Extracted from:

Build Talking Apps

Develop Voice-First Applications for Alexa

The Pragmatic Bookshelf

Raleigh, North Carolina

The Pragmatic Programmers

Build Talking Apps

Develop Voice-First Applications for Alexa

Craig Walls

edited by Jacquelyn Carter

Build Talking Apps

Develop Voice-First Applications for Alexa

Craig Walls

The Pragmatic Bookshelf

Raleigh, North Carolina

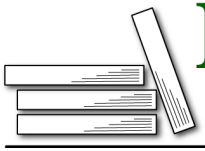

For our complete catalog of hands-on, practical, and Pragmatic content for software developers, please visit *https://pragprog.com*.

The team that produced this book includes:

CEO: Dave Rankin
COO: Janet Furlow
Managing Editor: Tammy Coron
Development Editor: Jacquelyn Carter
Copy Editor: Karen Galle
Indexing: Potomac Indexing, LLC
Layout: Gilson Graphics
Founders: Andy Hunt and Dave Thomas

For sales, volume licensing, and support, please contact *support@pragprog.com*.

For international rights, please contact *rights@pragprog.com*.

ISBN-13: 978-1-68050-725-6
Encoded using the finest acid-free high-entropy binary digits.
Book version: P1.0—April 2022

Acknowledgments

As a virtual assistant, Alexa can do some amazing things and can be a great help in many tasks. However, Alexa isn't much help when I say, "Alexa, help me write a book." That's why I'm super grateful to everyone who helped me make this book a reality.

First, I'd like to thank my editor, Jackie Carter. She guided me through the writing process and made sure I stayed on track. It's been great working with her again and I look forward to our next writing project.

Many thanks to all of the reviewers who provided feedback on the work in progress along the way to make this book the best it can be: Alessandro Bahgat, Jeff Blankenburg, Jesús Chamizo, Cornelia Davis, Jack Frosch, Srey Sang, Nate Schutta, Ken Sipe, Lalit Surampudi, Dzulkiflee Taib, Brian Tarbox, Erik Weibust, and Miles Woodroffe.

A shout out to my fellow speakers on the No Fluff Just Stuff conference tour as well as everyone who attended my sessions over the years. I have learned so much from all of you and knowing that I'll be standing at the front of the room sharing what I know pushes me to stay sharp and learn new things, including the material that went into this book.

As always, I'd like to thank the Phoenicians. Without their contribution, this book would not have been possible.

Finally, all of my love and thanks to my beautiful wife, Raymie, and to my amazing daughters, Maisy and Madi, for encouragement, inspiration, and your enduring love through yet another writing project.

Introduction

If you're reading this, then it's a safe bet that you're a software developer. If so, then what drew you to this craft?

Speaking for myself, I was drawn to working with computers early on by the way that computers were portrayed in the science fiction shows that I watched as a kid. In shows like *Star Trek*, the crew of the Enterprise could talk to the ship's computer, and the computer would answer with a spoken response. Will Robinson would talk to the robot on *Lost in Space*, and the robot would talk back, often warning him of danger. Fast-forward to modern-day science fiction, who wouldn't want a computer like Iron Man's Jarvis to assist with your work?

When I finally got to use a real computer, I was a bit disappointed that I couldn't talk to my Commodore VIC-20. And even if I tried, it wouldn't do anything or talk back in response. Where was my talking computer? Oh well, that's why it's called science *fiction*.

Nevertheless, I continued to pursue a career in software development as I observed that science fiction often eventually becomes science fact. The computers we have today are far more powerful than the writers of science fiction would have imagined. The tablet that Captain Picard used in his ready room on *Star Trek: The Next Generation* is now commonplace in tablet devices such as the iPad. Smart watches worn by many these days rival anything James Bond ever had on his wrist. And recently, voice assistants such as Alexa have given us a talking computer that's not entirely unlike Jarvis.

As a software developer, there are very few things more exciting than writing code that implements what was once science-fiction. That's exactly what we'll do in this book—write code that implements science fiction with Alexa and voice-first applications.

Who Should Read This Book?

This book is for software developers who are interested in creating voice-first user interfaces or adding voice to their existing applications. We'll be developing in JavaScript using Node.js, so some experience with the JavaScript language and tooling will be helpful, but even those new to or unfamiliar with JavaScript should be able to follow along (and maybe learn JavaScript in a fun way).

About This Book

This book will introduce you to the exciting topic of voice user interfaces, building applications known as "skills" for the Alexa platform.

We'll start where all good learning projects start, with a simple "Hello World" skill that we'll deploy to Alexa and talk to. Then, before diving in deep, you'll learn techniques for testing Alexa skills, with both semi-automated and completely automated tests.

Next, you'll learn how to parameterize our conversations with Alexa using *slots*. When the user leaves out important information or says something that doesn't make sense, you'll learn how to create dialogs where Alexa can elicit missing details and validate that the information given is acceptable.

Often, voice interfaces must collaborate with other applications. We'll explore options for integrating Alexa skills with external applications and APIs.

Next, we'll look at how to change how Alexa speaks and sounds, and even include special sound effects and music in her responses using the Speech Synthesis Markup Language (SSML). That will lead right into a discussion of how to add localization to our skills so that Alexa can speak in the user's own language (and even change her voice to include the language's accent).

Then we'll see how voice-first applications aren't necessarily voice-only applications, by creating visual experiences that complement Alexa's vocal responses when our skill is launched on screen-enabled devices such as Echo Show and Fire TV.

Following our exploration of visual interfaces, we'll see how skills can act in response to out-of-band events when the user isn't actively using the skill. This includes handling events that are triggered when the user modifies the skill's settings as well as notifying the user of events.

Next, you'll learn how to monetize skills by offering in-skill purchases, allowing the user to purchase goods, both tangible and virtual.

We'll then unleash our skill to the world by publishing it. You'll learn what the requirements for publication are as well as perform some last-minute fine tuning to ensure that our skill is ready for real-world use.

Finally, we'll have a look at a relatively new way to create rich conversational interactions using Alexa Conversations. Like Dialogs we learned about earlier in the book, Conversations involve a question and answer session between Alexa and the user to achieve some outcome. But where Conversations shine is in how they allow the user to change their answers and take tangents in the conversational flow.

Online Resources

The examples and source code shown in this book can be found under the source code link on the Pragmatic Bookshelf website.[1]

Please report any errors or suggestions using the errata link that is available on the Pragmatic Bookshelf website.[2]

If you like this book and it serves you well, I hope that you will let others know about it—your reviews really do help. Tweets and posts are a great way to help spread the word. You can find me on Twitter at @habuma, or you can tweet @pragprog directly.

Craig Walls
April 2022

1. https://pragprog.com/book/cwalexa/build-talking-apps
2. https://pragprog.com/book/cwalexa/build-talking-apps