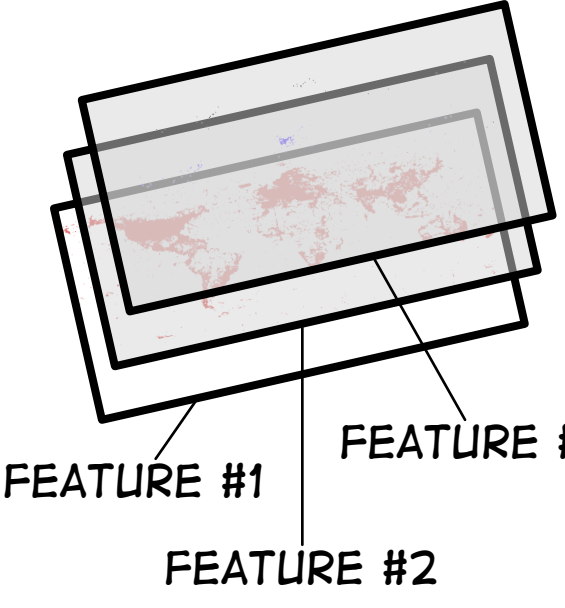
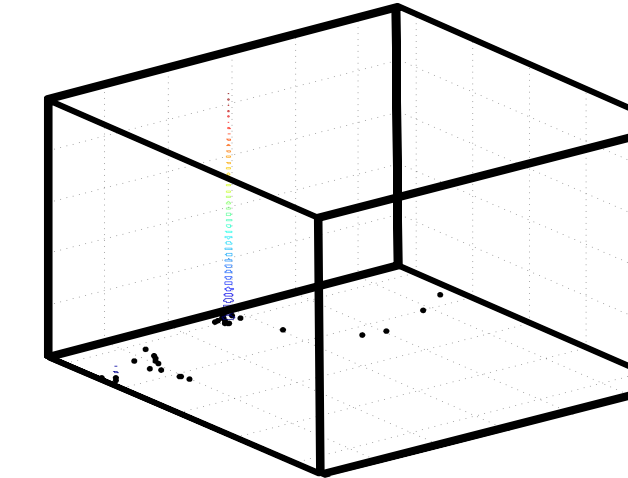


OVERALL APPROACH

MATCHING FEATURES PROVIDE SETS OF <LAT,LNG> PAIRS WHICH CAN BE SEEN AS DISCRETE SAMPLES DRAWN FROM A PDF OVER THE EARTH'S SURFACE

