

Design Patterns of Deep Learning with TensorFlow

*Building a customer hyper-personalisation
ecosystem using deep learning design patterns*

Thomas V Joseph



www.bpbonline.com

First Edition 2024

Copyright © BPB Publications, India


ISBN: 978-93-55516-497

All Rights Reserved. No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY

The information contained in this book is true to correct and the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.



Dedicated to

*To my wife **Anu**, my mother **Marykutty**
and my children **Joe** and **Tess***

About the Author

Thomas V Joseph is a seasoned professional with over 23 years of extensive experience. He currently leads the Data Science and Analytics function at Hogarth Worldwide, focusing on building creative analytics data products for multiple brands. He has diverse industry background spanning Telecom, Healthcare, Retail, Manufacturing, Edtech, Infrastructure, and Fintech, and is a known author of published works on AI/ML. He excels in developing AI/ML strategic plans, cutting-edge solutions, and building teams for technology competitiveness. His research interests include Autonomous vehicles and Causal AI, where he explores areas such as obstacle detection, road and lane detection, LiDAR-3D Point cloud, causal graphs, causal interventions, and do-calculus.

In addition to his technical expertise, Thomas possesses strong soft skills, including strategic leadership, product design thinking, and research and development. He is adept at enhancing business engagements for Data Science initiatives, partnering with customers to architect solutions, spearheading delivery teams, and leading R&D efforts to build new solutions.

About the Reviewer

Vijender Singh, a multi-cloud professional with over six years of expertise, holds an MSc with Distinction from Liverpool John Moores University, where his research centered on keyphrase extraction. He boasts an impressive collection of cloud certifications, including Google MLPE, five Azure certifications, two AWS certifications, and a TensorFlow certification. As a technical reviewer for numerous books, Vijender reflects his commitment to improving the future.

Acknowledgement

I count myself fortunate to have been granted the opportunity to embark on this book, and I attribute it to the Almighty. The Almighty's divine providence has endowed me with the discernment to acquire knowledge, wisdom to comprehend, and fortitude to progress in this endeavor.

I am deeply thankful to my beloved family for their unwavering love, support, and encouragement. Special thanks to my wife Anu, whose support and understanding have been my pillars of strength, and to my mother Marykutty, whose blessings have always inspired me. My children Joe and Tess have been a constant source of joy and motivation.

I am grateful to BPB Publications for their invaluable guidance and expertise in bringing this book to fruition. Their support and collaboration have been instrumental in shaping this project.

I extend my heartfelt appreciation to my colleagues and friends in various organizations where I have worked over the years. Their insights, feedback, and camaraderie have enriched my journey in the tech industry.

Last but not least, I extend my sincere thanks to all the readers who have shown interest in my book and supported me throughout this endeavor. Your encouragement and feedback have been immensely valuable, and I am truly grateful for your support.

Preface

The world of Artificial Intelligence (AI) has witnessed revolutionary changes in recent years, particularly with the emergence of Generative AI and Large Language Models. These advancements have opened up new possibilities for building innovative products and solutions. However, to leverage these technologies effectively and build successful products, it is crucial to have a strong understanding of the foundational aspects of AI.

This book on design patterns of deep learning is designed to provide readers with a comprehensive guide to understanding the basics of AI concepts, including Computer Vision and Natural Language Processing (NLP). Whether you are new to AI or an experienced practitioner looking to expand your knowledge, this book aims to equip you with the necessary skills to build robust and reliable enterprise applications using AI technologies.

In this book, you'll dive into fundamental topics like deep learning, explore design patterns in Computer Vision and NLP, and examine the components of Transformers, which is driving the Gen AI revolution. Additionally, you'll learn about best practices for building deploying enterprise applications, empowering you to develop robust AI-powered solutions. Each concept is reinforced with practical examples, enhancing your understanding of the material.

This book is intended for industry practitioners who are eager to explore the exciting field of AI and build enterprise applications that leverage the latest advancements in AI. Whether you are just starting out in AI or seeking to advance your career in enterprise development, this book will provide you with the knowledge and skills to succeed.

