Abstract

Artificial neural networks (ANNs) have very successfully been used in numerical simulations for a series of computational problems ranging from image classification/image recognition, speech recognition, time series analysis, game intelligence, and computational advertising to numerical approximations of partial differential equations (PDEs). Such numerical simulations suggest that ANNs have the capacity to very efficiently approximate high-dimensional functions and, especially, indicate that ANNs seem to admit the fundamental power to overcome the curse of dimensionality when approximating the high-dimensional functions appearing in the above named computational problems. There are a series of rigorous mathematical approximation results for ANNs in the scientific literature. Some of them prove convergence without convergence rates and some of these mathematical results even rigorously establish convergence rates but there are only a few special cases where mathematical results can rigorously explain the empirical success of ANNs when approximating high-dimensional functions. The key contribution of this article is to disclose that ANNs can efficiently approximate high-dimensional functions in the case of numerical approximations of Black-Scholes PDEs. More precisely, this work reveals that the number of required parameters of an ANN to approximate the solution of the Black-Scholes PDE grows at most polynomially in both the reciprocal of the prescribed approximation accuracy $\varepsilon > 0$ and the PDE dimension $d \in \mathbb{N}$. We thereby prove, for the first time, that ANNs do indeed overcome the curse of dimensionality in the numerical approximation of Black-Scholes PDEs.