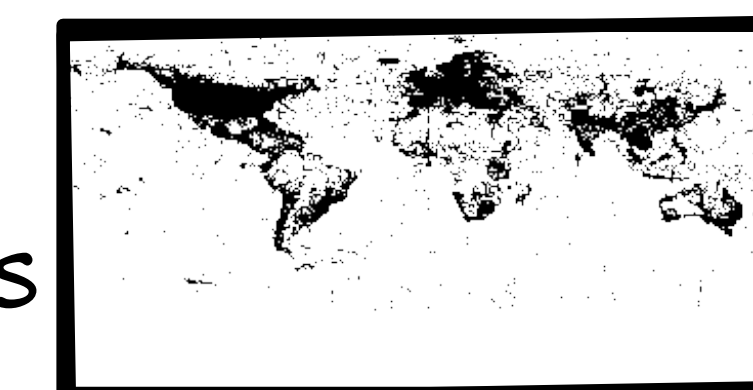


MEAN-SHIFT IS USED TO ESTIMATE THE MOST PROBABLE LOCATION FROM THE KERNEL DENSITY ESTIMATE OF THE PDF

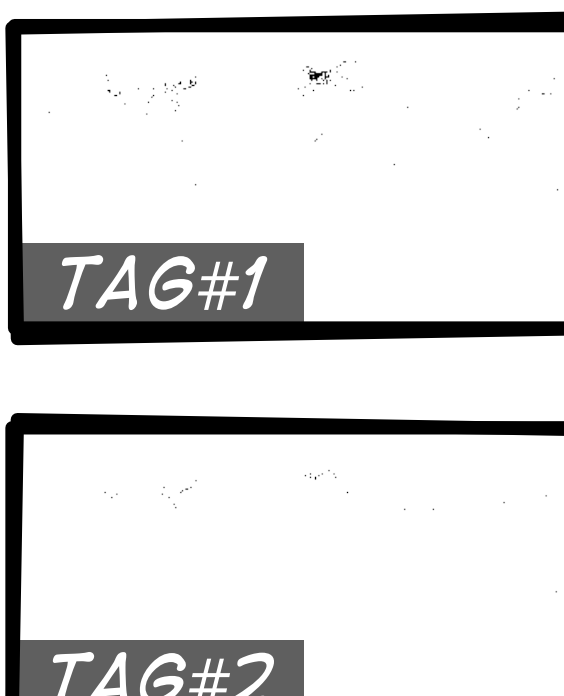


METADATA FEATURES

LOCATION PRIOR
A PRIOR FEATURE WAS CONSTRUCTED BASED ON WHERE PHOTOS ARE MOST LIKELY TO BE TAKEN

