

Quantization in Append-Only Collections



Salman Mohammed, Matt Crane, and Jimmy Lin

University of Waterloo

tl; dr – Quantization substantially increases retrieval efficiency without effectiveness loss, but requires proper parameters. In an append-only collection, we can use old time blocks to predict values for new time blocks.

Quantization

For static collections we can pre-compute term/document weights and then quantize floating point values into integers.

Why? Ranking becomes super fast integer additions!

Previous work: the number of bits used for quantization represents a tradeoff between effectiveness and efficiency.

Append-Only Collections

Quantization substantially improves efficiency in static collections, but most collections are not static!

What about append-only collections?

Like Tweets!




- Treat collection as a sequence of time blocks
- Learn quantization parameters from previous block


Experimental Setup

- Tweets2013 collection: 254 million tweets from Feb to March, 2013 – treat each week as a time block
- Topics MB III-225 from TREC Microblog Tracks for evaluation

Results



Much like prior work, quantization range impacts both efficiency and effectiveness.



Although for Tweets2013 the number of bits is much smaller than predicted as needed.

Finding: We can use previous blocks to predict parameters for newer blocks and achieve minimal effectiveness and efficiency impact.

