1
Introduction

Machine learning encompasses a broad range of problems ranging from detecting objects in images, finding documents relevant to a given query, or predicting the next element in a sequence, among countless others. Traditional approaches to these problems operate by collecting large, labeled datasets for training, uncovering informative features, and mining complex patterns that explain the association between features and labels. Typically, labels are regarded as an underlying 'truth' that should be predicted as accurately as possible.

Increasingly, though, there is a need to apply machine learning in settings where the 'correct' outcome is subjective, or otherwise depends on the context and characteristics of individual users. As we browse online for movies to watch, products to buy, or romantic partners to connect with, we are likely engaging with these new forms of *personalized* machine learning: Results are tailored to us specifically, based on the types of movies, products, or partners that *we specifically* are likely to engage with.

Much like traditional machine learning algorithms, *personalized* machine learning algorithms are at their heart essentially forms of pattern discovery. That is, predictions are made *for you* by analyzing the behavior of people *similar to you*. A recommendation such as 'people who liked this also liked' is perhaps the most simple example of this type of personalized pattern discovery:[1] based on the contextual attribute of a user liking a particular item, recommendations are extracted based on users who share this common preference. At the other end of the spectrum are complex deep learning approaches that learn 'black box' representations of users in order to make predictions, though these too at their heart rely on the intuition that 'similar' users (in terms of some complex representation) will have similar interaction patterns.

[1] Though strictly speaking maybe not one that we would call 'machine learning.'

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

1.1 Purpose of This Book

We seek to introduce *Personalized Machine Learning* by exploring a family of approaches used to solve the aforementioned problems, and construct a narrative around the common methods and design principles involved. We show that even in applications as diverse as song recommendation, heart-rate profiling, or fashion design, there is a common set of techniques around which personalized machine learning systems are built.

By introducing this underlying set of principles, the book is intended to teach readers how modern machine learning techniques can be improved by incorporating ideas from personalization and user modeling, and to guide readers in building machine learning systems where accurately modeling the users involved is key to success.

There is currently an abundance of models, datasets, and applications that seek to capture human dynamics or interactions. Examples pervade in diverse areas including web mining, recommender systems, fashion, dialog, and personalized health, among others. As such, there is an emerging set of techniques that are used to capture the dynamics of 'users' in each of these settings. This book is designed to act as a reference point to explain these techniques, and explore their common elements. As a starting point, we will begin the book (chaps. 2 and 3) with a primer of machine learning (and especially supervised learning) that will bring readers up-to-speed on the basic techniques required later. Although this introductory material is likely familiar to many readers, we have a particular focus on user-oriented datasets, and show that even with 'standard' machine learning techniques, there is considerable scope for building personalized systems through careful feature engineering strategies that capture relevant user characteristics.

Following this, our main introduction to personalized machine learning will be to explore recommender systems (chaps. 4 and 5). Recommendation technology has traditionally relied on personalization and user modeling, whether through simple similarity functions among users ('people like you also bought', etc.) or through more modern approaches involving temporal pattern mining or neural networks.

More recently, the need to account for personalization and to model users has spread into a variety of new areas of machine learning. Following our study of recommender systems, exploring personalization and user modeling in these new areas—and giving readers the tools they need to design personalized approaches in new settings—is the main goal of this book.

1.2 For Learners: What Is Covered, and What Is Not

Although this book is primarily intended as a guide to the specific topics of personalization, recommendation, and user modeling (etc.), it should also serve as a relatively gentle introduction to the topic of machine learning in general. Topics such as web mining and recommender systems serve as an ideal starting point for learners seeking a more 'application oriented' view of machine learning compared to what is typically covered in introductory machine learning texts.

Throughout the book, we focus on building examples on top of large, real-world datasets, and exploring techniques that are practical to implement in projects and exercises. Our particular focus guides us toward (and away from) certain topics, as we describe below.

Regressors, Classifiers, and the Learning Pipeline: We give a detailed introduction of the end-to-end machine learning process in Chapters 2 and 3, which (while condensed) should be suitable for learners with no background in machine learning. When introducing basic machine learning concepts in Chapters 2 and 3, we limit ourselves to linear regression and linear classification (logistic regression), since these serve as building blocks for the methods we develop later. Consequently, we ignore dozens of alternative regression and classification methods that are often the core of standard machine learning texts (though we briefly discuss the merits of alternatives in sec. 3.2).