I hope you find this book informative and valuable in your journey to becoming a proficient practitioner in the field of AI.

Throughout this book, we will navigate through a series of chapters, each dedicated to dissecting specific design patterns in deep learning and their applications in hyper-personalization

Chapter 1: Customer Hyper-personalization - In this chapter, we dive into the world of hyper-personalization, which is about giving tailored experiences to customers especially in the digital world. We start by understanding what hyper-personalization is and how it differs from traditional segmentation methods. Then, we explore how deep learning enables hyper-personalization through real-world examples like personalized fashion recommendations, video streaming, news, and music recommendation systems. Finally, we present the hyper-personalization enablement framework, providing organizations with a roadmap to leverage deep learning in enhancing user experiences.

Chapter 2: Introduction to Design Patterns and Neural Networks - This chapter explores neural networks and its design patterns, and also get acquainted with TensorFlow, our tool for implementing deep learning. As we build the groundwork for deep learning, we'll also introduce design patterns. These are like guiding principles that help us make our deep learning models even better. We'll delve into the thinking behind these models and how we can adjust them to fit the specific needs.

Chapter 3: Design Patterns in Visual Representation Learning - Further in this chapter we will enter the world of visual representation learning, a crucial aspect of enabling hyper-personalization. We'll unravel the process of converting real-world images into mathematical vectors, essential for understanding customer preferences. Using **convolutional neural networks (CNNs)**, we'll explore the fundamentals of visual representation learning, starting from basic CNN constituents to advanced design patterns like VGG and Resnets. Through practical exercises, we'll demonstrate how CNNs classify images and extract valuable visual features, setting the stage for personalized customer experiences.

Chapter 4: Design Patterns for Non-Visual Representation Learning - Along with visual representation there is lot of information that can be captured from product descriptions, which will aid in enabling hyper-personalization. This chapter deals with non-visual representation learning, an important aspect in enhancing hyper-personalization. From preprocessing techniques like tokenization to concepts like word embeddings and sequence-to-sequence models, we'll explore the design patterns of **Natural Language Processing (NLP)** models. We will also explore design patterns of attention mechanisms, laying the groundwork for revolutionary models like Transformers.

Chapter 5: Design Patterns of Transformers - In this chapter we will continue our discussion on non-visual representation learning by introducing the transformer model. We'll examine key design patterns including positional embedding, scaled dot-product attention, and multi-head attention, discussing about their significance in shaping the model's architecture. Additionally, we'll delve into building end-to-end Transformer models for applications like language translation.

Chapter 6 : Data Distribution Challenges and Strategies - Hyper-personalization is a transient phenomenon, which means customer preferences keep changing with context. For this reason, the data distributions involved in hyper-personalization, which consist of both unstructured and structured data sets, are set to have huge variance in distribution. In this chapter, we explore the importance of analysing data distributions in deep learning. Throughout this chapter, we address the challenges posed by data distributions and offer best practices for maximizing their utility. Additionally, we analyze different types of drifts and provide insights into strategies for monitoring and correcting drifts to ensure model accuracy and relevance.

Chapter 7 : Model Training Philosophies - The essence of deep learning is in learning the model parameters during the training process. This chapter is a comprehensive guide on best practices in model training, covering essential topics such as feedforward and backpropagation, data splitting and normalization, validation, and addressing overfitting. Additionally, it explores advanced techniques including checkpointing, early stopping, and the utilization of data generators. By the end of this chapter, readers will have a solid understanding of how to train models effectively and optimize their performance for real-world applications.

Chapter 8 : Hyperparameter Tuning - In this chapter, we'll delve into the crucial aspect of optimizing model performance through adjustment of hyperparameters. These parameters, distinct from model weights and biases, play a vital role in enhancing model training. Throughout the chapter, we'll explore best practices for hyperparameter tuning, including strategies such as weight initializations, hyperparameter search techniques like grid search and random search, learning rate schedulers, and regularization methods. By mastering these techniques, readers will be equipped to fine-tune their models effectively, achieving optimal accuracy within efficient timeframes.

