Introduction to Quiver Representations

Will Dana

July 9, 2019

In this class, all vector spaces are finite-dimensional and over $\mathbb{C}$, unless stated otherwise.

1 A Fundamental Idea of This Class

Many statements/theorems/problems in linear algebra can be expressed in terms of *decomposition*.

Example. Here's a theorem which makes all of linear algebra work:

Theorem. *Every vector space has a basis.*

Here's an equivalent statement of that theorem:

Theorem (Reprise)**.** *An n-dimensional vector space V is isomorphic to a direct sum*

$$V \cong \mathbb{C} \oplus \mathbb{C} \oplus \cdots \oplus \mathbb{C}$$

of n copies of $\mathbb{C}$.

Why are these the same thing? The value of having a basis $v_1, \ldots, v_n$ of V is that we can write any element of V in a unique way as a linear combination

$$a_1 v_1 + \ldots + a_n v_n$$

with complex numbers $a_1, \ldots, a_n$, and our isomorphism will pair this up with the tuple

$$(a_1, \ldots, a_n) \in \mathbb{C}^{\oplus n}.$$

The point demonstrated by this example is that the existence of a basis is what allows us to write vectors in a nice way—and this "writing vectors in a nice way" property corresponds to breaking our space up as a direct sum of simple pieces (in this case, copies of $\mathbb{C}$). In this class, we will break many other things up into simple pieces, and use that to say things about writing matrices in a nice way.

2 Classifying Matrices by Rank

2.1 The Classification

Theorem 1. *Let M be any $m \times n$ matrix. Then we can choose an invertible $m \times m$ matrix A and an invertible $n \times n$ matrix B such that*

$$AMB = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \end{pmatrix}$$

where the number of 1's is the rank of M.

Let's phrase this in terms of bases of abstract vector spaces, as this will make it easier to prove the theorem:

Theorem 2 (Theorem 1 Reprise)**.** *Let $f : V \to W$ be a linear map of vector spaces, of rank r. Then we can choose bases $v_1, \dots, v_n$ of V and $w_1, \dots, w_m$ of W such that*

$$f(v_i) = \begin{cases} w_i & 1 \le i \le r \\ 0 & r+1 \le i \le n \end{cases}$$

That is, there are bases of V and W such that the matrix of f is of the form above. The matrices A and B in Theorem 1 are the matrices used to change into these bases.

We can fit this into the "decomposition" framework by using the description of bases from section 1:

Theorem 3 (Theorem 1 Reprise Reprise)**.** *Let $f : V \to W$ be a linear map of rank r. Then we can express V and W as direct sums of 1-dimensional spaces*

$$V \cong \mathbb{C} \oplus \cdots \oplus \mathbb{C}$$
$$W \cong \mathbb{C} \oplus \cdots \oplus \mathbb{C}$$

such that f maps the first r summands of V into their counterparts in W by the identity, and is 0 on the others.

This is pretty clunky, though. (It will get a lot nicer once we have the language of quiver representations at our disposal.) Let's prove the second version.

Proof. By the rank-nullity theorem, $\ker(f)$ has dimension $n - r$. We choose a basis for $\ker(f)$, and for reasons which will become apparent almost immediately, we label it $v_{r+1}, \dots, v_n$. Then we break out an important fact from linear algebra:

Fact. *Any linearly independent set of vectors in a vector space can be completed to a basis.*

So we can complete our basis of $\ker(f)$ to a basis $v_1, \ldots, v_n$ of V.

Now for $1 \leq i \leq r$, we define $w_i := f(v_i)$. These are linearly independent: if

$$a_1 w_1 + \ldots + a_r w_r = 0$$

then

$$f(a_1 v_1 + \ldots + a_r v_r) = 0$$

and so $a_1 v_1 + \ldots + a_r v_r \in \ker(f)$. But any element of $\ker(f)$ can also be expressed as a combination $b_{r+1} v_{r+1} + \ldots + b_n v_n$; unless all a_i are 0, this violates the linear independence of the v_i.

Then we again use the above Fact to complete the w_i to a basis $w_1, \ldots, w_m$. The theorem is satisfied by construction. $\square$

Before moving on, let's look at the Fact we used in a little more detail. It, too, can be rephrased as a type of decomposition:

Fact (Reprise)**.** *Let $V' \subset V$ be a subspace. Then there is another vector space V'' such that $V \cong V' \oplus V''$ (and such that the inclusion map $V' \subset V$ is preserved by this isomorphism).*

Exercise 1. *Convince yourself that these two statements of the Fact are equivalent.*

Although this may seem like a very natural property, it's quite special to vector spaces. We refer to the property specified by the Reprise as **semisimplicity**.

Exercise 2. *Show that abelian groups do not have the semisimplicity property. That is, given an abelian group G and a subgroup G', it need not be the case that $G \cong G' \oplus G''$ for some G'' in a way that preserves the inclusion $G' \subset G$.*

It will also be helpful to note that if we have a decomposition $V \cong V' \oplus V''$, we can identify V'' with V/V'.

2.2 Implications of the Classification

In essence, this theorem says that as long as we look at the vector spaces V and W from the correct perspectives, there are only so many ways a linear map can behave.

Geometrically, the matrix in the original statement of Theorem 1 corresponds to projecting onto an r-dimensional subspace (by setting all but the first r coordinates to 0) and then including this subspace as a subspace of W (by appending 0s to the end or chopping off extraneous 0s if necessary). So the theorem says that *every* linear map is like this—the only things that change are which subspace we're projecting to and which subspace we're including into, and we can pick the coordinates of our spaces to accommodate that.

Algebraically, the theorem says that, if there is no other structure on the vector spaces involved, then the rank of a linear map is the only thing that really matters. Everything else is dependent on some arbitrary choice of bases. (If other maps get involved, this changes, because we need to tailor our bases to the map.)

Finally, we can make a slightly more nuanced statement than "every linear map falls into one of $\min(m, n)$ buckets". It is actually true that *most* matrices have full rank, since this happens as long as some polynomials[1] don't simultaneously vanish. Furthermore, any matrix can be expressed as a limit of full-rank ones. In topological terms, the full-rank matrices are *dense*, and everything else is icing on the matrix cake.

3 Classifying Matrices by Jordan Form

It's not always the case that the domain and codomain of a linear map can be manipulated independently. We might want to think of a square matrix as a transformation from a space to itself. In this case, changing the basis with change-of-basis matrix A transforms a matrix M by

$$M \mapsto AMA^{-1}$$

which limits our options substantially. But there is still a nice canonical form theorem in the vein of the last section:

Theorem 4 (Jordan Canonical Form)**.** *Let M be any $n \times n$ matrix. Then we can choose an invertible $n \times n$ matrix A such that AMA^{-1} has the block form*

$$AMA^{-1} = \begin{pmatrix} \mathbf{J_{i_1}}(\boldsymbol{\lambda_1}) & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \mathbf{J_{i_2}}(\boldsymbol{\lambda_2}) & \dots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{J_{i_k}}(\boldsymbol{\lambda_k}) \end{pmatrix}$$

where $\mathbf{J_i}(\lambda)$ is the $i \times i$ matrix

$$\begin{pmatrix} \lambda & 1 & 0 & \dots & 0 & 0 \\ 0 & \lambda & 1 & \dots & 0 & 0 \\ 0 & 0 & \lambda & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & \lambda & 1 \\ 0 & 0 & 0 & \dots & 0 & \lambda \end{pmatrix}$$

Once again, we can rephrase this in terms of a decomposition:

[1]Specifically, the determinants of the maximal submatrices.

Theorem 5 (Jordan Canonical Form Reprise)**.** *For any vector space V and linear map $f : V \to V$, we can write V as a direct sum*

$$V \cong V_1 \oplus \cdots \oplus V_k$$

such that f maps V_i into V_i, and such that the matrix of f when restricted to V_i, in an appropriate basis, has the form

$$\begin{pmatrix} \lambda & 1 & 0 & \dots & 0 & 0 \\ 0 & \lambda & 1 & \dots & 0 & 0 \\ 0 & 0 & \lambda & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & \lambda & 1 \\ 0 & 0 & 0 & \dots & 0 & \lambda \end{pmatrix}$$

for some λ.

These theorems are equivalent because of an important explicit description of direct sums:

Proposition 1. *Suppose $f_1 : V_1 \to W_1$ and $f_2 : V_2 \to W_2$ are linear maps represented by matrices M_1 and M_2. Then the map $f_1 \oplus f_2 : V_1 \oplus V_2 \to W_1 \oplus W_2$ is represented by the block matrix*

$$\begin{pmatrix} M_1 & \mathbf{0} \\ \mathbf{0} & M_2 \end{pmatrix}$$

Check this yourself if you're not convinced. It will be vital to understanding how we extract concrete statements about matrices from direct sum decompositions.

