Chapter 1

Introduction

"*Begin at the beginning and go on till you come to an end; then stop.* "

Lewis Carroll

This chapter introduces the key concepts and vocabulary used in the rest of the thesis. We start with introducing recommender systems in Section 1.1, followed by a brief description of the existing methodologies used for recommender systems in Section 1.2, followed by the limitations of existing methods in Section 1.4.1. Moving forward, we discuss our motivation for our work in Section 1.5. Next, we highlight the research gaps by defining various research questions in Section 1.6. Towards the end of the thesis, we summarise the main contributions of the thesis in Section 1.7. Finally, Section 1.8 presents the structure of the rest of the thesis.

1.1 Recommender System

Although automated recommendation generation emerged as an independent area of research in the 1990s, recommender systems became more popular in the late 1970s with the work proposed by [1] based on cognitive science. This work uses stereotypes to perform a user modeling task. The work following that, based on information retrieval and forecasting theory might be seen as forming the base for the recommendation system research [2].

Definition- *Recommender Systems (RS) are software tools and techniques used for providing suggestions for items to be of use to a user.*

A recommendation system is a technology that leverages data analysis, machine learning, and artificial intelligence techniques to suggest relevant items, products, content, or services to users. These systems assist users in navigating through vast amounts of information and provide personalized recommendations tailored to their preferences and needs [3]. A recommender system provides users with suggestions so the target user will choose relevant items from a large set of catalogues. Recommendation systems are significant because they may improve user experiences, increase engagement, and improve business performance across various professions. In a recommendation system, item suggestion for the target user has four key aspects:

- **Decide**: Predict ratings for the set of items so that the highest rating item will be preferred.
- **Discover**: Provide suggestions of unknown and relevant items to the users that will be appreciated.
- **Compare**: Provide a ranked list of suggestions or items.
- **Explore**: For target users, find similar items or similar users for a broad domain of item selection.

Recommender system research has strong emphasis on practical and commercial applications. There are various applications of recommender systems, e.g. E-commerce websites like Amazon, Music recommendation systems, e.g. Pandora and Spotify, video recommendation systems like YouTube, and many other domains like health, journal recommendation, and job recommendation.

More formally, the recommendation problem can be formulated as follows. Let U be the set of users available in the dataset and, similarly, I is the set of all the possible items for the recommendations. Let u be the utility function that measures the usefulness of an item i for a target user u. The utility function u can be defined as $U \times I \rightarrow R$, where R is the ordered set of available ratings. The values for the ratings are non-negative integer values within a specific range, ideally 1-5. Then, for the target user u, we select such items $i' \in I$ that maximize the relevance or utility for the user [4].

$$\forall i \in I,\ i'_i = \arg\max_{i \in I} u(i, I)) \tag{1.1}$$

In most recommender systems, the utility of an item is the rating value, or sometimes it is a binary value whether the user likes an item or not. A rating matrix for the movie recommendation system is shown in Table 1.1. In the Table 1.1, the rows represent user names and the column represents movie names and an entry is the rating value given by the respective user to the movie. The entry "-" indicates that the rating information for the respective user and the movie is unavailable. The task of the recommendation algorithm is to calculate the unrated item's rating value for the target user. The rating of the non-rated items can be calculated using various machine learning and deep learning algorithms, reinforcement learning, and multiple heuristics. Based on the rating estimation method, recommender systems are classified into collaborative filtering, content-based, and hybrid methods, further discussed in the section 1.2.

Table 1.1: A Fragment of User-Item Matrix for movie recommendation system

	Star war	KGF	Life of brain	Notorious
Alice	1	4	-	2
Bob	2	-	4	-
David	-	2	5	1

1.2 Classification of the Recommender System

There exist several classifications of recommender systems based on their mechanisms for recommendation generation. Adomavicius and Tuzhilin [4] described a survey for recommender systems and suggested three types of recommender systems based on their recommendation generation mechanism. The most common way to represent and specify the different types of recommender systems is as follows:

1. Content-based recommendation system
2. Collaborative filtering system
3. Hybrid recommendation system

Apart from these common recommendations system, there are also social filtering recommendation systems and context-aware recommendation systems. A music recommendation systems is an application of the recommendation system techniques. Music recommendation systems can also be seen as a subset of recommender system and they have a similar taxonomy to the recommendation system in general. Another broader classification is proposed by Burke and Robin [5]. They identify each recommendation method's input data and the algorithm used. They define five types of recommendation techniques:

1. Collaborative Filtering (CF)
2. Content-based recommendation system (CBF)
3. Hybrid recommendation system (HS)
4. Demographic based recommendation system (DRS)
5. Knowledge based recommendation system (KBRS)

A brief description of these techniques is described in the section 1.3.

