REVIEWER COMMENTS

"For a while, I have been looking for a book in Data Science using Python that would cover the most relevant technologies. Well, my search is over. A must-have book for any practitioner of this field. The machine learning chapter is a real winner!! The dynamic visualization is fantastic." —**Ramon Mata-Toledo, Professor, James Madison University**

"IBM Watson is an exciting chapter. I enjoyed running the code and using the Watson service. The code examples put together a lot of Watson services in a really nifty example. I enjoyed the OOP chapter—doctest unit testing is nice because you can have the test in the actual docstring so things are traveling together. The line-by-line explanations of the static and dynamic visualizations of the die rolling are just great." —**Daniel Chen, Data Scientist, Lander Analytics**

"A lucid exposition of the fundamentals of Python and Data Science. Excellent section on problem decomposition. Thanks for pointing out seeding the random number generator for reproducibility. I like the use of dictionary and set comprehensions for succinct programming. "List vs. Array Performance: Introducing %timeit" is convincing on why one should use ndarrays. Good defensive programming. Great section on Pandas Series and DataFrames—one of the clearest expositions that I have seen. The section on data wrangling is excellent. Natural Language Processing is an excellent chapter! I learned a tremendous amount going through it. Great exercises."
—**Shyamal Mitra, Senior Lecturer, University of Texas**

"My game programming students would appreciate these exercises." —**Pranshu Gupta, Assistant Professor, DeSales U.**

"I like the discussion of exceptions and tracebacks. I really liked the Data Mining Twitter chapter; it focused on a real data source, and brought in a lot of techniques for analysis (e.g., visualization, NLP). I like that the Python modules helped hide some of the complexity. Word clouds look cool." —**David Koop, Assistant Professor, U-Mass Dartmouth**

"I love the text! The right level for IT students. The examples are definitely a high point to this text. I love the quantity and quality of exercises. Avoiding heavy mathematics fits an IT program well." —**Dr. Irene Bruno, George Mason University**

"A great introduction to deep learning." —**Alison Sanchez, University of San Diego**

"I was very excited to see this textbook. I like its focus on data science and a general purpose language for writing useful data science programs. The data science portion distinguishes this book from most other introductory Python books."
—**Dr. Harvey Siy, University of Nebraska at Omaha**

"The collection of exercises is simply amazing. I've learned a lot in this review process, discovering the exciting field of AI. I liked the Deep Learning chapter, which left me amazed with the things that have already been achieved in this field. Many of the projects are really interesting." —**José Antonio González Seco, Consultant**

"An impressive hands-on approach to programming meant for exploration and experimentation."
—**Elizabeth Wickes, Lecturer, School of Information Sciences, University of Illinois at Urbana-Champaign**

"I was impressed at how easy it was to get started with NLP using Python. A meaningful overview of deep learning concepts, using Keras. I like the streaming example." —**David Koop, Assistant Professor, U-Mass Dartmouth**

"Really like the use of f-strings, instead of the older string-formatting methods. Seeing how easy TextBlob is compared to base NLTK was great. I never made word clouds with shapes before, but I can see this being a motivating example for people getting started with NLP. I'm enjoying the chapters in the latter parts of the book. They are really practical. I really enjoyed working through all the Big Data examples, especially the IoT ones." —**Daniel Chen, Data Scientist, Lander Analytics**

"A good overview of various neural networks with coding examples for classification problems for which neural networks are commonly used. The exercises in this chapter will give students insight into how changing the structure of neural networks and the amount of training/testing data affect performance. The Twitter examples covering trending topics, creating word clouds, and mapping the location of users are instructive and engaging. I like the real-world examples of data munging. Reviewing this book was enjoyable and even though I was fairly familiar with Python, I ended up learning a lot." —**Garrett Dancik, Associate Professor of Computer Science/Bioinformatics, Eastern Connecticut State University**

"I really liked the live input-output. The thing that I like most about this product is that it is a Deitel & Deitel book (I'm a big fan) that covers Python." —**Dr. Mark Pauley, University of Nebraska at Omaha**