Compressive neural representation of sparse, high-dimensional probabilities

xaq pitkow UNIVERSITY of ROCHESTER

sparse signal $\boldsymbol{s}$

$$y_i = \sum_j A_{ij} s_j$$

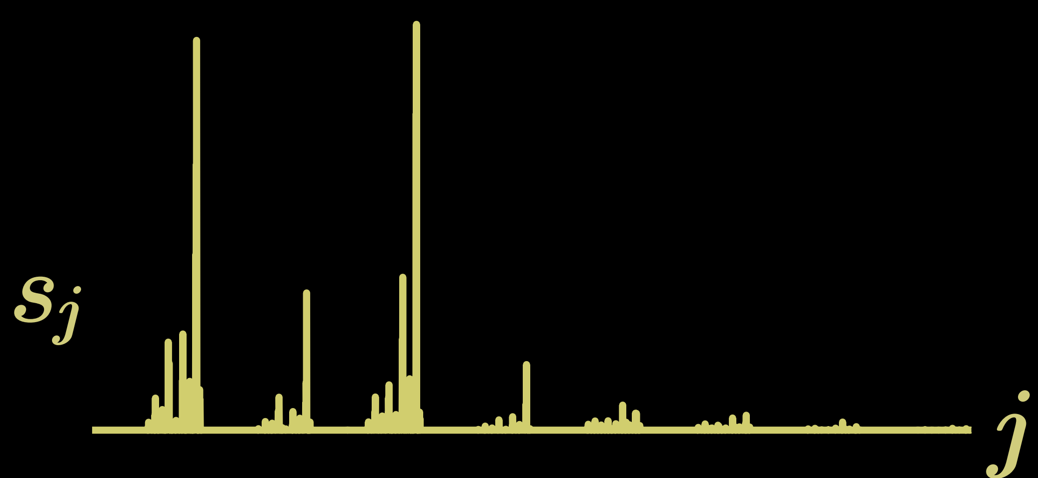

sparse probability $p(\boldsymbol{x})$

$$y_i = \sum_{\boldsymbol{x}} A_i(\boldsymbol{x}) p(\boldsymbol{x}) = \langle A_i \rangle$$

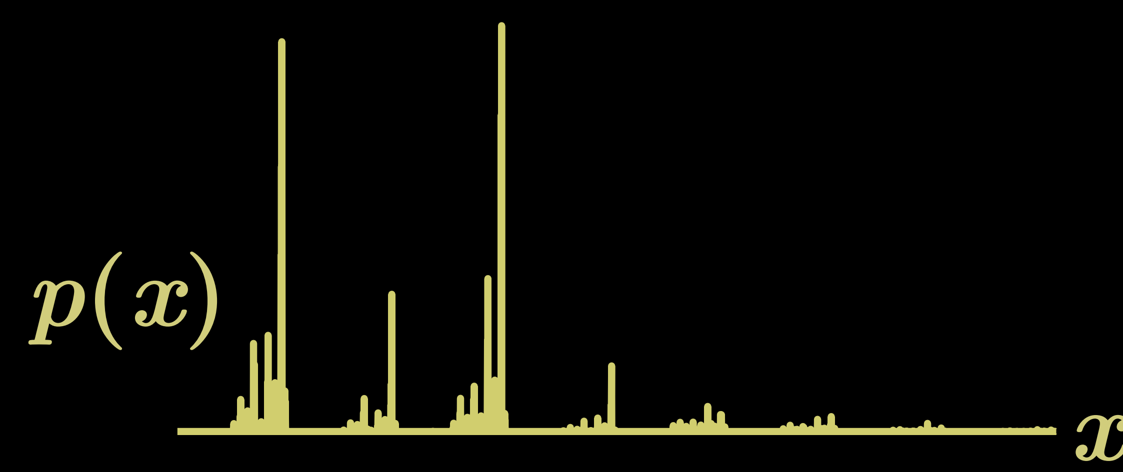

Information about $p(\boldsymbol{x})$ can be preserved by $\mathrm{O}(\boldsymbol{c}\,\boldsymbol{n}) \ll \mathrm{O}(2^{\boldsymbol{n}})$ measurements

Compressive Sensing GUARANTEE !