1.3 Recommendation Techniques

The core goal of the recommendation algorithms is to predict the valuable items for the target users. In order to predict these items, a recommendation system has to predict the utility value for the non-rated items for the target user; based on that value, the system decides the items' recommendation list for the target user.

1.3.1 Collaborative Filtering (CF)

Collaborative filtering is the most popular type of recommendation technique. Collaborative filtering (CF) is a process of filtering information using machine learning and deep learning methods using user similarities. The basic assumption of CF-based approaches is that similar users will have identical item preferences [6]. The advantage of using CF-based approaches is that these approaches will not need any descriptive information about the user and item. The only required information in CF is the user's past preference history in terms of items rating information [7, 8]. Generally, CF-based approaches are classified into two categories-

- Model-based
- Memory-based

Model-based Approaches: Model-based CF algorithms build models that include all the rating records of users' past preferences, and the estimated models predict ratings for the target users [9]. The two most popular methods for model-based CF algorithms are the clustering methods for collaborative filtering and the aspect-based model for rating prediction [10–12].

Memory-Based Approaches Memory-based CF approaches make rating predictions by combining ratings of selected users or relevant items for the target user [13]. The model's simplicity characterizes this approach since a minimal or no learning phase is involved. These approaches are also called neighbour methods, which are further divided into two sub-categories.

- User-based CF approach
- Item-based CF approach

User-based and item-based CF-based approaches are based on the concept of the nearest neighbourhood. In user-based collaborative filtering, the idea is to calculate the similarity between two users. For the target users, the preferences of similar users can serve as final recommendations [14]. In item-based collaborative filtering, the similarity between items is calculated based on the common consumption of items by the users. The recommendation generation is based on the idea of the user's past preferred items, and items similar to their last preference will be given as recommendations [15].

1.3.2 Content-based Recommendation System

In CBF approaches, the recommendations are generated from similar items a target user has liked. Indeed, the similarity of items is calculated from the items' content description and users' attribute preferences [16]. The recommendation process in CBF approaches is based on three steps: content analyzer, feature extraction, and filtering components. In content analysis, various feature extraction techniques are used for the content description of items. The profile analyzer module collects representative information about users' preferences and tries to generalize the data according to the user profile for recommendation generation. Further, filtering module exploits the user profile to suggest relevant items by matching the profile representation against recommended articles [17]. In music recommendation, CBF approaches are the most effective techniques because music is rich in content description, so using these content descriptions will lead to a good recommendation. Apart from user profile information, we use music content descriptions like music artist information, lyrics, audio data, MIDI data etc., for recommendation generation [18].

1.3.3 Hybrid Recommendation

Previously mentioned CBF and CF-based approaches have their advantages and disadvantages. So combining these techniques for recommendation generation will lead to a hybrid recommendation technique. The reason for combining these approaches is that "hybrid recommendation approaches can provide a more accurate recommendation than a single approach and the disadvantages of one method can be overcome by the other" [5, 19]. Literary surveys identify seven types of hybrid recommendation techniques listed below:

- **Weighted**: The different recommendation models' scores are combined numerically for final recommendation generation.
- **Switching**: The final recommendation is selected from various best-performing models.
- **Mixed**: This model will present final recommendations from various recommendation models in a combined form.
- **Feature Augmentation**: In this recommendation model, features generated from the model are used as a piece of input information for the second model from which the final recommendation will be generated.
- **Feature combination**: In this recommendation model, features are derived from different knowledge sources, and further, these sources are combined for a single recommendation generation model.
- **Cascade**: This recommendation process is more like a ranking process where recommenders are given some priority, with the lower priority ones breaking ties in the scoring of the higher ones.
- **Meta-level**: In this type of recommendation generation, the model generated by one of the recommendation systems are used as the input for other recommendation models.

1.3.4 Demographic-based recommendation system

The basic idea behind these recommendation techniques is that the recommendation should be generated based on the user's demographic information like country of origin, language, age, etc. These are simple yet effective recommendation methods because demographic features like language are important in music recommendation generation. In web-based recommendations, the demographic information of users is easily gathered from the country or region [20, 21].