Exercise 3. *Consider the map $f : \mathbb{C}^2 \to \mathbb{C}^2$ defined by*

$$\begin{pmatrix} \lambda & 1 \\ 0 & \lambda \end{pmatrix}$$

for some $\lambda \in \mathbb{C}$. Show that we cannot decompose $\mathbb{C}^2$ into a direct sum $V \oplus V'$ such that f maps V and V' into themselves. (In other words: we can't make a canny change of basis and break down Jordan blocks even further.)

3.1 Comparison and Contrastison

Jordan Canonical Form takes more effort to prove than the rank theorem in section 2. Also in contrast to that theorem, it depends on the fact that we're working over $\mathbb{C}$ (or more generally, an algebraically closed field). A similar theorem, known as **rational canonical form**, exists over general fields, but it's messier.

Looking at the statements of the theorems, the biggest difference is that maps between two vector spaces can be classified into finitely many canonical forms (parametrized by rank), while maps from a vector space into itself require infinitely many. At the same time, it's a fairly well-behaved infinity[2]: the

[2] Don't tell Susan or Steve I used this phrasing.

Jordan blocks are controlled by a single parameter, the value λ which appears on the diagonal. So this is still a reasonable classification of linear maps.

3.2 Looking Forward

We may encounter situations more complicated than just a single linear map. For example, what if we have two maps $f : V \to W$ and $g : W \to U$? Can we choose bases of V, W, and U such that f and g both have nice matrix expressions? And more generally, if we have some collection of vector spaces and maps between them, how can we choose the bases of our spaces so that all the matrices have a nice canonical form?

A priori, there's no reason why we should be able to do this, since in Theorem 1 we had to tailor our bases to a single map. Indeed, sometimes the answer to the last question above is just "we can't". But the language of quivers gives us a systematic way to understand the cases in which we can. Keep the examples of Theorem 1 and Jordan form in mind as we build up the more general theory.

4 Quiver Representations

4.1 Basic Definitions

A quiver is just a directed graph, which we allow to have loops and multiple edges. We give a more precise definition (and introduce the relevant notation) here:

Definition. *A **quiver** is given by two sets N (of vertices) and E (of edges), together with maps $t : E \to N$ ("tail") and $h : E \to N$ ("head") which give the endpoints of an edge.*

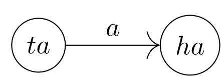

Figure 1: A simple quiver, with tail and head notation illustrated.

A **representation** V of a quiver is an assignment of a vector space $V(x)$ to every vertex x and a linear map $V(a) : V(ta) \to V(ha)$ to every edge a.

Examples.

- A representation of the quiver

  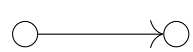

  amounts to a linear map between two vector spaces.

- A representation of

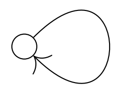

  is a single vector space together with a transformation from the space to itself.

- In the cases of the quivers

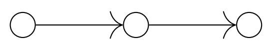

  and

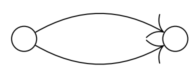

  representations are given in either case by a pair of linear maps. However, in the first case the codomain of one map is the domain of the other, while in the second case the two maps share domain and codomain.

- For any quiver without self-loops, we can define a family of representations S_x associated to the vertices. On vertices, S_x is given by

$$S_x(y) = \begin{cases} \mathbb{C} & y = x \\ 0 & \text{otherwise} \end{cases}$$

  Then since any edge connects two different vertices, one of which is associated to the 0 space, all of the maps in the representation must be 0.

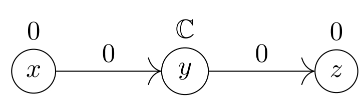

Figure 2: The representation S_y for a particular quiver.

It readily becomes apparent that quiver representations are kind of boring by themselves. As with any algebraic object, we're interested in studying maps between them. And as with any algebraic object, we want those maps to preserve the underlying structure in some way. For quiver representations, this structure is given by the maps within each representation.

Definition. *A **morphism** of representations V, W of a fixed quiver Q is given by a collection of maps $\varphi_x : V(x) \to W(x)$, as x ranges over the vertices of Q, such that for every edge a,*

$$\varphi_{ha} \circ V(a) = W(a) \circ \varphi_{ta}$$

This is best visualized using a **commutative diagram**:

$$\begin{array}{ccc} V(ta) & \xrightarrow{V(a)} & V(ha) \\ \Big\downarrow{\scriptstyle \varphi_{ta}} & & \Big\downarrow{\scriptstyle \varphi_{ha}} \\ W(ta) & \xrightarrow{W(a)} & W(ha) \end{array}$$

The condition for a collection of maps to be a morphism is summed up by saying that for every edge a, this diagram **commutes**: if we compose the arrows in either of the paths from the top left corner to the bottom right, we get the same result.

Exercise 4. *Define representations of*

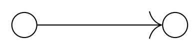

as follows: let V_1 be

$$\mathbb{C}^2 \xrightarrow{\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)} \mathbb{C}^2$$

and let V_2 be

$$\mathbb{C} \xrightarrow{\mathrm{id}} \mathbb{C}.$$

Show that mapping each copy of $\mathbb{C}^2$ to the corresponding copy of $\mathbb{C}$ by projecting to the first coordinate does not define a morphism of quiver representations. Then find an actual example of a morphism $V_1 \to V_2$.

We now proceed to define a flurry of notions analogous to important concepts in linear algebra.

Definition. *A **subrepresentation** W of a representation V is a collection of subspaces $W(x) \subset V(x)$ for each vertex x such that, for each edge, $V(a)$ maps $W(ta)$ into $W(ha)$. The restrictions of $V(a)$ to maps $W(ta) \to W(ha)$ make W a representation.*

Note that we can also start with a representation W whose vector spaces are subspaces of the $V(x)$, and this definition then says that the inclusion maps $W(x) \hookrightarrow V(x)$ define a morphism of representations.

Definition. *Given a representation V and subrepresentation W, the **quotient representation** V/W is defined by $(V/W)(x) = V(x)/W(x)$ for all vertices x, with $(V/W)(a) : V(ta)/W(ta) \to V(ha)/W(ha)$ the map induced by $V(a) : V(ta) \to V(ha)$ for any map a.*

Since W is a subrepresentation, $V(a)$ sends $W(ta)$ into $W(ha)$, and so the induced map $(V/W)(a)$ is well-defined.

Definition. *The **dimension vector** of a quiver representation is the function sending each vertex to the dimension of the space there.*

Example. Consider the representation

$$\mathbb{C}^2 \xrightarrow{\left(\begin{smallmatrix}1&1\\0&0\end{smallmatrix}\right)} \mathbb{C}^2 \xrightarrow{\left(\begin{smallmatrix}0&0\\1&1\end{smallmatrix}\right)} \mathbb{C}^2$$

Then a subrepresentation is given by

$$0 \xrightarrow{0} \mathbb{C}\begin{pmatrix}0\\1\end{pmatrix} \xrightarrow{\left(\begin{smallmatrix}0&0\\1&1\end{smallmatrix}\right)} \mathbb{C}^2$$

These representations have dimension vectors $(2,2,2)$ and $(0,1,2)$ (up to possible reordering of the vertices) respectively.

On the other hand, this representation does not have a subrepresentation with dimension vector $(2,1,0)$. Such a subrepresentation would have the form

$$\mathbb{C}^2 \xrightarrow{\left(\begin{smallmatrix}1&1\\0&0\end{smallmatrix}\right)} \mathbb{C}v \xrightarrow{0} 0$$

for some vector v; but the only possibility for v is $\left(\begin{smallmatrix}1\\0\end{smallmatrix}\right)$, and since that is not sent to 0 by $\left(\begin{smallmatrix}0&0\\1&1\end{smallmatrix}\right)$, this subrepresentation would not be compatible with the original one. So we already see that, while the definitions of these things carry over directly from the definitions for vector spaces, their behavior is a lot more subtle.

