SET AND FINITE AFFINE GEOMETRY

JULIA CARNEY

ABSTRACT. In this paper, we explore some basic structures in finite affine geometry through the lens of the popular card game SET. We begin with a brief introduction to the game itself. Next, we present the axioms of finite affine plane geometry. We prove that the SET deck is consistent with these axioms, and provide numerous concrete examples of how the deck can be used to visualize this geometric structure. We then use the SET deck as an introduction to finite affine geometry in higher dimensions. Finally, we introduce the notion of a cap and end with a brief discussion of maximal cap size in AG(n,3), which is currently an open problem for $n > 6$.

Contents

1. Introduction to the Game SET

In 1974, in the course of her work studying epilepsy in German shepherds, geneticist Marsha Jean Falco invented a system of images that she used to record patterns in her findings. This system would eventually lead Falco to invent the now-popular card game SET, a pattern-recognition game that is played with a special deck of cards featuring images similar to those of her record-keeping system. In this section we will introduce the structure and features of the SET deck and explain how the game itself is played.

As previously mentioned, the game SET is played using a special deck of cards. Each card displays an image that contains four different attributes, which we refer to as *number*, *shading*, *color*, and *shape*. For each of these four attributes, there are three possible values a card can take: the *number* of figures on the card can be **one, two,** or **three**, the *shading* of the figures can be **open, striped,** or **solid**, the *color* of the figures can be **red, green,** or **purple**, and the *shape* of the figures can be **diamonds, ovals,** or **squiggles.** A table of the four attributes and the three possible values for each is below for reference.

Date: September 1, 2021.

Attribute	Number	Shading	Color	Shape
Values	One Two Three	Open Striped Solid	Red Green Purple	Diamond Oval Squiggle

TABLE 1. Attributes and Values

Every possible combination of values appears in the deck exactly once; thus, there are $3^4 = 81$ cards altogether. Further, because each combination appears only once, we can refer to each card by its particular combination of attributes. For example, Figure 1 displays three different SET cards, which we can give names to: the first card is **one open green diamond**, the second is **two striped red ovals**, and the third is **three solid purple squiggles**.[1].

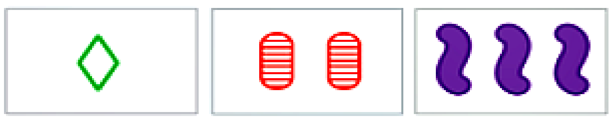

FIGURE 1. Three SET cards.

Gameplay is relatively simple, and any number of people can play. To begin the game, twelve cards are placed on the table, and players all simultaneously attempt to find SETs: collections of three cards that meet certain requirements, which we will define in a moment. When a player finds a SET, she calls out "SET!" and removes the cards from the table. The three cards are then replaced with three more cards from the deck, and gameplay proceeds as before. If at any point there are no SETs on the table (a scenario which we will return to in the final section of this paper), new cards are added to the table, three at a time, until a SET can be found. This time, when a SET has been found, new cards are not placed on the table unless there are still no SETs; cards are only replaced in order to bring the total number of cards on the table back up to twelve.

Play continues until the deck has been exhausted and there are no more SETs remaining on the table. The player who has collected the most SETs when the game ends is the winner.

So what exactly is a SET?

Definition 1.1. A SET is a collection of three cards such that for each of the four attributes, all three cards either share the same value or each have a different value.

For example, the three cards in Figure 1 form a SET: the cards contain all different numbers, all different shadings, all different colors, and all different shapes.

Figure 2 is another, perhaps more visually intuitive, example of a SET: the three cards contain all the same shading, all the same color, and all the same shape, but all different numbers.

Figure 3, however, is not a SET. Notice that the figures on the cards are all the same number, all the same shape, and all different shadings; these attributes all fit the requirements for a set. But color is the attribute where our SET check

[1]All images of SET cards courtesy of `https://geekandsundry.com/the-card-game-that-puzzled-mathematicians-for-decades`

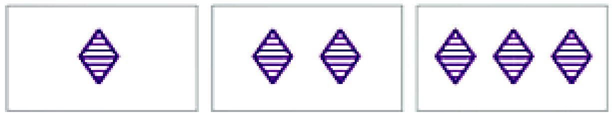

Figure 2. Another SET.

fails. Two cards are green and one is red, so the color is neither all the same nor all different; thus, these three cards are not a SET.

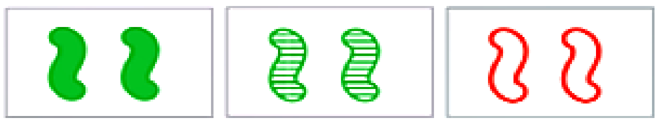

Figure 3. Not a SET!

New players who have not yet developed an intuitive eye for SETs will often find that it is much easier to identify when three cards do *not* form a SET than it is to identify when they do. Checking attribute by attribute to confirm that three cards are, in fact, a SET can be quite tedious at first; however, it is often relatively easy to spot when, for any particular attribute, two cards share a value and the third card does not. And when this is the case, as in Figure 3, the three cards cannot be a SET.

In addition to being a thoroughly enjoyable card game, SET is also an excellent framework for a variety of topics in higher-level mathematics. In the remainder of this paper, we will explore connections between the SET deck and a particular geometric structure called a finite affine plane.

