Math 168 Weekly assignment #9

About Networks and Frameworks

Reminder: it does not suffice to simply provide answers, you need to justify/explain why your answers are correct. And do not forget your sources and acknowledgments.

Recall that a *network* is a set V of "vertices" together with a set E of two-element subset of V, called the "edges." In class we define (I will use present tense throughout, because whether it should be future or past tense depends on when you are doing these problems) an *edge-weighted network* to be a network together with a function $l : E \to \mathbb{R}^+$, giving the "length" of each edge. If M is a metric space, then an *M-framework* (on a set of vertices V) is just a map $p : V \to M$. The key definition: we say that an M-framework *realizes* an edge-weighted network if for all $\{u, v\} \in E$, $d(p(u), p(v)) = l(\{u, v\})$. We will really only use the two cases in which M is either the Euclidean plane or Euclidean 3-space, in which case we just say refer to a "plane framework" or a "space framework." Important note: the "edges" of a Euclidean framework (i.e., the line segments connecting the images of vertices that are connected by edges in the network) are allowed to intersect, or even completely overlap; a vertex is allowed to lie on an edge it is not a part of, etc.

1. Choose any network you like that has at least one cycle containing four or more edges, and give an explicit example of an edge-weighted network based on that network, and a space framework that realizes that edge-weighted network.

An edge-weighted network is said to be *globally M-rigid* if for any two M-frameworks $p(), q()$ realizing it, there is a symmetry σ of M such that for all $v \in V$, $\sigma(p(v)) = q(v)$. That turns out to be a very strong condition – not many edge-weighted networks satisfy it. Hence, a more useful notion (because examples satisfying it come up much more frequently) is the following: An M-framework p realizing an edge-weighted network N is said to be *locally rigid*, or just *rigid*, if there is a a real number $\epsilon > 0$ such that if q is any other framework realizing N such that for all vertices v, $d(p(v), q(v)) < \epsilon$, then there is a symmetry σ of M such that for all $v \in V$, $\sigma(p(v)) = q(v)$. The introduction of the parameter epsilon captures the fact that a framework will "hold its shape" if there are no essentially different configurations "nearby" that framework.

2. Give an edge-weighted network with more than one edge which is both globally plane-rigid and globally space-rigid.

3. Give a plane framework and associated edge-weighted network which is locally rigid but such that the edge-weighted network is *not* globally plane-rigid. (Hint: the easiest example I can think of has exactly two realizations, and consists of two parts, each globally rigid by itself, but which share one edge and so they can fit together in two different ways. Remember to argue why the framework you designate is locally rigid.)

4. Give a space framework and associated edge-weighted network which is locally rigid, such that the edge-weighted network is not globally space-rigid. (Hint: will basically the same trick, adapted to three dimensions, work?)

Once we have a plane or space framework, we can think about it moving around in space. A *motion* of a framework is just an assignment of a direction (vector) $m(v)$ to each vertex v of the underlying network. However, some motions would tend to stretch or compress the edges. We're interested in motions that don't do this, at least not initially. They are called "infinitesimal motions" of the framework. The details of this definition are motivated in class, but we need the dot product (also called the "inner product") of two vectors. In two dimensions this is $(x, y) \cdot (u, v) = xu + yv$ and in three it is $(x, y, z) \cdot (u, v, w) = xu + yv + zw$. Then we say that a motion m of a framework is an infinitesimal motion for the underlying network if for all edges $\{u, v\}$ in the network, $(p(u) - p(v)) \cdot (m(u) - m(v)) = 0$. However, some of the infinitesimal motions of a network are not interesting because they come from a global infinitesimal motion of the entire space that always produces an infinitesimal motion of every network. In class we show that the infinitesimal motions of a Euclidean space are the constant motions, assigning the same vector to every point, and in the plane, infinitesimal rotations around a point (a, b) given by $g((x, y)) = (b - y, x - a)$, and in three dimensions, infinitesimal rotations around the line in direction d through the point c, given by $g(t) = (t - c) \times d$, where $\times$ is the usual "cross product" in three dimensions. If an infinitesimal motion m has the property that for all vertices u, $m(u) = g(p(u))$, then we say that m is a trivial infinitesimal motion. Finally, we say that a framework is infinitesimally rigid if it has no non-trivial infinitesimal motions. It has been proven that an infinitesimally rigid framework is rigid. On the other hand:

5. Show that the following plane framework is rigid, but not infinitesimally rigid. (Hint: since there are more than two vertices, any motion in which the motion of only one vertex u is non-zero must be non-trivial (explain why). And then to check such a "singly supported" motion is an infinitesimal motion, you only need to check the condition on the edges leaving that one vertex, which becomes $(p(u) - p(v)) \cdot m(u) = 0$, meaning that the direction of the motion of u is perpendicular to the direction of the edges connected to u. And don't forget to argue why the framework is actually rigid.) The underlying network has five vertices $\{0, 1, 2, 3, 4\}$. The edges of the network are $\{\{0, 1\}, \{0, 2\}, \{0, 3\}, \{1, 2\}, \{1, 3\}, \{2, 4\}, \{3, 4\}\}$. The edge weighting is $l(\{0, 1\}) = 1$,$l(\{0, 2\}) = l(\{0, 3\}) = \sqrt{2}, l(\{1, 2\}) = l(\{1, 3\}) = \sqrt{5}, l(\{2, 4\}) = l(\{3, 4\}) = 1$. And finally, the position map of the framework is given by

vertex	0	1	2	3	4
position	(0,1)	(0,2)	(-1,0)	(1,0)	(0,0)

.