User Representations and Dimensionality Reduction: Many of the techniques we explore when learning user representations are essentially forms of manifold learning (or dimensionality reduction), and borrow ideas from related topics such as matrix factorization (sec. 5.1). While readers should have some basic familiarity with linear algebra, we carefully avoid a linear algebra-heavy presentation of 'traditional' dimensionality reduction techniques: in terms of actual implementation, these have little in common with the methods we develop (though we discuss the connection to, for example, the singular value decomposition in sec. 5.1).

Deep Learning: Any discussion of 'modern' machine learning approaches necessitates a fairly broad discussion of deep learning. For example, we discuss multilayer perceptron-based recommendation in Section 5.5.2, sequence models based on recurrent neural networks in Chapter 7, and models of visual preferences based on convolutional neural networks in Chapter 9. However in doing so we are merely scratching the surface of deep learning-based personalization, and will largely refer readers elsewhere for in-depth discussion of specific architectures, or for a first-principles presentation of deep learning methods.

Offline versus Online Learning: We largely limit ourselves to traditional offline, supervised learning problems, that is, uncovering patterns and making predictions from historical collections of training data. Generally, we prefer this setting since it allows us to focus on methods that we can develop on top of real-world, publicly available datasets. Of course, in practice, when deploying predictive models, data may be obtained in a streaming setting and updates must be made in real time. This type of training regime is known as *online learning*, which we briefly cover in Section 5.7; we also avoid discussion of (e.g.) reinforcement learning algorithms, though mention their use briefly in settings such as conversational recommendation (e.g., sec. 8.4.4).

Bias, Consequences, and User Considerations: By design, our study of personalization is largely confined to machine learning approaches. That is, we are generally concerned with building predictive systems that can estimate—as accurately as possible—how a particular user will respond to a given stimulus. By doing so, we can estimate preferences, predict future activities, retrieve relevant items, and so on.

Of course, we are mindful of the dangers associated with 'black-box' approaches to machine learning, and want to avoid the pitfalls of blindly optimizing model accuracy, such as filter bubbles, unwanted biases, or simply a degraded user experience. In Chapter 10 we discuss these issues, as well as potential approaches to address them.

Again our discussion is mostly limited to machine learning solutions, that is, we investigate *algorithmic* approaches to correct for biases, increase recommendation diversity, and so on. We note that algorithmic solutions are only part of the picture, and that while having better algorithms is critical, it is also critical that those algorithms are appropriately *used*. Our presentation is complementary to a large body of work that explores personalization from the perspective of human computer interaction, or user interface design, where the primary concern is maximizing the quality of the user experience (ease of finding information, satisfaction, long-term engagement, etc.).

Implementation and Libraries: All code examples are presented in *Python*. While we assume a working familiarity with data processing, matrix libraries (etc.), further links on our online resources page (sec. 1.4) will help users with less familiarity. When discussing deep learning approaches, and more generally when fitting complex models, we base our implementations on *Tensorflow*, though these examples can easily be interchanged with alternate libraries (PyTorch, Theano, etc.).

While we focus on implementation, we largely avoid 'systems building' aspects of personalized machine learning, such as concerns around deploying machine learning models on distributed servers (etc.), though we discuss high-level libraries and implementation of best practices throughout the book.

1.3 For Instructors: Course and Content Outline

This book is inspired by my own experience teaching classes on recommender systems and web mining at UC San Diego. Courses on these topics have proved extremely popular and are often chosen as learners' first exposure to machine learning.

One reason this topic acts as a good first contact with the machine learning curriculum is that it has a somewhat lower bar for entry than many machine learning courses, including (e.g.) courses on deep learning, or even many 'introductory' machine learning classes. Partly this is due to the material being less dependent on deep and complex theory, and partly it is due to the ability to quickly build working solutions that are fairly representative of the state-of-the-art, rather than mere proofs-of-concept. As such, a focus of this book is to quickly build working solutions, and covering a wide breadth of approaches, rather than diving too deep into the theory behind any one approach. This approach can be useful in helping learners to understand the practical considerations behind building predictive systems based on user data, and is complementary to the more theoretical treatment given in most introductory texts.

Another feature that has made this material popular among learners is the ability to work quickly with large, real-world datasets. The ability to work with collections of user data from *Amazon*, *Google*, *Steam* (etc.) on applications that are representative of real use cases, has proved immensely valuable for students building their project portfolios or preparing for interviews. As such, each chapter is paired with project suggestions, each of which would be suitable as a major class project. These projects aim to synthesize the material from each chapter, with more focus on system building considerations, design choices, and thorough model evaluation.

