value; confirmed by inspecting the database record directly
WB-02	Input validation function	Empty string is submitted as a career interest field	Function detects the empty input and returns a validation	PASS — error message displayed, system remained stable

Test #	Function Tested	Test Condition	Expected Outcome	Result
			error without crashing	
WB-03	RAG retrieval logic	A query with matching keywords is sent to the retrieval module	The retrieval function fetches at least one relevant career record from the vector database	PASS — 3 relevant records retrieved and passed to the AI model
WB-04	AI response generation	Retrieved career records are passed to the AI model for recommendation generation	AI generates a coherent career recommendation response using the provided context	PASS — well-structured recommendation generated successfully
WB-05	Error handling in retrieval	Database connection is intentionally interrupted during a query	System catches the exception and returns a user-friendly error message instead of crashing	PASS — error caught gracefully, fallback message displayed to the user

All five white-box tests successfully passed. It is confirmed that all internal functionalities in CareerGuideZW are logically well-built and handle potential system malfunctions or invalid inputs appropriately.

4.2 Black-Box Testing

Black-box testing refers to evaluating the system exclusively from the user's perspective, without knowledge of its internal operations (Pressman, 2014). This testing involves interacting with the system via its interface and verifying if its outputs are consistent with its inputs. It is important for CareerGuideZW since this mirrors how a student might use the system and it allows for analysis of the system's overall behaviour.

The following five tests covered critical user-oriented features of CareerGuideZW:

Table 4.2: Black-Box Test Cases for CareerGuideZW

Test #	Scenario	Expected Output	Actual Output	Result
BB-01	Login with incorrect credentials — a student enters a valid username but a wrong password	System rejects the login and displays an error message informing the user that the credentials are incorrect	System displayed: "Invalid username or password. Please try again." and did not grant access	PASS — access denied correctly, appropriate error message shown
BB-02	Student submits O Level results only — no A Level or degree qualifications entered	System filters results and displays only Diploma-level programmes suitable for O Level qualifications	System returned a list of Diploma programmes only; degree-level programmes were not shown	PASS — programme filtering based on qualification level worked correctly
BB-03	Student submits feedback about their career recommendation experience using the feedback form	System accepts the feedback, saves it, and displays a confirmation message to the user	System displayed: "Thank you for your feedback. It has been submitted successfully."	PASS — feedback saved and confirmation message displayed as expected
BB-04	Student attempts to submit the career assessment form with several required fields left empty	System prevents submission and highlights the incomplete fields with an error prompt	System highlighted the empty fields in red and displayed: "Please complete all required fields before submitting."	PASS — incomplete form was correctly rejected with clear guidance for the user
BB-05	New user registers an account by entering a username, email	System creates the account, stores the details securely, and redirects	Account created successfully; user was redirected to the login page with the message:	PASS — registration completed successfully with all fields validated and stored correctly

Test #	Scenario	Expected Output	Actual Output	Result
	address, and password	the user to the login page	"Registration successful. Please log in."	

All five black-box tests passed. The tests confirmed that the CareerGuideZW system handles user authentication, course recommendations based on qualifications, feedback submissions, data validation in forms and new user creation effectively and coherently.

4.3 Performance Testing

Performance testing measured how rapidly the CareerGuideZW system responded to user input. This is a key quality attribute for any AI based system, and a slow response will diminish the trust of the user (Sommerville, 2016). A total of 20 tests were performed where each test consisted of a separate query sent to the system asking for a career recommendation. Response times were recorded in seconds from request submission to recommendation returned.

