The Way We Think About Data

ADRIAN COLYER

**HUMAN
INSPECTION OF BLACK-
BOX ML MODELS;
RECLAIMING
OWNERSHIP OF DATA**

The two papers I've chosen for this issue of *acmqueue* both challenge the way we think about and use data, though in very different ways.

In “Stop Explaining Black-box Machine-learning Models for High-stakes Decisions and Use Interpretable Models Instead,” Cynthia Rudin makes the case for models that can be inspected and interpreted by human experts. There are many occasions when understanding what a model is doing and how it reaches its decisions is sufficiently important that interpretability has to be taken into account as a key design objective for a system. The common wisdom is that turning away from deep models inevitably means sacrificing accuracy. But Rudin points out that's not necessarily so.

What if you never deployed a black-box model when an interpretable model exists with the same level of performance? That line of thinking leads to the follow-on question, how can you know the optimal performance an interpretable model can achieve? If this topic interests you, you can read more about that in The Morning Paper write-ups on CORELS (certifiably optimal rule lists; <https://blog.acolyer.org/2019/10/30/corels/>) and RiskSLIM (Risk-calibrated Supersparse Linear Integer Model; <https://blog.acolyer.org/2019/11/01/optimized-risk-scores/>).

The second paper, “Local-first Software: You Own Your Data, in Spite of the Cloud,” describes how to retain

sovereignty over your data. Can you combine the freedom and sense of ownership you used to have with applications writing open file formats to the local file system, with the ease of multidevice access and multiuser collaboration that comes from cloud services? Kleppmann et al. think so and set forth a compelling call for us to start building local-first applications. The only question is, how?

For the complete column, go to <https://queue.acm.org/TheMorningPaper/6/>

Adrian Colyer is a venture partner with Accel in London, where his job is to help find and build great technology companies across Europe and Israel. (If you're working on an interesting technology-related business, he would love to hear from you: acolyer@accel.com.) Prior to joining Accel, he spent more than 20 years in technical roles, including CTO at Pivotal, VMware, and SpringSource.

Copyright 2019 held by owner/author. Publication rights licensed to ACM. Reprinted with permission from <https://blog.acolyer.org>