An electrical circuit is a network of wires with current (electrical charge) flowing along each wire. Kirchhoff's current law says that at each junction in the circuit, the total current flowing into the junction equals the total current flowing out of the junction. Pictured below is a circuit of seven wires and three junctions. The currents (in milliamperes) are known on four of the wires and unknown on three of the wires.

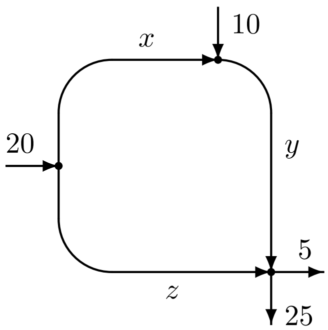

A. Write a system of linear equations for the unknown currents.

B. Solve that system to find the unknown currents.

You're a hydrologist studying the groundwater under Carleton's campus. For simplicity, let's assume that campus is utterly flat. However, due to variations in the kinds of soil, clay, and rock under campus, the surface of the groundwater is not flat, as the surface of a lake is. Instead, the depth of the water's surface below ground varies from place to place. For example, here are some depths (in m) that you measure in an evenly spaced grid on the Bald Spot:

1.21	1.16	1.11	1.06
1.25	1.19	1.13	1.07
1.29	1.22	1.15	1.08
1.33	1.25	1.17	1.09

Now you want to understand the groundwater in the area of Goodsell Hall. The problem is that you can't measure under Goodsell itself; you can only measure around it. You get

1.22	0.91	0.75	0.83
1.43	a	b	0.91
1.29	c	d	0.93
1.23	1.17	1.09	1.04

where a, b, c, and d are unknown, because they're under Goodsell. Water diffusion theory tells you that each depth measurement should be approximately equal the average of the four measurements adjacent to it.

A. Using this principle, write (but don't solve) a system of linear equations that can be used to find a, b, c, and d.

In chemistry, x molecules of hydrogen peroxide (H_2O_2) decompose into y molecules of water (H_2O) and z molecules of oxygen (O_2). A chemist might write the reaction as

$$xH_2O_2 \to yH_2O + zO_2.$$

The total number of hydrogen atoms has to match before and after the reaction, and the same goes for oxygen. What can x, y, and z be?

The World Wide Web consists of billions of web pages, with one-way connections, called links, among those pages. See the figure below for an example with only eight pages. When you use Google to search the Web, Google finds all pages relevant to your search and shows the most valuable ones first. But what does "valuable" mean? Well, some web sites pay Google to rank their pages highly. Let's ignore that advertising aspect of the system. Let's instead try to understand the algorithm that assigns baseline values to web pages.

Roughly, a page is valuable if many other pages link to it — especially if those pages are themselves valuable. More precisely, pages have values and links have values. A page gets value from the links coming into it; the page's value is simply the sum of the values of the incoming links. A link gets value from the page that it comes out of; a page's value is evenly shared among all the links that come out of it.

If this logic sounds circular, that's because it is kind of circular. You can't figure out the value of one web page in isolation; you have to figure out all of the values at once. The values are "coupled" to each other.

Actually, in one special situation you can figure out the value of a web page in isolation. Consider the web page at the top of the diagram. It has no links coming into it. So its value must be 0. It has one link going out, so the value of that link is $0/1 = 0$. So that link contributes no value to the web page that it points to.