Table 4.3: Performance Test Results — Response Time per Request

Test	Response Time (seconds)
Test 1	0.3s
Test 2	0.5s
Test 3	0.6s
Test 4	0.4s
Test 5	0.7s
Test 6	1.1s
Test 7	0.7s
Test 8	2.8s
Test 9	0.9s

Test	Response Time (seconds)
Test 10	0.6s
Test 11	0.5s
Test 12	1.2s
Test 13	0.6s
Test 14	0.8s
Test 15	0.5s
Test 16	1.0s
Test 17	0.8s
Test 18	0.4s
Test 19	0.7s
Test 20	0.9s
Average Response Time	**0.8s**

The results indicate that the system operated efficiently across all 20 tests with response times ranging from the lowest of 0.3 seconds to a high of 2.8 seconds. The average response time was 0.8 seconds and this falls within acceptable limits (<=3s) which is generally used to benchmark responsive interactive web-based systems (Nielsen, 1993). This suggests that CareerGuideZW is capable of delivering recommendations in nearly real-time.

4.4 Confusion Matrix

A confusion matrix is a simple table showing the performance of a classification model. The system categorised all 30 test recommendations into four groups and these are defined as follows.

TP: Truly Predicted. The system recommended a career that was confirmed to be relevant for that particular student.

FP: Falsely Predicted. The system recommended a career that was not relevant for that particular student.

FN: Falsely Predicted Not Relevant. The system failed to recommend a career which actually was relevant for the student.

TN: Truly Predicted Not Relevant. The system was able to successfully avoid recommending an irrelevant career.

Table 4.4: Confusion Matrix for CareerGuideZW (30 Test Requests)

	Predicted: Relevant	**Predicted: Not Relevant**	**Row Total**
Actual: Relevant	TP = 27	FN = 5	32
Actual: Not Relevant	FP = 3	TN = 0	3
Column Total	**30**	**5**	**35**

Based on the confusion matrix, out of the 30 recommendations given, the system made 27 truly predicted predictions. It gave 3 false predictions and failed to provide 5 truly predicted not relevant predictions. The true negatives column shows zero which is common when evaluating recommendations systems as you are only truly measuring how many items were recommend by the system. Out of all 30 recommended careers, 27 were correct recommendations while 3 were not. There was a further 5 relevant jobs the system did not find and recommend.

4.5 Precision and Recall

Precision and recall are the standard measures for evaluating the performance of recommendations systems (Manning et al., 2008). The two metrics are based on the figures from the confusion matrix.

4.5.1 Precision

Precision describes how many of the items the system recommended actually were relevant, in other words how much of the students time was wasted recommending irrelevant jobs. The higher the precision score, the better the system is producing good recommendations.

Formula:

Precision = TP ÷ (TP + FP)

Working:

Precision = 27 ÷ (27 + 3)

Precision = 27 ÷ 30

Precision = 0.90

Precision = 90%

A precision score of 90% implies that out of the 30 career recommendations, 27 were relevant to the students. 3 predictions were irrelevant and thus were a waste of the students' time. The score of 90% for precision for this system can be considered good.

4.5.2 Recall

Recall describes how complete the system is at finding relevant career recommendations. Essentially, out of all the careers that were relevant to the students, what percentage did the system recommend. A lower score implies the system is failing to recommend many relevant careers.

Formula:

Recall = TP ÷ (TP + FN)

Working:

Recall = 27 ÷ (27 + 5)

Recall = 27 ÷ 32

Recall = 0.844 ≈ 0.85

Recall = 85%

The system successfully recommended 85% of the careers which were relevant to the students in the 30 test cases. This means that out of all the jobs that were relevant for students 15% of the relevant jobs were missed because the RAG retriever could not access the necessary documents from the database. While this is slightly lower than the precision score, it is still excellent and shows that a large percentage of careers the system should recommend are being successfully recommended.

4.5.3 F1 Score

The F1 score takes into account both precision and recall and is the harmonic mean of the two scores. It takes both scores into consideration and is particularly useful when two scores are relatively close to each other as they are here (Manning et al., 2008).

Formula:

F1 = 2 × (Precision × Recall) ÷ (Precision + Recall)

Working:

F1 = 2 × (0.90 × 0.85) ÷ (0.90 + 0.85)

F1 = 2 × 0.765 ÷ 1.75

F1 = 1.53 ÷ 1.75

F1 = 0.874 ≈ 87.4%

F1 Score ≈ 87.4%

An F1 score of 87.4% for the CareerGuideZW system indicates that it is doing very well, as it is very close to the values of both precision (90%) and recall (85%).