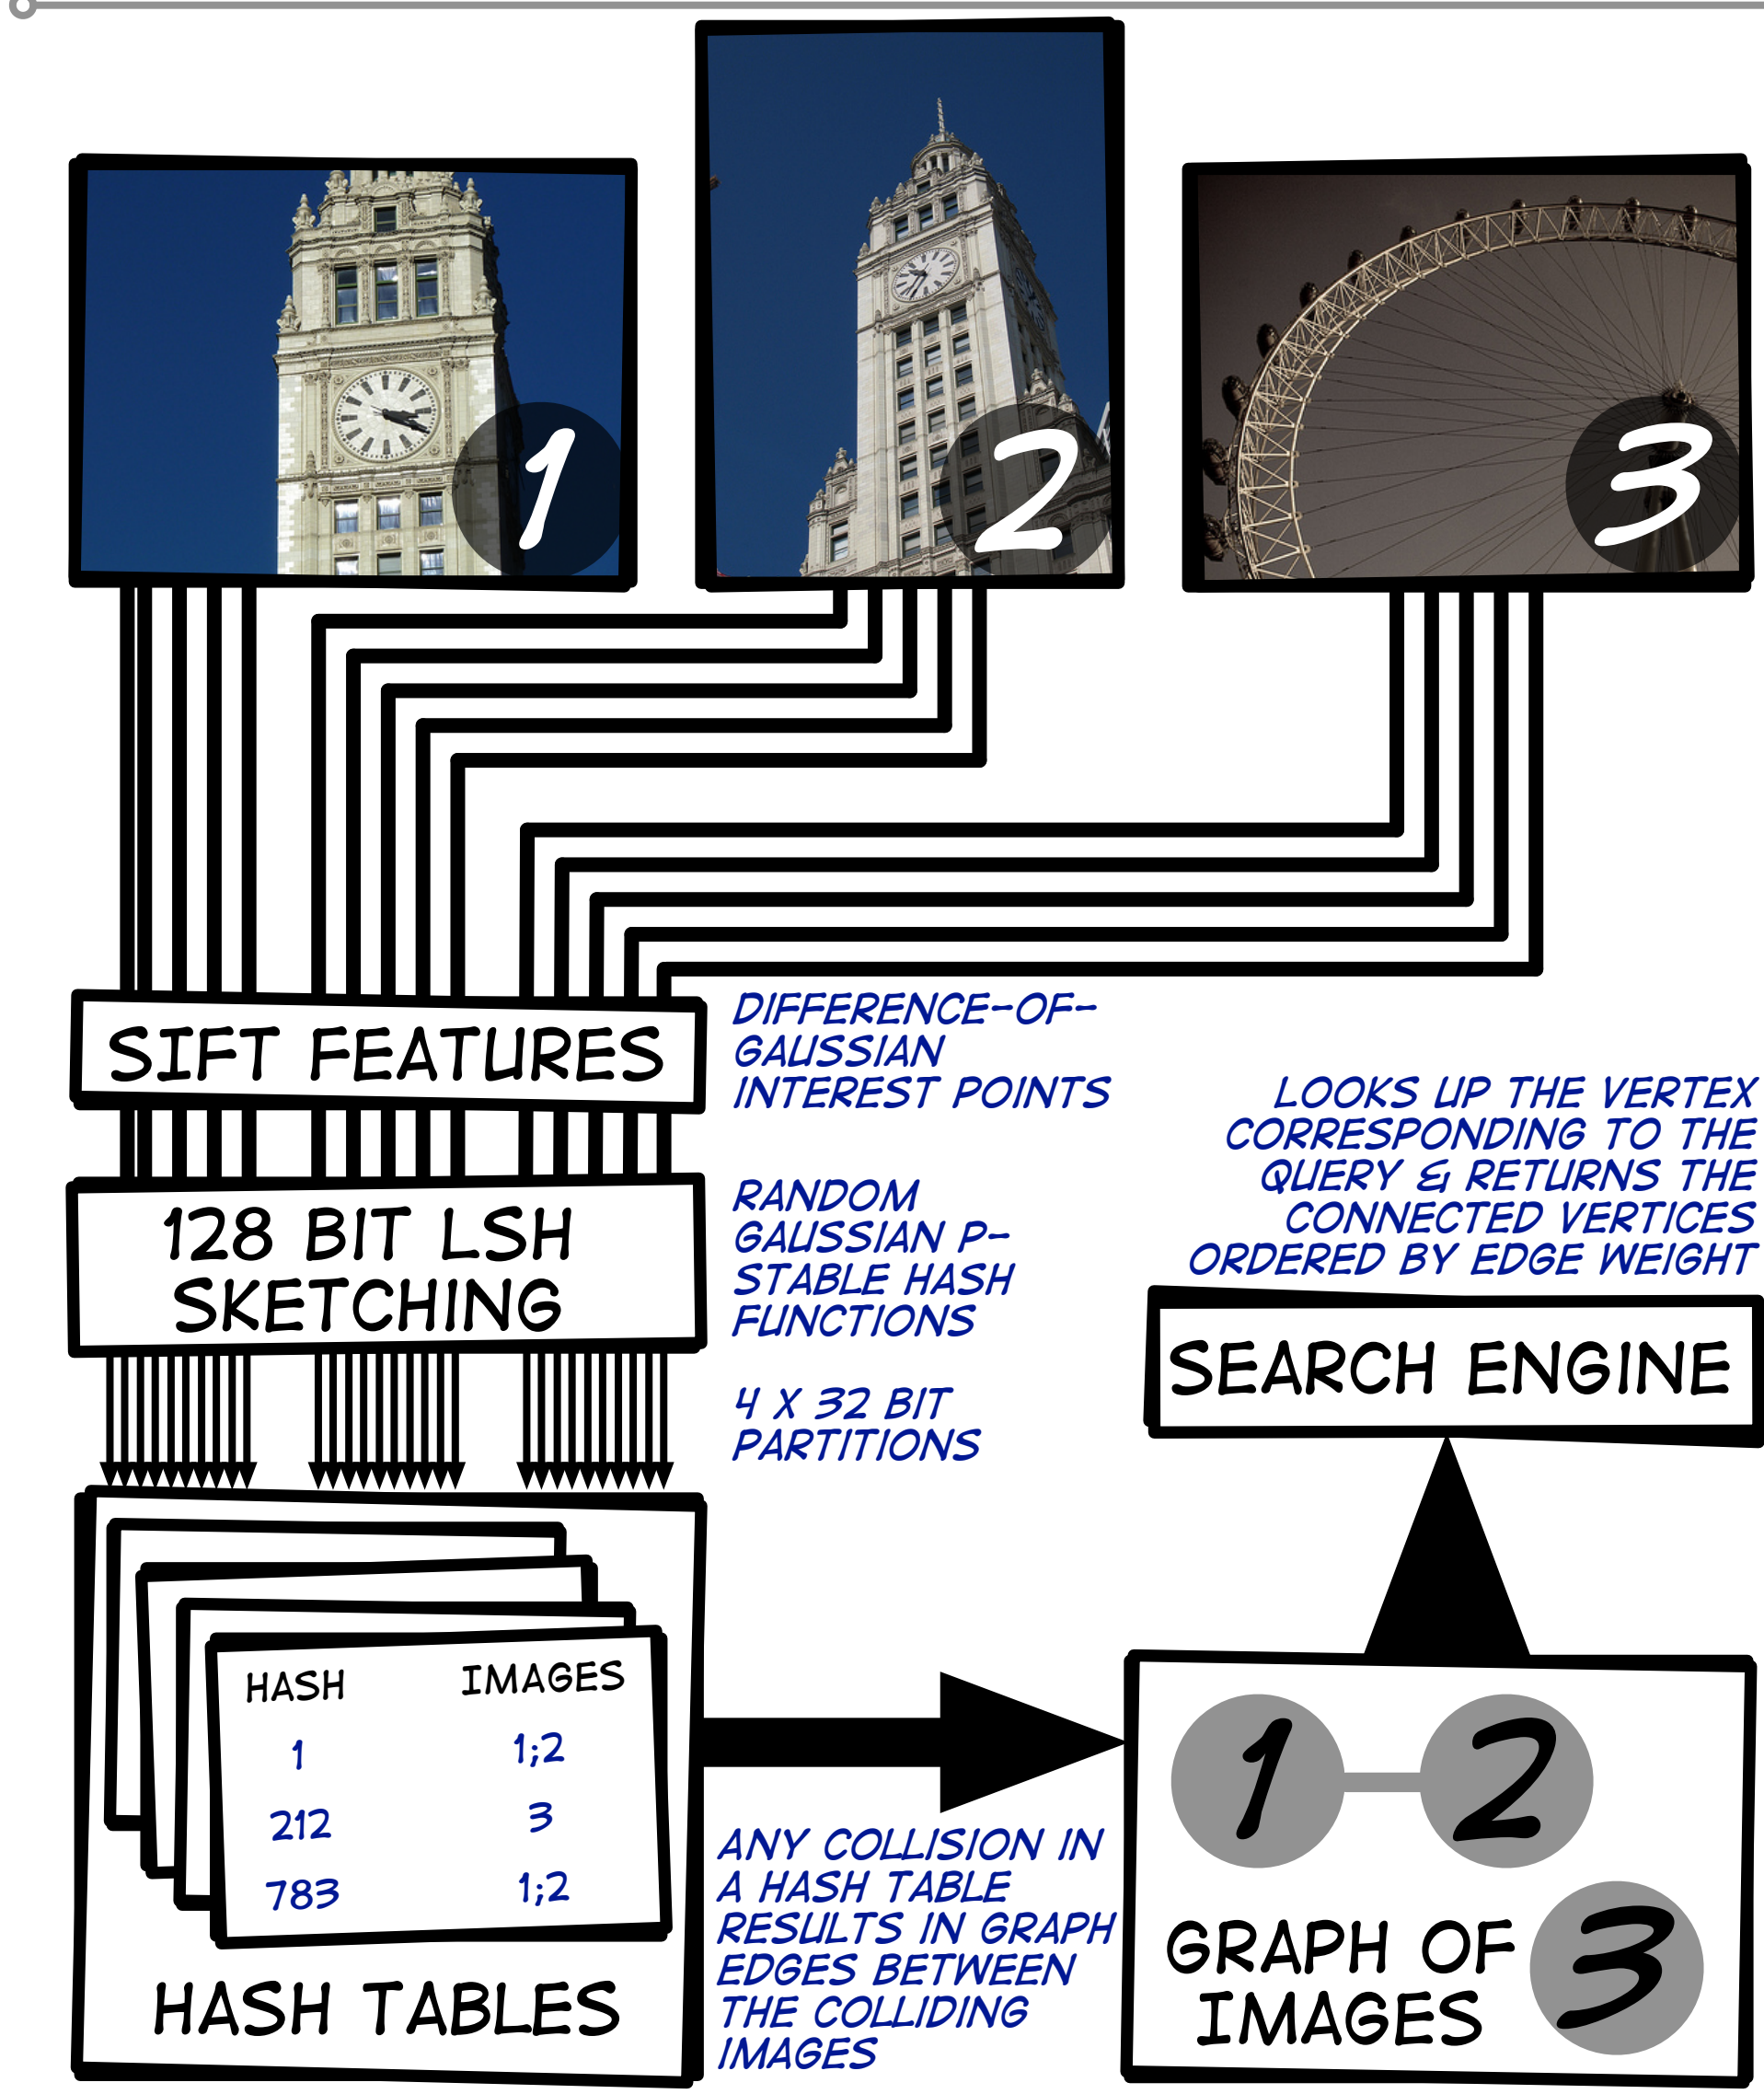


TAGS
EACH TAG IS TREATED AS A SEPARATE FEATURE AND PRODUCES A SET OF LIKELY LOCATIONS. ALL TAGS ARE CONSIDERED AS THEY GIVE SOME HINT, EVEN IF THEY ARE NOT RELATED TO A GEOGRAPHICAL PLACE - THE WORST CASE IS THAT THEY LOOK LIKE THE PRIOR.

