

Negative Sampling⁵

Given a pair (w, c) , can we determine if this came from our corpus or not? Model probabilistically as

$$p(D = 1 | w, c) = \sigma(\mathbf{w} \cdot \tilde{\mathbf{c}}) = \frac{1}{1 + \exp(-\mathbf{w} \cdot \tilde{\mathbf{c}})}.$$

Goal: maximize $p(D = 1 | w, c)$ for pairs (w, c) that **occur in the data**.

Also maximize $p(D = 0 | w, c_N)$ for (w, c_N) pairs where c_N is **sampled randomly** from the empirical unigram distribution.

⁵Tomas Mikolov et al. "Distributed Representations of Words and Phrases and their Compositionality". In: *Advances in Neural Information Processing Systems 26*. 2013, pp. 3111–3119.