



Cartification: from Similarities to Itemset Frequencies (Bart Goethals)

Suppose we are given a multi-dimensional dataset. For every point in the dataset, we create a transaction, or cart, in which we store the k -nearest neighbors of that point for one of the given dimensions. The resulting collection of carts can then be used to mine frequent itemsets; that is, sets of points that are frequently seen together in some dimensions. Experimentation shows that finding clusters, outliers, cluster centers, or even subspace clustering becomes easy on the cartified dataset using state-of-the-art techniques in mining interesting itemsets.

Bart Goethals is a professor at the University of Antwerp.