More on Exercise 3.3

We start with a reminder.

Exercise 3.3: Largest common point sets under ε-congruence Given two finite sets of points A and B in the plane, write a program that finds equally sized subsets $A' \subseteq A$ and $B' \subseteq B$ of maximal cardinality such that points in A' match points in B' under translation, up to distance at most some given ε. Namely, each point a in A' has a unique point b in the translated B' (and vice versa) such that the Euclidean distance between a and b is at most ε and A' is the largest cardinality such subset.

More technical details about the program (input/output files etc.) can be found at the course's site. Here we outline a possible direction for solution.

As we typically approach a variety of problems in computational geometry, the first step is to rephrase the problem in a "convenient" space where we can get helpful insights. In this case this is the space of all possible translations in the plane, which we denote by $\mathcal{T}$, and which is two dimensional. Each point $t = (x_t, y_t) \in \mathcal{T}$ represents a translation (of the set B) by the vector t.

Given a pair of points $a_i \in A$ and $b_j \in B$ we denote by D_{ij} all the translation vectors which when applied to b_j will make it ε-congruent to a_i, that is

$$D_{ij} = \{t \in \mathcal{T} | \rho(a_i, b_j + t) \leq \varepsilon\} ,$$

where ρ is the Euclidean distance between a pair of points. It is easily verified that D_{ij} is a disk in the translation space.

Let m and n denote the number of points in A and B respectively. Let $\mathcal{D}$ denote the collection of the mn disks D_{ij} induced by all pair of points one from A and one from B. Let C_{ij} be the boundary circle of D_{ij}, and let $\mathcal{C}$ be the collection of the mn circles C_{ij}. Our main object of study is the arrangement $\mathcal{A}(\mathcal{C})$ of circles in the translation space. (We will refer both to the disks and to the circles, as needed.)

Consider now the following question, which is a restricted version of the Exercise, where we fix the translation. Given a translation vector t_0, find equally sized subsets $A' \subseteq A$ and $B' \subseteq B$ of maximal cardinality such that points in A' match points in B' under translation *by the vector* t_0, up to distance at most ε.

The solution can be obtained as follows. We construct a bipartite graph: The nodes of the graph correspond to points of A and B. We connect two nodes corresponding to the points a_i and b_j iff $t_0 \in D_{ij}$. That is, we go over all the disks in $\mathcal{D}$, and for each disk that contains t_0 (either inside or on its boundary) we add the respective edge to the graph. It remains to find a maximum matching in this graph, which will yield the desired solution. We have to repeat this process for every translation vector $t \in \mathcal{T}$. A nice feature of the arrangement $\mathcal{A}(\mathcal{C})$ defined above is: *Given any fixed cell (vertex, edge, or face) of the arrangement, all the points inside it induce exactly the same bipartite graph.* This means that the arrangement $\mathcal{A}(\mathcal{C})$ gives us an effective discretization of the problem without giving up any information. Now it remains to go over all the cells of the arrangement, and for each cell sample one point, construct its relevant bipartite graph and proceed as in the single vector case described above. This process can be sped up by observing that between two adjacent cells the graph changes only a little.