Triangulations, rigid motions, and applications to representation theory

Asilata Bapat

July 1, 2022

1 Triangulations of points in a convex arrangement

Consider a configuration of n distinct points in $\mathbf{R}^2$, in a convex arrangement. By a *diagonal* or *edge* we mean a straight line segment joining two of the points. An edge can be either *external* (i.e., on the boundary of the convex hull), or *internal*.

Definition 1.1. A *triangulation* of a given convex configuration of n points in the plane is a maximal set of diagonals that are pairwise non-crossing. That is, no two diagonals in the set intersect in their interiors.

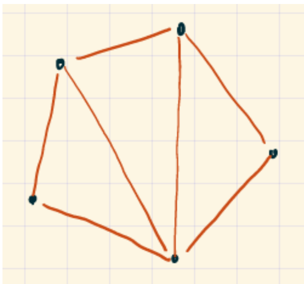

1.1 Combinatorics of triangulations

Given a convex arrangement of n points, the number of triangulations has a well-known formula.

Exercise 1.2. Find the number of possible triangulations of a convex arrangement of n points in the plane.

Note that the definition of a triangulation makes no reference to triangles!

Exercise 1.3. Show that a triangulation divides the convex hull of the configuration into triangular regions.

Using the previous exercise, we see that each internal edge in a triangulation has one triangular region on either side. That is, each internal edge forms the diagonal of a convex quadrilateral.

Definition 1.4. Let T be a triangulation of a fixed convex configuration of n points. Let a be an internal edge. The *edge-flip* of a in T is the unique edge a' that is the other diagonal of the convex quadrilateral formed by the two triangles bordering a.

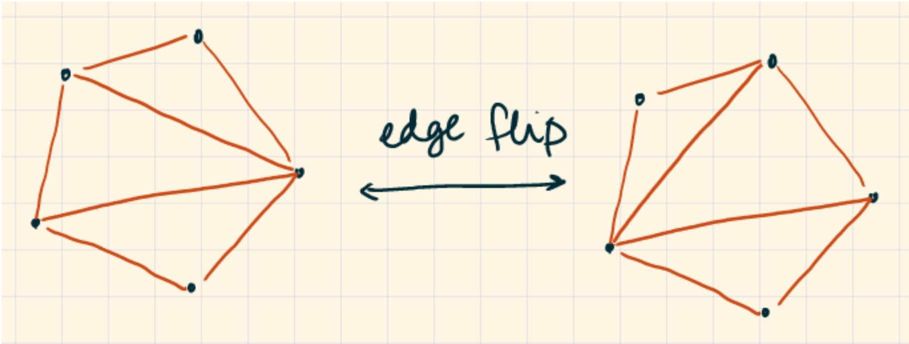

Note that the edge flip a' of a in T crosses a. It is not too hard to see that the set $(T \setminus \{a\}) \cup \{a'\}$ is also a triangulation.

Exercise 1.5. Let T be a triangulation, and let a be an internal edge. Prove that there is a unique edge a' such that $(T \setminus \{a\}) \cup \{a'\}$ is a triangulation, and that a' is precisely the edge flip of a in T.

We can also make the following observation.

Proposition 1.6. *Any triangulation of a convex n-gon has $2n-3$ edges and $n-2$ triangles.*

Proof. Consider any triangulation of a convex n-gon. Suppose that it has e edges and f triangular faces. Recall by Euler's formula for a planar graph, that we have

$$n - e + f = 1.$$

Moreover, each face has three edges, and each edge except for the external edges (of which there are n) borders two faces. We see that

$$3f = 2e - n.$$

It follows from the two equations that $e = 2n - 3$, and $f = n - 2$. □

We see that every triangulation has $(2n-3)-n = n-3$ internal edges, and they can each be flipped. So we can draw a graph whose vertices are the triangulations of a fixed n-gon, and where there is an edge between two triangulations if one is obtained from the other via an edge flip. It is clear that this graph is $(n-3)$-regular.

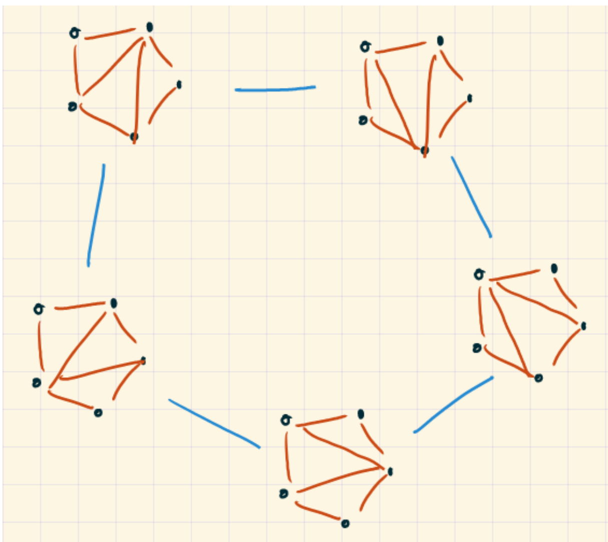

Exercise 1.7. Show that the *flip graph* of a convex n-gon (as described above) is connected. That is, one can reach a fixed triangulation from any triangulation via a sequence of flips.

The flip graph has a number of beautiful applications. For example, see this paper of Sleator, Tarjan, and Thurston [20]. More about triangulations in the direction of combinatorial geometry can be found in [8].

Moreover, the flip graph can be realised as the 1-skeleton of a polytope called the *associahedron*, or more specifically, the associahedron of type A. It is named as such because it can also be defined by vertices corresponding to the various ways to parenthesise a product of n terms, with edges coming from the associativity relation.

Exercise 1.8. Show that the triangulations of a convex n-gon are in bijection with the ways to parenthesise a product of $(n-1)$ terms, and that an edge flip corresponds to a change of parethesisation coming from an associativity relation.

There is a wealth of work on this topic, which has been generalised in several directions. For an introduction see, e.g. [9]. We explore one aspect of this in the next topic.

1.2 Connection to the type-A cluster algebra and the coordinate ring of a Grassmannian

1.2.1 The mutation rule

Fix a convex n-gon and a triangulation T. For each $ij \in T$, let x_{ij} be an indeterminate. To the vertex corresponding to the triangulation T in the flip graph, associate the set $C_T = \{x_{ij} \mid ij \in T\}$. This is called the *cluster* corresponding to T.

