

Hashing and similarity learning for tracking with the HL-LHC ATLAS detector

Moritz Kiehn¹, S. Amrouche², N. Calace¹, T. Golling², A. Salzburger¹

 1 CERN

² Université de Genève



Why new tracking techniques?

Tracking is the core of almost every particle physics experiment Tracking at the High-Luminosity LHC means pile-up 200 Combinatorial reconstruction approaches explode

 \Rightarrow Better algorithms, better physics!





A typical hash function

h(complex object) = unique number

used e.g. in hash maps (Python dict, std::unordered_map)





A perfect tracking hash function

h(hit) = track number

does not exists.



Hashing?

But, approximate grouping of hits

$$h($$
track 1 hit 0 $) =$ group x
 $h($ track 1 hit 1 $) =$ group x
 $h($ track 0 hit 1 $) =$ group x

. . .

can be constructed.



Hashing as bucketing strategy



Random projections to create item hash

Groups similar items/ neighbors for a given metric



General: approximate nearest neighbors



KD-tree R-tree VA-file Locality Sensitive Hashing Product Quantization

Our choice: https://github.com/spotify/annoy







Reconstruction quality



TrackML dataset: $t\bar{t}$, $\mu = 200$

Select buckets/neighborhoods of hits using randomly drawn hits as query points

 $\frac{\text{Efficiency} =}{\frac{\text{trackable particles in buckets}}{\text{total trackable particles}}}$

Find (almost) all tracks with sufficient queries



Possible improvement: metric learning

Ex.: initial space



Ex.: learned space



Neighbor definition depends on a metric (so far: angular distance)

Learn metric/transformation based on known grouping for max. separability

 \Rightarrow Better metric, better buckets



Metric learning for tracks



Uses Local Fisher Discriminant Analysis

Spatially close tracks (left) separated in transformed space (right)

Caveat: uses TrackML dataset with additional hit features Not easily translatable to other cases (yet)



Can this run on accelerators?

(Approximate) nearest neighbor search is a general problem Think: similar images, movies, people, ...

E.g. "Billion-scale similarity search with GPUs" arXiv:1702.08734 from Facebook, based on a different technique



CPU/GPU comparison



TrackML dataset $t\overline{t}$, $\mu = 200$

GPU NVidia Tesla K40m 12Gb RAM, 2880 CUDA Cores

Separate CPU/GPU implementations, same input data, same algorithm

 \rightarrow Consistent bucket definition



GPU performance





Running on ATLAS ITk

 $tar{t}$ with $\mu=$ 200 and ATLAS ITk geometry

Use initial method to select buckets of hits

Apply ATLAS reconstruction in buckets As a proxy for (future, optimized) local reconstruction

ATLAS-IDTR-2019-008









Reconstruction in bucket vs full event



clusters on tracks with full event reconstruction versus restricting the reconstruction to a bucket of 50 hits.

Diagonal: fully matched tracks

 $Below: \ bucket \ track \ is \ missing \\ hit(s)$

Above: bucket track contains additional hit(s)

ATLAS-IDTR-2019-008





Tracking via approximate-nearest-neighbor methods: hashing and similarity learning

General, detector-independent techniques

Promising results on TrackML and ATLAS ITK upgrade simulation dataset

Ongoing: work towards full reconstruction chain





home.cern

Reconstruction in extended buckets



As before, but buckets extended to contain all truth hits for contained tracks to emulate optimized track road.

Diagonal: fully matched tracks

Below: bucket track is missing hit(s)

Above: bucket track contains additional hit(s)

ATLAS-IDTR-2019-008