1.3.1 Course Plan and Overview

The content in this text is aimed at developing a quarter- or semester-long course, for students with some background in linear algebra, probability, and

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

data processing. After revising basic material in Chapters 2 and 3, Chapters 4 and 5 cover the core material upon which the remainder of the book builds. Chapter 6–9 are somewhat more orthogonal, such that components can be selected and combined as time or student background allows. A final chapter on bias, fairness, and the consequences of personalization (chap. 10) provides an opportunity to revisit earlier material through a new lens.

Each chapter is paired with homework and a project. Again the focus on these components is mainly on developing practical implementations, working with real data, and understanding the design choices involved, rather than testing theoretical concepts. Below we briefly summarize the material from each chapter:

Machine Learning Primer (chaps. 2 and 3) 2–3 weeks. Introduces the foundational concepts of machine learning, feature design, and evaluation, via a selection of datasets that capture user interactions. Exercises range from simple data manipulation to building a working machine learning pipeline (training, validation, etc.). Exercises are mainly concerned with feature design, including projects (Projects 1 and 2) that involve experimenting with activity data involving temporal and geographical dynamics.

Recommender Systems (chaps. 4 and 5) 2–3 weeks. Introduces the core set of techniques used for recommendation. Traditional heuristics are presented in Chapter 4 followed by machine learning approaches in Chapter 5. Recommender systems are used to develop the concept of a *user manifold* which is used throughout the following chapters to capture variation among users in several settings (sec. 1.7). Exercises are mainly focused on the basics of building practical recommendation approaches, and projects (Projects 3 and 4) are concerned with building an end-to-end recommendation pipeline for a book recommendation scenario.

Content and Structure in Recommender Systems (chap. 6) 1 week. Explores how to incorporate features (i.e., side information) into personalization (mostly recommendation) approaches, and explores personalization in settings with additional structure, such as socially aware recommendation and settings involving price dynamics. A particular focus is given to leveraging side-information in *cold-start* scenarios, where interaction histories are not yet available (sec. 6.2). Some of these content-aware approaches (such as factorization machines) are revisited later in the book when developing more complex models based on (e.g.) temporal or sequential dynamics. A project

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

(Project 5) consists of developing recommender systems for use in cold-start settings.

Temporal and Sequential Models (chap. 7) 1–2 weeks. We revise some of the basic approaches to temporal and sequential modeling, such as autoregression and Markov chains, and later develop more complex personalized approaches based on recurrent neural networks. The *Netflix Prize* (sec. 7.2.2) is presented as a case study to explore the basic design principles of temporal modeling. A project (Project 6) compares various approaches to temporal recommendation.

Personalized Models of Text (chap. 8) 1 week. After revising some of the basic predictive models of text (such as bag-of-words representations), we explore how text can be used to understand the dimensions of preferences. We revisit sequential modeling by exploring techniques that borrow from natural language to model interaction sequences. We also visit methods for text *generation*, which can be personalized in settings ranging from conversation to justification of machine predictions. A project (Project 7) consists of building personalized systems for document retrieval.

Personalized Models of Visual Data (chap. 9) 1 week. Explores applications involving visual data, ranging from personalized image search, to applications in fashion and design. A project (Project 8) consists of building visually aware recommendation systems for applications in fashion.

The Consequences of Personalized Machine Learning (chap. 10) 1 week. The final chapter explores the consequences and pitfalls of developing personalized machine learning systems. Examples include filter bubbles, extremification, and issues of bias and fairness. The chapter has a significant focus on applied case studies, and allows us to revisit several of the topics from previous chapters through a new lens. A project (Project 9) consists of improving recommendation approaches in terms of gender parity and other fairness objectives.

1.4 Online Resources

To help readers with exercises, projects, and to collect resources including datasets and additional reading materials, an online supplement is available to augment the material covered here with working code and examples:

https://cseweb.ucsd.edu/~jmcauley/pml/

The online supplement includes:

- Code examples covering the material in each chapter. These cover complete worked examples from which the code samples presented in each chapter are drawn. Additional code samples are included that correspond to various figures and examples presented throughout the book.
- Solutions to all exercises from each chapter.
- Links to datasets used in the book (as well as various other personalization datasets), including small, processed datasets useful to complete the exercises.
- Links to additional reading, mostly focused on introductory material useful to learners less familiar with some of the background material described in Section 1.2.