Suppose that $ik \in T$, and consider the edge flip of ik in T to form T'. This corresponds to replacing ik by jl in a quadrilateral $ijkl$ consisting of the two triangles bordering ik. Set $C_{T'}$ to be $(C_T \setminus \{x_{ik}\}) \cup \{x_{jl}\}$, where

$$x_{jl} = \frac{x_{ij}x_{kl} + x_{jk}x_{il}}{x_{ij}}.$$

This comes from the *Ptolemy relation* (of cyclic quadrilaterals):

$$x_{ik}x_{jl} = x_{ij}x_{kl} + x_{jk}x_{il}.$$

The operation of replacing C_T by $C_{T'}$ is called *mutation*. Let $\mathscr{F}$ be the field of rational functions in $\{x_{ij} \mid ij \in T\}$. Let $\mathscr{A}$ be the subring of $\mathscr{F}$ generated by the union of all possible successive mutations of C_T. This is the *cluster algebra* associated to the chosen convex n-gon (and the starting triangulation T). It is a special case of a much more general construction: see, e.g. [14, 22].

The construction does not make it clear whether the cluster corresponding to a particular triangulation T' depends on the sequence of flips. In fact, it does not.

Theorem 1.9. *Consider the flip graph of a convex n-gon, and fix a triangulation T and an initial cluster C_T. Then if T' is any other triangulation, the cluster associated to T' via a sequence of mutations from T is independent of the path chosen.*

The theorem above shows that this cluster algebra has finitely many clusters: i.e., it is *of finite type*. It is a special case of the classification of cluster algebras of finite type. These are classified by (yet again) the Dynkin diagrams of type A, D, and E — the original reference is [10].

Exercise 1.10. Verify for triangulations of a convex pentagon that you get finitely many clusters by repeated mutation.

1.2.2 The coordinate ring of $\mathrm{Gr}(2,n)$

The cluster algebra of type A has a "concrete" realisation via the coordinate ring of a Grassmannian. Fix a field $\mathbf{k}$. Recall that the *Grassmannian* $\mathrm{Gr}_{\mathbf{k}}(m,n)$ is a space that parameterises m-dimensional vector subspaces of a fixed n-dimensional vector space.

A point in $\mathrm{Gr}(m,n)$ is specified by a set of m linearly independent vectors $v_1,\dots,v_m \in \mathbf{k}^n$. Thus this point is specified by an $n \times m$ matrix:

$$H = \begin{pmatrix} \uparrow & \uparrow & \cdots & \uparrow \\ v_1 & v_2 & \cdots & v_m \\ \downarrow & \downarrow & \cdots & \downarrow \end{pmatrix}.$$

Conversely, any $n \times m$ matrix H of rank m specifies a point $[H]$ in $\mathrm{Gr}(m,n)$. Note that several such matrices may specify the same point, because the chosen m-dimensional subspace has many bases.

Exercise 1.11. Let H and H' be two $n \times m$ matrices of rank m. Show that $[H] = [H']$ if and only if there is an invertible $m \times m$ matrix A such that $HA = H'$.

Exercise 1.12. Check that the action of $\mathrm{GL}_m(\mathbf{k})$ on the space of $n \times m$ matrices of rank m by right multiplication is free.

We see that $\mathrm{Gr}(m,n)$ is the quotient of the space of $n \times m$ matrices of rank m by the action of $\mathrm{GL}_m(\mathbf{k})$.

Recall that an $n \times m$ matrix H has rank m if and only if some $m \times m$ minor has non-zero determinant. Further, if $A \in \mathrm{GL}_m(\mathbf{k})$, then the determinants of the $m \times m$ minors of HA are just the corresponding determinants from H, scaled by $\det A$.

Exercise 1.13. Check that right multiplication by an element of $\mathrm{GL}_m(\mathbf{k})$ scales the determinants of the $m \times m$ minors of an $n \times m$ matrix by $\det A$.

This motivates the *Plücker map* $\mathrm{Gr}(m,n) \to \mathbf{P}^{\binom{n}{m}-1}$, described as follows. The homogeneous coordinates of the projective space on the right hand side are indexed by a choice of m indices out of n, which we can think of as $(1 \le i_1 < i_2 < \cdots < i_m \le n)$. Given H such that $[H] \in \mathrm{Gr}(m,n)$, let $H_{i_1<\cdots<i_m}$ be the $m \times m$ minor of H corresponding to rows numbered $i_1,\dots,i_m$. Then the Plücker map is defined as

$$[H] \mapsto [\det(H_{i_1<\cdots<i_m})]_{1\le i_1<\cdots<i_m\le n}.$$

Exercise 1.14. Check that the Plücker map is well-defined and injective.

In fact, it is an embedding of algebraic varieties.

We now focus on $\mathrm{Gr}(2,n)$. Then a point $[H] \in \mathrm{Gr}(2,n)$ is specified by an $n \times 2$ matrix of rank 2, and the Plücker embedding looks like

$$[H] \mapsto \left[\det(H_{i<j})_{1\le i<j\le n}\right].$$

For $i < j$, set p_{ij} to be the (ij)th homogeneous coordinate of $\mathbf{P}^{\binom{n}{2}-1}$. For convenience, if $j < i$, set $p_{ij} = p_{ji}$.

Exercise 1.15. Check that every point in the image of the Plücker embedding satisfies the *Ptolemy relations*: if $1 \le a < b < c < d \le n$, then

$$p_{ac}p_{bd} = p_{ab}p_{cd} + p_{ad}p_{bc}.$$

Exercise 1.16. (*) Check that a point of $\mathbf{P}^{\binom{n}{2}-1}$ lies in the image of the Plücker embedding if and only if it satisfies all possible Ptolemy relations.

In fact, it follows from the exercises above that the homogeneous coordinate ring of $\mathrm{Gr}(2,n)$ (thought of as a projective algebraic variety) is the quotient of $\mathbf{k}[(p_{ij})_{1\leq i<j\leq n}]$ by the Ptolemy relations.

Note that if we think of the numbers from 1 to n as arranged in a convex n-gon, then the Ptolemy relation relates the Plücker coordinates of the two diagonals of a sub-quadrilateral $abcd$ to the Plücker coordinates of the sides.

Therefore the coordinate ring of $\mathrm{Gr}(2,n)$ has the structure of the cluster algebra described in the previous section, where the clusters are in bijection with the triangulations, and for a given triangulation T, consist of $\{p_{ij} \mid ij \in T\}$.

2 Non-convex arrangements: triangulations and pseudo-triangulations

2.1 Triangulations of not-necessarily-convex configurations

Consider a possibly non-convex (but generic) arrangement of n distinct points in the plane. A *triangulation* is once again a maximal collection of non-crossing edges. As before, a triangulation divides the convex hull of the configuration into triangular regions.

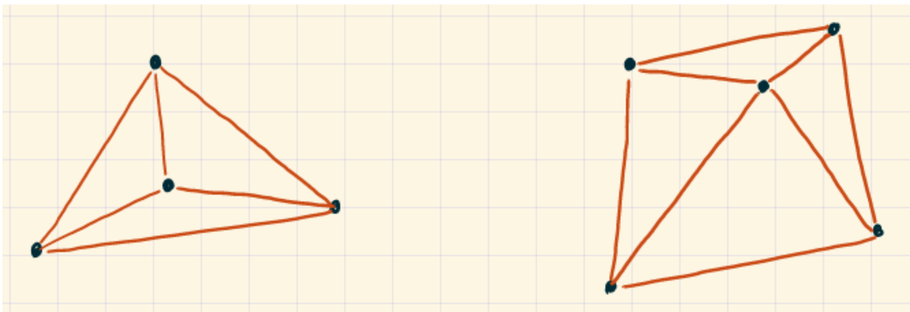

Once again, we can count the number of edges in a triangulation.