Exercise 5. *Determine what you get by quotienting by the subrepresentation above.*

Definition. *The **kernel** of a morphism of representations $\varphi : V \to W$ is a subrepresentation of V given by $(\ker(\varphi))(x) = \ker(\varphi_x)$ for all vertices x, and with maps given by the restrictions of the maps in V.*

Definition. *The **image** of a morphism of representations $\varphi : V \to W$ is a subrepresentation of W given by $(\operatorname{im}(\varphi))(x) = \operatorname{im}(\varphi_x)$ for all vertices x, and with maps given by the restrictions of the maps in W.*

Definition. *The **cokernel** of a morphism of representations $\varphi : V \to W$ is the quotient $W/\operatorname{im}(\varphi)$.*

Exercise 6. *Check that $\ker(\varphi)$ and $\operatorname{im}(\varphi)$ are indeed subrepresentations of V and W, respectively.*

This is a lot to deal with, but remember that all of these notions come down to applying the notions for vector spaces at each vertex. What's important is that we always be working with morphisms of quiver representations.

Definition. *A morphism of quiver representations $\varphi : V \to W$ is **injective** if all φ_x are injective, and **surjective** if all φ_x are surjective. It is an **isomorphism** if it has an inverse which is also a morphism of quiver representations. If there is an isomorphism between representations V and W, we say $V \cong W$ (V and W are isomorphic).*

Exercise 7. *Check that a morphism of representations is an isomorphism if and only if it is both injective and surjective.*

Let's take a minute to pick apart this notion of isomorphism of quiver representations. We say $V \cong W$ if there are invertible linear maps $\varphi_x : V_x \to W_x$ such that $\varphi_{ha} \circ V(a) = W(a) \circ \varphi_{ta}$ for each edge a. But since the maps are invertible, we can rephrase this condition slightly and say

$$W(a) = \varphi_{ha} \circ V(a) \circ \varphi_{ta}^{-1}$$

The key observation here is that, if $V(x)$ and $W(x)$ are both identified with $\mathbb{C}^n$ by choices of basis, then the map φ_x is a change of basis between them. The isomorphism relation above then says that the matrices defining the $W(a)$ are the matrices representing the $V(a)$ in this new choice of bases.

We can now rephrase our theorems from sections 2 and 3 in terms of quiver representations!

Theorem 6 (Theorem 1 Director's Cut)**.** *Any representation of*

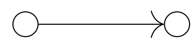

of dimension vector (m, n) *is isomorphic to one of* $\min(m, n)$ *different representations, given by the matrices in Theorem 1.*

Theorem 7 (Jordan Canonical Form Director's Cut)**.** *Any representation of*

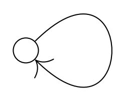

is isomorphic to one given by a matrix in Jordan canonical form.

And more generally: the problems of **finding canonical forms** for matrices which we started the class by considering are equivalent to problems of **classifying** quiver representations **up to isomorphism**.

4.2 Irreducibility and Indecomposability

This problem of classifying representations, as with many other classification problems, will become easier if we can break down our representations into component parts. So we'll define one more notion from vector spaces in this new context.

Definition. *For quiver representations* V *and* W*, the* ***direct sum*** $V \oplus W$ *is defined by* $(V \oplus W)(x) = V(x) \oplus W(x)$ *for every vertex* x*, and* $(V \oplus W)(a) = V(a) \oplus W(a) : V(ta) \oplus W(ta) \to V(ha) \oplus W(ha)$ *for each edge* a*.*

As with direct sums of single linear maps, this is easiest to visualize using block matrices.

Example. Define a representation V by

$$\mathbb{C}^2 \xrightarrow{(1\ 1)} \mathbb{C} \xrightarrow{2} \mathbb{C}$$

and a representation W by

$$\mathbb{C} \xrightarrow{-1} \mathbb{C} \xrightarrow{\left(\begin{smallmatrix} 3 \\ 2 \end{smallmatrix}\right)} \mathbb{C}^2$$

Then $V \oplus W$ is

$$\mathbb{C}^3 \xrightarrow{\left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 0 & -1 \end{smallmatrix}\right)} \mathbb{C}^2 \xrightarrow{\left(\begin{smallmatrix} 2 & 0 \\ 0 & 3 \\ 0 & 2 \end{smallmatrix}\right)} \mathbb{C}^3$$

Using direct sum, we'll want to break quiver representations down into constituent pieces to make classification easier. However, there's a slight snag here: what should our atomic pieces be? One option, which is of fundamental importance in other branches of representation theory (and in analogy with decomposing integers into prime numbers) is the notion of irreducibility:

Definition. *A quiver representation V is **irreducible** if it has no subrepresentations other than itself and the 0 representation.*

However, irreducible representations of quivers are frequently boring and unhelpful. This is best captured by the following theorem:

Definition. *An **acyclic** quiver is one whose edges do not form directed cycles.*

Theorem 8. *Let Q be an acyclic quiver. Then the only irreducible representations of Q are the representations S_x defined in the previous section.*

Exercise 8. *Prove this.*

Exercise 9. *By contrast, consider this quiver Q:*

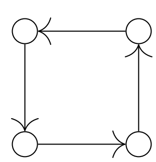

For $\lambda \neq 0$, define a representation V_λ by

$$\begin{array}{ccc} \mathbb{C} & \xleftarrow[\text{id}]{} & \mathbb{C} \\ {\scriptstyle \text{id}}\downarrow & & \uparrow{\scriptstyle \text{id}} \\ \mathbb{C} & \xrightarrow{\lambda} & \mathbb{C} \end{array}$$

(a) *Show that the V_λ are irreducible.*

(b) *Show that, for $\lambda \neq \mu$, $V_\lambda \ncong V_\mu$.*

(c) *Show that the only irreducible representations of Q are the S_x and V_λ.*

This theorem shows quite decisively that irreducible representations are not the right choice for breaking general representations into pieces. The maps of the representations S_x don't contain any information, because they're all 0! And in particular, the only representations we can form from irreducibles in an acyclic quiver by direct sum are those in which every map is 0.

Another way of stating this deficiency, looking back at the previous sections, is that quiver representations almost never satisfy the semisimplicity property: if $W \subset V$ is a subrepresentation, then it need not be the case that $V \cong W \oplus (V/W)$. If this were true, we could continue breaking off irreducible subrepresentations of any representation as direct summands, until we had expressed the entire thing as a direct sum of irreducibles—and we've seen that's not going to cover everything.

So we need to cast our net a little wider to find our building blocks:

Definition. *A quiver representation V is* ***indecomposable*** *if it cannot be written $V \cong V' \oplus V''$ for nontrivial representations V', V''.*

It's straightforward to show inductively from this definition (and the fact that we're working with finite-dimensional spaces) that indecomposables are the building block we need:

Proposition 2. *Any quiver representation is isomorphic to a direct sum of indecomposable representations.*

It's harder to show that we have some form of "unique factorization" which makes the breakdown of a representation into indecomposables well-defined, but this is also true (though we won't prove it here):

Theorem 9 (Krull-Remak-Schmidt Theorem)**.** *The decomposition of a quiver into indecomposables is unique up to isomorphism and permutation of factors. That is, if*

$$V_1 \oplus \cdots \oplus V_k \cong W_1 \oplus \cdots \oplus W_l$$

with all factors indecomposable, then $k = l$ and there is some permutation σ of the indices $1, \ldots, k$ such that $V_i \cong W_{\sigma(i)}$.

With this idea in hand, our original question of "what are the canonical forms we can place matrices in by changing bases?", which got transformed into "what are the isomorphism classes of representations of a particular quiver?", now reduces to "what are the *indecomposable* representations of a quiver?" Let's see how this informs our old friend Theorem 1.

4.3 Indecomposable Representations of One Non-Loop Edge

Theorem 10 (Theorem 1 Director's Cut Ultimate Edition)**.** *The indecomposable representations of*

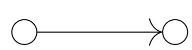

are (up to isomorphism)

$$\mathbb{C} \xrightarrow{0} 0$$
$$0 \xrightarrow{0} \mathbb{C}$$
$$\mathbb{C} \xrightarrow{\text{id}} \mathbb{C}$$

Proof. First, we check that these representations are indecomposable. For the first two, this is clear, as they are actually irreducible. The only way we could decompose the third one is as the sum of the first two; but this results in the representation $\mathbb{C} \xrightarrow{0} \mathbb{C}$, which is not isomorphic to $\mathbb{C} \xrightarrow{\text{id}} \mathbb{C}$.

Now, consider an arbitrary representation

$$\mathbb{C}^n \xrightarrow{f} \mathbb{C}^m$$

and suppose that it is indecomposable.