1.5 About the Author

I have been a Professor at UC San Diego since 2014, following postgraduate training at Stanford University, and undergraduate and graduate training in Australia. *Personalized Machine Learning* is the main theme of my research lab at UCSD. Our lab's research has pioneered the use of images and text in recommendation settings (e.g., McAuley et al. (2015); McAuley and Leskovec (2013a)), with applications including fashion design, personalized question answering, and interactive dialog systems. Our lab has also studied personalization outside of typical recommendation settings, such as developing personalized models of heart-rate profiles (Ni et al., 2019b), and systems for generating personalized recipes (Majumder et al., 2019).

Figure 1.1 The author

Our lab regularly collaborates with industry to develop state-of-the-art systems for personalized machine learning. We have worked on problems including visually aware recommendation with *Adobe* and *Pinterest*, understanding user budgets and personalized price dynamics with *Etsy* and *Microsoft*, and question-answering and dialog systems with *Microsoft* and *Amazon*. We will explore several of these approaches through case studies throughout the book.

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

1.6 Personalization in Everyday Life

Other than introducing the techniques underlying personalized machine learning systems, one of our goals in this book is to explore the wide range of practical applications where personalization is applied, to explore the history of the topic, and eventually to explore the associated risks and consequences.

Personalized machine learning is increasingly becoming pervasive to the point that most of us are likely to interact with personalized machine learning systems every day. Systems that generate playlists based on our listening habits, mark e-mails as ‘important,’ suggest products or advertisements based on our recent activities, rank our newsfeeds, or suggest new connections on social media, all personalize their predictions or outputs in some way. Techniques range from simple heuristics (e.g., we’re likely to become friends with somebody if we already share mutual friends), to complex algorithms that account for temporal patterns, or incorporate natural language and visual signals.

Below we will study a few common (and less common) scenarios in which personalization plays a key role, many of which will form the basis of case studies throughout this book.

1.6.1 Recommendation

Many of the examples we cover in this book will relate to *recommender systems*, and more broadly to modeling users’ interactions with data collected from the web. Part of the reason for this focus is opportunistic: user interaction datasets are widely available, allowing us to build models on top of real data.

Pedagogically, recommender systems are also appealing as an introduction to personalized machine learning as they allow us to quickly implement working systems that are close to the state-of-the-art. As we will see, even widely deployed systems turn out to be surprisingly straightforward, relying on simple heuristics and standard data structures (sec. 4.5).

Ultimately though, our main reason for studying recommender systems is because they are a fundamental tool for modeling *interactions between users and items*. The basic techniques developed when building recommender systems can be applied in a variety of other situations where we want to predict how a user will respond to some stimulus. Many of the settings we describe later build on this general theme.

Recommender systems represent perhaps the purest settings where *variation among individuals* captures a large fraction of the variability in a dataset. To build recommender systems we must understand the underlying *preferences*

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

of users and *properties* of items that explain why an item might be purchased by one user and not another. Users might vary due to subjective preferences, budgets, or demographic factors; both users and items might change over time due to social, temporal, or contextual factors (etc.).

Building on the techniques we develop for recommendation, we argue that there are countless settings where capturing variation among individuals is key to making meaningful predictions. In settings like personalized health, users may vary in terms of their physical characteristics, medical histories, or risk factors; or in settings involving natural language (or dialog), users may vary in terms of their writing styles, personalities, or their specific context.

Below we describe a few such examples, partly to highlight the wide range of settings where personalization is critical, but also to demonstrate the common set of ideas involved in modeling them.

1.6.2 Personalized Health

Beyond 'obvious' applications in electronic commerce or social media, personalization is increasingly playing a role in high-stakes and socially important problems. *Personalized health* is a key emerging domain for personalization: like recommendation, problems in health have the key characteristic that predictions are highly contextual and exhibit significant variation among individuals. Critically, when estimating symptoms, responses to medication, or heart-rate profiles, it would be impossible to make useful predictions *without* personalization.

Estimating what symptoms a patient will exhibit on their next hospital visit is a canonical task in personalized health, with applications in (e.g.) preventative treatment. This task closely resembles the settings we explore when developing recommender systems, given the goal to estimate patients' interactions with certain stimuli (symptoms) over time (Yang et al., 2014). As such, techniques for such tasks borrow ideas from recommender systems, especially temporal and sequential recommendation, as we develop in Chapter 7.