6. Give an example of a space network that is rigid but not infinitesimally rigid. You don't need to do a full-blown detailed demonstration of both facts, but give intuitive arguments for both parts (why rigid, why not infinitesimally rigid).

7. Let p be a plane framework with at least one edge, and pick any particular edge $\{u, v\}$. Show that p is infinitesimally rigid if and only if there is no non-zero infinitesimal motion m such that $m(u) = m(v) = 0$. (Intuitively, this says that a plane framework is infinitesimally rigid if and only if it becomes completely fixed if you "pin" two of its vertices connected by an edge.)

8. Intuitively, we can count the number of internal degrees of freedom of a framework by pinning two of its vertices, and then counting the number of coordinates of the other $m(w)$ values we have to specify before the entire infinitesimal motion is uniquely determined. A plane framework with n vertices and no edges has $2n - 3$ internal degrees of freedom, and typically each added edge reduces the number of degrees of freedom by one, but it does not have to (intuitively, it might be "redundant" with other edges in determining the configuration of the network). Formally, we can at least define "one internal degree of freedom": say that a plane framework has one internal degree of freedom if its network has an edge $\{u, v\}$, it has at least one non-zero infinitesimal motion satisfying $m(u) = m(v) = 0$,("$\{u, v\}$ is pinned") and if m_1 and m_2 are two distinct non-zero infinitesimal motions satisfying this condition that $\{u, v\}$ is pinned, then no non-zero coordinates of $m_1(w)$ and $m_2(w)$ coincide for any vertex w of the network. Give an example of a plane framework on five vertices with one internal degree of freedom, even though it has 7 edges. (That is interesting because (10-3)-7=0, so at least one of the edges did not reduce the number of internal degrees of freedom.) Can you describe why/how some edge or edges are redundant in determining the degrees of freedom?

We'd like to make the intuitive ideas of the previous example more definite. To do so, we need some basics of linear algebra. A *real vector space* is a commutative group V of "vectors," the binary operation and identity of which are written "+" and "0," respectively, together with a map from $\mathbb{R}\backslash\{0\}$ to automorphisms (= symmetries) of V; we write just rv for the automorphism corresponding to r applied to the vector v. We extend this operation to all of $\mathbb{R}$ by setting $0v = 0$ (where the left-hand 0 is a real number, and the right-hand 0 is the identity vector), and call that operation "scalar multiplication." Further, these operations must satisfy two axioms: $r(sv) = (rs)v$ and $(r + s)v = rv + sv$, where r and s are real numbers and v is any vector.

9. Explain why, based on the above definition, in any vector space if v and w are vectors, and r is a real number, then $r(v + w) = rv + rw$.

In any vector space, take any set of vectors S. The *span* of S is the set of vectors you can produce by *finite linear combinations* from S, namely $\{r_1v_1 + \cdots + r_nv_n | r_1, \ldots, r_n \in \mathbb{R}, v_1, \ldots, v_n \in S\}$. (The numbers $r_1, \ldots, r_n$ are called the *coefficients* of the linear combination.) The set S is called (*linearly*) *dependent* if there is a finite linear combination with coefficients not all equal to 0, such that this finite linear combination is the 0 vector, and (*linearly*) *independent* otherwise. We show in class that if S is linearly independent, then every vector in span(S) has a *unique* representation as a finite linear combination from S. If span(S) is the entire vector space, then we say that S is a *basis* for the vector space.

10. Show that if a vector space has a finite basis, then every basis for that vector space has the same number of elements. (We call this number the "dimension" of the vector space.)

11. If a subset of a vector space is closed under vector addition and scalar multiplication, then it becomes a vector space in its own right, and is called a vector subspace. Show that a vector subspace of a finite-dimensional vector space is also finite dimensional, with dimension less than or equal to the dimension of the original space.

If we pick a basis $\{v_1, \ldots, v_n\}$ for a finite-dimensional vector space, then the coefficients $r_1, \ldots, r_n$ of the unique linear combination $r_1 v_1 + \cdots + r_n v_n = v$ are called the *coordinates* of v for that basis. For any n, $\mathbb{R}^n$ is a vector space: we add elements (vectors) componentwise, and perform scalar multiplication by multiplying every component by the scalar. It has a "canonical" basis, namely $(1, 0, 0, \ldots, 0), (0, 1, 0, \ldots, 0), \ldots, (0, \ldots, 0, 1, 0), (0, \ldots, 0, 0, 1)$. With this basis, the coordinates of a vector are just its components. Finally, any vector space of dimension n is isomorphic to $\mathbb{R}^n$ by choosing a basis and then mapping a vector to its vector of coordinates for that basis.