Chapter 9 : Transfer Learning - In this chapter, we'll explore an alternative approach to representation learning for hyper-personalization: transfer learning. Unlike training models from scratch, transfer learning involves leveraging pre-trained models and fine-tuning them to adapt to specific tasks. This method allows us to achieve comparable results in significantly less time and with minimal data. We'll delve into the intricacies of transfer

learning, focusing on its application in representation learning. By harnessing transfer learning, we can efficiently capture meaningful patterns and features from data, enabling the creation of highly effective models for personalized recommendations tailored to user preferences.

Chapter 10 : Setting Up Data and Deployment Pipelines - This chapter explores the crucial process of deploying models to production efficiently. By structuring each step as a pipeline, from raw data ingestion to continuous model monitoring, we can streamline the deployment process. We'll introduce the **TensorFlow Extended Framework (TFX)** and demonstrate how it facilitates chaining together various stages of the deep learning lifecycle, including data ingestion, preprocessing, model training, and evaluation. By leveraging TFX components, you'll learn how to create comprehensive end-to-end machine learning pipelines tailored for model deployment.

Code Bundle and Coloured Images

Please follow the link to download the *Code Bundle* and the *Coloured Images* of the book:

<https://rebrand.ly/35ec1c>

The code bundle for the book is also hosted on GitHub at

<https://github.com/bpbpublications/Design-Patterns-of-Deep-Learning-with-TensorFlow>.

In case there's an update to the code, it will be updated on the existing GitHub repository.

We have code bundles from our rich catalogue of books and videos available at **<https://github.com/bpbpublications>**. Check them out!

Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

errata@bpbonline.com

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at www.bpbonline.com and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

business@bpbonline.com for more details.

At **www.bpbonline.com**, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on BPB books and eBooks.

Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



Table of Contents

1. Customer Hyper-personalization.....	1
Structure.....	3
Objectives	4
What is hyper-personalization?	4
Enabling hyper-personalization.....	5
Hyper-personalization vs. segmentation.....	6
Enabling hyper-personalization with deep learning	8
<i>Personalized fashion recommendations.....</i>	<i>8</i>
<i>Creation of style profiles</i>	<i>8</i>
<i>Fine-tuning of personalized style quotient.....</i>	<i>10</i>
<i>Feedback to refine personal profiles.....</i>	<i>11</i>
<i>Personalized request notes.....</i>	<i>12</i>
<i>Personalized video streaming.....</i>	<i>12</i>
<i>Personalized news recommendation systems.....</i>	<i>14</i>
<i>Personalized music recommendation systems.....</i>	<i>15</i>
The hyper-personalization enablement framework.....	17
Introduction to deep learning and design patterns.....	19
<i>Visual representation learning</i>	<i>19</i>
<i>Non-visual representation learning</i>	<i>19</i>
Data distribution challenges and strategies	20
<i>Model training philosophies.....</i>	<i>20</i>
Hyperparameters and its optimization.....	20
<i>Transfer learning</i>	<i>21</i>
<i>Deployment pipeline.....</i>	<i>21</i>
Conclusion.....	21
Points to remember	22
References.....	22
2. Introduction to Design Patterns and Neural Networks	23
Structure.....	23
Objectives	24

Introduction to design patterns.....	24
Design patterns for deep learning	25
Design patterns for deep learning life cycle.....	26
<i>Design patterns in the institutionalization phase.....</i>	<i>26</i>
Data generation system.....	27
Data featurizer.....	27
Inference engine.....	28
Output consumption/visualization system.....	28
Feedback layer.....	28
Output/feedback storage.....	28
<i>Design patterns in the development phase of deep learning models.....</i>	<i>29</i>
Business discovery.....	29
Data discovery.....	30
Feature engineering.....	31
Model spot checking and training.....	31
Model registry and deployment.....	31
<i>Design patterns in the conception of model architecture.....</i>	<i>32</i>
Input layer.....	33
Learner block.....	33
Task layer.....	33
Convolutional Neural Networks.....	33
Introduction to neural networks.....	35
Exploring and implementing neural networks.....	37
Introduction to TensorFlow.....	38
Loading and preprocessing the data set.....	40
Preprocessing the image.....	42
Defining the model architecture.....	43
Forward pass with TensorFlow.....	45
Sequential class implementation.....	45
Back propagation.....	50
Calculate loss.....	52
Calculating gradients.....	53
Update process.....	55
Initiating the training process.....	56