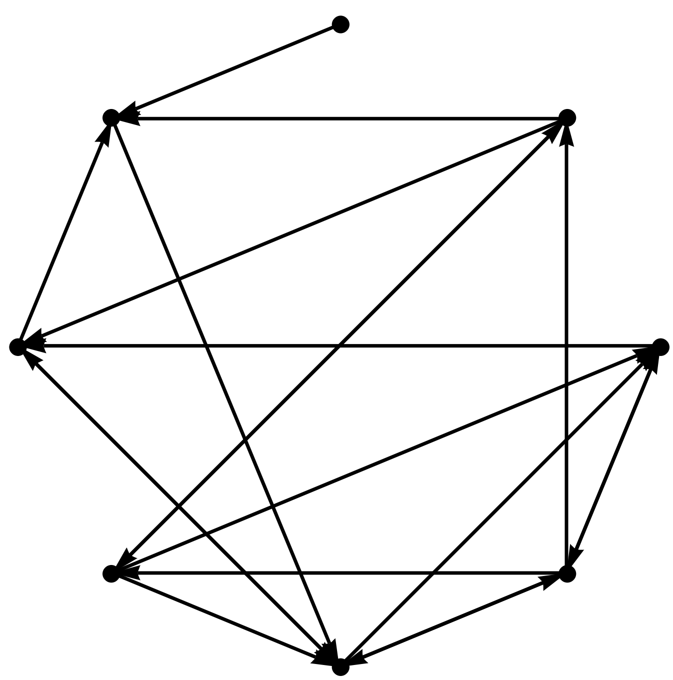

A. In the figure above, number the pages from 1 to 8, going clockwise from the top. Let x_i be the value of page i. Write (but don't solve) a system of equations, whose solution tells us the values of all of the pages.

We wish to find a cubic polynomial $y = ax^3 + bx^2 + cx + d$ that passes through the points $(-1, 3)$, $(0, 1)$, $(1, 3)$, and $(2, 4)$. That is, we wish to find a, b, c, and d.

A. Write (but don't solve) a system of equations for the unknown coefficients a, b, c, d.

B. What if there were a fifth data point, that we wished the cubic polynomial to hit? Discuss thoroughly.

In the Democratic Republic of Pretendland (DRP) there are n major cities, which are called $C_1, \ldots, C_n$ (in decreasing order of population). Some of the cities are connected by non-stop high-speed trains; let A be the $n \times n$ adjacency matrix of this train network. Some of the cities are connected by non-stop airplane flights; let B be the $n \times n$ adjacency matrix of this airplane network. Just so we're clear: There is one set of cities here, but there are two separate graphs built on them, with different adjacency matrices.

A. You are planning a trip to the DRP in which you'll be traveling to many of its major cities. How is the matrix $A + B$ useful to you?

B. Mathematically, what is the difference between $(A + B)^2$ and $A^2 + B^2$? Be sure to show every step of your work.

C. Practically, what is the difference between $(A + B)^2$ and $A^2 + B^2$, to your trip?

Let S, M, H, E, O denote the concentrations of a certain pollutant in lakes Superior, Michigan, Huron, Erie, and Ontario, respectively. All of the lakes already contain some of the pollutant. Each year, more pollutant is added to Lake Superior by industrial activity along its shores; also, the pollutant mixes among the lakes as they slowly drain eastward into the Atlantic Ocean. While researching the topic you come across the following discrete dynamical system model for the dispersion of the pollutant:

$$\begin{bmatrix} S_{k+1} \\ M_{k+1} \\ H_{k+1} \\ E_{k+1} \\ O_{k+1} \end{bmatrix} = \begin{bmatrix} 1.1 & 0.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 0.8 & 0.0 & 0.0 & 0.0 \\ 0.1 & 0.1 & 0.7 & 0.0 & 0.0 \\ 0.0 & 0.0 & 0.2 & 0.5 & 0.0 \\ 0.0 & 0.0 & 0.0 & 0.2 & 0.6 \end{bmatrix} \begin{bmatrix} S_k \\ M_k \\ H_k \\ E_k \\ O_k \end{bmatrix}.$$