1.3.5 Knowledge-based recommendation system

A knowledge-based recommendation system uses a user's past preferences with some specific information or answers to queries made to the user. The users are prompted with queries which they answer. For example, in a music recommendation system, the recommendation for the target user may be generated for a specific artist or for a particular genre as given by the user [22, 23].

1.4 Music Recommendation System

Music is one of the most popular media available over the Internet to entertain users. Music is one of the best ways to express ideas and emotions through sounds and lyrics. There is a vast amount of digital data for music, and their associated data representations are available, which describe music at various semantics levels [24]. Typically music data are available in multiple digital formats like lyrics, MIDI, audio data, video, etc. Music can also be categorized into genres, such as pop, rock, jazz, blues, folk, etc. Nowadays, the availability of digital music is very abundant compared to the previous era, so sorting out all this digital music is time-consuming and causes information fatigue. A music recommender system addresses this problem by automatically searching music libraries and suggesting suitable songs for users [25]. A music recommendation system is a particular kind of recommendation system that concentrates on recommending songs, albums, playlists, or artists to users based on their musical tastes, listening history, and other pertinent criteria. These systems examine user data and employ several methods to produce personalised music recommendations that help listeners find new music and improve the listening experience. Many applications are available based on music recommendations to suggest to users their customized playlists, like Spotify [26], Pandora [25], Apple Music [27] etc. A music recommender system is meant to help the users to filter relevant songs from a wide range of music available over the Internet. A good music recommender system should be able to give appropriate music recommendations to a user based on their preferences automatically. Automatic music recommendation is a challenge for researchers to develop algorithms that detect users' moods and priorities with minimal user interaction. Also, creating such a recommender system will also provide an excellent opportunity for the industry to aggregate users interested in music [18]. In order to analyze and compare various songs, feature extraction, which involves extracting significant characteristics from audio signals and text information from lyrics, is an essential component of music recommendation systems. These extracted features provide valuable insights into songs' musical content and enable the system to make accurate recommendations based on user preferences. Here are some commonly used feature extraction techniques in music recommendation systems:

- Spectral Features: In a music recommendation system, spectral features like the Short-Time Fourier Transform (STFT) and Mel Frequency Cepstral Coefficients (MFCCs) are used. A music recommendation system looks into these features. Short-Time Fourier Transform (STFT) is used to examine the frequency content of the audio signal. In order to represent the spectral content of the song, it determines the magnitude of the frequencies present in the signal at each time window. The audio signal's spectral properties are captured by MFCCs, simulating how sound is perceived by the human auditory system. They are frequently used to represent the timbral characteristics of songs in music analysis.
- Rhythm and Tempo Features: These techniques identify a song's rhythmic structure and tempo. Tempo analysis determines the music's speed or tempo, while beat tracking algorithms can identify the beats and rhythmic patterns. Recommending songs with similar rhythmic qualities is made easier with the aid of rhythm and tempo features.
- Pitch and Harmony Features: Pitch analyzes the audio signal's frequency content

to identify the song's pitch or musical notes. Chroma features represent the distribution of pitch classes or musical chords in a song, ignoring the octave information. They provide a condensed representation of the harmonic content of the music, facilitating comparisons and similarity calculations.

- Timbral and Texture Features: Timbral and texture features in music describe tone colour, sound quality, and texture characteristics. These characteristics accurately represent the distinctive qualities that set one musical instrument or sound source apart from another. Timbral and texture features are important because they offer information about the sonic characteristics of songs, allowing for more individualized and nuanced recommendations.

These are merely a handful of instances of feature extraction methods used in music recommendation systems. The precise set of features used may vary depending on the system requirements and the properties of the music dataset. In order to give users personalized recommendations based on their musical preferences, recommendation systems can compare and analyze songs.

1.4.1 Limitations of Music Recommender System

The current state of the research and development in recommendation systems has undoubtedly contributed to user satisfaction and increased business revenue. Still, many open issues and limitations of the recommender system are present that challenge the limit of usefulness of the recommendation systems. This section discusses some well-known limitations of the music recommendation systems.

- **Cold-start Problem**: One of the main problems of certain kinds of general recommender systems or music recommendation systems is the cold-start problem [28, 29]. The recommender system based on collaborative filtering, where the recommendations are generated based on the similarities of users or items, face such a problem. When a new user or item is introduced into the system, the system will not generate recommendations for these users or items due to insufficient data [30, 31]. The cold-start problem is also known as the new user problem or new item problem.