Building the model using Functional API	59
Custom training method using gradient tape.....	62
Overfitting and regularization	68
Conclusion.....	69
Points to remember	70
References.....	70
3. Design Patterns in Visual Representation Learning.....	71
Introduction.....	71
Structure.....	71
Objectives	72
Hyper-personalization and the importance of visual representation learning.....	72
Visual representation learning.....	73
Visual representation learning using CNNs.....	75
<i>Kernels/filters</i>	76
<i>Padding</i>	77
<i>Stride</i>	79
<i>Activation</i>	80
<i>Pooling</i>	81
<i>Pooling operations at the intermediate layers</i>	82
<i>Pooling operations at the final layers</i>	82
Batch normalization layers	83
Fully connected layers	84
CNN represented in TensorFlow	85
Operational complexities of CNNs.....	87
<i>Complexity in normal convolution operation</i>	87
<i>Depth-wise convolutions</i>	88
<i>Point-wise convolutions</i>	89
<i>Spatial separable convolutions</i>	90
<i>Depth-wise separable convolutions</i>	90
Properties of CNNs that differentiate them from fully connected networks	91
<i>Sparse interactions</i>	91
<i>Parameter sharing</i>	93
<i>Equivariant representation</i>	93

<i>Translation invariance</i>	94
Design patterns of CNN architecture.....	95
<i>Design patterns of major CNN architectures</i>	96
<i>Deep networks</i>	96
<i>Wide networks</i>	99
<i>Inception network</i>	100
Networks for resource constrained devices	101
<i>MobileNet</i>	102
Visual representation learning for hyper-personalization using CNNs	103
Conclusion.....	107
Points to remember	107
4. Design Patterns for Non-Visual Representation Learning	109
Introduction.....	109
Structure.....	110
Objectives	110
Design patterns for NLP model	110
Design patterns in data processing for NLP.....	111
<i>Tokenization</i>	112
<i>Character level tokenization</i>	112
<i>Word-level tokenization</i>	114
<i>Subword-level tokenization</i>	116
Word embeddings	118
Common design patterns for NLP models.....	121
<i>Sequence to sequence model patterns</i>	122
<i>Recurrent neural networks</i>	123
<i>Stacked RNN</i>	124
<i>Bidirectional RNNs</i>	125
<i>Long short term memory architecture</i>	127
<i>Patterns in the application of RNNs</i>	130
<i>Sequence Labelling architecture pattern</i>	130
<i>Sequence classification pattern</i>	135
<i>Language generation pattern</i>	138
<i>Encoder-decoder pattern</i>	143

Attention mechanism.....	147
<i>Bahdanau attention</i>	148
<i>Encoder</i>	149
<i>Attention mechanism</i>	150
<i>Attention score</i>	152
<i>Attention weights</i>	154
<i>Context vector</i>	154
<i>RNN with attention mechanism</i>	154
Conclusion.....	164
Points to remember	165
5. Design Patterns for Transformers	167
Introduction.....	167
Structure.....	168
Objectives	168
Design patterns of transformer	168
Word embedding and positional embedding	170
Scaled dot product attention.....	174
Multi-head attention	177
Encoder	181
Decoder.....	188
<i>Decoder inputs</i>	188
<i>Masked multi-head attention</i>	190
The complete transformer model.....	196
Masks for encoder and decoder	197
Transformer training	200
<i>Downloading and preparing the data</i>	200
<i>Start and end tokens</i>	201
<i>Tokenization</i>	201
Sequence padding	203
Training operations	204
Conclusion.....	210
Points to remember	211