Proposition 2.1. *Consider a fixed configuration of n distinct points in the plane, such that m of the points lie in the interior of the convex hull of the configuration. Then the number of edges in any triangulation of the configuration is $2n-3+m$, and the number of triangular regions is $n-2+m$.*

Proof. Suppose we have e edges and f triangular faces. We apply Euler's formula once again:

$$n-e+f=1.$$

Once again, each triangular face bounds three edges. Every edge except for the external edges bounds two faces. There are exactly $n-m$ external edges, and so we have

$$3f=2e-(n-m).$$

Solving this system completes the proof. □

However, we can immediately see that not every internal edge in a triangulation has an edge flip. This is because the quadrilateral formed by the triangles on either side need not be convex. We can, however, form the flip graph as before.

Exercise 2.2. Find an example that shows that the flip graph of triangulations is not necessarily regular.

Exercise 2.3. (*) Prove that the flip graph of triangulations of any generic arrangement of n points in the plane is connected.

For more details about this exercise see, e.g. [13].

We could attempt to impose regularity on the flip graph by including curved arcs. To do this, we could define a triangulation to using (isotopy classes of) curved arcs that begin and end at two of the points and whose interior does not intersect any point in the configuration. Then a triangulation would be a maximal collection of such arcs that divides the convex hull into regions whose interiors are topologically interiors of triangles.

Exercise 2.4. Make the above construction precise, and define flips. Is the flip graph regular? Are there finitely many triangulations?

2.2 Pseudotriangulations

We choose to extend the story of triangulations to non-convex arrangements using a slightly different gadget.

Definition 2.5. A *pseudotriangle* is a non-crossing m-gon in the plane with the following properties:

1. its convex hull is a triangle, and
2. each vertex other than the three on the convex hull has an interior reflex angle.

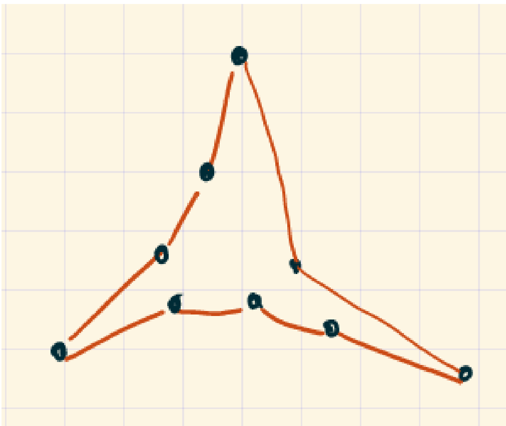

We will instead consider subdivisions of non-convex arrangements into pseudotriangles.

Definition 2.6. A pseudotriangulation of a non-convex generic arrangement of n points is a collection of edges that divides the convex hull into pseudotriangles.

Note that a triangle is always a pseudotriangle, and so a triangulation is always a pseudotriangulation. So what did we achieve? Let us see. More details can be found in [19].

2.2.1 Edge and face counts

Consider a pseudotriangulation on n points that has e edges and f interior faces. As before, we have

$$n - e + f = 1.$$

This time, a face no longer necessarily has three edges. However, it does have three (convex) angles. The total number of convex angles is then $3f$. Let r be the number of reflex angles, including the ones on the outside of the convex hull. The number of angles (both convex and reflex) at a vertex is exactly the number of edges emanating from that vertex. The total number of all the angles is therefore twice the number of edges (each edge is counted once at each endpoint). We have

$$3f + r = 2e.$$

Since we also have $3f = 3 + 3e - 3n$, we see that

$$3 + 3e - 3n + r = 2e,$$

which gives

$$e = 3n - r - 3$$

$f = 3n - r - 2$.

Note that having $r = n$ gives the same edge count as in the convex case, namely $2n - 3$! We now impose this condition. Every external vertex always has a reflex angle (the one on the outside). So now stipulate that every internal vertex should also have a reflex angle.

Definition 2.7. We say that a vertex is *pointed* in a given pseudotriangulation if it has a reflex angle. A *pointed pseudotriangulation* (or ppt) is a pseudotriangulation in which every vertex is pointed.

The calculations above show that every ppt has exactly $2n - 3$ edges and $n - 2$ pseudotriangular faces. As a corollary, we see that ppts are exactly the pseudotriangulations with the minimal number of edges.

Exercise 2.8. (*) Consider a configuration of n points in the plane in general position. Show that a maximal collection of edges that are non-crossing and such that every vertex is pointed is necessarily a ppt.

2.2.2 Flips

Note that a pseudotriangle may more generally be defined as a non-crossing polygon with exactly three internal convex angles. This accounts for degenerate cases. Similarly, say that a pseudo-k-gon is a non-crossing polygon with exactly k internal convex angles.

Definition 2.9. Let P be a pseudo-k-gon. A diagonal d that lies in the interior of P is called a *bitangent* if the graph $P \cup \{d\}$ is pointed.

Exercise 2.10. Prove that pseudotriangles have no bitangents.

Exercise 2.11. (*) Prove that a pseudoquadrilateral has exactly two bitangents.

Exercise 2.12. Show that if e is any internal edge of a ppt, then removing this edge creates a pseudoquadrilateral. Check that e is a bitangent of this pseudoquadrilateral, and that inserting the other bitangent e' also produces a ppt. Conclude that every internal edge of a ppt has a unique flip.

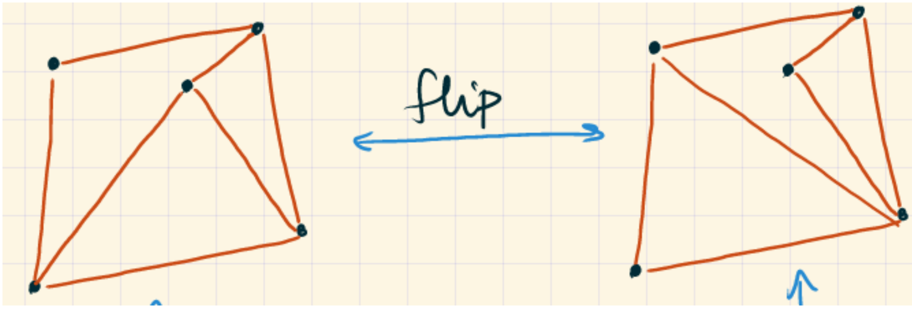

2.2.3 PPTs on non-generic configurations

Consider a possibly non-generic configuration. The most extreme case of this is when all n points lie on a single line.