Beyond estimating patient symptoms, personalized machine learning techniques can be adapted to related tasks ranging from estimating the duration of surgical procedures (Ng et al., 2017), modeling the progression of heart-rate sequences in response to physical stimuli (Ni et al., 2019b), or estimating the distribution kinetics of drugs (such as anesthetics) (Ingrande et al., 2020). Modeling such problems requires understanding the characteristics of patients or physicians (and the interactions between them). Techniques range from simple regression (e.g., to predict surgery duration) to recurrent neural networks (e.g., to forecast heart-rate profiles).

Many problems in personalized health also depend upon natural language data, for example, modeling the characteristics of clinical notes or generating reports based on radiology images (Ni et al., 2020). Such applications build on techniques for personalized natural language processing and generation, as we develop in Chapter 8.

These techniques span the different 'types' of personalized learning systems (see sec. 1.7): some systems leverage traditional machine learning techniques, in which 'personalization' merely means extracting features that capture the relevant properties about users (or patients, physicians, etc.); others use complex deep-learning approaches, in which the underlying dimensions that capture patterns in behavior are harder to interpret.

1.6.3 Computational Social Science

Often the goal of modeling user data is not merely to *predict* future events or interactions, but to *understand* the underlying dynamics at play. Using machine learning and data-driven approaches to understand the underlying dynamics of human behavior from large datasets is one of the main goals of *computational social science*.

Likewise, for many of the models we develop, our goals are as much about building more accurate predictors as they are about understanding social or behavioral dynamics. When we develop regressors to predict content success on *reddit* (sec. 2.6.1), our main goal is to disentangle what factors lead to success, such as community dynamics, titles, submission times, and so on. Or, when building recommender systems our goals are to understand and interpret the underlying preference dimensions that guide users' decisions, and what causes those preferences to change over time, including how users acquire tastes, develop nostalgia for old items, or simply respond to changes in a user interface.

Finally, as we begin to explore the ethical consequences of personalization (which we introduce in sec. 1.8), we will underline the point that accurate prediction is rarely a desirable goal in and of itself. In Chapter 10, we will examine the long-term effects on users who interact with personalized systems: this includes studying what factors drive users to extreme content, and how to algorithmically mitigate such undesirable outcomes.

1.6.4 Language Generation, Personalized Dialog, and Interactive Agents

Finally, given the new modalities via which people interact with predictive systems, there are new demands for personalization.

For example, personalization is critical in a broad range of settings involving natural language. User-generated language data exhibit substantial variability due to differences in writing style, subjectivity, and so on. When dealing with such data, non-personalized models may struggle with this nuance. For example, automated systems for dialog, whether in task-oriented settings or for open-domain 'chit-chat,' can benefit from personalization, in order to generate responses that are more personalized or empathetic to the tone or context of individual users (Majumder et al., 2020).

We will see several instances of personalized language modeling throughout the book: language models are increasingly important to explain or interpret machine predictions (sec. 8.4.3), to facilitate new modalities of interaction with predictive systems (such as conversation, sec. 8.4.4), and to develop new kinds of assistive tools, for example, to help users respond to e-mail (sec. 8.5).

1.7 Techniques for Personalization

As mentioned in Section 1.1, one of the goals of this book is to establish a common narrative around the tools and techniques used to design personalized machine learning systems. Although we have shown that such systems are applied in domains as diverse as online commerce to personalized health, we find that the techniques used to implement these models follow a few common paradigms.

1.7.1 User Representations as Manifolds

One of the main ideas we will revisit throughout this book—and which allows us to adapt ideas from recommender systems to other types of machine learning—is that of a *user manifold*. That is, most of the personalized methods we will explore will involve *representations* of users that describe the common patterns of variation in their activities and interactions.

In the case of recommender systems, this 'user manifold' will be a vector that describes the principal dimensions that explain variance among user preferences (fig. 1.2). For example, we might discover that the principal dimensions that explain variance in preferences in a movie recommendation setting center around certain genres, actors, or special effects. Throughout the book, we will revisit the idea of user manifolds, as a general-purpose means of capturing common patterns of variation among users. Some examples include:

- In Chapter 5, we will use low-dimensional user representations to describe the dimensions of preferences and activities, which can be used to recommend items that users are likely to interact with.

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

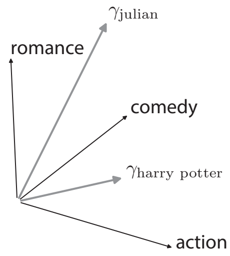

Figure 1.2 The basic idea behind recommender systems, and various other types of personalized machine learning, is to represent users by *low-dimensional manifolds* that describe the patterns of variance among their interactions. In a recommendation setting, a low-dimensional user vector might describe my *preferences* while a low-dimensional item vector describes an item's *properties*; compatible users and items have vectors that point in the same direction (chap. 5).