6. Data Distribution Challenges and Strategies.....	213
Introduction.....	213
Structure.....	213
Objectives	214
Data distributions.....	214
Covariate shift.....	215
Label shift	217
Concept drift	217
Detecting and correcting data drift.....	219
Conclusion.....	223
Points to remember	224
References.....	224
7. Model Training Philosophies	225
Introduction.....	225
Structure.....	225
Objectives	225
Introduction to deep learning model training	226
The model training process.....	227
Data transformation processes	228
<i>Data normalization</i>	<i>229</i>
<i>ImageDataGenerator class</i>	<i>232</i>
<i>Z-score normalization</i>	<i>233</i>
<i>Log normalization</i>	<i>234</i>
<i>Data augmentation.....</i>	<i>235</i>
<i>Flips</i>	<i>235</i>
<i>Shift</i>	<i>237</i>
<i>Rotation</i>	<i>238</i>
<i>Brightness.....</i>	<i>238</i>
<i>Zoom.....</i>	<i>239</i>
<i>Hyperparameters</i>	<i>240</i>
<i>Hyperparameters controlling flow of data.....</i>	<i>240</i>
<i>Hyperparameters controlling training duration</i>	<i>242</i>
<i>Hyperparameters controlling convergence of models.....</i>	<i>243</i>

<i>Optimization algorithms</i>	243
<i>Optimization processes</i>	244
<i>Gradient descent</i>	244
<i>Other optimization algorithms</i>	246
<i>Training metrics monitoring</i>	249
<i>Model baselining processes</i>	252
Conclusion.....	255
Points to remember.....	256
8. Hyperparameter Tuning	257
Introduction.....	257
Structure.....	257
Objectives.....	257
Introduction to hyperparameters.....	258
Hyperparameter tuning process.....	258
<i>Grid search</i>	260
<i>Random search</i>	262
<i>Bayesian optimization</i>	263
<i>Gradient-based optimization</i>	264
<i>Ensemble-based approaches</i>	264
Deep learning life cycle and hyperparameters.....	265
<i>Data preprocessing hyperparameters</i>	266
<i>Model architecture hyperparameters</i>	270
<i>Weight initialization hyperparameters</i>	274
<i>Weight initialization and tuning with Tensorflow</i>	276
<i>Model optimization hyperparameters</i>	280
<i>Tuning of regularizers</i>	284
Conclusion.....	287
Points to remember.....	288
Reference.....	288
9. Transfer Learning	289
Introduction.....	289
Structure.....	289
Objectives.....	290

Introduction to transfer learning.....	290
Design patterns for transfer learning	290
<i>Base model and its prediction</i>	291
<i>Feature extraction</i>	294
<i>Fine-tuning of base model</i>	299
Data pipelines for transfer learning.....	305
Transfer learning for NLP tasks	309
Conclusion.....	312
Points to remember	313
References.....	313
10. Setting Up Data and Deployment Pipelines	315
Introduction.....	315
Structure.....	315
Objectives	316
Production Deployment Framework	316
Production deployment at scale.....	317
TensorFlow extended framework	319
Data preparation.....	319
TFX framework for feature engineering	323
Data ingestion	326
Data validation	327
Data transformations	328
Training.....	331
Model evaluation.....	338
Model deployment.....	340
Linking all the components in a pipeline	342
Conclusion.....	347
Points to remember	348
Index	349-355

CHAPTER 1

Customer Hyper-personalization

It is the holiday season, and I log in to Amazon and am welcomed with a banner telling me about great savings on train tickets:




Figure 1.1: Amazon banner

I was not surprised to see this ad as over the past two weeks, I have been buying multiple train tickets for some of my weekend trips. Needless to say, I have not been using the Amazon Pay facility and transacted only with my credit card. Amazon must have sensed an opportunity here based on my recent purchase behavior, explaining the banner on savings on train tickets.