We can extend the definition of pseudotriangulations to this case. We now consider isotopy classes of arcs that start and end at two points in the configuration, and which does not contain any point of the configuration in its interior. We say that an arc is "allowed" if it has representatives in its isotopy class that are arbitrarily close to the straight line segment between its endpoints. We will focus on allowed isotopy classes, that is, isotopy classes of allowed arcs.

The unsigned angle between two allowed isotopy classes $[a_1]$ and $[a_2]$ (at a common endpoint) is the limit of the unsigned angles between tangents to differentiable representatives a_1 and a_2 at that endpoint, as a_1 and a_2 approach straight-line segments. In particular, we can now have angles with values 0 and 2π.

We say that a collection of allowed isotopy classes is non-crossing if each class has a representative such that no two of these representatives intersect.

Exercise 2.13. Convince yourself that you can make the above construction precise.

Exercise 2.14. Classify all pseudotriangles such that all vertices lie on a straight line.

Now a pointed pseudotriangulation is exactly the same as before.

Definition 2.15. A pointed pseudotriangulation on a possibly degenerate configuration of n points is a maximal non-crossing collection of allowed isotopy classes.

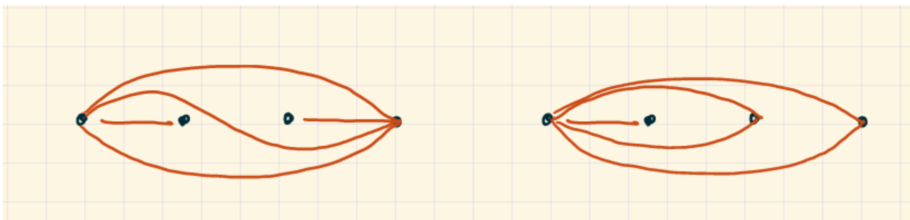

2.3 Assorted exercises

Exercise 2.16. (Open question.) Find a closed formula for the number of ppts for a configuration of n arranged in a line.

Exercise 2.17. Check that every internal edge of a ppt in a possibly degenerate configuration is flippable.

Exercise 2.18. (Open-ended.) Find families of configurations of points with a large number of ppts.

Exercise 2.19. (**) Show that among all possible configurations of n distinct points, a convex configuration has the minimum number of ppts.

The proof of this exercise is the subject of the paper [1].

Exercise 2.20. (Open question?) Show that among all possible configurations of n distinct points, a totally degenerate configuration has the maximum number of ppts.

3 Rigidity and infinitesimal rigidity

A *framework* is specified by a set of points $\{p_1, \ldots, p_n\} \in \mathbf{R}^2$, together with a collection of undirected edges of the form $ij = \{i, j\} \in E$. The edge ij is thought to be a straight and rigid rod connecting p_i and p_j.

Let ℓ_{ij} be the (Euclidean) length of the edge $p_i p_j$. The *configuration space* of the framework is the real solution set of the system of equations

$$(x_i - x_j)^2 + (y_i - y_j)^2 = \ell_{ij}^2, \quad ij \in E.$$

A *(differentiable) rigid motion* of the framework is a (differentiable) function

$$P\colon (-\epsilon, \epsilon) \to (\mathbf{R}^2)^n$$

for some $\epsilon > 0$, satisfying

$$(p_i(t) - p_j(t))^2 = \ell_{ij}^2, \quad ij \in E.$$

Note that rotations and translations (i.e., isometries of the plane) are always rigid motions. A framework is *rigid* if it is has no rigid motions other than isometries of the plane.

An *infinitesimal rigid motion* of the configuration is a set of "velocity vectors" $\{v_1, \ldots, v_n\}$ satisfying

$$\langle p_i - p_j, v_i - v_j \rangle = 0, \quad ij \in E.$$

Note that any differentiable rigid motion gives rise to an infinitesimal rigid motion, by taking the derivative at the point 0.

Let G be a framework with vertex set P of size n, and edge set E of size e. We have an "infinitesimal rigidity map"

$$M_G\colon (\mathbf{R}^2)^n \to \mathbf{R}^e$$

defined as follows. A point in the source space is a tuple of the form $(v_1, \ldots, v_n)$, where each v_i is in $\mathbf{R}^2$. It is sent to the element $(\langle p_i - p_j, v_i - v_j \rangle)_{ij \in E}$.

Note that the kernel of this matrix represents the infinitesimal rigid motions, and is at least three-dimensional (as it includes the isometries).

Definition 3.1. Consider a framework with vertex set P and edge set E. A *self-stress* of a framework is a function $\omega\colon E \to \mathbf{R}$ satisfying the following. For each $i \in P$, we have

$$\sum_{ij \in E} \omega_{ij}(p_i - p_j) = 0.$$

Clearly, the set of all self-stresses forms a vector space.

Exercise 3.2. Prove that a function $\omega\colon E \to \mathbf{R}$ (thought of as an element of $\mathbf{R}^e$) is a self-stress of G if and only if it is orthogonal to the image of M_G.

Now suppose that G is a framework with infinitesimal rigidity matrix M_G. Recall that $\ker(M_G)$ consists of the infinitesimal rigid motions. Let $\dim(\ker(M_G)) = d+3$, so that G has d non-trivial infinitesimal motions. Let s be the dimension of the space of self-stresses of G.

Observe that

$$e = \mathrm{rk}(M_\mathrm{G}) + \dim(\mathrm{im}(M_\mathrm{G})^\perp) = (2n-d-3)+s = (2n-3)+(s-d). \tag{1}$$

In particular, ppts can be thought of as frameworks, and the number of edges in a ppt is $2n-3$. We conclude that the number of non-trivial infinitesimal motions of a ppt is precisely the dimension of the space of self-stresses.

In fact, we can independently compute the self-stresses of frameworks using the Maxwell–Cremona correspondence, which dates back to the 1800s [7, 15, 16] .

Consider a planar graph G with straight edges, thought of as a framework. Let P be the vertex set of G. A *3d lifting* of G is a (height) function $h: P \to \mathbf{R}$ with the following property. Consider a face of G with vertices $\{p_i\}$. Then the points $\{(p_i, h(p_i))\}$ are all coplanar.

A *Maxwell lifting* of G is simply a 3d lifting in which all vertices of the outer face have height zero.

Exercise 3.3. Check that the set of all Maxwell liftings forms a vector space.

Theorem 3.4 (Maxwell–Cremona correspondence)**.** *The following vector spaces are isomorphic.*

1. *The space of self-stresses of G.*

2. *The space of Maxwell liftings of G.*

The proof of this theorem is via an intermediate gadget called *reciprocal diagrams*. Let G be a planar framework with vertex set P and edge set E. A *reciprocal diagram* of G is a (not necessarily planar) framework G' with the following properties:

1. The vertex set $P(G')$ of G' is in indexed by the set of faces of G, including the outer face. That is, $P(G') = \{p_F \mid F \text{ a face of } G\}$. (Note that the points p_F need not be distinct.)