You could compute the eigensystem of the 5×5 matrix using Mathematica's `Eigensystem` function. To save you time, I'll just tell you that the eigenvalues are

$$\lambda_1 = 1.1, \quad \lambda_2 = 0.8, \quad \lambda_3 = 0.7, \quad \lambda_4 = 0.6, \quad \lambda_5 = 0.5,$$

and their respective eigenvectors are approximately

$$\vec{v}_1 = \begin{bmatrix} 0.97 \\ 0.00 \\ 0.24 \\ 0.08 \\ 0.03 \end{bmatrix}, \quad \vec{v}_2 = \begin{bmatrix} 0.00 \\ 0.59 \\ 0.59 \\ 0.39 \\ 0.39 \end{bmatrix}, \quad \vec{v}_3 = \begin{bmatrix} 0.00 \\ 0.00 \\ 0.41 \\ 0.41 \\ 0.82 \end{bmatrix}, \quad \vec{v}_4 = \begin{bmatrix} 0.00 \\ 0.00 \\ 0.00 \\ 0.00 \\ 1.00 \end{bmatrix}, \quad \vec{v}_5 = \begin{bmatrix} 0.00 \\ 0.00 \\ 0.00 \\ 0.45 \\ -0.89 \end{bmatrix}.$$

A. Explain why the numbers in the third row of the 5×5 matrix are reasonable.

B. Based on this model, what is the long-term trend for the distribution of the pollutant among the Great Lakes? For example, which two lakes will end up most polluted? How many times as polluted, will the most polluted lake be, as the second-most-polluted lake?

After your trip to the DRP, you return to your job as an aerospace engineer. You are working on a spacecraft that, at a crucial point in its maneuvers, rotates in two stages. First the spacecraft rotates about the z-axis of space through an angle α. Three minutes later it rotates about the y-axis of space through an angle of β.

A. Find a matrix that expresses the combined effect of the z-axis rotation followed by the y-axis rotation.

In 3D computer graphics applications such as video games, an object is described as a list of polygons. Each polygon is described by its vertices, which are points in $\mathbb{R}^3$. For a variety of reasons, it is also useful to give each polygon a *normal vector* $\vec{n}$ — meaning a nonzero vector, perpendicular to the polygon, that points "out" from the object. For example, the cube below left is made up of six square sides, each with a normal vector (although only three of the sides are visible).

Frequently (as in, millions of times per second) we wish to transform our polygons by a 3×3 matrix A. Transforming the vertices of a polygon is simple; we just apply A to each vertex. Transforming the normal vector is more subtle; if we just apply A to the normal vector $\vec{n}$, then the resulting vector $A\vec{n}$ may no longer be perpendicular to the polygon; see the transformed cube below in the middle. To produce transformed normals that are perpendicular to their transformed polygons, we must transform them in a different way; see below right.

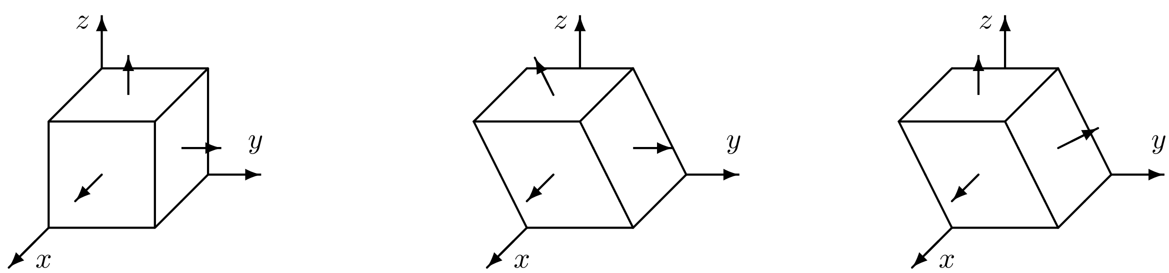

A. Show that if $\vec{n}$ is perpendicular to a given polygon, then $\left(A^{-1}\right)^\top \vec{n}$ is perpendicular to the transformed polygon. In other words, when vertices transform by A, normal vectors transform by $\left(A^{-1}\right)^\top$.