I scroll further down and see the Harry Potter movie recommended and a nudge to continue watching the latest James Bond movie, as shown in the following image:

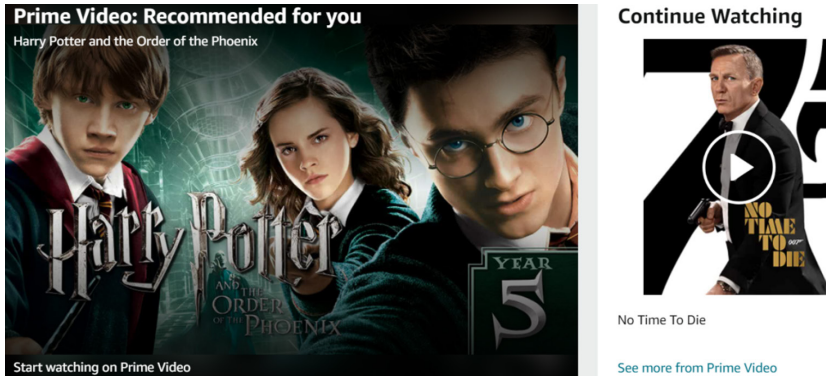


Figure 1.2: Amazon personalized recommendations

The nudge for the James Bond movie was expected as I had stopped that movie after 10 minutes. However, the recommendation of the Harry Potter movie was rather perplexing as I had not watched any Harry Potter movies earlier. Then I realized that a couple of weeks ago, I had bought an illustrated edition of a Harry Potter book from Amazon:

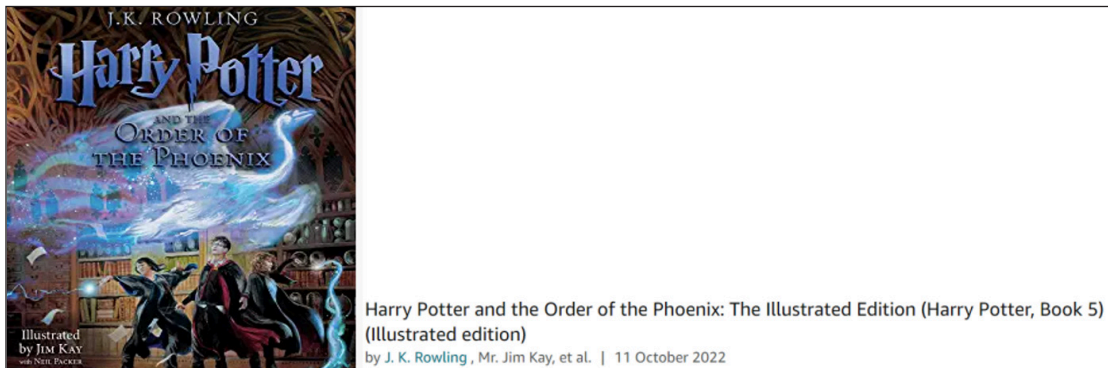


Figure 1.3: Amazon recommendations

Scrolling further down, I see the deals of the day, enticing me with some amazing discounts:




Figure 1.4: Amazon recommendations

These deals are the aftermath of my search for a new laptop over the past 2-3 weeks. All these deals that appeared on top of the portal had me thinking about how well the recommended products aligned with my needs and propensities. This is not a scenario unique to me; this will be the case of every individual who transacts regularly on e-commerce platforms like Amazon.

Welcome to the world of hyper-personalization!

In this chapter, we will lay the foundation for the concept of hyper-personalization. We will also be showcasing different industry use cases on how hyper-personalization is enabled through the confluence of different technologies, which have deep learning at their hearts.

Structure

In this chapter, we will discuss the following topics:

- What is hyper-personalization?
- Enabling hyper-personalization
- Hyper-personalization v/s segmentation
- Enabling hyper-personalization with deep learning
 - Personalized fashion recommendations
 - Personalized video streaming
 - Personalized news recommendation systems
 - Personalized music recommendation systems
- The hyper-personalization enablement framework
- Introduction to deep learning and design patterns
 - Visual representational learning
 - Non-visual representational learning
- Data distribution challenges and strategies
 - Model training philosophies
- Hyper parameters and its optimization
 - Transfer learning
 - Deployment pipeline

Objectives

Upon completion of this chapter, you will acquire a comprehensive understanding of several key concepts pertaining to hyper-personalization and its significance in the digital age. The chapter will help you explore why hyper-personalization has become increasingly relevant in today's technologically advanced era, where individuals expect tailored experiences that cater to their unique preferences and needs. Moreover, you will delve into the technology enablers that facilitate hyper-personalization, with a specific focus on the role of deep learning.

