

## Letter from the Editor-in-Chief

### TCDE Chair Election

**Voting is IMPORTANT.** Every two years, the Technical Committee on Data Engineering has an election for a new TCDE chair. The chair is the person who is empowered and entrusted with the authority to appoint members of the TCDE Executive Committee. The Executive Committee has full authority in managing the Technical Committee. So when you vote, you are choosing someone who is critical to the on-going operation of the Technical Committee. Elections have consequences.

Page two of this issue has a letter from the nominating committee describing their work and naming the candidates. Also on page two, Brookes Little from the IEEE Computer Society describes how you can vote. This is followed on pages three and four with candidate statements and biographies. Brookes describes how you can vote, but only you can exercise that function. Please do.

### The Current Issue

Machine learning (ML) is a huge deal, driven by the lure of finding the proverbial needle in the haystack insight that is transformative. Also driven by the fear that someone else will find the needle first. So this is a hugely competitive space, with implications for business, science, education, indeed a vast array of areas. This issue should be of interest to any number of database technologists, be they practitioners or researchers, data scientists or data platform specialists.

This is more than simply taking a look at Map/Reduce. That is pure infrastructure, and now only part of the platform space that includes distributed databases as well as more specialized infrastructures. And platform work only captures part of what is going on. Like OLAP, one needs more than just a platform, one needs to bring machine learning methods into play. Inferencing, probability, model building, categorizing, and more can be usefully integrated into ML systems. While data intensive ML related work has gone on for a number of years, this is nonetheless a young area with enormous potential for high impact work that may span the better part of many careers.

Thus, ML is a great topic for the Data Engineering Bulletin. Chris Jermaine is the editor for this introductory issue on ML, but I am sure there will be more ML Bulletin issues in the future. The current issue spans platforms, models, scaling, ML approaches, and more. Chris has assembled a diverse collection of articles in this area that can function as that can function as an introduction to the area, and the start of a journey to an in-depth understanding of ML technology and its uses. I want to thank Chris for this fine result, and for some additional sweat in dealing with Bulletin formatting difficulties.

David Lomet  
Microsoft Corporation