0.1 Friend Trends

This puzzle is based on the following anecdote concerning a Hungarian sociologist and his observations of circles of friends among children.

> "In the 1950s, a Hungarian sociologist S. Szalai studied friendship relationships between children. He observed that in any group of around 20 children, he was able to find four children who were mutual friends, or four children such that no two of them were friends. Before drawing any sociological conclusions, Szalai consulted three eminent mathematicians in Hungary at that time: Erdos, Turan and Sos. A brief discussion revealed that indeed this is a mathematical phenomenon rather than a sociological one. For any symmetric relation R on at least 18 elements, there is a subset S of 4 elements such that R contains either all pairs in S or none of them. This fact is a special case of Ramsey's theorem proved in 1930, the foundation of Ramsey theory which developed later into a rich area of combinatorics."
>
> (Quoted from lecture notes by MIT Prof. Jacob Fox.)

The puzzle we now present follows the same idea but with some smaller numbers. Specifically, we are interested in investigating the smallest size of a group of people that *necessitates* a subgroup of three people that are all mutually friends or all mutually enemies.

> Assume that amongst a group of people, any two of them are either friends or enemies, and that these are the only possible relationships (i.e. no acquaintances or frenemies or anything like that). Take a group of four people and try to assign a designation of friend/enemy to each pair so that there are *no* groups of three people that are all friends or all enemies. Can you do this with a group of five people? How about six? Seven? Ten? Twenty? Try to identify a cutoff number for the size of the group where you can be *guaranteed* to find a subgroup of three people that are all friends or all enemies.

Think carefully about this before turning the page and reading our solution.

Did you figure it out? This is a very tricky puzzle, so don't feel bad if you struggled with finding a solution. In fact, we think that investigating this puzzle is just as important as actually finding an answer, because there are several ways to approach this puzzle and it's always interesting to see how different people interpret the puzzle.

Let's start by discussing how to even write down/draw/talk about this situation. For any of the questions posed by this puzzle, we are meant to consider a group of people with a certain size and think about the relationship between any two people in the group. To tackle this puzzle, we will need a way to represent all of these relationships in an efficient and easily-interpretable way, so that we can verify the desired property about the subgroups of size three. Specifically, we want to easily identify whether or not there are any subgroups of size three that are *homogeneous*, in the sense that the three people are *all* friends or *all* enemies. From now on, we will refer to this as the "homegeneity property" to be more concise. How could we do this? We could number the people in the group, then write out a list of all pairs of numbers and label each pair with F (friend) or E (enemy). Let's try doing that for a group of four people:

$$12F \quad 13E \quad 14F \quad 23F \quad 24E \quad 34F$$

Does this friend/enemy group satisfy the subgroup property? It's not so easy to verify, is it? For one, the numbering makes it difficult to find subgroups of size three, and to verify the property, we need to check *all* such subgroups to make sure they are not EEE or FFF. Perhaps we should find a better way of *representing* the information of the puzzle before attempting a solution. Can you think of a more visually pleasing way of representing whether two people are friends or enemies, for *all* possible pairs in the group? Specifically, we would like to have a relatively efficient way of looking for subgroups of size three and recognizing whether they are all friends or all enemies.

Let's try representing each person in the group as a single dot and connecting two people with a type of line dependent on whether those two are friends or enemies. For example, let's connect two people that are friends with a blue line and two people that are enemies with a red line (and remember that any two people are either friends or enemies and nothing else, so each pair of dots must have *some* colored line between them). Now, what would we be looking for to verify the homogeneity property? We want three dots (three people) so that all of the lines between them are either blue (all friends) or red (all enemies). That's right—we're looking for monochromatic triangles! (Note: we want the vertices of the triangle to be one of the original dots we drew; that is, we don't want a vertex to be a place where two lines cross. Also, *monochromatic* comes from the Greek words *monos* and *khroma*, meaning "one" and "color", respectively.) This representation is much easier to interpret visually and makes checking for a solution much faster. Try drawing the dots-and-lines diagram that matches the list we gave above. Did you find any monochromatic triangles? What happens if you move your four dots around on the paper? Did that affect the solution? Did it make it any easier to look at? How else could you draw the lines?

Based on the diagram that you just drew, it would seem that we have addressed the question regarding four people: we have an arrangement of friends and enemies so that there are no subgroups of size three that are either all friends or all enemies. Can you find another arrangement? How can you be sure that it's a *different* arrangement than the one we've already seen? How many different arrangements are there that satisfy the homogeneity property? Now, try drawing an arrangment that *does not* satisfy the homogeneity property. What does that look like? How many of these arrangements are there?