2. The vertex corresponding to the outer face lies at the origin.

3. Let $e = pq$ be an edge of G that separates faces A and B of G. Then $\vec{p_A p_B}$ is parallel to $\vec{pq}$. That is, there are real numbers $\{\alpha_e : e \in E\}$ that satisfy the following. Suppose that A lies to the left of the oriented edge $\vec{pq}$, and B lies to the right.

$$(p_A - p_B) = \alpha_e \cdot (p-q).$$

Exercise 3.5. Check that the reciprocal diagrams of G form a vector space.

Exercise 3.6. Consider a reciprocal diagram of G with edge set E. Show that the real numbers $\{\alpha_e : e \in E\}$ satisfying the second condition in the definition of the reciprocal diagram give a self-stress of G. Conversely, show that a self-stress of G gives rise to a reciprocal diagram.

We can now relate the reciprocal diagrams to the space of Maxwell liftings.

Exercise 3.7. Consider a Maxwell lifting $h: P \to \mathbf{R}$ of G. Use the fact that each face has a well-defined normal vector (after fixing an orientation) to show that h gives rise to a reciprocal diagram of G. Conversely, show that each reciprocal diagram produces a Maxwell lifting.

In order to compute the self-stresses of a ppt, it is therefore sufficient to compute its Maxwell liftings.

Exercise 3.8. Let h be a lifting of a planar framework G. Let F be a face of G and p be a vertex of G that is a reflex vertex of F. Prove that if h has a global maximum (or global minimum) at p, then h is constant on F.

Now consider a ppt, and suppose it has a nontrivial Maxwell lifting h. The convex hull vertices have height zero, so some interior vertex must have nonzero height. Suppose (WLOG) that p is an interior vertex such that $h(p)$ is non-negative, and a global maximum of h. The vertex p is necessarily a reflex vertex of some pseudotriangular face, and therefore every vertex of this face has the same (positive) height. We can then argue the same for every other vertex of this face, and continuing by induction, we conclude that every vertex p has the same height. In particular, the global maximum is zero. (The same can be argued about global minima.) In other words, ppts have no non-trivial Maxwell liftings! As a consequence, the dimension of the space of self-stresses of a ppt is zero. Therefore ppts have no non-trivial infinitesimal motions, i.e., they are infinitesimally rigid.

Exercise 3.9. Construct an example of a framework that is rigid but not infinitesimally rigid.

Equation (1) implies that if a framework is infinitesimally rigid, it has $2n-3+s$ edges, where s is the dimension of the self-stresses. We conclude that the minimal number of edges in an infinitesimally rigid framework is exactly $2n-3$.

Also note that removing one or more external edges of a ppt framework does not increase the dimension of the space of self-stresses: indeed, any non-trivial self stress of the configuration without the external edges can be extended to the whole framework by zero on the external edges that were removed. We conclude the following.

Proposition 3.10. *A ppt framework with a single external edge removed has a unique non-trivial infinitesimal motion.*

Making heavy use of this proposition, Rote–Santos–Streinu construct in [18] a polytope whose vertices are in one-to-one correspondence with ppts of a fixed point configuration. This is a direct generalisation of the associahedron that we saw earlier.

3.1 Expansive motions

We say that an infinitesimal rigid motion is *expansive* if we have

$$\langle p_i - p_j, v_i - v_j \rangle \geq 0$$

for every pair $\{i, j\}$. We say that a rigid motion is expansive at time 0 if its derivative evaluated at 0 is an infinitesimal expansive rigid motion.

In fact, the previous proposition has a better version, due to Streinu [21].

Theorem 3.11 (Streinu)**.** *The unique nontrivial infinitesimal motion of a ppt framework with a single external edge removed is expansive in one direction and contracting in the other. Furthermore, there is a rigid motion whose derivative is precisely the unique infinitesimally expansive rigid motion.*

3.2 Carpenter's rule problem (and solution)

Ppts were used by Streinu [21] to give an elegant solution to the *Carpenter's rule problem*, which asks the following. (The problem had been solved by a different method very shortly before Streinu's solution, in [6].)

Question 3.12 (Carpenter's rule problem)**.** Consider a non-crossing polygonal chain in the plane. Is it always possible to convexify this polygon via rigid motions in the plane?

Here is a sketch of Streinu's solution algorithm. Consider a non-crossing polygonal chain as a framework.

- If it is already convex, we are done.
- Otherwise, virtually fill in enough edges (external and internal) to form a ppt.
- Note that at least one external edge is virtual, so choose one and delete it.
- Now expand the mechanism along the unique expansive direction, until a collinearity occurs.
  - If the collinearity is between two adjacent edges of the polygon, "freeze" the joint, disregarding the middle vertex, and continue.
  - If the collinearity is about to occur either between two adjacent virtual edges or between a virtual edge and an adjacent framework edge, perform an appropriate flip so that we avoid collinearity.
  - If one of the edges aligns with the missing external edge, we still have a ppt without an external edge, so simply continue an expansive motion.

Streinu proved that this algorithm is well-defined and that it terminates, resulting in a convex position.

4 Configuration space and the braid group action on arcs

Let $\mathrm{Conf}_n(\mathbb{R}^2)$ be the configuration space of n ordered distinct points in $\mathbf{R}^2$, namely:

$$\mathrm{Conf}_n(\mathbb{R}^2) = \{(p_1, \dots, p_n) \mid p_i \in \mathbb{R}^2,\, p_i \neq p_j \text{ for } i \neq j\}.$$

Let $\mathrm{UConf}_n(\mathbb{R}^2)$ be the configuration space of n unordered distinct points in $\mathbb{R}^2$, namely:

$$\mathrm{UConf}_n(\mathbb{R}^2) = \mathrm{Conf}_n(\mathbb{R}^2)/S_n.$$

Recall that the n-strand Artin braid group B_n is defined as follows.

Definition 4.1. The n-strand (Artin) braid group is a group defined by generators

$$\sigma_1, \dots, \sigma_{n-1},$$

and relations

$$\sigma_i\sigma_j\sigma_i = \sigma_j\sigma_i\sigma_j \text{ if } |i-j| = 1;$$
$$\sigma_i\sigma_j = \sigma_j\sigma_i \text{ if } |i-j| \neq 1.$$

It is closely related to the symmetric group S_n.

Exercise 4.2. Find a surjective homomorphism $B_n \to S_n$ and describe its kernel. This kernel is called the *pure braid group* PB_n.

Elements of the braid group may be visually represented as braid diagrams (up to isotopy), where multiplication is represented by stacking two braid diagrams together and then "pulling tight".