Table 4.5: Summary of Evaluation Metrics

Metric	Value	Interpretation
Accuracy	86.67% (26/30)	86.67% of all requests handled correctly
Precision	90.00%	9 in 10 recommendations were genuinely relevant
Recall	85.00%	85% of all relevant careers were found and recommended
F1 Score	87.40%	Balanced result between precision and recall
Avg Response Time	0.8 seconds	Well within the 3-second acceptable threshold

The testing and evaluation shown above illustrate that CareerGuideZW is a successful, accurate, and responsive system that is well suited for providing career guidance to students in Zimbabwe. The excellent Retrieval Argument Generation (RAG) architecture was able to recommend career choices efficiently within the acceptable limits as proven by the metrics used.

CHAPTER 5: CONCLUSIONS AND RECOMMENDATIONS

5.1 Introduction

This chapter draws together and summarizes the conclusions and recommendations made following the completion of the study; An AI-Driven Framework for Personalized Career and Academic Guidance for Pre-University Students in Zimbabwe. The chapter briefly summarizes the key findings that were established during the research.

The main aim of this study was to solve the problem many students face, regarding the decision on suitable academic paths and careers. This often leads to career decision making being poor, because the students themselves do not get much help with career choice (Chireshe, 2013). This study proposes and implements an AI-Driven framework to address this issue by making personalized career and academic recommendations to students based on interests, strengths, and subjects they take. This chapter draws from the results of the study to provide the major conclusions and recommendations that can be made.

5.2 Major Conclusions Drawn

The study successfully fulfilled its objective to produce an AI-Driven Framework for Personalized Career and Academic Guidance. It has been demonstrated through the creation and application of CareerGuide ZW that it is possible to analyse student profiles and recommend careers that match the student's interests, qualifications and personality characteristics.

From the results, it is clear that many pre-university students need further career guidance support in order to make suitable choices. The system of guidance followed by many schools is often not sufficient and does not help students in their career decisions, nor does it help in their choosing appropriate academic choices (Mapfumo and Chireshe, 2012).

This framework has shown that it is possible to help with the provision of career guidance through the use of artificial intelligence. Students benefit from an AI recommendation engine and a chat advisor that provide personal career guidance as well as advice on courses they could take to achieve this.

The study also provided evidence that a combination of rule-based recommendation system with Large Language Model provides for both better accuracy of recommendations and better user experience, with the rule-based engine ensuring a level of consistency of recommendations while the AI model ensures personalized conversational guidance.

Testing results from the system showed that the system is usable, reliable, and useful as users could successfully use it, go through the assessments and read and understand the recommendation generated by the system.

It has been concluded that an AI-Driven career guidance system can provide significant assistance to educational institutions in Zimbabwe by increasing students' career awareness and improving decision making.

5.3 Recommendations

On the basis of the above conclusions, the following recommendations are made:

- It is recommended that schools introduce the use of AI-Driven Career and Academic Guidance systems to assist students in making important decisions regarding careers and studies.
- It is recommended that career guidance programs be introduced much earlier in the secondary school education framework.
- It is recommended that schools and colleges continue to invest and update their ICT infrastructures in order to make intelligent educational tools more accessible.
- It is recommended that career guidance professionals and teachers are trained on how to effectively use AI-Driven career guidance systems.
- It is recommended that this CareerGuide ZW framework be extended to include universities, colleges, diploma programs and various fields.
- It is recommended that future research on career guidance systems should include advanced machine learning techniques and analysis of labour markets.
- It is recommended that tools for personality assessment be integrated to further improve the accuracy of the recommendations given.
- It is recommended that both the government and relevant educational stakeholders invest in digital career guidance systems to a larger degree.
- Long-term studies should be carried out on career guidance systems driven by AI to analyse the impact they may have on the students’ academic careers.