- **Sparsity**: This is a sub-problem of the cold-start problem for CF-based systems, but it can also occur in other kinds of recommender systems. Sparsity is when the number of available ratings is much lower than all possible ratings, which is particularly required for accurate recommendation generation [32,33]. In general, information on user-item rating value is much less than the available user-item interaction pairs. The high sparsity in the system is responsible for the unreliable recommendation generation for the user and low rating coverage in the system.
- **Grey-sheep Problem**: When a group of users has specific tastes or some unusual preferences, then recommendation generation for these users is not convenient in the CF-based approach. The grey-sheep problem exists if users have low correlation coefficients with other users or who have unique or unconventional preferences that don't align with the majority. It is also a variation of the cold start problem in recommendation systems [34].
- **Diversity**: In the CF-based approach, recommendations are generated based on similarity, but generating a similar recommendation every time will reduce the system's relevance. This is known as a diversity issue in the recommendation system. For example, generating similar or repetitive music recommendations in a music recommender system will make the system monotonous [35]. Diversity in recommendation systems refers to including a broad range of items or content in the suggested recommendations.
- **Scalability**: This problem impacts our real-life scenario because the exponential growth of data over the web leads to scalability issues in real-time-based recommendation handling of the growing data and the potential to accommodate these data to the system in recommendation generation.

- **Popularity Bias Problem**: Popularity bias is a significant concern in the music recommendation system or any other general recommendation problem. Popularity bias occurs when recommender systems recommend popular items to users frequently compared to less popular items [36].

1.5 Motivation

In a music recommendation system, users' past interaction history and content information are used for recommendation generation. The user information (e.g., personality, emotional states, activities, social environment) and the the music information (e.g., music audio content, novelty, diversity) influences the nature of recommendations. Earlier recommendations were generated using basic collaborative filtering methods, assuming that like-minded people prefer similar recommendations. One of the problem with these systems is the cold start users and items where the recommendation generation for the new user and item is tricky. To overcome the limitations of collaborative filtering by using content information of users and items called a content-based recommendation system. In music, content information is audio, video, lyrics, artist information, etc. Researchers use this content information as additional input information with the user's preference history for better recommendation generation. We require lots of data for input processing for an accurate recommendation prediction, so using audio and video information requires lots of space and computational resources [37]. This thesis aims to address such problems.

The goal of our thesis is two-fold. Besides considering only accuracy for our recommendation system, we try to include diversity in recommendation generation. The reasons for including diversity in recommendation models are listed below:

- Earlier approaches suffer from drawbacks like redundancy where recommending items based on the user's past interaction history limits their interaction domains, i.e., a reduction of diversity in the model.

- Diversity in the recommendation provides different but interrelated perspectives that will affect the recommendation system's quality and user satisfaction by generating new and exciting items in the recommendation list for the target user.

1.6 Research Goals

The main objective of the research contribution presented in this thesis is to adopt some techniques to develop a music recommendation system with improved accuracy and diversity in recommendation generation. We mainly focus on two aspects of music recommendation systems during as part of the work leading up to this thesis. First is the enhanced performance of music recommendations, and in the second we focus on the diversification of the recommendation system. There are many issues with the current music recommendation systems which still poses challenges. We explore some problems and try to find answers to some research problems related to these issues in our work. Our research goals can be listed as:

- **RG1**: *How to handle the cold-start problem for an accurate music recommendation system using MIDI data?*

  To deal with a cold-start issue, i.e., recommendation generation for a new user and a new song in the music recommendation system. We proposed an MSA-SRec model to overcome this issue, where music's content information (MIDI) is used with users' past interaction histories.

- **RG2**: *How to handle cold start problem with limited use of personalized data in the music recommendation system?*

  In recommendation systems research, a lot of demographic data has been used to address the cold-start problem, which creates some privacy concerns in recommendation generation. The usage of song content information will significantly aid in resolving this problem and improve recommendation accuracy. Our proposed model MSA-SRec uses a combination of MIDI data and users' interaction history to address this concern.

- **RG3**: *How to handle the sparsity issue in music recommendation system?*

  Data sparsity is a significant issue in music recommendation systems where we have many missing users and item interaction information, which will cause lost entries in the user-item interaction matrix. So to avoid sparsity in our model, instead of considering only the user-item interaction matrix, we used other information, like user and item content information, in recommendation generation.