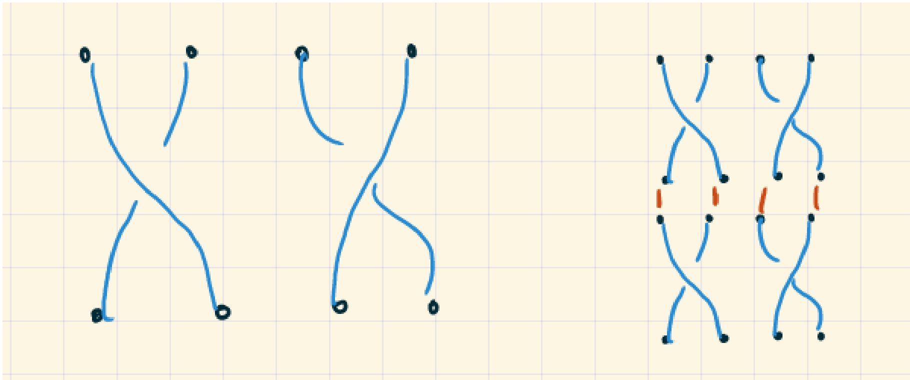

Note that a loop in the space $\mathrm{UConf}_n(\mathbb{R}^2)$ is a path from a chosen configuration to itself, up to reordering of the points. It is easy to see that such a loop corresponds to a braid diagram.

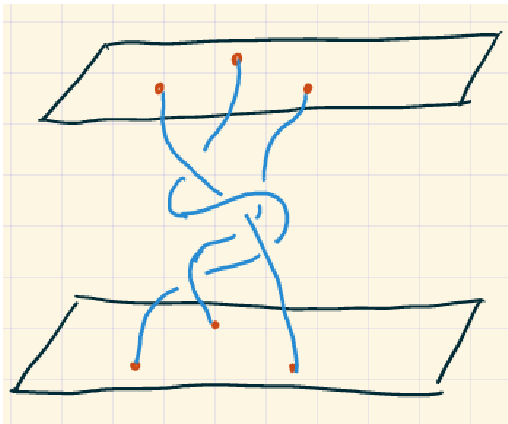

In fact, the following theorem is well-known.

Theorem 4.3. *The fundamental group of* $\mathrm{UConf}_n(\mathbb{R}^2)$ *is exactly* B_n*. The fundamental group of* $\mathrm{Conf}_n(\mathbb{R}^2)$ *is exactly* PB_n*.*

In particular, consider two distinct points in a chosen configuration. Consider a loop that keeps all other points fixed, but swaps the positions of these two points "clockwise". This is called a *Dehn half-twist* around the arc joining these two points, and it is an element of the fundamental group (i.e. the braid group).

Exercise 4.4. Show that the fundamental group of $\mathrm{UConf}_n(\mathbb{R}^2)$ is generated by Dehn half-twists. What are some minimal generating sets of Dehn half-twists?

Exercise 4.5. (Open-ended) Find a finite generating set for the fundamental group of $\mathrm{Conf}_n(\mathbb{R}^2)$.

Now fix an element of $\mathrm{UConf}_n(\mathbb{R}^2)$, that is, a configuration of (unordered) points in $\mathbb{R}^2$. Consider the collection of arcs of this configuration up to isotopy. Enumerate the positions of the points as $p_1, \ldots, p_n$, say with increasing x-coordinate for convenience. Then B_n acts on the collection of arcs via Dehn half-twists in the segments $p_1p_2, \ldots, p_{n-1}p_n$ acting as the standard generators, respectively.

We say that a collection of arcs $a_1, \ldots, a_{n-1}$ forms an A_{n-1}-chain if the arcs are non-crossing, and if a_i begins at position p_i and ends at position p_{i+1}. Clearly, B_n acts on A_{n-1}-chains in the same way as above.

Exercise 4.6.

1. Prove that the support of any arc is a subset of a ppt minus an external edge.

2. Prove that the same is true for any A_n-chain of arcs.

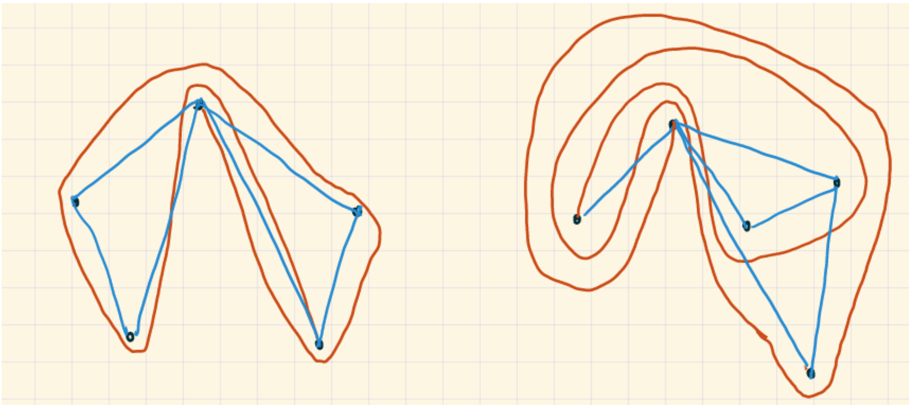

Exercise 4.7. Fix a convex configuration, and also fix a triangulation minus an external edge. Find an algorithm to draw an arc whose support is precisely this fixed framework.

Exercise 4.8. (I don't know a nice proof of this statement!) Fix a ppt on a configuration of n points (not necessarily generic), and delete an external edge. Prove (or better, find an algorithm) that there exists an arc whose support is exactly the fixed framework.

5 Bridgeland stability conditions on a triangulated category

5.1 A few (only a few!) words about triangulated categories

Neeman's book [17] is an excellent reference for triangulated categories. I will not give the definition here in the interest of time, but will rather stick to a couple of examples.

One important example is the following. Let A be an algebra over a fixed field, and consider the (additive) category of projective finite-dimensional A-modules. Call this $\mathrm{Proj}(A)$. The homotopy category of $\mathrm{Proj}(A)$, denoted $K(\mathrm{Proj}(A))$ is the category whose objects are cochain complexes of projective finite-dimensional A-modules, and whose morphisms are chain maps up to chain homotopy. This is a triangulated category.

Exercise 5.1. (If you know all the definitions! In any case, *) Check this.

Another important example to keep in mind is that of the (bounded) derived category $\mathscr{D}(\mathscr{A})$ of a fixed abelian category $\mathscr{A}$.

Consider cochain complexes of objects in $\mathscr{A}$, and cochain maps. Recall that a cochain map is a *quasi-isomorphism* if it induces an isomorphism on cohomology in every degree. A lightning definition of $\mathscr{D}(\mathscr{A})$ (resp. $\mathscr{D}^b(\mathscr{A})$) is as follows: it is the category whose objects are (bounded) cochain complexes of objects in $\mathscr{A}$, and whose morphisms are cochain maps where quasi-isomorphisms have been formally inverted. This is a triangulated category.