2. Finite Affine Geometry

In order to understand the mathematical structure of the SET deck, we must first introduce the axioms of finite affine geometry. In this section we will define a structure called a finite affine plane and prove a handful of results that allow us to quantitatively describe these planes. In the followings section, we will return to the SET game and explain how the two are related.

Affine geometry is a generalization of Euclidean geometry that focuses only on the relationship between points and lines, ignoring other notions such as distances and angles. We begin with a formal definition.

Definition 2.1. An *affine plane* is a system of points and lines satisfying the following three axioms.

Axiom 1: There exist at least three non-collinear points.

Axiom 2: Two points uniquely determine a line.

Axiom 3: Given any line l and point p not on l, there exists exactly one line containing p that does not intersect l. We say that this line is *parallel* to l. We also say that a line is parallel to itself.

A finite affine plane, then, is an affine plane containing a finite number of points and lines. Note that we can see from these axioms that standard two-dimensional Euclidean space is an example of an affine plane, but not a finite one.

Given this definition, we would like to know if there is any way to quantitatively describe such planes. For example, can a finite affine plane contain any number of points and lines, or do finite affine planes only exist in certain sizes? What is the relationship between the number of points and the number of lines in a plane? How might we envision these structures geometrically?

We begin to answer these questions by proving a result that allows us to classify finite affine planes in a convenient way.

Theorem 2.2. *Every line in a particular finite affine plane contains the same number of points.*

Proof. Let k and l be two distinct lines in a finite affine plane. We know from our axioms in Definition 2.1 that k and l must either be parallel, or else intersect at some point p. (Observe that two distinct lines can intersect at at most one point; otherwise they would be the same line by Axiom 2.) Let us first consider the case where k and l intersect at point p.

Axiom 2 gives us that every line must contain at least two points. Thus, since k and l intersect at point p, there exists at least one other point q on k that does not intersect l. By Axiom 3, there exists a single line j through q that does not intersect l. j must also contain at least one other point s. Further, s cannot lie on k; otherwise s and q would both lie on k, and then k and j would be identical by Axiom 2 and thus j would intersect l.

Thus, we have that there exists a point s in our plane that does not lie on either k or l. Now let h be any line through s.

It must be the case that h intersects at least one of k and l; we know that j is the only line through s that is parallel to l, and j intersects k, so there does not exist a line through s that is parallel to both k and l. Then if h intersects only one of the two lines, it is by definition parallel to the other. Thus, h must either be parallel to one of k or l but not the other, or intersect both k and l.

Now, by Axiom 3, we have that there exists exactly one line through s that is parallel to l, and we know that this line intersects k at one point. Similarly, there exists exactly one line through s that is parallel to k, and this intersects l at one point. Further, we have just proved that every other line through s intersects both k and l; thus, every one of these lines contains exactly one point of k and one point of l. These two facts establish a one-to-one correspondence between points of k and points of l. Therefore, the two lines must have the same number of points, as desired.

Now let us consider the case where k and l are parallel. Take a point p on k and a point q on l, and consider the line j containing p and q. Then j intersects both k and l. Thus, by the previous part of the proof, we can establish a one-to-one correspondence between points of j and points of k, and between points of j and points of l. Thus, k and l must also have the same number of points, and we are done. □

Hence, since all lines in a finite affine plane must contain the same number of points, we can classify finite affine planes according to the number of points on each line.

Definition 2.3. A *finite affine plane of order n* is an affine plane where every line contains exactly n points.

We write this as AG$(2, n)$. AG stands for affine geometry, the 2 designates a plane, i.e. a two-dimensional space, and the n designates the order of the plane. (Later on in this paper, we will also explore affine geometry in dimensions higher than two.)

Next, we would like to know if we can quantify the number of points and lines in AG$(2, n)$. The following theorem and its corollaries will allow us to do exactly that. Before introducing the theorem, however, we need a lemma, which will allow us to say that lines parallel to the same line are also parallel to each other.

Lemma 2.4. *Parallelism in a finite affine plane is an equivalence relation.*

Proof. First, note that by the definition given in Axiom 3, parallelism is already reflexive and symmetric. Thus, we just need to show that it is transitive.

Let j, k, and l be three distinct lines in AG$(2, n)$. Let j and k both be parallel to l. Suppose for the sake of contradiction that j and k are not parallel to each other. Then they must intersect at some point p. Note that p cannot lie on l, since p lies on both j and k, and neither of these lines intersects l. But then we have two distinct lines through a point not on l that are both parallel to l, which contradicts Axiom 3. Therefore j and k must be parallel to each other. Thus, parallelism is transitive and is therefore an equivalence relation as desired. □

We are now ready to prove the theorem that will allow us to quantitatively describe the finite affine plane of order n.

Theorem 2.5. *There exist exactly n^2 points in $AG(2, n)$.*

Proof. By Axiom 1, we have that there exist at least three non-collinear points; call two of these points p_1 and p_2, and call the third q_1. Since we are in AG$(2, n)$, the unique line p containing p_1 and p_2 exists and contains n distinct points. Call these points $p_1, \ldots, p_n$. Since our initial three points are non-collinear, q_1 does not lie on p, so by Axiom 3 there exists a unique line q containing q_1 that is parallel to p. q also contains n distinct points; call these points $q_1, \ldots, q_n$. Further, since q is parallel to p, all the q_ks and p_ks, where $1 \leq k \leq n$, are distinct from each other.