12. Show that the map described just above really is an isomorphism; in other words, show that it is bijective and preserves vector sums and scalar products.

In general, a structure-preserving map of vector spaces (that preserves vector addition and scalar multiplication) is called a *linear* map. Given a basis for each of the source and target vector spaces, say of dimension n and m, respectively, such a map can always be represented by a *matrix* with m rows and n columns: the first column consists of the coordinates of the image of the first basis vector of the source, the second column consists of the coordinates of the image of the second basis vector, and so on. You can then calculate the value of the map on any vector by writing the coefficients of that vector in a column, and performing ordinary multiplication of the matrix by the column vector, to yield the coefficients of the image of the original vector. Let V and W be finite-dimensional vector spaces, and $l : V \to W$ a linear map. Note that if $l(v) = 0$ and $l(u) = 0$, then $l(u + v) = 0$ and for any real r, $l(rv) = 0$. So the collection of vectors which map to 0 is a vector subspace of V, called the *kernel* of l, which by problem 11 has a dimension less than or equal to that of V. Conversely, if w and t are two vectors in the image of l,then $t + w$ and rw are both in the image of l, so the image is a vector subspace of W. The dimension of the image is called the *rank* of l, and a key fact from linear algebra (the "rank-nullity law") that we need to use is that the dimension of the kernel of l plus the rank of l is always equal to the dimension of V. If M is a matrix, then it represents a linear map, and we also call the rank of this linear map the *rank* of M.

The big payoff here is that there is a linear map R called the *rigidity* map from the vector space of all motions of a framework to $\mathbb{R}^e$ where e is the number of edges of the underlying network, such that the infinitesimal motions of the framework are exactly the kernel of R. The map is directly related to the condition $(p(u) - p(v)) \cdot (m(u) - m(v)) = 0$ to be an infinitesimal motion. The matrix for R is very easy to write down. It has $2n$ or $3n$ columns (depending on whether this is a plane or a space framework) and e rows. The first two or three columns correspond to coordinates of the first vertex of the network, the next two or three to coordinates of the second vertex, and so on. The ith row pertains just to the ith edge. And if the ith edge is $\{v_j, v_k\}$,then the ith row consists of all 0s except that in the columns for the jth vertex, you put the coordinates of $p(v_j) - p(v_k)$ and in the columns of the kth vertex, you put the coordinates of $p(v_k) - p(v_j)$. Now you can investigate this matrix with any linear algebra software. For example, you can compute its rank, and then use the rank-nullity law mentioned above to determine the dimension of the kernel of R, from which you can immediately read off whether it is infinitesimally rigid (recall that the dimension of the trivial infinitesimal motions of the plane is 3, and of space is 6). Free packages you might use to do the rank computation include, but are not limited to, R (http://www.r-project.org/), Octave (http://www.gnu.org/software/octave/), Euler (http://euler.rene-grothmann.de/), Maxima (http://maxima.sourceforge.net/), and Sage (http://www.sagemath.org/).

13. Consider the following modified framework from problem 5. Compute the rank of its rigidity map. Is this framework rigid? The underlying network is the same as in problem 5, but the edge-weighting is slightly different: $l(\{0,1\}) = 1, l(\{0,2\}) = l(\{0,3\}) = \sqrt{2}, l(\{1,2\}) = l(\{1,3\}) = \sqrt{5}, l(\{2,4\}) = l(\{3,4\}) = \sqrt{2}$. And finally, the position map of the framework is given by

vertex	0	1	2	3	4
position	(0,1)	(0,2)	(-1,0)	(1,0)	(0,-1)

.

What's interesting about edges $\{2,4\}$ and $\{3,4\}$ in the underlying network, particularly in light of your conclusions about rigidity?

14. Consider the following space framework. Compute the dimension of the kernel of its rigidity map, and hence its internal number of degrees of infinitesimal freedom (which by the above is just the dimension of the kernel minus six). Do you think this framework is rigid?

The underlying network has five vertices $\{0, 1, 2, 3, 4\}$. The edges of the network are $\{\{0, 1\}, \{0, 2\}, \{0, 3\}, \{1, 2\}, \{1, 3\}, \{1, 4\}, \{2, 3\}, \{2, 4\}, \{3, 4\}\}$. The edge weighting is $l(\{0, 1\}) = l(\{0, 2\}) = l(\{0, 3\}) = l(\{1, 2\}) = l(\{1, 3\}) = l(\{2, 3\}) = 1,\ l(\{1, 4\}) = l(\{2, 4\}) = l(\{3, 4\}) = 1/\sqrt{3}$.

And finally, the position map of the framework is given by

vertex	0	1	2	3	4
position	$(0, 0, \sqrt{2/3})$	$(1/\sqrt{3}, 0, 0)$	$(-1/\sqrt{12}, 1/2, 0)$	$(-1/\sqrt{12}, -1/2, 0)$	$(0, 0, 0)$

.

15. What software package did you use to compute the ranks/dimensions in problems 13 and 14? Had you used the package prior to this course? Was it easy/hard to use? (Give some detail about your experience with the package.)