- **RG4**: *Explore various applications of theories and methodologies from IR (Information Retrieval) to include diversity in the recommendation system.*

  Generating only accurate and similar recommendations from users' past preferences will make the system monotonous. This is known as a diversity problem in the recommendation system. We explore many available research proposals for diversity inclusion in the recommendation system. In prior research, diversity in the recommendation system has often been achieved using post-processing methods, where the final recommendation list is reshuffled according to the dissimilarity context. We put forth two models for diversifying recommendation systems using machine learning and deep learning methods. Instead of depending solely on post-processing techniques, we looked into adding pre-trained models for diversification in the recommendation system.

- **RG5**: *How to maintain accuracy-diversity trade-off in the recommendation system?*

  In a recommendation system, while including diversity, the model's accuracy suffers due to the accuracy diversity tradeoff, a significant issue for recommendation systems. We were able to achieve diversity in recommendations in the ClusDR model. However, accuracy remained a concern in our model. To address this, we propose the Div-HERec model, which is based on a graph neural network (GNN), and allows us to get a better trade-off between accuracy and diversity.

1.7 Contributions

This thesis is committed to developing an effective and diverse music recommendation system. The main contributions of this thesis are related to the design, implementation, and evaluation of the performance of recommender systems, where we have addressed several existing issues using novel models and methods. The contribution of this thesis is twofold. First, To improve the recommendation system's accuracy and address the cold start issue in music recommendation systems, one could start exploring music content information, such as MIDI. Second, we'll attempt to provide a solution for the diversity problem with recommendation generation. We have proposed three novel models for recommendation generation. Our first model is for an accurate music recommendation system, and our other two models are for a recommendation system that takes diversity into account. A detailed description of our main contributions is briefly discussed below in terms of the chapters of this thesis:

1.7.1 Deep Content-based Music Recommendation System

This chapter describes a model for music recommendation generation using content information of music for better recommendation generation. We proposed a novel deep content-based music recommendation system using deep learning models. Earlier in the music recommendation system, the user's past interaction history was considered for recommendation generation, which led to a cold start problem. To solve this issue, content-based music recommendation has been proposed, which uses the content information of music for recommendation generation. Music's audio, video, or artist information has been used for recommendation generation in other state of the art methods. The novelty of our proposed approach is the exploration of new content information using MIDI (Musical Instrument Digital Interface), a compressed version of audio information used for recommendation generation. The following methodology for music recommendation is based on a multimodal concept where we used lyrics information for recommendation generation, apart from MIDI information. The critical contributions of this work are as follows:

- The key contribution of this model is to mitigate the cold-start problem of recommendation systems.
- We integrate the content information of the song in terms of MIDI data and user-item interaction data to resolve the cold-start recommendation problem and reduce reliance on personalization in the recommendation system.
- We transform MIDI data into multidimensional time-series data to use it as a piece of sequential information for use in deep learning models and also manage the dynamicity of users' preferences.
- We also discuss the importance of MIDI as content information in music recommendation systems. We present a comparative results that prove that MIDI leads to accurate and efficient sytems, compared to other music metadata.
- We also presented a comprehensive experiment over a real-world dataset, i.e., the million song dataset, to evaluate the performance of our proposed model MSA-SRec and multimodal music recommendation system. The proposed model outperforms other state-of-the-art methods for recommendation systems.

1.7.2 Clus-DR: Cluster-based Diversity Recommendation System

The primary objective of our proposed thesis includes some principled approaches to evaluating and enhancing the diversity of the recommendation model. We also tried to balance the tradeoff between accuracy and diversity. The improvement of such fundamental dimensions of the usefulness of recommendations has to consider how users explore and perceive recommendations, the problems that novelty and diversity solve, and the causes of such issues. We have pursued the following research goals and proposed a novel methodology for diversity enhancement in the recommendation system. The critical contribution of our proposed approach is as follows:

- To deal with the diversity issue in the recommendation system, we proposed a method called Clus-DR, which is a pre-trained model for diverse recommendation generation.
- We propose an individual diversity score (IDS) and an individual feature score (IFS) for each user, calculated from users' past preferences. These scores are two decisive criteria for each user's inclusion in their recommendation list.
- Next, we include a cluster-based algorithm to group users. The basic strategy for user clustering is to quantify more diverse users using their past interaction histories. We used two clustering strategies for users: (1) individual diversity of the user and (2) using a content feature of the item.
- Our proposed Clus-DR model involves a tunable hyper-parameter (ILD-Score/IFS-Score) that users can use to adjust the diversity levels of their recommendation list based on their past item interactions. The goal of model training for target users with more aggregate diversity is to include interaction domains among diverse users to increase the target user's exploration domain.
- We also present extensive experiments on three real-world datasets to demonstrate our Clus-DR model's effectiveness with various baseline algorithms for diversification in the recommendation system.