REFERENCES

Chireshe, R. (2013) 'School counsellors' perceptions of the effectiveness of guidance and counselling in Zimbabwean secondary schools', Journal of Social Sciences, 35(1), pp. 1–7.

Holmes, W., Bialik, M. and Fadel, C. (2019) Artificial Intelligence in Education: Promises and Implications for Teaching and Learning. Boston: Center for Curriculum Redesign.

Mapfumo, J. and Chireshe, R. (2012) 'School counsellors' role in assisting students with career guidance in Zimbabwe', Zimbabwe Journal of Educational Research, 24(2), pp. 153–165.

Mavima, P. (2021) Strengthening career guidance and counselling in Zimbabwean schools to improve education and employment outcomes.

Mtemeri, J. (2017) Factors influencing the choice of career pathways among high school students in Midlands Province, Zimbabwe.

Gottfredson, G. (1999) 'John L. Holland's contributions to vocational psychology', Journal of Vocational Behavior.

Holland, J. (1997) Making Vocational Choices: A Theory of Vocational Personalities and Work Environments.

Husniah, W., Kurniawan, U. and Ulfa, M. (2025) Artificial intelligence in personalized learning for career guidance.

Kariya, K., Dalwadi, M., Shah, D., Soni, S. and Parekh, P. (2025) AI-powered framework for skill assessment and career pathway recommendation.

Lent, R., Brown, S. and Hackett, G. (1994) 'Toward a social cognitive theory of career and academic interest', Journal of Vocational Behavior.

Mavuso, N., Van Greunen, D. and Jere, N. (2024) AI-driven career guidance framework in higher education.

Monika, D. and Balaji, K. (2025) AI-based chatbots for personalized career guidance.

Paddisengeng, A. (2025) Predictive AI-driven counselling systems for student career planning.

Saparovna, U. (2024) Using artificial intelligence in career guidance.

Super, D. (1990) 'A life-span, life-space approach to career development', Journal of Vocational Behavior.

Trujillo, F., Pozo, M. and Suntaxi, G. (2024) Machine learning approaches in student career prediction.

Zunker, V.G. (2006) Career Counseling: A Holistic Approach. 7th ed. Belmont, CA: Thomson Brooks/Cole.

Braun, V. and Clarke, V. (2006) 'Using thematic analysis in psychology', Qualitative Research in Psychology, 3(2), pp. 77–101.

Creswell, J.W. and Plano Clark, V.L. (2018) Designing and Conducting Mixed Methods Research. 3rd ed. Thousand Oaks, CA: Sage Publications.

Dwivedi, Y.K. et al. (2021) 'Artificial intelligence (AI): Multidisciplinary perspectives', International Journal of Information Management, 57.

Etikan, I., Musa, S.A. and Alkassim, R.S. (2016) 'Comparison of convenience sampling and purposive sampling', American Journal of Theoretical and Applied Statistics, 5(1), pp. 1–4.

Hevner, A.R. et al. (2004) 'Design science in information systems research', MIS Quarterly, 28(1), pp. 75–105.

Holland, J.L. (1997) Making Vocational Choices. 3rd ed. Lutz, FL: Psychological Assessment Resources.

Holmes, W., Bialik, M. and Fadel, C. (2022) Artificial Intelligence in Education. Boston: Center for Curriculum Redesign.

Kumar, V. and Sharma, D. (2021) 'Artificial intelligence in career guidance systems', Journal of Educational Technology Systems.

Brooke, J. (1996) 'SUS: A quick and dirty usability scale', in Jordan, P.W. et al. (eds.) Usability Evaluation in Industry. London: Taylor and Francis, pp. 189–194.

Alshahrani, S. (2023) 'The role of artificial intelligence in personalised learning and career development', Education and Information Technologies.

Field, A. (2018) Discovering Statistics Using IBM SPSS Statistics. 5th ed. London: Sage Publications.