Scanning the Issue

Point of View:
A Retrospective and Prospective View of Approximate Computing
by W. Liu, F. Lombardi, and M. Shulte

Computing systems have been facing severe technology challenges in recent years with regard to power consumption, circuit reliability, and high performance. For many years, the issues of power consumption and performance have been addressed with the use of technology scaling. However, as Dennard's scaling tends toward an end, it has become difficult to further improve the performance under the same power constraints. In addition to power, reliability also becomes a critical issue when the feature size of the complementary metal–oxide–semiconductor (CMOS) technology is reduced below 7 nm. Thus, ensuring the complete accuracy of the signal has become increasingly challenging in recent years.

The month's regular papers issue highlights the use of adversarial learning for defenses against attacks on deep neural network classifiers and the use of combinatorial optimization for GUI design.

As it happens, there are a number of pervasive computing applications, which are inherently error-tolerant and, in general, require acceptable—not accurate—results. Approximate computing, a technique that returns possibly inaccurate results rather than a guaranteed accurate result, is an excellent choice for such applications as it offers high energy efficiency at an acceptable level of output. In this article, the authors take a closer look at approximate computing and offer their perspective on the future of this promising technique/methodology. They begin by defining approximate computing and introducing the variations of the technique, before discussing its application at different levels (hardware and software). Current challenges are then outlined in detail, and the authors conclude by offering an outlook into the future.

Adversarial Learning Targeting Deep Neural Network Classification:
A Comprehensive Review of Defenses Against Attacks
by D. J. Miller, Z. Xiang, and G. Kesidis

Machine learning (ML)-based systems—particularly, in recent years, deep neural networks (DNNs)—have found broad applications ranging from military, industrial, medical, multimedia/Web, and scientific (including genomics) to even the political, social science, and legal arenas. As their integration into the modern world's infrastructure continues to grow, they become evermore enticing targets for adversaries, including individual hackers, criminal organizations, as well as government intelligence services, which may seek to "break" them. Thus, adversarial learning (AL)—a technique used for devising attacks against ML systems as well as defenses against such attacks—has become a popular topic over the past decade.

This article provides a timely survey of contemporary attacks against supervised classifiers. In particular, the focus is on such attacks, and defenses against them, for DNNs, which have been a recent target due to their state-of-the-art performance in many applications. The article begins by laying out some background material and then proceeds to survey recent work on test-time evasion (TTE), data poisoning (DP), backdoor DP, and reverse engineering (RE) attacks and particularly defenses against the same. Reviewed works are assessed technically, including identifying any issues/limitations, required hyperparameters, needed computational complexity, as well as the performance measures evaluated and the obtained quality. This article then moves on to provide novel insights, supported by experimental results that challenge conventional AL wisdom and that target a number of key unresolved issues. This article also discusses attacks on

Digital Object Identifier 10.1109/JPROC.2020.2975522

the privacy of training data. Benchmark comparisons of several defenses against TTE, RE, and backdoor DP attacks on images are then provided. This article concludes with a discussion of continuing research directions, including the supreme challenge of detecting attacks whose goal is not to alter classification decisions, but rather simply to embed, without detection, "fake news" or other false content.

Combinatorial Optimization of Graphical User Interface Designs
by A. Oulasvirta, N. R. Dayama, M. Shiripour, M. John, and A. Karrenbauer

The graphical user interface (GUI) has become the prime means for interacting with computing systems. It leverages human perceptual and motor capabilities for elementary tasks such as command exploration and invocation, information search, and multitasking. For designing a GUI, numerous interconnected decisions must be made such that the outcome strikes a balance between human factors and technical objectives. Normally, design choices are specified manually and coded within the software by professional designers and developers. Among the various computational techniques used for choosing the GUI design and interactions, combinatorial optimization is distinguished by its algorithmic capacity, controllability, and generalizability. Its potential to complement the human designers' work lies in its capability to search for large numbers of possible designs, a target that might otherwise be out of reach. In comparison to formal methods, such as logic, combinatorial optimization offers an effective but flexible way of expressing design knowledge and objectives in a computable manner. Compared to data-driven approaches based on machine learning, such as artificial neural networks, combinatorial optimization allows direct and meaningful control of design outcomes via specific design objectives.

Despite these benefits, applications of this technique were limited to keyboards and widget layouts until 15 years ago. The obstacle has been the mathematical definition of design tasks, on the one hand, and the lack of objective functions that capture essential aspects of human behavior, on the other.

This article presents definitions of layout design problems as integer programming tasks, a coherent formalism that permits identification of problem types, analysis of their complexity, and exploitation of known algorithmic solutions. It then surveys advances in formulating evaluative functions for common design-goal foci such as user performance and experience. The convergence of these two advances has expanded the range of solvable problems. Approaches to practical deployment are outlined with a wide spectrum of applications. The article concludes by discussing the position of this application area within optimization and human–computer interaction research and outlines challenges for future work.

Scanning Our Past: Electrostatic Telegraphy—1753–1816
by A. Allerhand

In this month's history article, the author traces the developments in electrical transmission between the years 1600 and 1823. William Gilbert's discovery in 1600 that substances besides amber attracted light objects after being rubbed led to the invention of electrostatic generators. These were used for experiments that contributed to the gradual understanding of electrical phenomena. Starting in 1672, electrostatic generators in Europe also enabled the transmission of electric charges over increasing distances, while the Leyden jar led to the proposal and invention of various electrostatic telegraph systems in the second half of the 18th century. Francis Ronalds' telegraph of 1816 contained elements used in battery-powered, commercial telegraphs; in 1823, he also explained the phenomenon of retardation of the electric current by induction in underground wires, a problem that plagued cable telegraphy into the 20th century. ■