


LINEAR TIME ALGORITHMS FOR PAIRWISE STATISTICS (Ram, Lee, March, Gray, '09)


Use kd-tree for finding nearest neighbor in expected $O(\log N)$ (Freidman, Bentley, Finkel, '77)



Hierarchical tree for $O(N \log N)$ force calculation (Barnes & Hut, '86)

Fast Multipole method for efficient almost $O(N)$ particle simulation (Greengard & Rokhlin, '87)


Combine these ideas to get super efficient *dual-tree* algorithms for pairwise statistics (Gray, Moore, '00)



Cover trees can answer a nearest neighbor query with a rigorous $O(\log N)$ runtime bound (Beygelzimer, Kakade, Langford, '06)

In this paper, we prove that dual-tree nearest neighbor with cover tree is actually $O(N)$...

Dual-tree nearest neighbor with cover trees might be $O(N)$



... and so is dual-tree approximate kernel summation with cover trees and dual-tree approx N-body potential calculation.



Yes! We have Linear-Time Algorithms for Pairwise Statistics