

Is Quantized ANN Search Cursed?


Case Study of Quantifying Search and Index Quality

Gylfi Þór Guðmundsson / Reykjavik University, Iceland / gylfig@ru.is

Björn Þór Jónsson / Reykjavik University, Iceland / bjorn@ru.is




Goal: Understand and quantify the quality of deep index structures for large-scale high-dimensional search ... **when it fails!**



Search Can Fail in Three Ways


Representation Failure (RF)

No good representative available in index due to curse of dimensionality




Index Hierarchy Failure (IHF)

Index structure fails to identify the most similar representative



Segment Search Failure (SSF)

Search fails despite reaching the correct segment

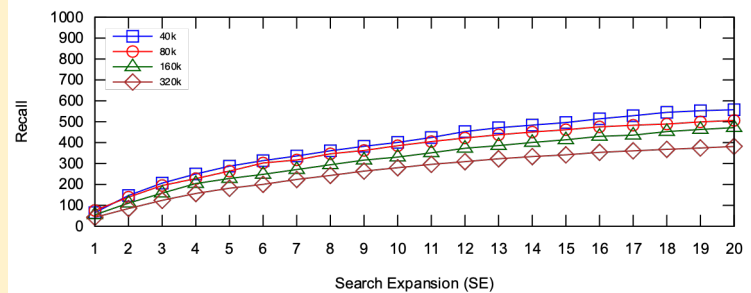


Key Contributions and Findings

Using **only the index**, we can:

- Quantify the different failures
- Estimate recall

We show that **RF** is the **primary source of failures** in eCP



Improving eCP's simple index structure will not repair RF failures