Let's move on and think about a group of five people. Our diagram changes because we have five dots now, and this means there are more lines to draw. Still, we are looking to fill in all of the connections with a blue or red line and make sure there are no monochromatic triangles. Is this possible? (Hint: try arranging the dots into a regular pentagon shape and then filling in the lines.) Try to draw this a few times and see if any of your arrangements work. It may also help to draw in a few lines randomly and then guide your choices from there on out by making sure that you don't create any triangles when you add a new line.

Did you figure it out? Turn the page to see how we did it . . .

Here is our arrangment of red/blue lines amongst five dots that completely avoids the homogeneity property:

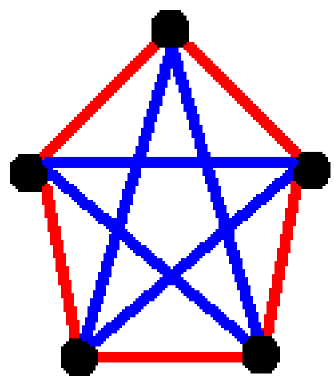

Notice the elegant symmetry of the figure: all of the red lines are on the outside of the pentagon, and all of the blue lines are in the interior of the shape. The reason this works is that any triangle using three dots as vertices must use either two outside lines and one inside line, or two inside lines and one outside line. (Think about that: why couldn't we use three inside lines or three outside lines to make a triangle?) This *guarantees* that any triangle we draw will use two differently-colored lines, so this figure does *not* have the homogeneity property! Of course, we could look at all possible triangles inside the diagram and make sure none of them use one color. How many such triangles are there? How quickly could you check all of them by hand? Is it easier to do that, or to notice the inside/outside property that we mentioned above?

Perhaps you found a solution that doesn't look like the drawing we have. How can you tell whether it's actually a different figure? How many blue and red lines are there in your diagram? In ours? Try redrawing your figure by moving the dots around but maintaining the relationships between the dots (i.e. the color of line drawn between any two of them). Can you make your figure look like ours? What do you think this says about the number of solutions to this puzzle?

Okay, now we're ready to think about what happens when we have six people. In terms of the dots and lines, we're looking to draw all possible connections between six dots with either blue or red lines and ensure that there are no triangles with the same line type. Before you start drawing, try to think about the solutions to this puzzle when we were working with four and five dots. What did those solutions look like? What was the number of lines that we had to fill in? How many will we have to draw this time? Can we try to make this figure look like the solution for five dots? It sometimes helps to think about how a solution to a current puzzle might be similar to previous work. Now, try to draw this figure and see what happens.

Did it work? Why not? Where did you run into trouble? How many lines could you draw before you were *forced* to make a monochromatic triangle? That is, how many lines could you fit into the figure before the next one you drew would make a monochromatic triangle, no matter whether it was blue or red? These are just tangential questions, in a way, to solving this particular puzzle, but they're worth thinking about because they're interesting in their own right and they may guide us toward a solution for this puzzle or a generalization thereof.

The situation we face now is interesting because it's of the opposite nature to the kind of situations we faced before. With four and five dots, we wanted to show that it was *possible* to arrange all of the lines to make no monochromatic triangles. To show that, we just had to *do* it! Exhibiting a *particular* figure with the desired property was sufficient to show that it was possible to achieve the property we wanted. With six dots, though, it seems like it is *not* possible to arrange the lines so that there are no monochromatic triangles. How can we prove that this is a fact? It is tempting to say that we should just look at *all possible* arrangements of the lines and argue that there is *at least* one monochromatic triangle in every single one of them. Is this feasible? How many arrangements of the lines are there? How could we easily show that there *is* a monochromatic triangle in any given diagram? Remember how we did this with the figure with five dots? We noticed that any triangle would have to use at least one line from the outside and at least one line from the inside, which guarantees right away that any triangle has two types of lines. Could we do the same thing here, and identify some property that *guarantees* a triangle?

The issue is that there are *too many* possible arrangments of the lines in the diagram with six dots for us to check all of them by hand! There are 15 lines to be drawn, each of which could be either blue or red, so it seems like there are 2^{15} possible arrangements. This is a big number! (In actuality, there are slightly fewer possibilities because some of them are equivalent in some sense; more technically, they are called "*isomorphic*". We will discuss these concepts in more detail later on.)

