

**Peter Richtarik**

University of Edinburgh, United Kingdom

peter.richtarik@ed.ac.uk

We design a novel stochastic dual coordinate ascent (SDCA) method for minimizing an L2-regularized empirical loss function. The method operates by updating a random **set** of coordinates of the dual problem at every iteration. We allow for an **arbitrary probability law** (sampling) to govern the choice of the set of coordinates to be updated in an iteration and show how this enters the complexity bound. By varying the sampling we obtain SDCA with importance sampling, mini-batch SDCA and distributed SDCA as special cases. In a statistically interesting regime for the choice of the regularization parameter, we obtain a linear speedup up to the square root of the number of coordinates (examples). However, our method also enjoys further **data-dependent speedup**. Lastly, unlike traditional analysis of SDCA, in our analysis we control the decrease of the duality gap directly.

*Joint work with Zheng Qu (Edinburgh), Tong Zhang (Baidu).*