First, using the semisimplicity property of vector spaces, we can split $\mathbb{C}^n$ up as $\ker(f) \oplus V$ for some space V. Since f just sends $\ker(f)$ to 0, the effect of f on $\mathbb{C}^n$ is determined entirely by its effect on V. Thus we can split up our representation as

$$(\ker(f) \xrightarrow{0} 0) \oplus (V \xrightarrow{\overline{f}} \mathbb{C}^m)$$

Since we assumed it was indecomposable, we must either have $\ker(f) = 0$ (so f is injective) or $m = 0$. In the latter case, our representation is just

$$\mathbb{C}^n \xrightarrow{0} 0$$

and unless $n = 1$, this splits up as a direct sum of copies of the first representation on our list.

So assume f is injective. Then using the semisimplicity property on the other side of the map, we have that $\mathbb{C}^m \cong \text{im}(f) \oplus W$ for some space W. Again, since f only takes values in $\text{im}(f)$, the W factor is more or less irrelevant to the representation, and we can split the whole thing as

$$(\mathbb{C}^n \xrightarrow{f} \text{im}(f)) \oplus (0 \xrightarrow{0} W)$$

Then because our representation is indecomposable, either $W = 0$ (so $\text{im}(f)$ is all of $\mathbb{C}^m$ and f is surjective) or $n = 0$. In the latter case, our representation is just

$$0 \xrightarrow{0} \mathbb{C}^m$$

and unless $m = 1$, this splits up as a direct sum of copies of the second representation on our list.

Thus it remains to consider the case that $f : \mathbb{C}^n \to \mathbb{C}^m$ is an isomorphism. But then we can split $\mathbb{C}^n$ up as a direct sum $\mathbb{C} \oplus \cdots \oplus \mathbb{C}$, and under the isomorphism these are all carried to copies of $\mathbb{C}$ in a similar decomposition of $\mathbb{C}^m$.[3]

[3]But it's important to note: this might not be the decomposition that's implied by the notation $\mathbb{C}^m$. This distinction is where the "change of basis" in the original theorem is hiding.

Unless $n = m = 1$, we can then break our representation up as a direct sum of copies of the third representation on our list. This completes our classification. □

How does this relatively abstract statement translate back into the concrete version with matrices we started with? Let's think about how we build up a matrix representation of the map by summing copies of the three indecomposable representations available:

- $(\mathbb{C} \xrightarrow{\text{id}} \mathbb{C})^{\oplus n}$ is represented by the $n \times n$ identity matrix. We can also think of building up this matrix by iteratively adding copies of the representation, which corresponds to building up a block matrix by appending to the diagonal 1×1 blocks consisting of the number 1.
- Adding copies of $\mathbb{C} \to 0$ corresponds to adding columns consisting entirely of 0s, since it increases the dimension of the domain space, but the extra coordinates are irrelevant to the map.
- Likewise, adding copies of $0 \to \mathbb{C}$ corresponds to adding rows consisting entirely of 0s.

In this way, we build up a matrix in precisely the canonical form of Theorem 1.

4.4 Indecomposable Representations of Two Composed Edges

Here we look at a slightly more complicated case, different from the motivating examples we looked at in the first few sections. However, its behavior is still fairly nice, and still allows us to make a statement about canonical forms of matrices.

Theorem 11. *The indecomposable representations of*

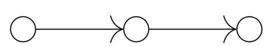

are

$$\begin{array}{ll}
\mathbb{C} \xrightarrow{0} 0 \xrightarrow{0} 0 & \quad 0 \xrightarrow{0} 0 \xrightarrow{0} \mathbb{C} \\
\mathbb{C} \xrightarrow{\text{id}} \mathbb{C} \xrightarrow{0} 0 & \quad 0 \xrightarrow{0} \mathbb{C} \xrightarrow{\text{id}} \mathbb{C} \\
0 \xrightarrow{0} \mathbb{C} \xrightarrow{0} 0 & \quad \mathbb{C} \xrightarrow{\text{id}} \mathbb{C} \xrightarrow{\text{id}} \mathbb{C}
\end{array}$$

Proof. Left as an exercise, but keep in mind that it is somewhat difficult and fiddly. You need to use the semisimplicity property of vector spaces a lot. (If I have time, I may include the proof as an appendix to these notes.) □

How can we transform this classification into a concrete statement about matrices? Once again, we consider how adding copies of the 6 indecomposable representations to our direct sum builds up a pair of block matrices. We denote the two matrices we build up this way as A (the left map) and B (the right one).

- $(\mathbb{C} \xrightarrow{\text{id}} \mathbb{C} \xrightarrow{\text{id}} \mathbb{C})^{\oplus n_1}$ is represented by two $n_1 \times n_1$ identity matrices.
- Adding $(\mathbb{C} \xrightarrow{\text{id}} \mathbb{C} \xrightarrow{0} 0)^{\oplus n_2}$ appends a further $n_2 \times n_2$ identity matrix to A, while adding n_2 columns of 0s to B.
- Likewise, adding $(0 \xrightarrow{0} \mathbb{C} \xrightarrow{\text{id}} \mathbb{C})^{\oplus n_3}$ appends an $n_3 \times n_3$ identity matrix to B, while adding n_3 rows of 0s to A.
- Adding some number of copies of $(0 \xrightarrow{0} \mathbb{C} \xrightarrow{0} 0)$ adds that many rows of 0s to A and columns of 0s to B.
- Likewise, adding some number of copies of $(\mathbb{C} \xrightarrow{0} 0 \xrightarrow{0} 0)$ adds that many columns of 0s to A.
- Likewise, adding some number of copies of $(0 \xrightarrow{0} 0 \xrightarrow{0} \mathbb{C})$ adds that many rows of 0s to B.

Putting together the matrix we've assembled in this way gives us the following theorem:

Theorem 12. *Let $f : U \to V$, $g : V \to W$ be morphisms of vector spaces. Then we can choose bases of U, V, and W such that f and g are represented by matrices of the block form*

$$\begin{pmatrix} \mathbf{I_{r\times r}} & \mathbf{0_{r\times s}} & \mathbf{0_{r\times v}} \\ \mathbf{0_{s\times r}} & \mathbf{I_{s\times s}} & \mathbf{0_{s\times v}} \\ \mathbf{0_{t\times r}} & \mathbf{0_{t\times s}} & \mathbf{0_{t\times v}} \\ \mathbf{0_{u\times r}} & \mathbf{0_{u\times s}} & \mathbf{0_{u\times v}} \end{pmatrix}$$

and

$$\begin{pmatrix} \mathbf{I_{r\times r}} & \mathbf{0_{r\times s}} & \mathbf{0_{r\times t}} & \mathbf{0_{r\times u}} \\ \mathbf{0_{t\times r}} & \mathbf{0_{t\times s}} & \mathbf{I_{t\times t}} & \mathbf{0_{t\times u}} \\ \mathbf{0_{w\times r}} & \mathbf{0_{w\times s}} & \mathbf{0_{w\times t}} & \mathbf{0_{w\times u}} \end{pmatrix}$$

for nonnegative integers r, s, t, u, v, w.

While the notation here gets a bit complicated, the forms are still relatively simple.

Exercise 10. *Show that the isomorphism class of a representation of*

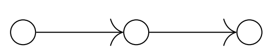

is completely determined by the dimensions of the three spaces, the ranks of the two maps, and the rank of their composition.

From these examples, one might get the impression that indecomposable representations, at least of acyclic quivers, must consist of spaces with dimension 0 or 1. (We've seen that the indecomposable representations of a loop can have any dimension, as they are given by Jordan blocks.) But this does not have to be the case.

Exercise 11. *Show that the representation of*

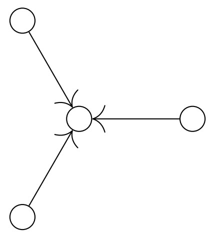

defined by

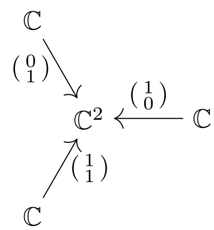

is indecomposable.

More generally, as more arrows are added to our quiver, the property of being a direct sum of two representations becomes harder to satisfy. As a result, indecomposable representations become more plentiful, and so the problem of classification gets harder. In the rest of the class, we'll single out the quivers for which classification is realistically doable, and we'll outline how the classification works when it is.

5 Gabriel's Theorem and the Finite Type Classification

Definition. *Say that a quiver is of* ***finite type*** *if it has finitely many indecomposable representations.*

Example. From the results above, we see that the quivers

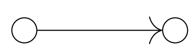

and

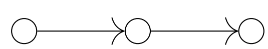

are of finite type.

On the other hand,

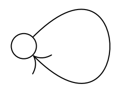

is not of finite type, since there are infinitely many possible Jordan blocks, depending on choices of eigenvalue and dimension.