We need to be more clever with our argument so that we can prove a property of *any* diagram without drawing a particular one. That is, we need to find some fact, some property that is true of all of the possible diagrams with six dots, that will still allow us to deduce that there *must* be a triangle. One way to approach this is to think about drawing the lines in one small section of the diagram. Specifically, let's take any of the six dots and consider the five lines coming out from that dot. How many are blue and how many are red? This is partially a trick question; because we're not considering any *particular* diagram but trying to find a fact about *all possible* diagrams, we can't answer that question too specifically. Here's what we *can* say, though: there must be at least three blue lines *or* at least three red lines. Do you see why this is true? The only way this *wouldn't* be true is if there were two (or fewer) blue lines *and* two (or fewer) red lines leaving this particular dot, for a total of four (or fewer) lines, but we know that all possible connections must be drawn, so there should be five! (This argument is an example of a concept known as the **Pigeonhole Principle**. The idea is that we can't place five objects, of two different colors, into two different boxes without putting three objects of one color into one box. This is an incredibly useful strategy with these types of problems, and we will examine the principle in much greater detail later on in Chapter 6.)

So where do we stand? We started with *any* of the figures with six dots and all lines filled in, and focused on one particular dot; coming out from this dot, there must be three blue lines or three red lines. It could be either color, so we can't just assume it's blue and follow an argument that way; we can do that, but we have to come back to this point afterwards and see what would change if the three lines were red. So let's do that: let's examine all of the figures so that there are three red lines exiting this particular dot. Where can we go from there? We haven't yet assumed anything about the other lines in the figure, so let's look at what those could be. Examine the picture below to see what line colors we have assumed exist so far:

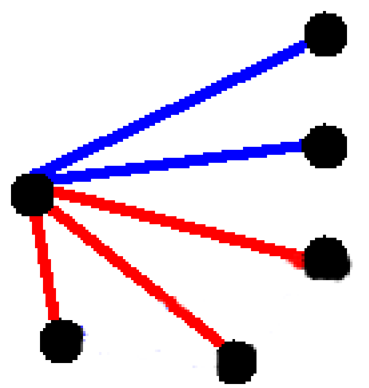

Now, what lines could be added to this diagram that would avoid making a triangle of one color amongst three dots? We can't *necessarily* make any assumptions about the lines coming out from the two dots that are isolated in the picture, so let's focus on the three dots on the bottom. What color could the lines among those be? Well, if any of them are red, that would form a monochromatic triangle between the two endpoints of that line and the original dot we focused on! That would be a problem. Okay, the only way to avoid that is to make all of those lines blue. But that would make a blue triangle among those three dots! Wow, it looks like we *cannot avoid* making a monochromatic triangle no matter what we do!

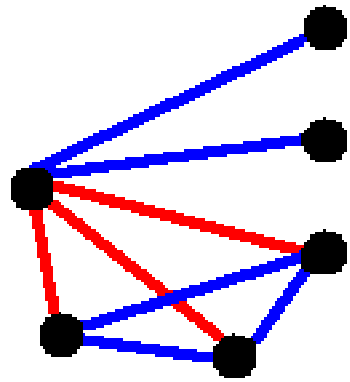

Let's go back to our Pigeonhole argument and reassess the situation. What if the three lines of the same type that were guaranteed by the argument were blue instead of red? Well, nothing would really change, right? We would still be stuck, in terms of adding new lines among the three dots on the bottom of the figure:

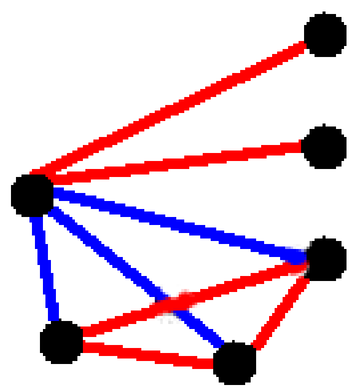

If we include any blue lines, that forms a monochromatic triangle with the original dot, and if we make all of them red, that forms a monchromatic triangle there! In this sense, the two arguments we followed after the Pigeonhole argument were *identical*; if we were to replace every instance of the word "blue" with "red" and vice versa, we would have the same argument. Sometimes, mathematicians will use this situation to shorten a proof and just say that "without loss of generality, the three lines are red." This is usually taken to mean that if we were to make the other choice instead (i.e. if the lines were blue) then the further argument would have an identical structure, mathematically, so we can save space and time by not writing the same words over again. This is so common, in fact, that you might sometimes see this phrase, "without loss of generality", abbreviated as WOLOG or WLOG.