Now, by Axiom 2, there exists a unique line l containing p_1 and q_1. Further, this line cannot contain any of the other p_ks or q_ks; otherwise l would be intersecting a distinct line at multiple points, which we know is not possible (see proof of Theorem 2.2.) Thus, since l contains n distinct points, we have the existence of $n - 2$ new points distinct from the p_ks and q_ks. Call these points $l_3, l_4, \ldots, l_n$.

By Axiom 3, there exists a unique line s_{l_3} containing l_3 that is parallel to p. By Lemma 2.4, since p is parallel to q, this new line is parallel to q as well. Further, by Axiom 2, s_{l_3} cannot contain any points of l besides l_3. Thus, including l_3, s_{l_3} contains n points that are distinct from the p_ks and q_ks. Similarly, there exists a unique line s_{l_4} containing l_4 that is parallel to p, q, and s_{l_3}. Thus, this line contains exactly n points, including l_4, that are distinct from any points on p, q, or s_{l_3}.

We can repeat this process for all the remaining points of l. When this is done, the total number of points we have found, including $l_3, l_4, \ldots, l_n$, will be $n(n - 2)$. Adding this to the n points of p and the n points of q, we have $n(n - 2) + 2n = n^2$ points in total. Now we just need to verify that there are no other points in the plane, and we will have achieved the desired result.

Suppose for the sake of contradiction that there exists a point c in AG$(2, n)$ that is distinct from the n^2 points we have already found. By Axiom 2 there exists a unique line j containing c and p_1. This means that j is not parallel to p, since it intersects p at p_1. Then by our lemma, j cannot be parallel to any of $q, s_{l_3}, s_{l_4}, \dots, s_{l_n}$; otherwise j would also be parallel to p. Therefore j must intersect each of these lines. Thus, j contains at least $n+1$ points: c itself, p_1, and one point from each of $q, s_{l_3}, s_{l_4}, \dots, s_{l_n}$. But this is a contradiction; a line in $\mathrm{AG}(2,n)$ must contain exactly n points. This means that c cannot be an element of $\mathrm{AG}(2,n)$. Therefore $\mathrm{AG}(2,n)$ contains exactly n^2 points, as desired. □

Once we have determined the number of points in the plane, the number of lines easily follows.

Corollary 2.6. *Each point in $AG(2,n)$ is contained in exactly $n+1$ lines, and there exist exactly n^2+n lines in total.*

Proof. Let p be a point in $\mathrm{AG}(2,n)$. By Theorem 2.5, there are n^2-1 other points in the plane. By Axiom 2, each of these points lies on a unique line through p; in addition, each of these lines through p contains $n-1$ other points. Thus, there are $(n^2-1)(n-1) = n+1$ lines through p. Then, since there are n^2 total points, $n+1$ lines through every point, and n points on every line, we have $n^2(n+1)/n = n(n+1) = n^2+n$ total lines in the plane. □

Another useful quantitative result that follows immediately from Theorem 2.4 is the number of parallel classes in $\mathrm{AG}(2,n)$.

Corollary 2.7. *In $AG(2,n)$, there exist exactly $n+1$ parallel classes, each containing n lines.*

Proof. By Theorem 2.5, we have that each point in $\mathrm{AG}(2,n)$ is contained in exactly $n+1$ lines. By Axiom 3, each parallel class contains exactly one line through any given point. Thus, there must also be exactly $n+1$ parallel classes. Then, since there are $n^2+n = n(n+1)$ total lines, there must be $n(n+1)/(n+1) = n$ lines in each class. □

These results provide us with a reasonably thorough description of finite affine planes; our next step will be to visualize these planes, which is where the SET deck comes back into play. Before moving on, though, one final question that is worth mentioning is whether it is possible to construct an affine plane of any order, or whether $\mathrm{AG}(2,n)$ only exists for particular values of n. This is currently an open question. It has been proven that an affine plane of order p exists for all primes p, and all known finite affine planes have orders that are prime powers. Further, finite affine planes have been proven *not* to exist for many different orders. We will not delve into the specifics of these proofs here, but it is worth noting that this remains an interesting topic for potential future research.

3. SET and the Axioms of Finite Affine Geometry

We now return to the SET deck. As the reader may have guessed by now, it is possible to build a finite affine plane out of SET cards! In this section, we will prove that certain subsets of the SET deck are isomorphic to $\mathrm{AG}(2,3)$, and we will work through examples of a few ways in which one can go about constructing such a subset. We will then use the SET deck to explore $\mathrm{AG}(3,3)$ and $\mathrm{AG}(4,3)$.

Recall that Definition 2.1 defines affine planes using three axioms about points and lines. We can translate these axioms into the language of SET by thinking of cards as points and SETs as lines. (Note that since a SET contains three points, we will be working with finite geometries of order three.) Thus, our goal is to show that SET satisfies the three axioms. That is, we want to show the following:

- There exist at least three SET cards that do not form a SET;
- Two SET cards uniquely determine a SET;
- Given a SET and a card not in the SET, there exists a unique SET that contains the card and is parallel to the original SET.