B. What happens in the special case when A is a rotation? Explain in detail.

In quantum physics, the state of the universe is represented as a vector $\vec{v}$ of very high dimension d. Depending on the specific assumptions that you're using, the dimension d might be infinite, or it might be merely gigantic, like $2^{10^{89}}$. Usually, quantum physicists don't study the entire universe at once, but rather a small part of the universe, such as a single electron. In any case, the state of the system under consideration is a vector of some dimension d. Let's assume that d is finite.

The state of the system changes with time, so $\vec{v}$ is actually a function of time t. We write this function as $\vec{v}(t)$. At time 0, which let's say is the start of the experiment, the state is a vector $\vec{v}(0)$. At time 1 the state is a vector $\vec{v}(1)$. And so on.

Moreover, the change in the state from time 0 to any other time t is a linear transformation. In other words, there exists a time-dependent matrix $U(t)$ such that $\vec{v}(t) = U(t)\vec{v}(0)$.

A. What is the size of $U(t)$? I mean, how many rows and columns does it have?

Moreover, quantum physics says that information is conserved. So if I tell you what $\vec{v}(t)$ is, you should (in principle) be able to work backwards to what $\vec{v}(0)$ was. For example, it is not possible for two different values of $\vec{v}(0)$ to produce the same value of $\vec{v}(1)$.

B. What does the information conservation part tell us about $U(t)$?

Let $\vec{e}_1, \vec{e}_2, \ldots, \vec{e}_d$ be unit vectors along the coordinate axes. So $\vec{e}_i$ is a column with a 1 in the ith entry and 0s elsewhere. When the state $\vec{v}$ is *measured*, it *collapses* to one of the $\vec{e}_i$. Which one? It's random. For $i = 1, 2, \ldots, d$, the probability that $\vec{v}$ collapses to $\vec{e}_i$ is v_i^2.

C. In any random event, such as the collapse of a state $\vec{v}$, the probabilities must sum to 1. What does this fact imply about the vector $\vec{v}$?

D. The matrix U needs to preserve the property that you just enunciated in Part C, for all vectors $\vec{v}$. What does this fact imply about U?

By the way, in actual quantum theory the vectors $\vec{v}$ and the matrices U are made up of complex numbers rather than real numbers. That fact forces adjustments, but only small ones, to the description above. Also, physicists use different notation. Instead of writing $\vec{v}(t) = U(t)\vec{v}(0)$, they might write $|\psi(t)\rangle = U(t)|\psi(0)\rangle$. But that is a cosmetic difference.

In an earlier problem, I mentioned that normal vectors are important in computer graphics. This problem explains how normals are used in computing how light reflects off a surface.

The diagram below shows a triangle (viewed edge-on, so that it looks like a line segment) with normal vector $\vec{n}$. The vector $\vec{\ell}$ points toward a light source. A light ray comes from the light source, hits the marked point on the triangle, and bounces off in the reflected direction $\vec{r}$. Given $\vec{n}$ and $\vec{\ell}$, we want to compute $\vec{r}$.

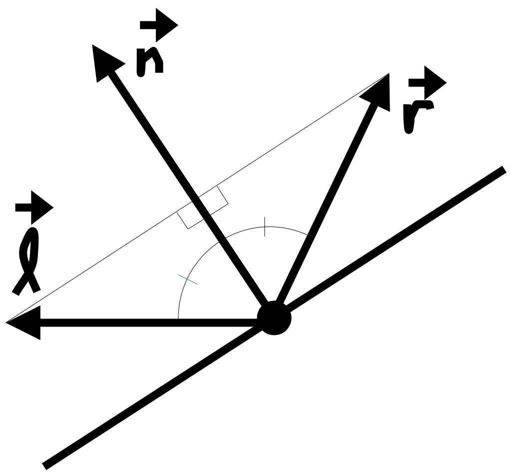

Assume $\vec{n}$ has length 1. Do not assume that $\vec{\ell}$ has length 1. Your answers should work in any dimension (although in practice these calculations are done in 3D or 4D).