Exercise 5.2. (If you know all the definitions! In any case, *) Check this.

Exercise 5.3. Play around with the construction of the bounded derived category for some manageable (but not necessarily trivial) abelian categories. For example, you could try:

1. taking $\mathscr{A}$ to be the category of finite-dimensional vector spaces;
2. taking $\mathscr{A}$ to be the category of finite-dimensional modules over the algebra $k[\epsilon]/\epsilon^2$.
3. taking $\mathscr{A}$ to be the category of finite-dimensional modules over the algebra $k[\epsilon]/\epsilon^n$.

As a warm-up, try to write down some non-trivial cochain complexes and some cochain maps on objects in the categories mentioned above. Compute cohomology: can you find two cochain complexes that are not chain homotopic, that have the same cohomology?

5.2 A crash course on bounded t-structures

Triangulated categories are not usually abelian, but sometimes they can be "filtered by abelian categories".

5.2.1 Bounded t-structures

A particularly nice way this might happen is if the category supports a *bounded t-structure*.

Definition 5.4. Let $\mathscr{C}$ be a triangulated category. A *bounded t-structure* on $\mathscr{C}$ is determined by a full abelian subcategory $\mathscr{A} \subset \mathscr{C}$, satisfying the following property. Every object $X \in \mathscr{C}$ has a unique finite filtration (called the *cohomology filtration with respect to $\mathscr{A}$*):

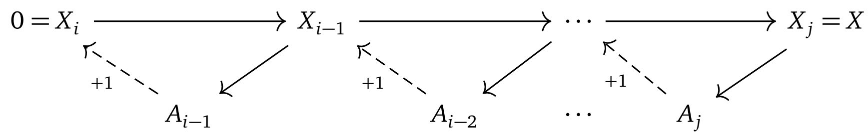

such that $A_m \in \mathscr{A}[m]$.

The subcategory $\mathscr{A} \subset \mathscr{C}$ is called the *heart* of this bounded t-structure. For the original reference and a more general definition of t-structures, see [4].

In the presence of a bounded t-structure, working with a triangulated category often (basically) reduces to working with the heart, and then extending everything via the unique filtration. We adopt this point of view for convenience.

Note, of course, that the same triangulated category $\mathscr{C}$ can (and often does) have many different bounded t-structures. The interplay between the different t-structures can have rich combinatorial consequences. See, e.g. [12].

5.2.2 Extended exercise: a non-standard t-structure

Let $\mathscr{A}$ be the abelian category whose objects are diagrams

$$\begin{CD} V_t \\ @VV{M}V \\ V_h \end{CD}$$

where V_t and V_h are complex vector spaces and M is a linear map from V_t to V_h. A morphism is a commutative diagram as follows.

$$\begin{CD} V_t @>{M}>> V_h \\ @VV{f_t}V @VV{f_h}V \\ W_t @>{N}>> W_h \end{CD}$$

(Some of you may recognise this as the category of representations of a certain quiver.)

Exercise 5.5. Convince yourself that this category is, in fact, abelian.

Exercise 5.6. Shaw that up to isomorphism, there are exactly two simple (or irreducible) objects in this category, namely $S_1 = \mathbb{C} \xrightarrow{0} 0$ and $S_2 = 0 \xrightarrow{0} \mathbb{C}$.

Next consider the bounded derived category $\mathscr{C} = \mathscr{D}(\mathscr{A})$. The objects of $\mathscr{C}$ can be thought of as diagrams

$$V^\bullet = \begin{array}{c} V_t^\bullet \\ \big\downarrow{\scriptstyle M^\bullet} \\ V_h^\bullet \end{array},$$

where $V_h^\bullet$ and $V_t^\bullet$ are cochain complexes of vector spaces, and $M^\bullet$ is a cochain map.

Set $\mathscr{B}$ to be the collection of objects whose cohomologies are supported in the shaded region shown below:

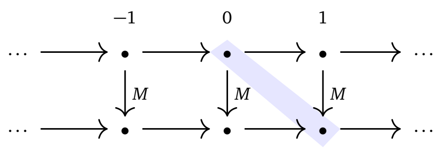

Explicitly, $V^\bullet \in \mathscr{B}$ if:

$$H^i(V_t^\bullet) = 0 \text{ for all } i \neq 0, \quad H^i(V_h^\bullet) = 0 \text{ for all } i \neq 1.$$

Exercise 5.7. Describe the category $\mathscr{B}$ as explicitly as possible. For example, what objects generate this category?

Exercise 5.8. (*) Show that $\mathscr{B}$ is the heart of a bounded t-structure on $\mathscr{C}$.

Exercise 5.9. Show that the categories $\mathscr{A}$ and $\mathscr{B}$ are not equivalent.

Exercise 5.10. Is it true that $\mathscr{C} = D^b(\mathscr{B})$?

5.3 A crash course on (Bridgeland) stability conditions

Consider a triangulated category $\mathscr{C}$ that has a bounded t-structure with heart $\mathscr{A}$. Recall that in suitably nice cases, objects of $\mathscr{A}$ (or indeed, any suitably nice abelian category) have Jordan–Hölder filtrations, whose subquotients are simple. It is therefore natural to try and understand objects of $\mathscr{A}$ via their (simple) composition factors. However, this point of view is necessarily limited for a few reasons, including:

1. the JH series itself is not unique, so it is hard to get a handle on a canonical JH series;
2. the JH series is fundamentally an "abelian" construction, because it relies on simple objects, and the notion of which objects are simple changes dramatically if you alter the t-structure; and
3. once you fix a heart, you don't have a choice of which objects are simple.

A stability condition is a piece of extra data on $\mathscr{C}$ that simultaneously remedies all of the problems above. Once again, we focus on the case when $\mathscr{C}$ has a chosen bounded t-structure with heart $\mathscr{A}$. (However, there is an independent, equivalent definition that does not involve any choices of t-structure. See [5, Definition 5.1 and Proposition 5.3])

Definition 5.11. Let $\mathscr{C}$ be a triangulated category with a bounded t-structure whose heart is $\mathscr{A}$. A *Bridgeland stability condition* on $\mathscr{C}$ (with respect to the heart $\mathscr{A}$) consists of a group homomorphism

$$Z\colon K(\mathscr{A}) \to \mathbb{C},$$

with the following properties.

1. For every $A \in \mathscr{A}$, we have $Z(A) \in \mathbb{H}$.
2. An object $A \in \mathscr{A}$ is called *semistable* if whenever $0 \neq B \subset A$, we have $\arg(B) < \arg(A)$.
3. Every object $X \in \mathscr{A}$ has a *unique* finite filtration

$$0 = X_0 \subset X_1 \subset \cdots \subset X_n = X$$

   whose successive subquotients $A_i = X_i/X_{i-1}$ are semistable of decreasing argument: $\arg(A_1) > \arg(A_2) \cdots > \arg(A_n)$.

