

ETSVL007- 2017- VLSI Extreme Learning Machine: A Design Space Exploration

Abstract

In this paper, we describe a compact low-power high-performance hardware implementation of extreme learning machine for machine learning applications. Mismatches in current mirrors are used to perform the vector-matrix multiplication that forms the first stage of this classifier and is the most computationally intensive. Both regression and classification (on UCI data sets) are demonstrated and a design space tradeoff between speed, power, and accuracy is explored. Our results indicate that for a wide set of problems, σVT in the range of 15-25 mV gives optimal results. An input weight matrix rotation method to extend the input dimension and hidden layer size beyond the physical limits imposed by the chip is also described. This allows us to overcome a major limit imposed on most hardware machine learners. The chip is implemented in a 0.35- μm CMOS process and occupies a die area of around 5 mm \times 5 mm. Operating from a 1 V power supply, it achieves an energy efficiency of 0.47 pJ/MAC at a classification rate of 31.6 kHz.