1.7.3 Graph-based Diverse Recommendation System

In this section, we proposed a model for the diversity problem of a recommendation system where improving diversity in personalized recommendations while maintaining accuracy is challenging [38]. While introducing diversity in the model, the model will avoid suggesting items similar to their last preferences [39]. The diversity in earlier models has been introduced using post-processing and optimization methods which will suffer from the accuracy-diversity tradeoff. To maintain this tradeoff, various approaches have been proposed [40–44]. The general framework for most of the methods for diversity inclusion in the system is to include similar as well as dissimilar users while modelling the interaction of the target user. We believe that to acquire an acceptable accuracy-diversity tradeoff in the system, we need an approach to capture the irregular and complex user-item interactions that are most relevant and somehow understated for the target user but not irrelevant. We proposed a methodology based on the metapath2vec node representation method and graph colouring algorithm to better address this tradeoff. The path selection for users in our proposed model is based on the graph colouring algorithm. Unlike the traditional metapaths algorithm, where the target user's neighbour is chosen randomly, we select a neighbour of the target user using a graph colouring algorithm. The main contributions of this work can be summarized as follows.

- We propose a novel accuracy-diversity trade-off framework for recommendation systems using HIN. The model operates on a diverse embedding representation of user and item using the metapath2vec algorithm.
- We propose a novel Div-HERec framework to improve the recommendations model's accuracy and diversity. As far as we are aware, we are the first to include GNN with a graph colouring algorithm on HIN for diversified recommendations generation.
- The proposed model includes pairwise preferences of all users for accurate and diverse item recommendations. A user's diverse pair is formed using a graph colouring algorithm over the user-user weighted homogeneous graph. After that, we use a metapath algorithm to generate diverse embedding representations of users and items in a heterogeneous graph neural network.
- We convert the implicit feedback information into pairwise preferences of users, so our proposed approach can be adapted to other domains, as we are only using

the user-item interaction history for recommendation generation.