Being finite type means that the essential properties of a quiver representation are conveyed by a finite collection of natural numbers giving the number of times each indecomposable representation shows up—much like how the rank in Theorem 1 is the only basis-independent property of a linear map, or how the select few ranks in Exercise 10 are the only basis-independent properties of a chained pair of maps. It also just makes the classification easy to describe. As such, it's an exceptionally nice feature and we would like to know which quivers have it!

Theorem 13 (Gabriel's Theorem)**.** *A connected quiver is finite type if and only if the graph obtained by forgetting the orientations of its edges is in one of the following:*

A_n :

D_n :

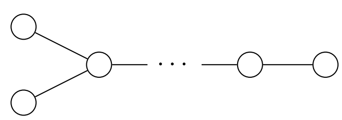

E_6 :

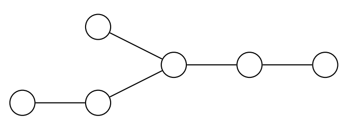

E_7 :

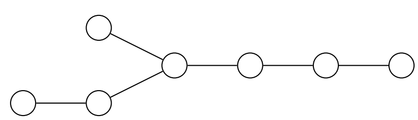

E_8 :

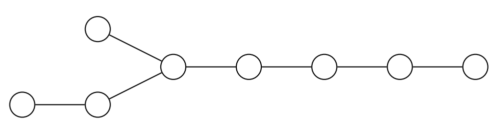

I find this theorem a little bit depressing, because it indicates that finite type is a very special property that only a few quivers have. (And the general story of classifying quiver representations is even more depressing, as we'll touch on later.)

But let's not be preoccupied with the scarcity of finite type quivers, and instead look on the bright side: these finite type quivers are cool! The graphs shown here are known as the **(type ADE or simply laced) Dynkin diagrams**, and they classify a remarkable variety of mathematical objects, such as:

- Lie groups and Lie algebras
- certain types of singularities on algebraic surfaces
- cluster algebras of finite mutation type (as seen in Véronique's class in week 1!)

If you're curious to see the Dynkin diagrams show up once more, you should check out Kevin's class on root systems next week.

We'll spend the next day or two proving the "only if" direction of this theorem: that a quiver of finite type means it must correspond to one of the Dynkin diagrams. The "if" direction is more technical, but proves a much stronger result, which allows us to classify all the indecomposable representations of any Dynkin quiver in a beautiful way. We'll outline the proof of that direction with whatever time remains in class.

5.1 Proof: Finite Type Quivers Must Be Dynkin

We set up this direction of the theorem by phrasing the idea of isomorphism of quiver representations in a slightly different way. As above, let Q be a quiver, and let N and E be respectively its vertex and edge sets.

We then fix a choice of vector spaces at the vertices of Q, which essentially amounts to fixing a dimension vector α, and consider all representations of Q given by maps between these vector spaces. Once we choose a basis of each space, these maps are just defined by matrices. Specifically, the map associated to the edge a is given by an $\alpha(ha) \times \alpha(ta)$ matrix. Thus we can identify the collection of all quiver representations on a fixed collection of spaces with dimension vector α, a collection we denote Rep_α, with the Cartesian product

$$\mathrm{Rep}_\alpha := \prod_{a \in E} \mathrm{Mat}_{\alpha(ha) \times \alpha(ta)}$$

where $\mathrm{Mat}_{m \times n}$ denotes the set of $m \times n$ matrices. For example, the set of possible choices of linear maps that fill in the representation

$$\mathbb{C}^n \xrightarrow{?} \mathbb{C}^m \xrightarrow{?} \mathbb{C}^l$$

is

$$\mathrm{Mat}_{m \times n} \times \mathrm{Mat}_{l \times m}\,.$$

What does it mean for two elements $V, W \in \mathrm{Rep}_\alpha$ to represent isomorphic quiver representations? It means that for each vertex x, there is an invertible map $\varphi_x : \mathbb{C}^{\alpha(x)} \to \mathbb{C}^{\alpha(x)}$—equivalently, an element of the group $\mathrm{GL}_{\alpha(x)}$ of invertible $\alpha(x) \times \alpha(x)$ matrices—such that $\varphi_{ha} \circ V(a) \circ \varphi_{ta}^{-1} = W(a)$ for every edge a.

But now let's look at this question from the other direction: rather than asking "how do we tell when two representations are isomorphic?", let's ask "given a particular representation, how do we track down all the ones that are isomorphic to it?" After all, if we're interested in classifying the isomorphism classes of quiver representations, it would be nice to have a way to start from a single representation and describe its class.

With this in mind, define the direct product of groups

$$\mathrm{GL}_\alpha := \prod_{x \in N} \mathrm{GL}_\alpha(x)$$

Then we can hit an element $(M_a \mid a \in E) \in \mathrm{Rep}_\alpha$ with an element $(A_x \mid x \in N) \in \mathrm{GL}_\alpha$ as follows:

$$(A_x \mid x \in N) \cdot (M_a \mid a \in E) = (A_{ha} M_a A_{ta}^{-1} \mid a \in E)$$

Essentially by the definition of isomorphism, this produces an isomorphic representation, and all representations in Rep_α isomorphic to the one we started with arise in this way. For a representation $V \in \mathrm{Rep}_\alpha$, we thus denote its isomorphism class in Rep_α by $\mathrm{GL}_\alpha \cdot V$. (If you know about group actions, you may recognize this as the **orbit** of V under the action of the group GL_α.)

Exercise 12. *Consider the quiver*

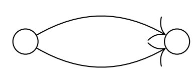

and let $\alpha = (n, m)$. What are Rep_α and GL_α? How does an element of GL_α act on Rep_α?

The question we'll investigate now becomes: are the isomorphism classes $\mathrm{GL}_\alpha \cdot V$ "big enough" that we can fill up Rep_α with finitely many of them? To proceed, I'll have to black-box some statements about dimension, which is going to be how we measure the size of these things. These statements are contained in the following white box:

(1) The dimension of Rep_α is $\sum_{a\in E}\alpha(ha)\alpha(ta)$, since a representation in Rep_α is specified by choosing entries to fill out an $\alpha(ha)\times\alpha(ta)$ matrix for each edge a.

(2) The dimension of GL_α is $\sum_{x\in N}\alpha(x)^2$. Each element of GL_α is specified by choosing entries to fill out an $\alpha(x)\times\alpha(x)$ matrix for each vertex x. While we do need to make each of the matrices invertible, almost all matrices are invertible, so that doesn't affect the dimension.

(3) If an n-dimensional group G is acting on an m-dimensional space R, but some k-dimensional subgroup has no effect on the space, then the dimension of an orbit $G\cdot r$ for $r\in R$ is at most $n-k$, because the k-dimensional part of G is irrelevant.

   – In particular, the dimension of the isomorphism class $\text{GL}_\alpha\cdot V$ is at most $\dim(\text{GL}_\alpha)-1$. This is because there is a 1-dimensional family of elements of GL_α which acts trivially on every representation: the elements $\lambda\cdot\text{id}$ in which every matrix acts by some fixed scalar multiple of the identity. (Check this if you're not convinced!)

(4) You can't take a union of finitely many things with dimension $\leq n$ and get something with dimension $> n$.

Lemma 1. *If the quiver Q is of finite type, then for any $\alpha: N\to\mathbb{Z}_{\geq 0}$,*

$$\sum_{x\in N}\alpha(x)^2-\sum_{a\in E}\alpha(ha)\alpha(ta)\geq 1$$

Proof. Suppose that the quiver Q is of finite type. Then in particular, for any fixed dimension vector α, Rep_α falls apart into finitely many isomorphism classes $\text{GL}_\alpha\cdot V$, since we can only build up so many representations of dimension α with a finite list of indecomposable ingredients. Points (1) and (4) from The Box then imply that one of the isomorphism classes must have dimension

$$\dim(\text{Rep}_\alpha)=\sum_{a\in E}\alpha(ha)\alpha(ta).$$

On the other hand, points (2) and (3) imply that the dimension of an isomorphism class can be at most

$$\dim(\text{GL}_\alpha)-1=\left(\sum_{x\in N}\alpha(x)^2\right)-1$$

Putting these together, we get that the inequality

$$\sum_{a\in E}\alpha(ha)\alpha(ta)\leq\left(\sum_{x\in N}\alpha(x)^2\right)-1$$

must hold for every possible dimension vector α, from which the statement of the lemma follows. $\square$

So we have a nice-looking combinatorial condition which is necessary for a quiver to be of finite type! We can simplify it a bit further with some linear algebra, as shown in the following exercise:

Exercise 13. *Suppose Q is a quiver with vertices N indexed by $\{1, \ldots, n\}$. Define an $n \times n$ matrix C (called the **Cartan matrix**) as follows:*

$$c_{ij} = \begin{cases} 2 - 2(\text{\# of loops at } i) & i = j \\ -(\text{\# of edges between } i \text{ and } j\text{, disregarding direction}) & i \neq j \end{cases}$$

Show that the inequality in the above lemma can be restated as

$$\alpha^T C \alpha \geq 1$$

for all $\alpha \in (\mathbb{Z}_{\geq 0})^n$.

