

## Preface

Michael L. Brodie, Computer Systems and Artificial Intelligence Laboratory,  
Massachusetts Institute of Technology

### The ACM A.M. Turing Award

This is the first in a series of books launched by the Association for Computing Machinery (ACM) to honor and celebrate the achievements of the ACM A.M. Turing Award recipients.

The A.M. Turing Award is the ACM's most prestigious technical award, given for major contributions of lasting importance to computing.<sup>1</sup>

Sometimes referred to as the "Nobel Prize of computing," the A.M. Turing Award was named in honor of Alan Mathison Turing (1912–1954), a British mathematician and computer scientist. Alan Turing made fundamental advances in computer architecture, algorithms, formalization of computing, and artificial intelligence. In 1950, Turing developed the now-prophetic "Turing test," used to determine if a computer appears to exhibit intelligent behavior indistinguishable from that of a human. He was also instrumental in British code-breaking work during World War II.

Created in 1966, the A.M. Turing Award has been conferred for 51 years to 66 recipients spanning all areas of computing. Since 2014, it has been accompanied by a \$1,000,000 prize. "The Turing Award is now funded at the monetary level of the world's most prestigious cultural and scientific awards and prizes," said ACM President Alexander Wolf. "With the generous support of Google, we can celebrate the mainstream role of computing in transforming the world and the way we communicate, conduct business, and access entertainment. We can also commemorate the pioneering, fundamental contributions of our ACM Turing Award recipients in advancing computing as a science and a profession."<sup>2</sup>

### Celebrating Michael Stonebraker's 2014 ACM A.M. Turing Award

This book celebrates Michael Stonebraker's accomplishments that led to his 2014 ACM A.M. Turing Award "*For fundamental contributions to the concepts and practices underlying modern database systems*".<sup>[1]</sup>

When Barbra Liskov, Turing Award committee chair, informed Mike that he had been awarded the 2014 Turing Award, he "... teared up. The recognition and validation for my lifetime work was incredibly gratifying." <sup>[2]</sup>

The book describes, for the broad computing community, the unique nature, significance, and impact of Mike's achievements in advancing modern database systems over more than forty years. Today, data is considered the world's most valuable resource,<sup>3</sup> whether it is in the tens of millions of databases used to manage the world's businesses and governments, in the billions of databases in our smartphones and watches, or residing elsewhere, as yet unmanaged, awaiting the elusive next generation of database systems. Every one of the millions or billions of databases includes features that are celebrated by the 2014 Turing Award and are described in this book.

*Why should I care about databases? What is a database? What is data management? What is a database management system (DBMS)?* These are just some of the questions that this book answers, in describing the development of data management through the achievements of Mike Stonebraker and his over 200 collaborators. In reading the stories in this book, you will discover core data management

---

<sup>1</sup> [https://amturing.acm.org/award\\_winners/stonebraker\\_1172121.cfm](https://amturing.acm.org/award_winners/stonebraker_1172121.cfm)

<sup>2</sup> <https://amturing.acm.org/prize-news.cfm>

<sup>3</sup> *The Economist*, May 6, 2017

concepts that were developed over the two greatest eras—so far—of data management technology. *Why do we need database systems at all? What concepts were added? Where did those concepts come from? What were the drivers? How did they evolve? What failed and why? What is the practice of database systems? And, why do those achievements warrant a Turing Award?*

The focus is on Michael Stonebraker, the 2014 Turing Award winner. But the achievements that the award honors are not just those of one person, no matter how remarkable he may be. The achievements are also due to hundreds of collaborators—researchers, students, engineers, coders, company founders and backers, partners, and yes, even marketing and sales people. *Did all of the ideas come from Mike?* Read on, especially Mike’s chapter “Where Do Good Ideas Come from and How to Exploit Them.”

I have had the great privilege of working with more than my fair share of Turing Award recipients starting as an undergraduate taking complexity theory from Steve Cook of  $P=NP$  fame. No two Turing Award winners are alike in topic, approach, methods, or personality. All are remarkably idiosyncratic. Mike is, to say the least, idiosyncratic, as you will discover in these pages.

This book answers questions, like those in *italics*, in 30 stories, each by storytellers who were at the center of the story. The stories involve technical concepts, projects, people, prototype systems, failures, lucky accidents, crazy risks, startups, products, venture capital, and lots of applications that drove Mike Stonebraker’s achievements and career. Even if you have no interest in databases at all,<sup>4</sup> you’ll gain insights into the birth and evolution of Turing Award-worthy achievements from the perspectives of 53 remarkable computer scientists and professionals.