We have already seen that the first axiom is true of the deck as a whole; we will see that it is also true within the particular subsets of the deck considered in this section. We can also prove that the second axiom holds for any two cards in the deck, which of course includes any two cards in any subset of the deck.

Theorem 3.1. *Two SET cards determine a unique SET.*

Proof. Let A and B be any two SET cards; we will identify the unique card C that completes a SET with A and B. Consider the values for each attribute of A and B. For each of the four attributes, A and B must either share the same value or have different values. If A and B share the same value for a particular attribute, then C must also have that value for that attribute. If A and B have different values for a particular attribute, then C must have the third possible value for that attribute. Thus, we can determine the one necessary value for each attribute of C in order for C to form a SET with A and B. This gives us the unique identity of C. $\square$

For example, let us consider the two cards in Figure 4, which we will refer to as A and B.

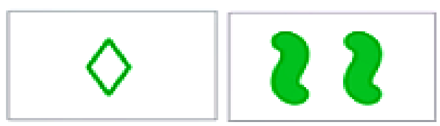

FIGURE 4. Card A and Card B.

Looking at each attribute, we can use the definition of a SET to determine the unique card C that completes the SET.

Number: A contains one figure and Card B contains two. These two numbers are different from each other; therefore C must take the third possible value and contain three figures.

Shading: A is open and B is solid. Therefore C must be striped.

Color: A and B are both green. Therefore C must also be green.

Shape: A contains diamonds and B contains squiggles. Therefore C must contain ovals.

We have thus arrived at the identity of C: **three striped green ovals.**

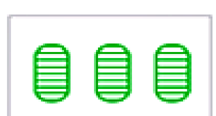

FIGURE 5. Card C.

Now all that remains to be shown is the third axiom. Here we must pause and consider what it means to be parallel in the context of SET. We begin with a definition.

Definition 3.2. Two SETs are *parallel* if every attribute that takes all the same value in one SET also takes all the same value in the other, and every attribute that takes all different values in one SET also takes all different values in the other. Further, there must exist some ordering of the three cards such that the values for those attributes that are different cycle in the same order in both SETs.

A few clarifying remarks on this definition:

First, note that when the value of an attribute is all the same in one SET, we do *not* require the second SET to take the same particular value for that attribute as the first SET. We only require that the value of that attribute be all the same *within the second SET*. For example, if we have a SET where all the cards are purple, then the parallel SET must also be all the same color, but that color need not necessarily be purple.

Second, when we speak about the order in which values for an attribute cycle within a SET, we are referring to a property of the particular physical ordering of the cards for that SET, rather than a property that is uniquely defined for the SET itself. For example, consider the SET in Figure 6. In this SET, shape is one of the attributes that is different. We have arranged the SET in Figure 6 so that the cycle of values is **diamond** $\rightarrow$ **squiggle** $\rightarrow$ **oval** $\rightarrow$ **diamond** etc.

FIGURE 6. One particular ordering of a SET.

However, we could just as well have placed the cards in a different order such that the values cycle differently, and they would still be the same SET. For example, in Figure 7, the cycle of values for the shape attribute is **squiggle** $\rightarrow$ **diamond** $\rightarrow$ **oval** $\rightarrow$ **squiggle** etc. But this is clearly the same SET as in Figure 6.

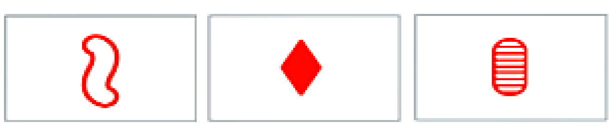

FIGURE 7. The same SET, in a different order.

How do we deal with this? The fact that cyclic ordering of values is a property of a particular physical layout of the cards has the potential to become problematic for a definition of parallelism, as we would like parallelism to be a property of the SETs themselves, rather than the particular order in which the cards are placed. Thus, we avoid this issue in our definition by only requiring that there exist *some* ordering of the cards for which the attributes cycle in the same way.

We are now ready to prove that SET satisfies the third axiom of finite geometry.

Theorem 3.3. *Given a SET and a card not in the SET, there is a unique SET containing that card that is parallel to the first SET.*

Proof. Begin by fixing an ordering for the first SET. Call the cards in the SET A, B, and C, respectively, and call the card not in the SET D. We will now identify the unique cards E and F that form a SET containing D that is parallel to the SET containing A, B, and C.

First, consider those attributes, if there are any, for which the values are all the same in the first SET. By our definition of parallelism, the value for these attributes must also be all the same within the second SET. Thus, whatever value D takes for these attributes, E and F must also take.

Now consider those attributes for which the values are all different in the first SET. By our definition of parallelism, the values for these attributes must cycle in the same order within the second SET as they do in the first. Thus, we can determine the values that E and F must take for those attributes as follows. First, note what value D takes for that particular attribute. Since A, B, and C all take different values for that attribute, D must share a value with one of them. If D takes the same value as A, then let E take the same value as B and F take the same value as C. Similarly, if D takes the same value as B, let E take the same value as C and let F take the same value as A. Finally, if D takes the same value as C, let E take the same value as A and let F take the same value as B. In each of these cases, then, the values for the attribute in question will cycle in the same order in both SETs, as desired.