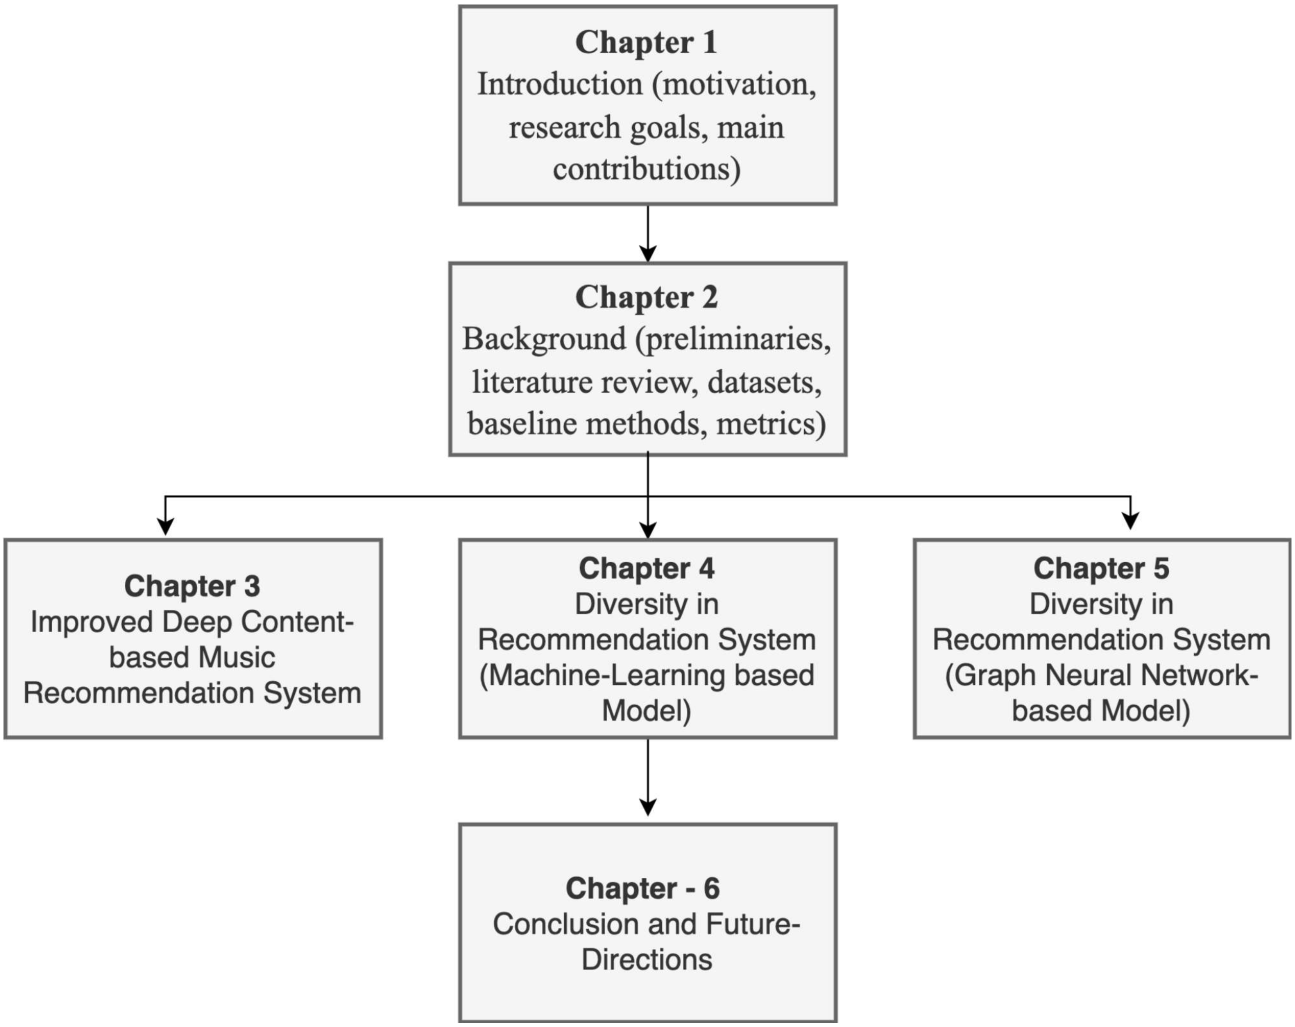

Figure 1.1: Illustration of thesis structure

1.8 Structure of the Thesis

The thesis consists of six chapters. Figure 1.1 provides an overview of the thesis and interconnection among chapters.

- **Chapter 1**: This chapter gives a brief introduction to recommender systems and their types. We also discuss motivation, research goals, and our novel contributions to the recommender system research.

- **Chapter 2**: This chapter is focused on the background and literature survey of the music recommender system and a state-of-the-art discussion on the diversification methodologies for the recommendation system. We also include theoretical knowledge for our research and different evaluation methods used in our proposed models. At the end, we have a detailed description of the datasets used in our model's performance evaluation.
- **Chapter 3**: This chapter explains our first contribution to deep content-based music recommendation using MIDI and song lyrics information using two different models. These proposed models are for an improved music recommendation system. It also describes the methodology for transforming MIDI and lyrics information into a deep learning model.
- **Chapter 4**: The model's accuracy when generating recommendations is primarily considered in Chapter 3. Ideally, accurate and diverse recommendation generation would be well-balanced. In Chapter 4, we began including diversification in the recommendation generation, but our model's precision suffered while considering diversity. In this chapter, we propose pre-trained approaches based on machine learning as a way to diversify the recommendation system.
- **Chapter 5**: In Chapter 4, diversification is considered in the recommendation system and the model's accuracy is less importance while generating recommendations for the target user. To overcome this, in Chapter 5, we proposed a model for maintaining a good tradeoff between accuracy and diversity. We proposed a model based on GNN that would increase the model's diversity while trying to maintain its accuracy. Earlier in this thesis, we started by mainly exploring music, but simultaneously we explored some ideas, which led to a consideration of a generalised framework for diversification in a recommendation that can be applicable to other domains as well.
- **Chapter 6**: This Chapter mentions the conclusions derived from work on this

thesis, along with some future work directions.