
Preface

This text is in part an outgrowth of my MIT graduate course *Digital Speech Signal Processing*, which I have taught since the Fall of 1990, and in part a result of my research at MIT Lincoln Laboratory. As such, principles are never too distant from practice; theory is often followed by applications, both past and present. This text is also an outgrowth of my childhood wonder in the blending of signal and symbol processing, sound, and technology. I first felt this fascination in communicating with two cans coupled by twine, in playing with a toy Morse code, and in adventuring through old ham radio equipment in my family's basement. My goals in this book are to provide an intensive tutorial on the principles of discrete-time speech signal processing, to describe the state-of-the-art in speech signal processing research and its applications, and to pass on to the reader my continued wonder for this rapidly evolving field.

The text consists of fourteen chapters that are outlined in detail in Chapter 1. The “theory” component of the book falls within Chapters 2–11, while Chapters 12–14 consist primarily of the application areas of speech coding and enhancement, and speaker recognition. Other applications are introduced throughout Chapters 2–11, such as speech modification, noise reduction, signal restoration, and dynamic range compression. A broader range of topics that include speech and language recognition is not covered; to do so would result in a survey book that does not fill the current need in this field. The style of the text is to show not only when speech modeling and processing methods succeed, but also to describe limitations of the methods. This style makes the reader question established ideas and reveals where advancement is needed. An important tenet in this book is that anomaly in observation is crucial for advancement; as reflected by the late philosopher Thomas Kuhn: “Discovery commences with the awareness of anomaly, i.e., with the recognition that nature has somehow violated the paradigm-induced expectations that govern normal science.”¹

The text body is strongly supplemented with examples and exercises. Each exercise set contains a number of MATLAB problems that provide hands-on experience with speech signals and processing methods. Scripts, workspaces, and signals, required for the MATLAB exercises, are located on the Prentice Hall companion website (<http://www.phptr.com/quatieri/>). Also on this website are audio demonstrations that illustrate a variety of principles and applications

¹ T. Kuhn, *The Structure of Scientific Revolution*, University of Chicago Press, 1970.