- In Chapter 8, user representations can describe the topics users tend to discuss (e.g., when writing reviews), or individual characteristics of their writing styles.
- In Chapter 9, user representations will describe the visual dimensions that users are interested in, allowing us to rank, recommend, or generate images in a personalized way.
- Throughout various case studies, user representations will capture characteristics ranging from dietary preferences (sec. 8.4.2), fitness profiles (sec. 7.8), social trust (sec. 6.4.1), or fashion choices (sec. 9.3).

1.7.2 Contextual Personalization and Model-Based Personalization

Although this book will predominantly cover methods that explicitly model user terms (as earlier), we will also cover a variety of models that deliberately avoid doing so.

Starting with simple approaches such as 'people who bought X also bought Y,' many classical approaches for (e.g.) recommendation *leverage user data, but do not include explicit parameters* (i.e., a 'model') associated with a user. However, such models are still *personalized*, in the sense that different predictions will be made for each individual based on how they interact with the system. Simple machine learning techniques, such as those we develop in Chapters 2 and 3, where users are represented by a few carefully engineered features, also follow this paradigm.

https://doi.org/10.1017/9781009003971.001 Published online by Cambridge University Press

We will distinguish between these two classes of approach using the terms *model-based* and *contextual* personalization. *Model-based* approaches learn an explicit set of parameters associated with each user, such as the 'user manifolds' described earlier (and in fig. 1.2); these models are typically intended to capture the predominant patterns of variation among users in a system, usually in terms of a low-dimensional vector. In contrast, *contextual* (also sometimes called 'memory-based,' as in chap. 5) approaches extract features from users' histories of recent interactions.

There are several settings in which contextual personalization may be preferable to explicitly modeling a user. When developing simple recommender systems in Chapter 4, and even more trivial personalized models in Chapters 2 and 3, we see that personalization can often be achieved with simple heuristics, or hand-crafted features or similarity measures. Such approaches may be desirable for a number of reasons: simple models may be more interpretable (and therefore preferable to expose to a user compared to 'black-box' predictions); or, we may lack adequate training data to learn complex representations from scratch.

1.8 The Ethics and Consequences of Personalization

Along with the increasing ubiquity of personalized machine learning systems, there is a growing awareness of the risks associated with personalization. Some of these issues have reached mainstream awareness, such as the idea that personalized recommendations can trap users in 'filter bubbles,' while other issues are considerably more subtle. For instance, considering the specific case of recommender systems, a naively implemented model can introduce issues including:

Filter Bubbles: Roughly speaking, recommendation algorithms rely on identifying specific item characteristics that are preferred by each user, and recommending items that most closely represent those characteristics. Without care, even a user with broad interests may be recommended only a narrow set of items that closely mimic their prior interactions.

Extremification: Likewise, a system that identifies features that a user is interested in may identify items that are most representative of those features, for example, a user who likes action movies may be recommended movies with *a lot* of action; in contexts such as social media and news recommendation this can lead to users being exposed to increasingly extreme content (the relationship between this and the previous issue is explained in chap. 10).

Concentration: Similar to the previous phenomenon, a user who has diverse interests may receive recommendations that only follow their most predominant interest (sec. 10.2). In aggregate, this may lead to to a small set of items being over-represented among all users' recommendations.

Bias: Given that recommenders (and many other personalized models) ultimately work by identifying common patterns of user behavior, users in the 'long-tail' whose preferences do not follow the predominant trends may receive sub-par recommendations.

Along with a rising awareness of these issues has come a set of techniques designed to mitigate them. These techniques borrow ideas from the broader field of fair and unbiased machine learning, whereby learning algorithms are adapted so as not to propagate (or not to exacerbate) biases in training data, though the fairness goals are often quite different. Diversification techniques can be used to ensure that predictions or recommendations balance relevance with novelty, diversity, or serendipity; related techniques seek to better 'calibrate' personalized machine learning systems by ensuring that predicted outputs are balanced in terms of categories, features, or the distribution over recommended items (sec. 10.3). Such techniques can mitigate filter bubbles by ensuring that model outputs are not highly concentrated around a few items, and more qualitatively can increase the overall novelty or 'interestingness' of model outputs. Other techniques follow more directly from fair and unbiased machine learning, ensuring that the performance of personalized models is not degraded for users belonging to underrepresented groups, or who have niche preferences (sec. 10.7).