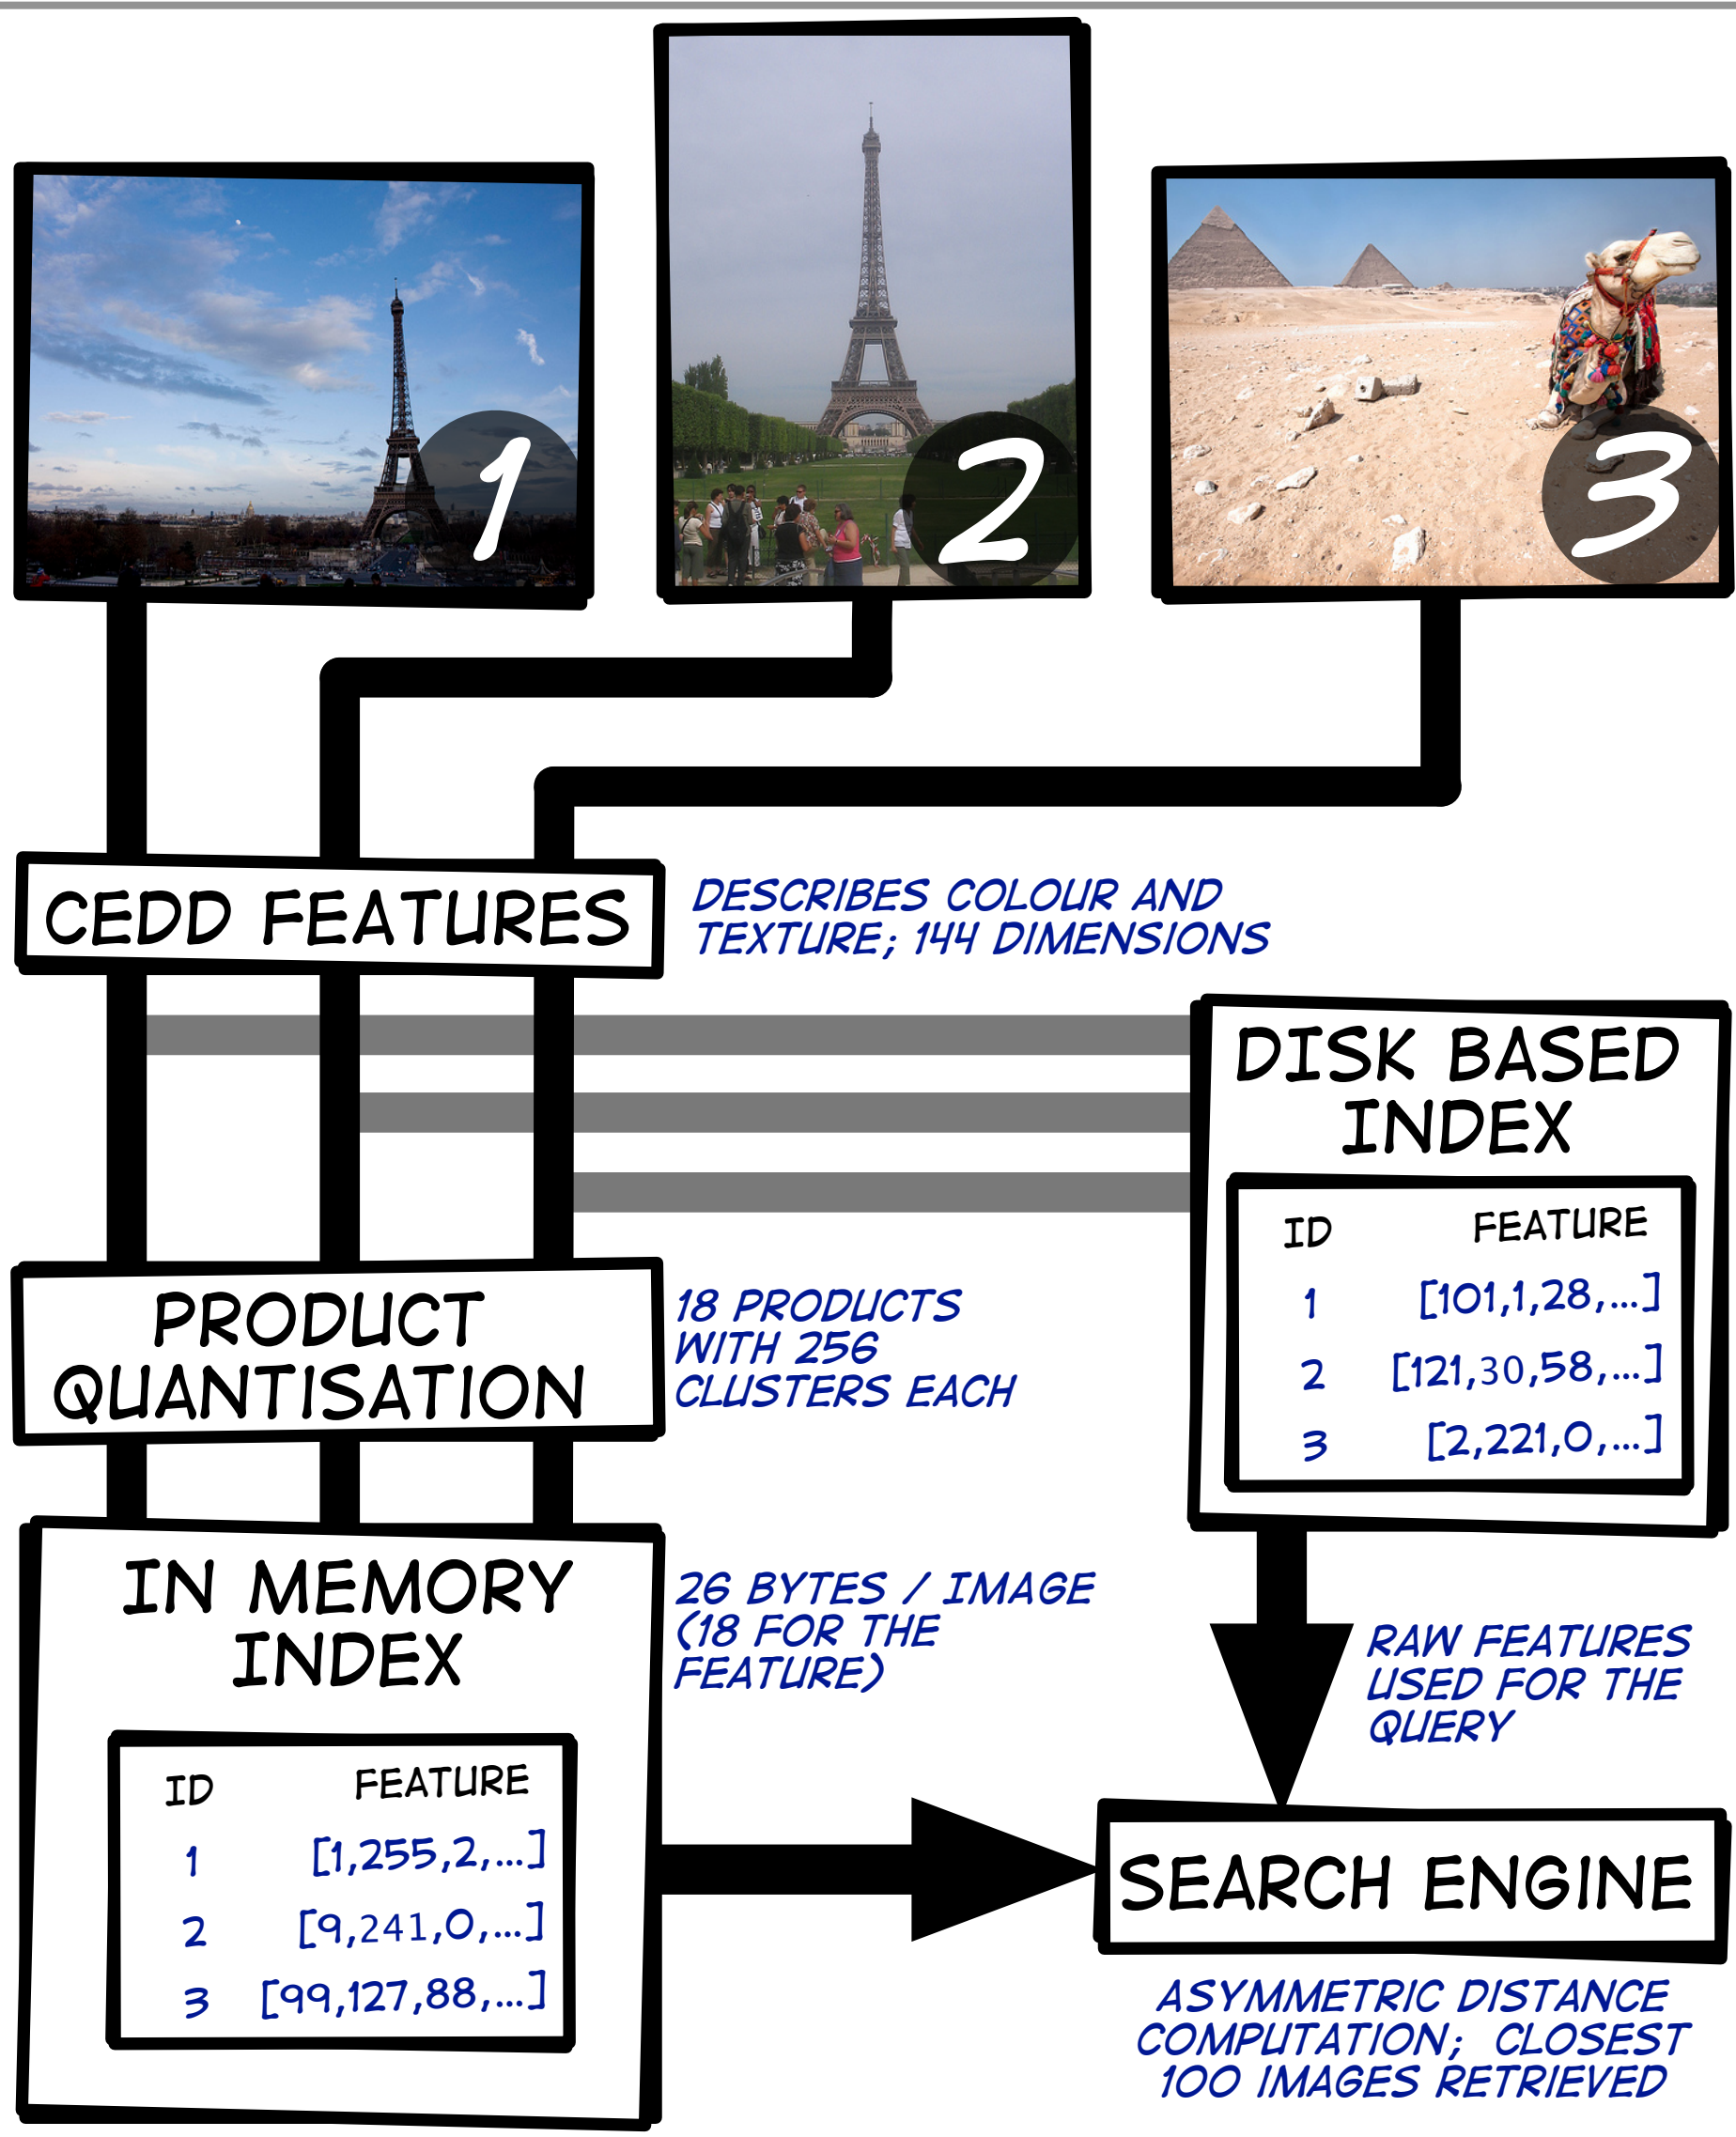


VISUAL FEATURES

LSH-SIFT: HIGH PRECISION/LOW RECALL SEARCH



PQ-CEDD: LOW PRECISION/HIGH RECALL SEARCH



NORMALISATION

ALL FEATURES ARE FORCED INTO PRODUCING EXACTLY 1000 <LAT,LNG> PAIRS BY SUB/SUPER SAMPLING. THIS GIVES EACH FEATURE AN EQUAL WEIGHT. SAMPLING IS UNIFORM FOR THE PRIOR AND TAGS FEATURES, BUT BASED ON THE SCORE RETURNED BY THE SEARCH ENGINE FOR THE VISUAL FEATURES.