A. Using vector projections, show that

$$\vec{r} = 2(\vec{\ell} \cdot \vec{n})\vec{n} - \vec{\ell}.$$

B. Algebraically check that, for the value of $\vec{r}$ just given, $\vec{r}$ has the same length as $\vec{\ell}$.

Let $\vec{w}$ be a unit 3D vector, and let α be a number. A rotation of 3D space, about the axis $\vec{w}$, through the angle α, should be a linear transformation. In the 1810s, Rodrigues found a succinct and computationally efficient matrix representation of this linear transformation. Rodrigues's discovery is still used heavily in robotics, aeronautics, computer graphics, crystallography, and other fields where the rotation of 3D objects is important. Let

$$W = \begin{bmatrix} 0 & -w_3 & w_2 \\ w_3 & 0 & -w_1 \\ -w_2 & w_1 & 0 \end{bmatrix}.$$

Rodrigues discovered that the matrix

$$U = I + (\sin\alpha)W + (1 - \cos\alpha)W^2$$

is rotation about $\vec{w}$ through the angle α. (There are actually two such rotations. One is "left-handed" and the other is "right-handed". Rodrigues's is the right-handed one. I can explain what that means if you ask me, but you don't need to know for the following questions.)

In the following questions, you will partially check that Rodrigues was right.

A. Algebraically show that $U\vec{w} = \vec{w}$. Geometrically, what does this mean?

B. Algebraically show, for all vectors $\vec{v}$, that $\|U\vec{v}\| = \|\vec{v}\|$.

C. Pick any one vector $\vec{v}$ that is perpendicular to $\vec{w}$, and show that the angle between $\vec{v}$ and $U\vec{v}$ is α. (Hint: The columns of W are perpendicular to $\vec{w}$.)

Suppose that you are a physical chemist studying the radioactive decay of a certain element into another element. You measure the quantity q (in moles) of the first element at various times t (in seconds, maybe). Here is your observed data set.

t	0	1	2	3	4	5	6
q	1.284	1.160	1.119	0.860	0.858	0.787	0.692

Based on the data and your knowledge of physics, you suspect that q should decay exponentially with t. This phrase means that $q = ae^{bt}$ for some as-yet-unknown constants a and b. You would like to find the "right" values of a and b.

The first obstacle is that no curve $q = ae^{bt}$ will perfectly fit your observations, because of unavoidable "noise" in the data. So you'll have to do statistics rather than just math.

The second obstacle is that you'd prefer to do linear regression, as opposed to some more complicated technique, but your proposed model for q does not depend linearly on the parameters a and b. In this case, there's an easy fix. Let $z = \log q$ and $c = \log a$ (where log means the "natural" or "base-e" logarithm). Then the equation $q = ae^{bt}$ becomes $z = c + bt$. Here, z depends linearly on c and b, so you can use linear regression to discover c and b, and then compute $a = e^c$.

You are welcome to do the following work by hand with the assistance of a calculator. You might find it easier to see the section "An exercise" near the bottom of the Mathematica notebook leastSquares.nb on our course web site. That notebook section is set up specifically for this problem. Either way you do it, show enough work that another human could reproduce your work. Reproducibility is important to the scientific process.

A. Express this curve-fitting problem in the form $T\vec{c} = \vec{z}$, where T and $\vec{z}$ are known and $\vec{c}$ is unknown. Explicitly show T, $\vec{c}$, and $\vec{z}$.

B. What is the matrix P that represents projection onto the column space of T?

C. What is the matrix that you multiply by $\vec{z}$ to get the least-squares solution $\vec{c}$, meaning the solution to the modified problem $T\vec{c} = P\vec{z}$?

D. So what are a and b? (So that you can check your work, I'll tell you that they are approximately $a = 1.288$ and $b = -0.103$. But I want to see your work that leads up to the answer.)