Say a quiver Q is a **subquiver** of a quiver Q' if its vertices are some subset of Q''s vertices and its edges are some subset of the edges between those vertices. Then the criterion just established interacts with the subquiver relation in a nice way.

Lemma 2. *Suppose Q is a subquiver of Q'. Let C_Q and $C_{Q'}$ be the Cartan matrices of Q and Q', respectively. Suppose there exists a dimension vector α on the vertices of Q such that $\alpha^T C_Q \alpha < 1$. Then the same is true for $C_{Q'}$. In particular, Q' is not of finite type.*

Proof. Extend α to a function α' on the vertices of Q' by defining it to be 0 everywhere else. Then we check the condition ${\alpha'}^T C_{Q'} \alpha' < 1$ in its original form stated in Lemma 1. Since we're not adding any nonzero values at the vertices,

$$\sum_{x \in N} \alpha(x)^2$$

cannot increase, and since all of the edges of Q are present in Q',

$$\sum_{a \in E} \alpha(ha)\alpha(ta)$$

cannot decrease. Thus ${\alpha'}^T C_{Q'} \alpha' \leq \alpha^T C_Q \alpha < 1$. $\square$

Lemma 3. *Suppose we can label the vertices of a quiver Q with nonnegative integers such that, at any vertex x, the sum of the labels of adjacent vertices (disregarding edge direction) equals twice the label at x. Then neither Q nor any quiver containing Q as a subquiver is of finite type.*

Proof. The condition described in the lemma is exactly what it means for a vector α to satisfy $C\alpha = 0$. (Check this!) Then $\alpha^T C\alpha = 0 < 1$ and the condition in the previous exercise is not satisfied. Lemma 2 then implies the rest of the statement. □

Together, these lemmas give us a strategy for figuring out which quivers can possibly be finite type:

- Consider a handful of graphs which are "minimally more complicated" than Dynkin diagrams. Then these graph should not be finite type, but they should also appear as subgraphs of other non-Dynkin graphs.
- On each of these minimally more complicated graphs, find a labeling with the property given in Lemma 3.

So what are these "minimally non-Dynkin" graphs going to be? These are known as the "extended Dynkin diagrams", and each of them is defined by adding a new vertex to the Dynkin diagrams as follows:

$\widehat{A}_n$:

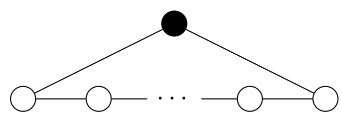

$\widehat{D}_n$:

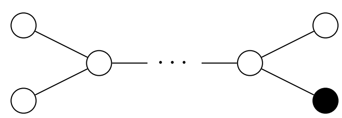

$\widehat{E}_6$:

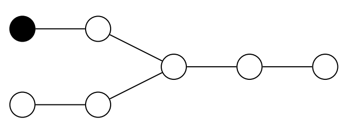

$\widehat{E}_7$:

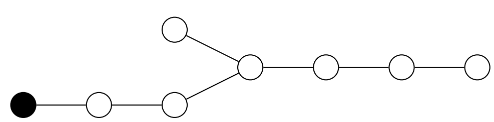

$\widehat{E}_8$:

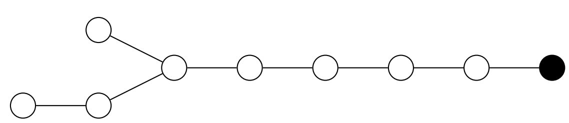

Now we exhibit labelings on these graphs which satisfy the condition of Lemma 3:

$\widehat{A}_n$:

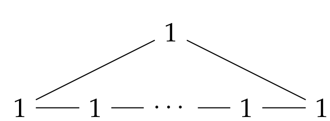

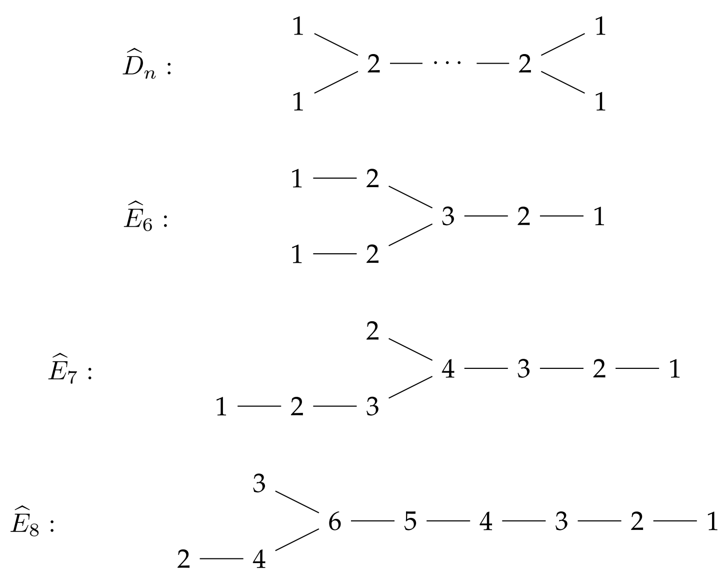

These may seem like they've been pulled out of thin air, but it helps to attempt to produce such a labeling on the Dynkin diagrams, and see how the extra vertices appearing in the extended diagrams appear out of necessity for such a labeling to work.

We now have everything we need to finish the first direction of Gabriel's theorem:

Lemma 4. *A connected graph which does not contain any of the extended Dynkin diagrams as subgraphs is a Dynkin diagram.*

Proof. Suppose we have a connected graph which does not contain any of the extended Dynkin diagrams. Then:

- Since it does not contain any of the $\widehat{A}_n$, it has no cycles, and so it is a tree.
- Since it does not contain $\widehat{D}_3$:

  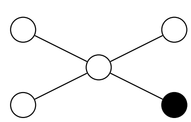

  it does not have any vertices of degree ≥ 4.
- Since it does not contain any of the other $\widehat{D}_n$'s, there can be at most one vertex of degree 3 (two vertices of degree 3 and the path between them will form a $\widehat{D}_n$). So either all vertices are of degree at most 2 (in which case the graph is an A_n) or it has three branches containing p, q, and r edges branching off from the lone vertex of degree 3.

- Since it does not contain $\widehat{E}_6$, one of the branch lengths (without loss of generality, p) must be 1.
- Since it does not contain $\widehat{E}_7$, one of the other two branch lengths (without loss of generality, q) must be less than 3. If $p = q = 1$, our graph is a D_n, so suppose $q = 2$.
- Finally, since the graph does not contain $\widehat{E}_8$, $r < 5$. Going through the various possibilities for r then produces E_6, E_7, and E_8.

□

With this, the first direction of the proof of Gabriel's theorem is done. By Lemma 3 and the reasoning below it, no quiver of finite type can contain any of the extended Dynkin diagrams as a subgraph. Thus the only quivers that can be of finite type are those given by orienting the Dynkin diagrams.

5.2 Tame and Wild Types

What does the above reasoning tell us about the possibility of classifying representations of non-finite quivers? Nothing good.

Recall that for a quiver Q with Cartan matrix C_Q, the expression

$$\alpha^T C_Q \alpha - 1$$

gives the difference between the dimension of the whole representation space and the maximum dimension of a single isomorphism class. If Q is based on a Dynkin graph, then $\alpha^T C_Q \alpha \geq 1$ for all dimension vectors α, and so this formula doesn't rule out being finite type.

It is still true that, for the extended Dynkin diagrams,

$$\alpha^T C_Q \alpha \geq 0$$

for all α. (We say that C_Q is **positive semidefinite**.) Additionally, the only α for which $\alpha^T C_Q \alpha = 0$ are scalar multiples of the labelings we gave in the above proof.

Exercise 14 (optional). *Prove this. (A potentially helpful fact: a symmetric matrix is positive semidefinite if and only if all of its eigenvalues are nonnegative).*

Then in particular,

$$\alpha^T C_Q \alpha - 1 \geq -1$$

for all α. So it's still possible that the dimension of an isomorphism class lags only 1 behind the dimension of Rep_α for any α. If this is the case, then we should expect that, although infinite, the indecomposable representations of a particular dimension vector should lie in a "1-dimensional family" which fills out all of Rep_α. And indeed, this is exactly what happens. We see an example of this in the representations of

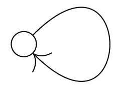

since the indecomposable representations of a given dimension vector are the Jordan blocks of that size, which are parametrized by a single eigenvalue.