Random perceptrons

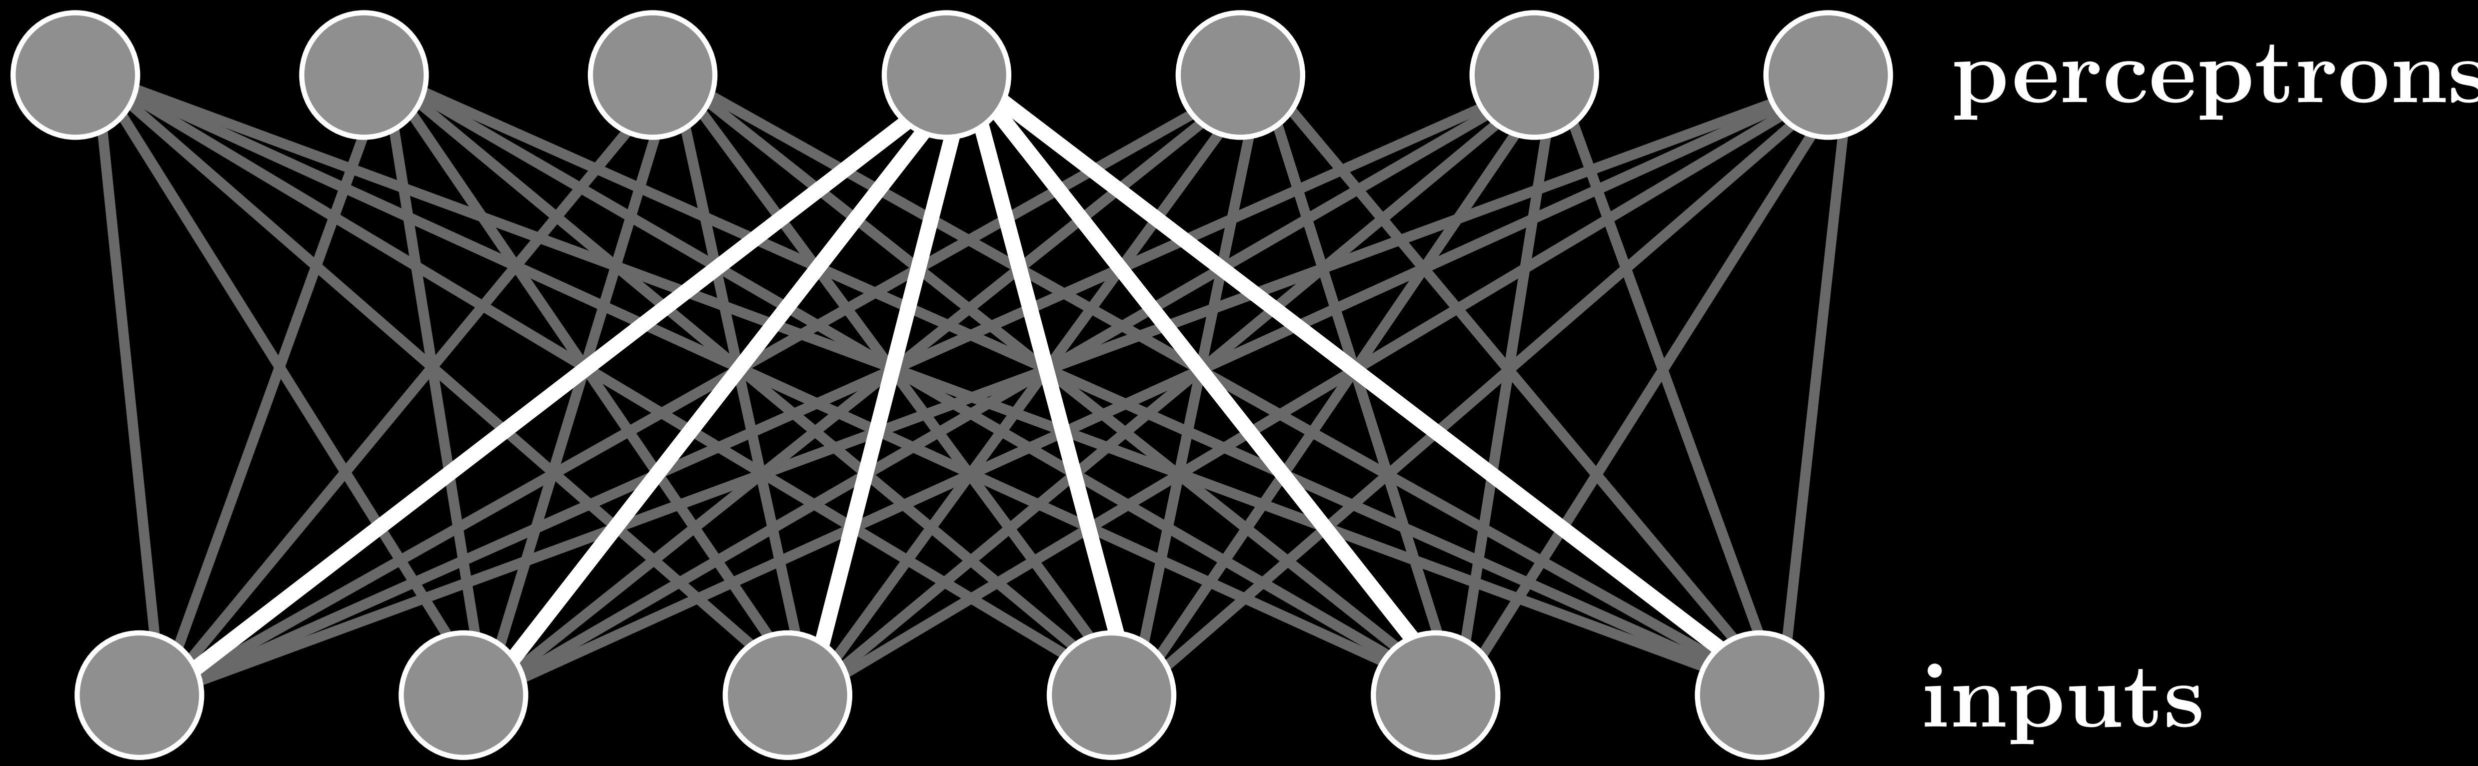

Sensing matrix $A_i(\boldsymbol{x}) = \mathrm{sgn}\left(\sum_j W_{ij} x_j\right)$