What have we accomplished so far? We produced *specific* diagrams that showed we could arrange the lines among four and five dots so that we can avoid a monochromatic triangle, and we argued that *any* diagram of lines with six dots *must* have a monochromatic triangle. In terms of the friends/enemies formulation of the puzzle, this means that any group of six people *must* have a subgroup of three people that are all friends or all enemies. Notice how helpful it was to recast the puzzle into this dots/lines formulation; it allowed us to completely forget about the social context of the problem (which can be distracting, in a way) and let us simplify our terminology and notation (we went from labeling pairs of people as "friends" or "enemies" to simply drawing one line between two dots). This is a very useful strategy: extracting the inherent *structure* of a puzzle—the underworkings, the relationships between the elements, how they interact, etc.—and rewriting everything in terms of just those parts. This can make the puzzle easier to understand and tackle, and it can guide us into devising better notation. What if we had continued to solve this puzzle with the $13F, 23E, \dots$ notation? That may have eventually worked, but it would have been much harder!

One of this puzzle's original questions was to identify a *cutoff* number so that any larger group of people would necessarily have this subgroup property. Do you think we have accomplished that? Have we identified a cutoff? Could six be that number? Why or why not? In any group of seven people, there's a smaller group of six people, and then our work above proves that there must be three mutual friends/enemies among that group! Certainly, this works for any group of people of size larger than six, so this must be that precise cutoff point we were looking for. This is an analogous result to the one mentioned in the original statement of the puzzle, where the Hungarian sociologist noticed this phenomenon for subgroups of size four. That problem is much more difficult to solve, so we handled a smaller, simpler case here. Both of these results are related to a larger class of problems known as **Ramsey Theory**. This branch of combinatorics and graph theory works with identifying these kind of "cutoff points" where, as the size of some structure (like a group of people) grows and grows, there is eventually a point where we can be *guaranteed* to find a subgroup with a certain property (three mutual friends/enemies). What was at first thought to be a sociological phenomenon turned out to be a rigorous mathematical fact. How about that!

Let's pose some interesting and related questions before moving on. What if we had been looking for homogeneous groups of a different size, like four or five or twelve? Certainly, we would have to have a larger group of people, overall, to *guarantee* finding such a subgroup. Can we always do this? That is, given any desired subgroup size, can we identify a cutoff point in the way that we did here? Can you figure out how to prove that such a cutoff point *must* exist, even without finding the particular number? Furthermore, what if we allowed for a third possibility: friend, enemy, unacquainted. Could we answer similar questions about homogeneous groups? These are all questions related to Ramsey Theory and some of them are quite difficult to answer and took mathematicians many years to address (and many are still open problems) so don't be discouraged if you don't make any progress on these questions. We believe that even attempting to answer them and thinking about the issues therein is meaningful and beneficial enough, in itself.

Lesson from this puzzle: This puzzle brought up several difficulties. First, we had to find a way to interpret the puzzle in a meaningful way so that we could even address the questions it asked, and this involved coming up with appropriate *notation* to represent the elements of the puzzle. This is an important part of mathematical problem-solving, particularly for a puzzle like this that doesn't incorporate the notation and visualization as part of the problem statement. Second, to identify six as the group cutoff size, we had to somehow prove that something *is not* possible, but the number of possible configurations to examine was way too large to examine each one individually. This happens frequently, particularly in problems related to computer science and algorithms. To address this, we had to employ a strategy far more clever than mere brute force, and it's not always clear what strategy that should be. Here, we essentially started to try to fill in the lines as if it was going to work out, and then realized that we reached a point that was impossible to fix. Proving that something is possible can amount to just showing an example of that phenomenon (which we did with the groups of size four and five), but proving something is *impossible* can be much trickier and require some context-dependent ingenuity. Lastly, we saw that it can be interesting to think of questions closely-related to the puzzle at hand that simply tweak one or more of the conditions of the problem. What if we look for larger subgroups? What if we allow for more types of lines? How does this change the results? Exploring the boundaries of puzzles by changing the conditions like this can lead to new mathematical discoveries and techniques, and keeps mathematicians actively searching for new knowledge and ways of sharing that knowledge.