What does this *actually do*?

It turns out this is nothing that new! Levy and Goldberg[6] show that SGNS is **implicitly factorizing** the matrix

$$M_{i,j} = \mathbf{w_i} \cdot \tilde{\mathbf{c}}_\mathbf{j} = PMI(w_i, c_j) - \log k = \log \frac{p(w_i, c_j)}{p(w_i)p(c_j)} - \log k$$

using an objective that weighs deviations in more frequent (w, c) pairs more strongly than less frequent ones.

Thus...

[6] Omer Levy and Yoav Goldberg. "Neural Word Embedding as Implicit Matrix Factorization". In: *Advances in Neural Information Processing Systems 27*. 2014, pp. 2177–2185.