GEONAMES BOOSTING
TAGS FEATURES THAT MATCHED A GENOMES PLACE OR ALTERNATE PLACE NAME HAD THEIR NUMBER OF SAMPLES DOUBLED TO 2000, EFFECTIVELY INCREASING THEIR POWER.

RUN CONFIGURATIONS

RUN	PRIOR	TAGS	PQ-CEDD	SIFT-LSH	GEONAMES
1	✓	✓	✓	✓	
2	✓		✓	✓	
3	✓	✓			
4*	✓	✓		✓	
5	✓	✓	✓	✓	✓

* RUN 4 USES A LARGER DATASET OF GEOTAGGED FLICKR IMAGES CRAWLED INDEPENDENTLY. ALL IMAGES FROM THE SET OF USERS FROM THE QUERY IMAGES WERE REMOVED FROM THIS DATASET TO KEEP TO THE SPIRIT OF THE TASK.

RESULTS

	1	2	3	4	5
NO. EST. WITHIN 1KM	20.6%	0.3%	23.2%	26.2%	23.7%
NO. EST. WITHIN 10KM	31.5%	0.6%	37.7%	40.9%	38.5%
NO. EST. WITHIN 100KM	36.1%	0.1%	43.8%	47.4%	44.2%
NO. EST. WITHIN 500KM	42.0%	3.5%	51.0%	53.5%	49.9%
NO. EST. WITHIN 1000KM	47.2%	7.4%	56.7%	58.4%	54.5%
MEDIAN ERROR IN KM	1352	6898	451	254	540
LINEAR CORRELATION OF ERROR	0.16	0.06	0.37	0.37	0.04