Additionally, you will be introduced to the concept of design patterns and their importance in the realm of deep learning. Design patterns provide reusable solutions to common design problems, enabling developers to build robust and scalable deep learning models efficiently.

What is hyper-personalization?

Hyper-personalization is a marketing phenomenon that derives meaningful insights about individuals from real-time interactions across multiple digital touch points. At the heart of hyper-personalization is the user transaction data, which includes buying behavior data like the product purchased, its specifications, browsing patterns, reviews of products; contextual data like time of day, day of week, geographic location; seasonal data and demographic data like age, gender and so on. These disparate datasets are used to extract trends, patterns, and meaningful information about customers, driving the recommendation process.

Personalization is the flavor of marketing strategies of companies like Amazon and Netflix. Right from the customized website tailored to each individual's taste to offers, prices, and even product designs, everything is driven by personalization. Amazon claims that nearly 35% of its sales comes from personalized recommendations, and nearly 56% customers are likely to convert to repeat buyers.[1]

Companies like Amazon and Netflix have reaped several benefits from hyper-personalization, some of which are as follows:

- **Better conversion rates:** The enhanced user engagement levels and better need fulfilment that come with hyper-personalization result in better conversion rates and better revenue realization.
- **Enhanced user engagement:** Hyper-personalization enhances the user experience across different touch points, which, in turn, leads to rise in engagement with the brand.
- **The power of knowing:** Hyper-personalization enables in-depth understanding of customer needs and aids the fulfilment of those needs. This results in better customer experience and more loyal customers.

- **Getting the timing right:** Hyper-personalization provides insights into the time/ periods when customers are more likely to engage with brands. This enables the targeting of customers at the right time and place with the right message.
- **Enabling long-term relationships:** When customer engagement happens at a personal level, customers get interested in the brand and feel more engaged. This results in long-term relationships with brands.

Enabling hyper-personalization

While hyper-personalization has been around for many years, what we are witnessing today is a higher level of adoption for personalized services. This has been enabled by the proliferation of technology and tools that make hyper-personalization possible. The explosion of connected devices and pervasiveness of apps in our day-to-day life have made contextual data on various touchpoints of the customer journey quite accessible. Extracting insights from these disparate customer datasets, thereby enabling hyper personalization at scale and speed, requires advanced technology landscape. At the heart of the technology landscape, enabling hyper-personalization at scale is machine learning and specifically, deep learning. There is the confluence of many deep learning techniques like computer vision, natural language processing, deep reinforcement learning and neural graph networks fueling hyper-personalization at scale. We will be exploring some of the deep learning techniques and their roles in enabling hyper-personalization in this book.

First of all, let us remind you that this book is not about hyper-personalization. This book is about deep learning, specifically about deep learning design patterns. You might now be wondering why we started off with the discussions on hyper-personalization then. Well, the answer is to get the context for learning design patterns of deep learning. The deep learning techniques, which we will discuss moving forward, will be in the context of enabling hyper-personalization. Therefore, it is important to understand hyper-personalization. This chapter will be focused on hyper personalization, and by the end of this chapter, we would connect the dots between hyper-personalization and deep learning, specifically looking at how hyper-personalization will be enabled by deep learning. In the subsequent chapters, we will get into the specifics of deep learning design patterns.

Now, let us get back to the focus area of this book, deep learning, specifically design patterns of deep learning. This book will not approach deep learning the way it is done in a traditional introductory book on deep learning. A traditional book lays a lot of focus on the what and how of deep learning, with explanations of the mathematical aspects of deep learning. In contrast, it tries to assimilate the experience, thumb rules, tips, and tricks that a deep learning practitioner will employ when applying to practical problems. That, in essence, is what design patterns attempt to address. The term design pattern has evolved from the field of architecture, where it entailed capturing of best practices, codifying experiences, and documenting proven practices. In this book, we will be looking at deep learning from the perspective of evolving design patterns , particularly when building hyper-personalization models.