Thus, we are able to determine the value of each attribute of E and F. Further, we can verify that this process will result in the same two cards no matter what initial ordering we choose for A, B, and C; the work is omitted here as it is identical to the proof above. Therefore E and F are uniquely determined by A, B, C, and D, as desired. □

Once again, let's walk through an example of this process. Consider the following SET and card not in the SET, which we will refer to as A, B, C, and D, respectively.

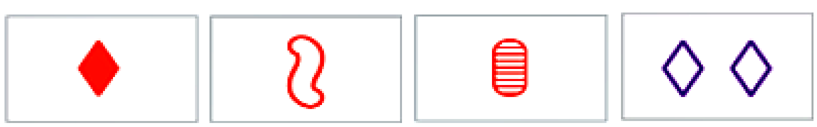

FIGURE 8. Cards A, B, C, and D. Note that A, B, and C form a SET.

Looking at each attribute and using our definition of parallelism, we can now determine the identities of two cards E and F, which will complete the parallel set containing D.

Number: Our first SET shares the same value for number, so our second SET must also share the same value for number. D has two figures on it; therefore E and F must also have two figures.

Shading: Our first SET takes all different values for shading. D is open, like B. Thus, E must be striped like C, and F must be solid like A.

Color: Our first SET shares the same value for color, so our second SET must do the same. D is purple; therefore E and F must also be purple.

Shape: Our first SET takes all different values for shape. D contains diamonds, like A. Thus, E must contain squiggles like B, and F must contain ovals like C.

Thus, we have arrived at the identities of E and F: **two striped purple squiggles** and **two solid purple ovals**, respectively. An interested reader can verify that we will indeed arrive at the same two cards regardless of what order we initially choose for A, B, and C.

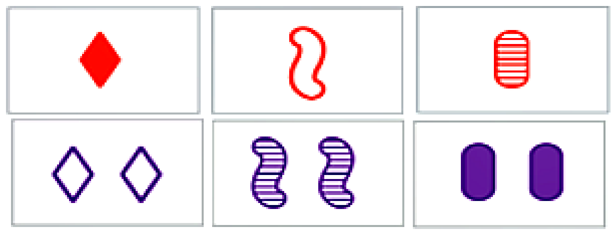

FIGURE 9. Our two parallel SETs.

4. Higher Dimensions

In the previous section, we successfully showed that the SET cards are consistent with the axioms of finite affine geometry. However, we know from Section 2 that a finite affine plane of order three contains exactly nine points and twelve lines; the full SET deck, with its 81 cards, is much larger than that. Clearly, then, the full SET deck is not just a finite affine plane. Two questions then arise: first, how do we find a subset of the deck that is isomorphic to AG $(2, 3)$? And second, if the whole deck is not a plane, what is it? In this section, we address these two questions. We begin by demonstrating how to construct a plane from three non-collinear SET cards. We then use the rest of the deck to introduce finite affine geometry in higher dimensions.

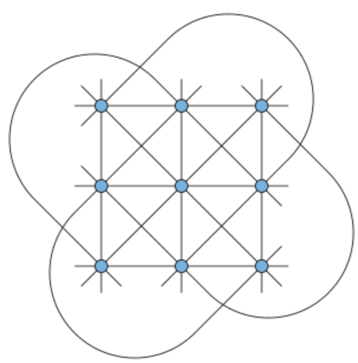

FIGURE 10. A template for AG$(2, 3)$.

First, for ease of visualization, we introduce a template for AG$(2, 3)$ in Figure 10; we will use this to place SET cards such that we can easily identify lines and parallel classes. Here the blue dots represent points, and any three points that are connected by a single straight or curved line represent a line. We can see that this diagram contains nine points and twelve lines, as desired. Note the ways in which this visual differs from standard Euclidean geometry; three points that make up a line do not need to lie along the straight lines we are accustomed to!

Now that we have a convenient way of visualizing the finite affine plane, we can create an array of nine SET cards such that all the SETs in the plane lie along the lines in the diagram. There are a few ways to do this. One easy (and aesthetically appealing!) method is based on the way in which the SET deck is constructed. We know that the deck consists of 81 cards, representing all possible combinations of three values across four attributes. Thus, we can reduce down to a nine-card structure that preserves the same mathematical properties as the whole deck by simply fixing values for any two of the attributes and selecting the nine cards that take the chosen values for both of those attributes. From there, it is relatively uncomplicated to place the cards so that all the SETs align with the lines in our template; just arrange the cards such that each horizontal line corresponds with a particular value of one of the attributes that takes different values, and each vertical line corresponds with a particular value of the other attribute that takes different values.

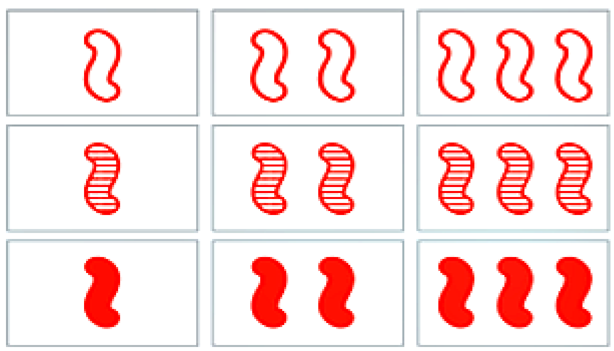

Figure 11. A plane with two fixed attributes, color and shape.