We say that quivers which give extended Dynkin diagrams when orientation is ignored are of **tame type**. There are infinitely many representations in some dimensions, but they're controlled by a single parameter, which isn't all that bad.

But now, suppose there exist dimension vectors α such that $\alpha^T C_Q \alpha < 0$.

Exercise 15. *(a) Show that for each of the graphs below, there is a dimension vector α such that $\alpha^T C_Q \alpha < 0$.*

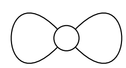

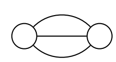

$\widehat{\widehat{A}}_n$:

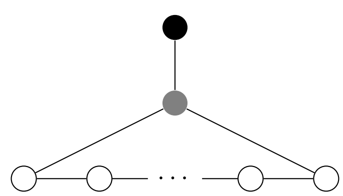

$\widehat{\widehat{D}}_n$:

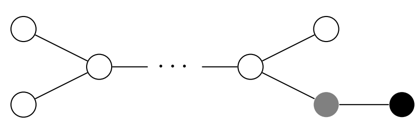

$\widehat{\widehat{E}}_6$:

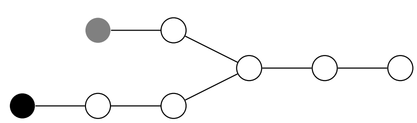

$\hat{\hat{E}}_7$:

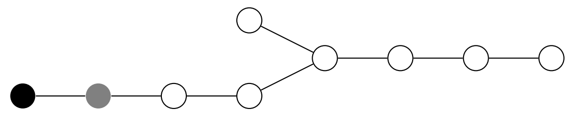

$\hat{\hat{E}}_8$:

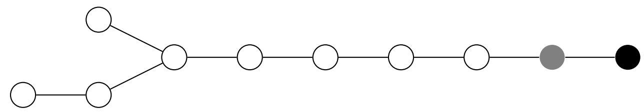

(b) Show that the only connected graphs not containing any of the above graphs as a subgraph are the Dynkin and extended Dynkin graphs. Conclude that every quiver Q with an underlying graph other than a Dynkin or extended Dynkin diagram admits some dimension vector α with $\alpha^T C_Q \alpha < 0$.

In this case, then we can find α such that $\alpha^T C_Q \alpha$ is an arbitrarily large negative number, just by scaling: scaling α by a constant k will scale $\alpha^T C_Q \alpha$ by k^2.

But this means that, as α gets larger, the gap between the dimension of an isomorphism class and the dimension of Rep_α does as well—and so increasingly complicated systems of indecomposables will be needed to account for all of the isomorphism classes. It is a generally held mathematical opinion that attempting to classify all the representations for such quivers is a hopeless task: they are said to be **wild type**. And as we've seen, most quivers (for example, any graph with a cycle and another edge) are wild! There are still interesting things we can say about their representation theory, but complete classification is not a realistic option.

5.3 Outline: Dynkin Quivers are Finite Type

The other direction of Gabriel's theorem is a bit more technical, but it actually provides an explicit way to determine all the indecomposable representations of a given Dynkin quiver. We outline the tools involved here.

5.3.1 Reflection Functors

Our goal is to construct indecomposable representations of the Dynkin quivers. An important piece of information to jump-start this is that, for any particular quiver (without loops), we already know a handful of indecomposable representations: the irreducible representations S_x, which are given by $\mathbb{C}$ at some vertex and 0 everywhere else. Our strategy will be to start with those representations and construct the other ones by transforming them. Then we can prove that this gives us *all* of the indecomposable representations by starting with an arbitrary indecomposable representation, running this process of transformation in reverse, and showing that we end up back at one of the S_x.

Our tools for transforming old indecomposable representations into new ones are called **reflection functors**. One subtle aspect of this tool is that it actually changes the quiver we're representing slightly.

Definition. *Let Q be a quiver and let x be one of its vertices. Then $\sigma_x(Q)$ is the quiver with the same vertex set and the same edges disregarding direction, but with the orientation of each edge incident to x reversed.*

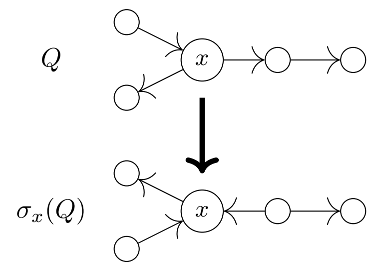

Figure 3: Like this.

We'll be specifically interested in the case that all of the edges incident to x point into it (x is a **sink**) or out of it (x is a **source**). Applying σ_x when x is a sink turns it into a source, and vice versa.

Now consider the case that x is a sink, and let V be a representation of Q. We define a representation $C_x^+(V)$ as follows:

- For vertices $y \neq x$, $C_x^+(V)(y) = V(y)$.
- For edges a not incident to x, $C_x^+(V)(a) = V(a)$.
- Adding together all of the arrows a pointing into x gives a single map of vector spaces
$$f := \bigoplus_{y \xrightarrow{a} x} V(a) : \bigoplus_{y \xrightarrow{a} x} V(y) \to V(x)$$
and we define
$$C_x^+(V)(x) = \ker(f)$$
- Since $C_x^+(V)(x) = \ker(f)$ is a subspace of $\bigoplus_{y \xrightarrow{a} x} V(y)$, we can define the maps $C_x^+(V)(a) : C_x^+(V)(x) \to V(y)$ to be the projections from this direct sum onto its summands.

This is a rather arcane-looking definition, but let's see what it gives us when we apply it to a representation we considered above.

Example. Let V be

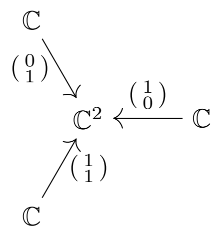

Let x be the sink in the center of the quiver. To find $C_x^+(V)(x)$, we first add together the three maps pointing into it, and get a map $f : \mathbb{C}^3 \to \mathbb{C}^2$ given by the matrix

$$\begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}$$

This matrix is full-rank, and so its kernel is 1-dimensional, spanned by the vector

$$v = \begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}$$

When we project the vector cv onto the three copies of $\mathbb{C}$ we started with, we get c, c, and $-c$ respectively. So with v as the basis of $\ker(f)$, the representation $C_x^+(V)$ is

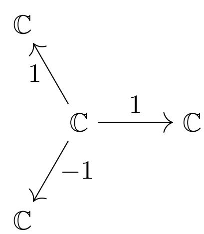

Now, suppose x is a source and V is a representation. We'll do pretty much the same construction, but with all the arrows in reverse. Define a representation $C_x^-(V)$ as follows:

- For vertices $y \neq x$, $C_x^-(V)(y) = V(y)$, and for edges a not incident to x, $C_x^+(V)(a) = V(a)$, same as before.
- Adding together all of the arrows a pointing out of x gives a single map of vector spaces

  $$g := \bigoplus_{x \xrightarrow{a} y} V(a) : V(x) \to \bigoplus_{x \xrightarrow{a} y} V(y)$$

  and we define

  $$C_x^-(V)(x) = \operatorname{coker}(g).$$

  (Recall—or just call—that for a map $F : V \to W$ of vector spaces, $\operatorname{coker}(F)$ is the quotient space $W/\operatorname{im}(F)$.)

- Since $C_x^-(V)(x) = \operatorname{coker}(g)$ is a quotient of $\bigoplus_{x \xrightarrow{a} y} V(y)$, we can define the maps $C_x^-(V)(a) : V(y) \to C_x^-(V)(x)$ to be the inclusions of the summands $V(y)$, followed by projection onto the quotient.

This is a slightly more arcane-looking definition, because quotients are annoying, but let's see what it gives us when we apply it to the representation we just got by applying C_x^+. Let W be

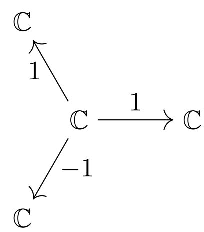

and let x be the source in the center of the quiver. To find $C_x^-(V)(x)$, we first add together the three maps pointing out of it, and get a map $g : \mathbb{C} \to \mathbb{C}^3$ given by the matrix

$$\begin{pmatrix} 1 \\ 1 \\ -1 \end{pmatrix}.$$

Then a basis of $\operatorname{coker}(g)$ is given by

$$e_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, e_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}$$

The inclusion maps of the three coordinate vectors send them to e_1, e_2, and $e_3 = e_1 + e_2 \mod \operatorname{im}(g)$. So in this basis, the representation $C_x^-(W)$ is

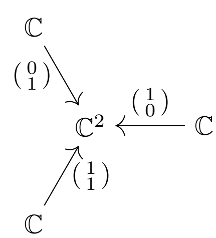

which is the representation we started with! What a remarkable coincidence.

