In theory, the linear autoencoder EASE is one of the most capable collaborative filtering recommenders for large item domains with sparse user-item feedback. However, the model's weights are determined by the inverse of a matrix of dimension equal to the item set size. This inverse matrix is generally dense, and for large item sets, the computed weight matrix might be too large to store in memory during inference. Consequently, scaling the model beyond tens of thousands of items quickly becomes very expensive.

We propose a modification of EASE called SANSA to alleviate the issue. SANSA approximates the weights of EASE with prescribed density via an end-to-end sparse training procedure. To find a method capable of computing the sparse approximation efficiently, we investigate approaches for constructing sparse approximate inverse preconditioners. We select a method fitting for very large SPD problems with general sparsity patterns. The training procedure is robust and finds a good approximation of EASE even on datasets with dense item relations. Moreover, as the number of items in datasets grows, SANSA achieves unparalleled efficiency, even compared to EASE's previous state-of-the-art modification focused on scalability. Consequently, SANSA effortlessly scales the concept of EASE to millions of items.