

# The Scourge of Sparsity

In the real world, **sparsity** hurts us:

- Low vocabulary coverage in training → high OOV rate in applications (poor generalization)
- “one-hot” representations → *huge* parameter counts
- information about one word *completely unused* for another!

Thus, a good representation must:

- reduce parameter space,
- improve generalization, and
- somehow “transfer” or “share” knowledge between words