Figure 11 is one example of this. Here, color and shape are fixed; all the cards are red squiggles. Thus, number and shading are the variable attributes. Each horizontal line contains a particular value for shading (open, striped, and solid, respectively), and each vertical line contains a particular value for number (one, two, and three, respectively.) Readers can verify that there are in fact twelve SETs present, each one corresponding to one of the lines in our diagram.

This approach to constructing a plane has the advantage of being quite visually appealing; since two attributes are fixed, it is much easier to check that all the lines are SETs. This plane also allows us to relatively easily identify the four parallel classes of three lines. One class consists of the three horizontal SETs, and one consists of the three vertical SETs. The other two each consist of one diagonal SET through the center card and two of the SETs along curved lines in the diagram. Interested readers can verify that within each class, the three SETs do indeed meet the definition of parallelism.

However, there is another, perhaps more interesting, way to construct a plane out of SET cards. Recall that the finite plane axioms only give us the existence of three non-collinear points; and all the work we did in Section 2 constructing and quantifying the plane was done using only those three axioms. Thus, we know that three non-collinear SET cards must uniquely determine a plane. Let us now demonstrate how to determine what that plane looks like.

We begin with three non-collinear points, i.e. three cards that do not form a SET.

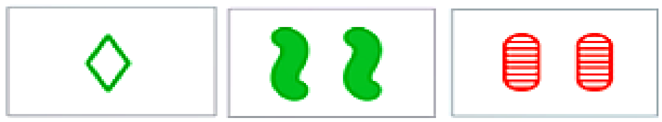

Figure 12. Three cards that do not form a SET.

We know from our diagram where the lines, or SETs, in the plane ought to be. Thus, since these three cards are not collinear, we can begin to form the plane by placing the cards in any three positions that do not all lie along the same line. Interested readers can verify that any non-collinear configuration of the initial three cards will eventually result in the same nine cards and twelve lines, although the physical layout will be different. We will choose the configuration shown in Figure 13.

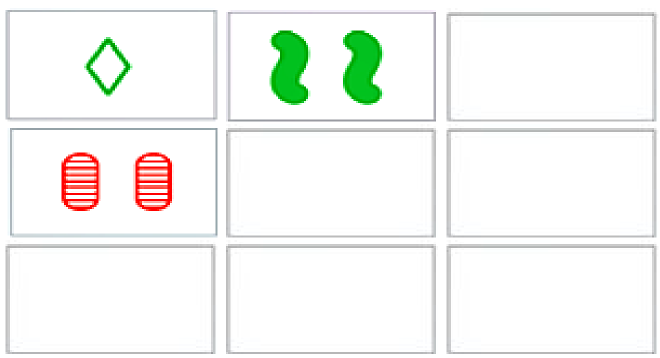

Figure 13. Beginning the plane.

Once again, we know where we want our SETs to be, and further, we know that two cards uniquely determine a SET. Thus, we can begin to fill in the plane by completing SETs where we already have two cards placed along the same line. Referring back to our diagram we see that three of the lines currently have two out of the three positions filled; thus, we can determine the cards needed to fill the third position on each of these lines. The first horizontal line contains **one open green diamond** and **two solid green squiggles,** so we can complete it with **three striped green ovals.** The first vertical line contains **one open green diamond** and **two striped red ovals**, so we can complete it with **three solid purple squiggles.** And finally, one of the four lines that loops around the plane contains **two striped red ovals** and **two solid green squiggles,** so we can complete it with **two open purple diamonds.** Figure 14 shows us the plane after filling in each of these positions.

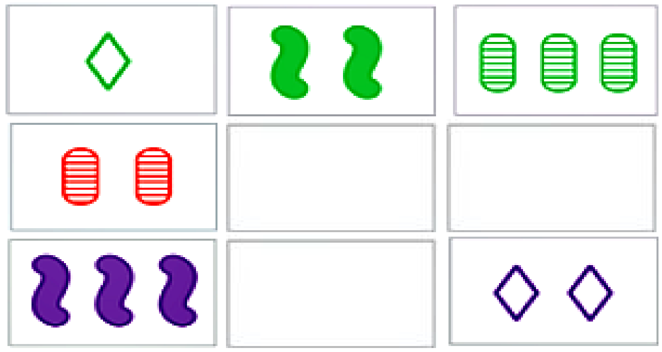

FIGURE 14. Filling in the plane.

We continue this process of filling in SETs until the whole plane is complete, as shown in Figure 15.

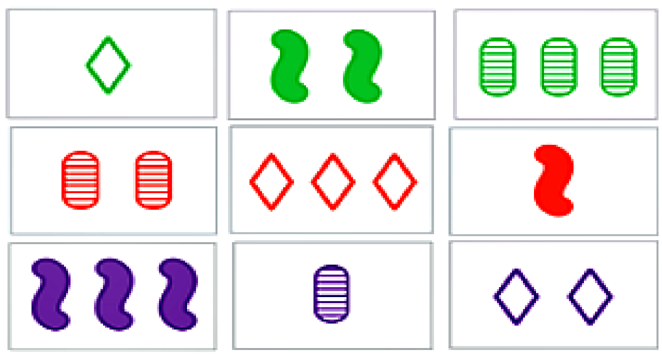

FIGURE 15. The completed plane, AG$(2, 3)$.