Theorem 14 (This Is, of Course, Not a Coincidence)**.** *(a) Suppose Q is a quiver with a sink x and V is a representation which does not have the irreducible representation S_x as a direct summand. Then $C_x^-(C_x^+(V)) \cong V$.*

(b) Suppose Q is a quiver with a source x and V is a representation which does not have the irreducible representation S_x as a direct summand. Then $C_x^+(C_x^-(V)) \cong V$.

Proof. Suppose x is a sink.

First, where is this requirement that S_x not be a summand coming from? The key observation to make here is that $C_x^+(S_x) = 0$. More generally,

Lemma 5. *A representation V has a direct summand isomorphic to S_x if and only if the map*

$$f = \bigoplus_{y \xrightarrow{a} x} V(a) : \bigoplus_{y \xrightarrow{a} x} V(y) \to V(x)$$

is not surjective.

Proof. If V has a direct summand isomorphic to S_x, then none of the maps $V(a) : V(y) \to V(x)$ will take values in the part of $V(x)$ coming from S_x, since S_x contributes nothing to the other vertices.

Conversely, suppose V is not surjective. Then choose some $v \in V(x)$ which is not contained in the image of f; since it does not appear in the image of any $V(a) : V(y) \to V(x)$, we can split a direct summand off of V which is given by the span of v at x and 0 everywhere else.

□

Now suppose V does not have S_x as a direct summand, so that

$$f : \bigoplus_{y \xrightarrow{a} x} V(y) \to V(x)$$

is surjective. By definition,

$$C_x^+(V)(x) = \ker(f).$$

Then the key observation to make is that

$$\bigoplus_{y \xrightarrow{a} x} C_x^+(V)(y) = \bigoplus_{y \xrightarrow{a} x} V(y)$$

since C_x^+ leaves all the spaces away from x untouched, and so by the definition of C_x^-,

$$C_x^-(C_x^+(V))(x) = \frac{\bigoplus_{y \xrightarrow{a} x} C_x^+(V)(y)}{C_x^+(V)(x)} = \frac{\bigoplus_{y \xrightarrow{a} x} V(y)}{\ker(f)}$$

But now, by the first isomorphism theorem, reducing f mod $\ker(f)$ gives an isomorphism

$$\frac{\bigoplus_{y \xrightarrow{a} x} V(y)}{\ker(f)} \cong \mathrm{im}(f) = V(x)$$

where we use here the fact that f is surjective.

Checking that this isomorphism at x, together with the identity maps at all other vertices, actually gives us an isomorphism of quiver representations, is routine but fiddly and is left as an exercise.

Likewise, part (b) of the theorem follows the same path as part (a)—just replace all instances of "surjective" with "injective", "kernel" with "cokernel" (and vice versa), and reverse all the arrows. This is also left as an exercise. □

So C_x^+ and C_x^- are almost inverses to each other. Using one more diagram-heavy but easily proven fact, we'll see how they can tell us about indecomposable representations in particular.

Exercise 16. *Show that C_x^+ and C_x^- preserve direct sums, in the sense that*

$$C_x^{\pm}\left(\bigoplus_i V_i\right) = \bigoplus_i C_x^{\pm}(V_i).$$

Theorem 15. *If V is an indecomposable representation which is not S_x, then $C_x^{\pm}(V)$ is also indecomposable.*

Proof. If $C_x^+(V)$ admitted a nontrivial direct sum decomposition, then by the above exercise, applying C_x^- would give a nontrivial direct sum decomposition of $C_x^-(C_x^+(V)) \cong V$. The only way this wouldn't be an issue is if one of the summands of $C_x^+(V)$ was S_x, and vanished when we applied C_x^-. However, by the definition of $C_x^+(V)$, the map

$$C_x^+(V)(x) \to \bigoplus_{y \xrightarrow{a} x} C_x^+(V)(y)$$

is injective (as it is just inclusion of a kernel), and so by the version of Lemma 5 which applies to C_x^- (with "surjective" replaced by "injective" and arrows reversed) $C_x^+(V)$ cannot have S_x as a direct summand. The same kind of reasoning works for C_x^-. □

When we were classifying indecomposable representations of quivers earlier, we started with a list pulled out of nowhere. The reflection functors give us a way to figure some part of that list out.

Exercise 17. *Pick a quiver and an indecomposable representation on it. Apply some reflection functors and see what you get.*

In particular, if Q is a tree (in particular in particular, if it's a Dynkin quiver), then it's possible to order its vertices $x_1, \ldots, x_n$, such that x_1 is a sink in Q, x_2 is a sink in $\sigma_{x_1}(Q)$, x_3 is a sink in $\sigma_{x_2}(\sigma_{x_1}(Q))$, and so forth. We can then apply the reflection functors at $x_1, \ldots, x_n$ in sequence; the result will be a representation of the quiver we started with, because we reversed each arrow exactly twice, once at each of its endpoints. Given an ordering that does this, we get a mapping from representations of Q to other representations of Q by applying $C_{x_n}^+ C_{x_{n-1}}^+ \cdots C_{x_1}^+$. We denote this composition C^+ and call it the **Coxeter functor**.

There's no reason that this should give us all the representations, and in general it doesn't. But for a Dynkin diagram, it does! The way we track this is by examining how the dimension vector of the representation changes.

Exercise 18. *Suppose that V is a representation of Q with dimension vector α, and let x be a sink. Suppose further that V does not have a direct summand isomorphic to S_x. Then $C_x^+(V)$ has dimension vector $\sigma_x\alpha$ given by*

$$\sigma_x\alpha(z) = \begin{cases} \alpha(z) & \neq x \\ \left(\sum_{y \xrightarrow{a} x} \alpha(y)\right) - \alpha(x) & z = x \end{cases}$$

The same formula, but summing over $y \xleftarrow{a} x$, holds for the dimension vector of $C_x^-(V)$ when x is a source.

The rest of the proof is showing that every indecomposable representation of a Dynkin quiver is obtained by applying reflection functors to some S_x. We don't have time to cover this in detail, but here is a rough description of how the proof goes:

- First, we consider the set of possible dimension vectors that can arise from applying reflection functors to S_x—that is, by applying the transformation in the above exercise to the standard basis vectors. We refer to these dimension vectors as **roots**. (This terminology comes from the theory of root systems—take Kevin's class next week if you're interested!)
- Then, we make use of the fact that we're working with a Dynkin diagram. We showed above that for any non-Dynkin diagram, there exists a dimension vector α such that $\alpha^T C_Q \alpha < 1$. In contrast, it is actually true that when Q is a Dynkin diagram, $\alpha^T C_Q \alpha \geq 1$ for any nonzero dimension vector α. In fact, $x^T C_Q x > 0$ for any nonzero vector x; we say that C_Q is **positive definite**.
- The map $\alpha^T C_Q \alpha$ interacts with the reflection operations σ_x in a nice way: $(\sigma_x\alpha)^T C_Q(\sigma_x\alpha) = \alpha^T C_Q \alpha$. In particular, $\alpha^T C_Q \alpha = 2$ for every root α.
- Using a compactness argument and the positive definiteness stated above, one shows that there are finitely many roots. (This works by analogy with the fact that there are only finitely many points in space with integer coordinates and length 2.)
- Once we know that there are finitely many roots, we need to show that every indecomposable representation arises from applying reflection functors to S_x. To do this, start with an indecomposable representation and repeatedly hit it with a Coxeter functor. We also track what happens to its dimension vector when we apply the corresponding σ_x transformations from the above exercise; by a small technical argument, one shows that this must eventually produce a vector with a negative component.
- But whenever we apply a reflection functor to an indecomposable representation other than S_x, we get another indecomposable representation, and σ_x will give us its dimension vector, which can't have a negative component! So the only way this can happen is if, in repeatedly applying

reflection functors to the representation we started with, we eventually got S_x. The last reflection functor can't have been at x, since by the argument at the end of the proof of Theorem 15, reflecting at x can't produce S_x. So we can undo all of the reflection functors we applied with the reverse reflection, starting at S_x, and get our indecomposable representation back!