## **Making Databases Work: The Pragmatic Wisdom of Michael Stonebraker**

The theme of this book is modern database *systems*. The 2014 A.M. Turing Award was conferred *For fundamental contributions to the concepts and practices underlying modern database systems*. It is 1 of only 4 Turing Awards given for databases, and 1 of only 2 out of 51 given for computer systems.

Mike addressed the systems theme in his Turing Award lecture (typically intended to summarize Turing-worthy achievements) in terms of the challenges that he faced and the approach he took to systems research, in four steps. “The first was to try to explain why system software is so hard to build, and why good teams screw it up on a regular basis. Second, it takes real perseverance to ‘stick it out’ and make something actually work. The third was to talk about the start-up experience, and why venture capitalists usually deserve their reputation as ‘land sharks.’ Lastly, it is clear that luck plays a significant role in successful startups, and I wanted to explain that. The overarching theme was to use a significant physical challenge as a metaphor for system software development. Over the years, the physical challenge has varied between our cross-country bike ride in 1988, and my climbing all 48 4,000-foot mountains in New Hampshire.” [2]

This description contains the seeds of answers to the previous *italicized* questions that are elaborated throughout the book.

The computer systems theme is pursued in the book by stories told from the research perspective: *What were the core database concepts? How did they develop? Why were they significant?* And stories told from the computer systems perspective: *What are the development or engineering challenges? What challenges arise in implementing a research idea? How are they overcome? Do essential research contributions arise from systems engineering? As you read these stories ask yourself: What is the relationship between research and systems engineering? Why build prototype systems at all? Having proven concepts in research and in a prototype system, why build a product?* (Spoiler alert: While money plays a significant role, it was by no means the goal.)

---

<sup>4</sup> Is that possible?

## Acknowledging 52 Remarkable Contributors

This book is a collection of 30 stories written by Mike<sup>5</sup> and 51 of his collaborators: 33 world-leading database researchers, 12 world-class systems engineers, 4 business partners, and 2 editors.

It is my great pleasure to acknowledge the fascinating contributions of all of these remarkable people. They responded with enthusiasm to recount their collaborations with Mike, looking for the essential contributions and how they emerged, all mixed with concern for accurately remembering the crucial facts for you, the reader—in some cases reaching back four decades. *What was important? What seemed to matter versus what really mattered?* Each contributor, like Mike, is idiosyncratic and strongly opinionated, as you will see. Their achievements reflect the state of the technology and data management demands of the time. Everyone had to be reminded to reflect disagreements with Mike (showing the normal give-and-take of computing research and product development), as well as to state why Mike's contributions warranted the Turing Award. Interestingly, few authors felt comfortable praising Mike, perhaps reflecting the personalities of computer scientists.

Mike can be intimidating. He has made a career of making bold, black-and-white statements to challenge and to inspire himself and the database community to greater accomplishments, as Phil Bernstein recounts so well.<sup>6</sup> It's a sign of maturity for a Ph.D., postdoc, or collaborator of any kind to stand up and refute Mike—and such a pleasure to experience, by Mike included. You will see, in each story, the efforts of the contributors to pass this benchmark.

A theme of Mike's career has been to question *conventional wisdom*, specifically as it ages and as new challenges and concepts arise or as it is undermined by poor practices. The most obvious example is Mike's claim that "one-size-does-not-fit-all" in databases, which is a complete contradiction of the claims of the Elephants, Mike's affectionate term for the DBMSs that dominate their market. Yet Mike was the chief proponent of "one-size-fits-all" in the relational database era. It has been fascinating to watch Mike's contributions become conventional wisdom, which Mike then questions toward the next level of achievement.

If you are an aspiring researcher, engineer, or entrepreneur you might read these stories to find these turning points as practice to tilt at your own computer-science windmills, to spur yourself to your next step of innovation and achievement.

### Janice Brown, Our Amanuensis

My greatest acknowledgement, for her contributions to this book, is for Janice L. Brown, technology writer/editor, startup consultant, and frequent Stonebraker collaborator, of Janice Brown & Associates, Inc. In many ways this is Janice's book. (If you dream about the book, it's yours.) Janice was our amanuensis: editor, copywriter, enthusiast, critic,<sup>7</sup> and berger du chats (cat herder) extraordinaire et malheureusement, très nécessaire.

---

<sup>5</sup> Mike wrote some of the most interesting chapters; he did not review other chapters so as not to influence the authors' voices.

<sup>6</sup> See Chapter 2: Leadership and Advocacy.

<sup>7</sup> "Truly a great story, yet perhaps you didn't really mean that; how about ..."

## Additional References

- [1] [2016 ACM A.M. Turing Award Citation and Biography](#)
- [2] Stonebraker, M., [What It's Like to Win the Turing Award](#), Blog@CACM, Communications of the ACM, July 24, 2015