And we are done! We can refer back to our diagram and verify that there are indeed twelve SETs here, each lying along one of the twelve lines. This plane is not as visually intuitive as our first example, but it demonstrates that three non-collinear points do in fact completely determine a unique SET plane.

We now return to the deck as a whole. As our first plane example demonstrated, there seems to be a correspondence of some sort between attributes in SET and dimensions in geometry. The natural question, then, is if we can use more of the SET deck to expand our finite geometry into more than two dimensions. For example, previously we constructed a plane by fixing values for two attributes and allowing only two attributes to vary. What if we fix values for only one attribute and allow the other three to vary? Let us examine one example of this and see what patterns we can find.

Figure 16 below has one attribute fixed, color; it contains all the SET cards that are red. We can see our plane from Figure 11 contained within this larger array, on the left side. We can also identify two very similar planes. The nine cards on the right side of this array form a plane that is nearly identical to Figure 11, except that it consists of cards with diamonds instead of cards with squiggles. Similarly, the nine cards in the middle form a plane that is nearly identical to the other two, but consisting of all the cards with ovals.

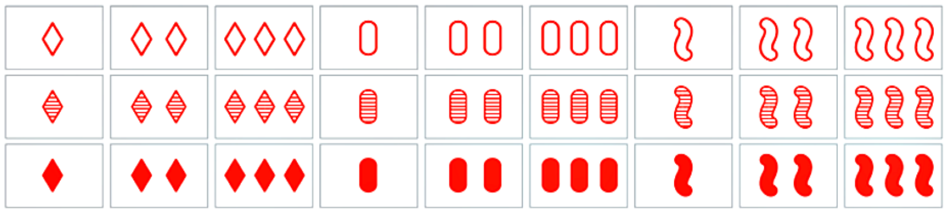

FIGURE 16. A hyperplane, AG$(3,3)$.

Further inspection reveals that these are not the only planes and SETs in this configuration of cards. Note that the three planes we just identified were identical in every value except shape; we can find similar groups of three planes for the other variable attributes. For example, consider the three rows of our array. We can see that the top row of cards forms a plane consisting of all the cards whose shading is open. The middle row forms a nearly identical plane, except that here, the shading is striped; similarly, the bottom row forms a nearly identical plane whose shading is solid. We can also see that the first, fourth, and seventh columns of the array form a plane consisting of all the cards whose number is one; the second, fifth, and eighth columns form a nearly identical plane, except that these cards' number is two; and the third, sixth, and ninth columns form a final nearly identical plane consisting of cards whose number is three.

These are just a few examples of planes present in this structure. In fact, we can verify that for any initial partition of this structure into three planes, we can find 36 more planes, for a total of 39 in the whole structure! We do this by choosing one of the twelve SETs in the first of the initial planes, then choosing one of the three SETs in the second of the initial planes that is parallel to this SET. The plane containing those six cards will then be uniquely completed by one of the SETs in the final of the initial planes. We omit this work here; this is a finite structure and so we can simply check manually that this process does in fact work in all cases, but doing so is quite time-consuming. Readers who own a SET deck may want to try building this configuration in three dimensions, with the initial three planes stacked on top of each other; this makes it much easier to visualize where the planes are located.

This structure, as readers may have guessed, models a three-dimensional analog of the finite affine plane; we call this a *hyperplane*, or AG$(3,3)$. Formally defining finite geometry in higher dimensions is somewhat technical, and can be done most precisely using finite fields; thus, we will not attempt to do so here.[2] This is one great benefit of the SET deck; we can use it to physically build these abstract structures and visualize the ways in which they work geometrically, even without a great deal of mathematical background.

It is also worth noting that this is not the only way to build a hyperplane. Recall that, when we were working in two dimensions, we could find a plane either by reducing the number of attributes from the full deck, or by taking any three

[2]See [2] for an algebraic treatment of this.

non-collinear cards and filling in the SETs determined by those cards. The same turns out to be true for any four non-coplanar cards in this space! We can use three cards to determine a plane, then use the fourth card to determine the placement of cards to form the rest of the SETs and planes in AG$(3,3)$.

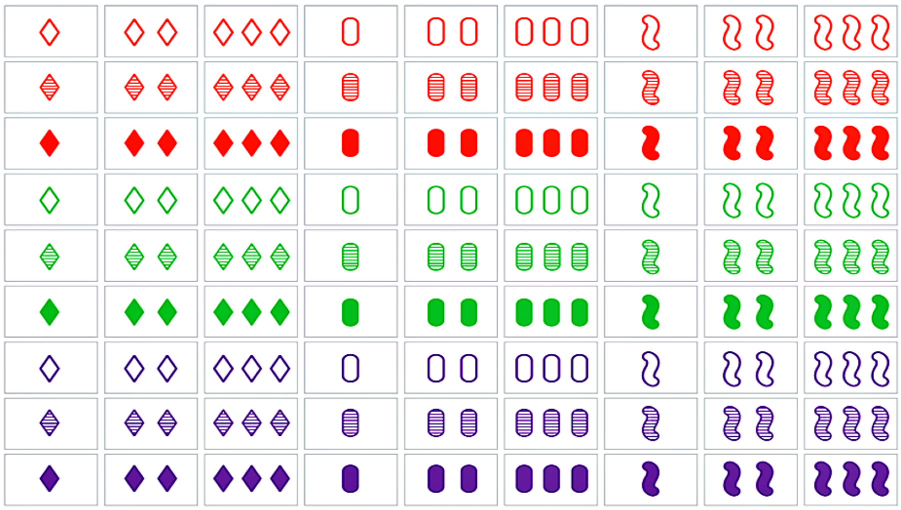

FIGURE 17. The full SET deck, AG$(4,3)$.