Some references for Bridgeland stability conditions: [3, 5].

Exercise 5.12. Write down one (or more!) explicit Bridgeland stability conditions on each of the hearts $\mathscr{A}$ and $\mathscr{B}$ in the extended exercise above.

6 The story of the zigzag category type A

Fix $n \in \mathbf{N}$. We work with a triangulated category $\mathscr{C} = \mathscr{C}_n$, which we call the zigzag category of type A_n. A precise definition can be found in e.g. [2, Section 2.3]. We will not dwell on the precise definitions here, but rather present the properties and flavour of this category.

- $\mathscr{C}$ is generated by objects P_i, where i ranges from 1 to n.
- The generating objects satisfy

$$\mathrm{Hom}^m(P_i, P_i) = \begin{cases} \mathbf{k} & m = 0, 2 \\ 0 & \text{otherwise.} \end{cases}$$

$$\mathrm{Hom}^m(P_i, P_j) = \begin{cases} \mathbf{k} & m = 1 \text{ and } |i - j| = 1, \\ 0 & \text{otherwise.} \end{cases}$$

  Note that the unique map of degree zero (up to scaling) from P_i to P_i is just the identity map. The unique map of degree two (up to scaling) from P_i to P_i is called the *loop map*.

- $\mathscr{C}$ has a bounded t-structure that we call the *standard t-structure* on $\mathscr{C}$. The heart is generated by complexes in which all maps have degree ≤ 1; that is, which have no loop maps in them. The simple objects in the heart are precisely the P_i.
- The Grothendieck group of $\mathscr{C}$ is isomorphic to the root lattice of type A, and carries an action of the symmetric group S_{n+1}.
- This action lifts to an action of the braid group B_{n+1} on $\mathscr{C}_n$.

Consider at first a configuration of n points arranged in a straight line. Khovanov–Seidel [11] establish a B_n-equivariant bijection between spherical objects of $\mathscr{C}_n$ and isotopy classes of arcs in the configuration that begin and end at points in the configuration but do not intersect any point of the configuration in their interior.

Now vary the configuration. The following results are joint work with Anand Deopurkar and Anthony Licata. We observe the following.

1. A stability condition with respect to the standard heart is simply a configuration of points.
2. Any other stability condition is specified by a configuration of points together with an A_n chain of arcs joining successive points.
3. The semistable objects in a standard stability condition prescribed by the given configuration are precisely the (isotopy classes of) allowed arcs in the configuration. The Harder–Narasimhan filtration pieces of a given spherical object can be read off from its support.
4. This procedure realises spherical objects as rational points of (the geometric realisation of) a simplicial complex whose maximal simplices are ppts with an external edge removed.
5. This simplicial complex is topologically a sphere, and the spherical objects of the category give a dense subset of this sphere.
6. An enhanced version of Streinu's algorithm for the carpenter's rule problem gives an explicit contraction of the space of stability conditions of $\mathscr{C}_n$ to a point, giving a new proof of its contractibility.

References

[1] O. Aichholzer, F. Aurenhammer, H. Krasser, and B. Speckmann, *Convexity minimizes pseudo-triangulations*, Computational Geometry, 28 (2004), p. 3–10.

[2] A. Bapat, A. Deopurkar, and A. M. Licata, *Spherical objects and stability conditions on 2-Calabi–Yau quiver categories*, (2021). https://asilata.github.io/assets/papers/stability-algorithm.pdf.

[3] A. Bayer, *A tour to stability conditions on derived categories*. https://www.maths.ed.ac.uk/~abayer/dc-lecture-notes.pdf, 2011.

[4] A. A. Beĭlinson, J. Bernstein, and P. Deligne, *Faisceaux pervers*, in Analysis and topology on singular spaces, I (Luminy, 1981), vol. 100 of Astérisque, Soc. Math. France, Paris, 1982, pp. 5–171.

[5] T. Bridgeland, *Stability conditions on triangulated categories*, Ann. of Math. (2), 166 (2007), pp. 317–345.

[6] R. Connelly, E. Demaine, and G. Rote, *Straightening polygonal arcs and convexifying polygonal cycles*, Proceedings 41st Annual Symposium on Foundations of Computer Science, (2000).

[7] L. Cremona, *Graphical statics. Two treatises on the graphical calculus and reciprocal figures in graphical statics. Translated by Th. H. Beare.* Oxford. Clarendon Press. [Nature XLIV. 221-222.] (1890)., 1890.

[8] S. Felsner, *Geometric Graphs and Arrangements*, Vieweg+Teubner Verlag, 2004.

[9] S. Fomin and N. Reading, *Root systems and generalized associahedra*, IAS/Park City Mathematics Series, (2007), p. 63–131.

[10] S. Fomin and A. Zelevinsky, *Cluster algebras. II. Finite type classification*, Invent. Math., 154 (2003), pp. 63–121.

[11] M. KHOVANOV AND P. SEIDEL, *Quivers, Floer cohomology, and braid group actions*, J. Amer. Math. Soc., 15 (2002), pp. 203–271.

[12] A. KING AND Y. QIU, *Exchange graphs and Ext quivers*, Adv. Math., 285 (2015), pp. 1106–1154.

[13] C. L. LAWSON, *Transforming triangulations*, Discrete Mathematics, 3 (1972), p. 365–372.

[14] B. R. MARSH, *Lecture notes on cluster algebras*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2013.

[15] J. C. MAXWELL, *On reciprocal figures and diagrams of forces*, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 27 (1864), p. 250–261.

[16] ———, *I.—on reciprocal figures, frames, and diagrams of forces*, Transactions of the Royal Society of Edinburgh, 26 (1870), p. 1–40.

[17] A. NEEMAN, *Triangulated categories*, (2001).

[18] G. ROTE, F. SANTOS, AND I. STREINU, *Expansive motions and the polytope of pointed pseudo-triangulations*, in Discrete and computational geometry. The Goodman-Pollack Festschrift, Berlin: Springer, 2003, pp. 699–736.

[19] G. ROTE, F. SANTOS, AND I. STREINU, *Pseudo-triangulations—a survey*, in Surveys on discrete and computational geometry, vol. 453 of Contemp. Math., Amer. Math. Soc., Providence, RI, 2008, pp. 343–410.

[20] D. D. SLEATOR, R. E. TARJAN, AND W. P. THURSTON, *Rotation distance, triangulations, and hyperbolic geometry*, Journal of the American Mathematical Society, 1 (1988), p. 647–681.

[21] I. STREINU, *Pseudo-Triangulations, Rigidity, and Motion Planning*, Discrete Computational Geometry, 34 (2005), p. 587–635.

[22] L. K. WILLIAMS, *Cluster algebras: an introduction*, (2012).