Verify information by L1 reconstruction*

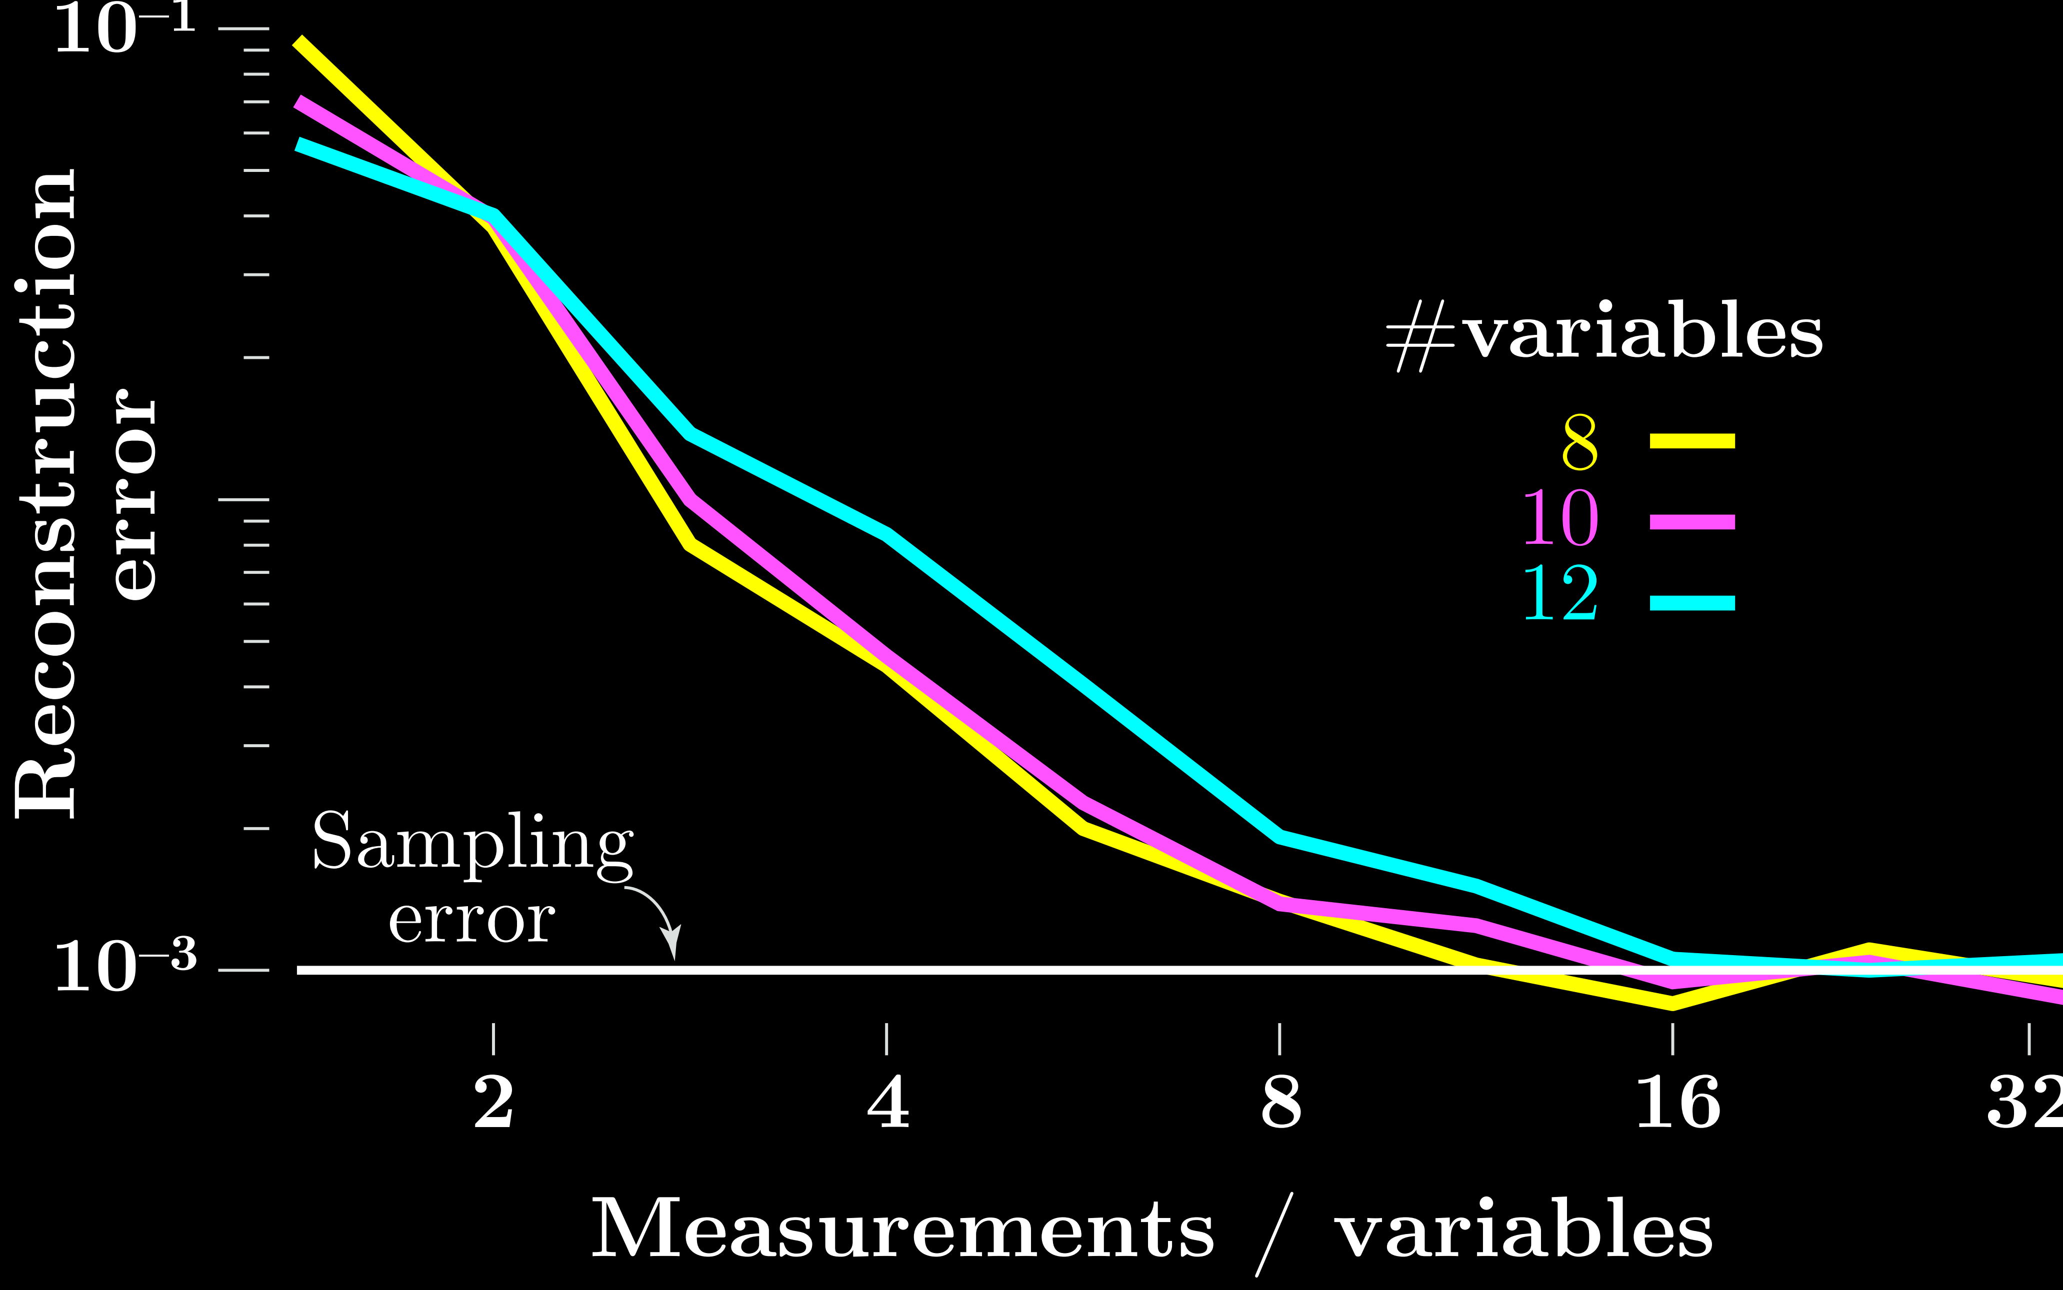

* $|p(\boldsymbol{x})|_{\mathrm{L1}}=1$! L1-minimization generally loses a little probability mass.

Compression preserves important relationships