Similarly, we can use the entire SET deck to construct AG$(4,3)$. This time it is not necessary to determine which cards belong in the subset, since we are using the whole deck. However, we can still arrange the cards in a way that allows us to more easily visualize the structure of the geometry. As with the previous dimensions, we can arrange the deck neatly by attribute and value, which allows us to intuitively visualize SETs, planes, and hyperplanes. Figure 17 shows one such nicely organized layout of the deck. Or, as in the two-and-three-dimensional cases, we can use three non-collinear cards to determine a plane, add another card to determine a hyperplane, and add yet another card to determine the final four-dimensional structure. And, although it would quickly become rather unwieldy, it is possible to create a SET deck with even more attributes and use it to model spaces in even higher dimensions!

5. Maximal Caps

Finally, we conclude this paper with a brief suggestion for future research: the topic of *maximal cap size.* This is an open problem in finite affine geometry with a clear parallel in SET gameplay; and once again, the SET deck proves to be a helpful physical model of the geometric structures being studied.

As mentioned in the first section of this paper, it is sometimes the case during a game of SET that there are no SETs on the table. When this occurs, more cards are added until a SET can be found. A question that eventually arises for any SET player is how long this process can continue. That is, how many cards need to be on the table before we can guarantee that there is a SET present? Or, in the language of finite affine geometry, what is the maximum number of points we can have in a subset of AG$(4, 3)$ without there being any complete lines in the subset? Such collections of points are called *caps.*

Definition 5.1. An *n-cap,* also referred to as simply a *cap*, is a subset of points in AG$(n, 3)$ that does not contain any lines.

For example, the four cards in Figure 18 form an example of a maximal 2−cap, taking our SET plane from figure 15 as AG$(2, 3)$. There are no SETs present; however, we can verify that adding any other card from the plane will complete a SET, and further, any collection of five cards from this plane will also contain a SET.

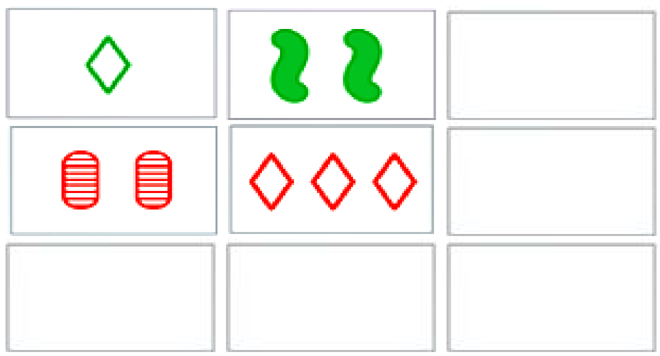

Figure 18. A maximal 2−cap.

The maximal size of an n-cap is known for n equal to 2, 3, 4, 5, and 6.[3] In four dimensions, which we know corresponds to the SET deck, the maximal size of a cap is 20. Thus, once there are 21 cards on the table, there is guaranteed to be a SET present somewhere.

Proving maximal cap size becomes much more complicated as the dimension increases; further, several patterns in the structure of maximal caps that appear to hold in earlier dimensions break down in higher ones. Thus, this remains a fruitful topic for future research.

Acknowledgments

First, I would like to thank my mentors, Iris Yunxuan Li and Cindy Tan, for providing accountability and support. Next, special thanks goes to Tony Brooks for moral support and to Chris Bolognese for sparking my interest in this particular topic many years ago. I'd also like to extend my gratitude to everyone who gave a talk or otherwise participated in the program this summer; thank you for sharing your knowledge and your time. Finally, thanks to Peter May for organizing an excellent REU program this summer!

[3]For proofs and visual examples of caps in multiple dimensions, see [3] and [4].

References

[1] Liz McMahon, Gary Gordon, Hannah Gordon, and Rebecca Gordon. *The Joy of Set.* Princeton University Press, 2017.

[2] Hillary Fairbanks. *The Game SET as* $\mathbf{F}_3^4$. University of California Davis, 2014.

[3] Charlotte Chan. *SETs and Anti-SETs: The Math Behind the Game of SET*. Princeton University, 2010.

[4] Benjamin Lent Davis and Diane MacLagan. *The Card Game SET*. The Mathematical Intelligencer, 2003.

[5] Peter J. Cameron. *Combinatorics: Topics, Techniques, Algorithms.* Cambridge University Press, 1994.

[6] Images of SET cards taken from `https://geekandsundry.com/the-card-game-that-puzzled-mathematicians-for-decades/`

[7] Affine plane diagram take from `https://en.wikipedia.org/wiki/Affine_plane_(incidence_geometry)`