7. ** Given $|X| = n$ and distinct nonempty subsets $A_1, \dots, A_n \subset X$. Show that there is an $x \in X$ such that the sets $A_i \backslash \{x\}$ are still distinct.*

I will give 3 proofs, 2 with linear algebra and 1 with graph theory.

Proof 1: We want to show that given an $n \times n$ matrix A with distinct columns, there is a row such that if we delete it, the resulting $(n-1) \times n$ matrix still has distinct columns.

First suppose $\det(A) \neq 0$. Then consider expanding the determinant on a column j, ie

$$\det(A) = \sum_{i=1}^{n} (-1)^{(i-1)(j-1)} a_{ij} \det(A_{ij}),$$

where a_{ij} is the entry of A on the ith row and jth column, and A_{ij} is the submatrix that we get after deleting the ith row and jth column. Because $\det(A) \neq 0$, there must be an i such that $a_{ij} = 1$ and $\det(A_{ij}) \neq 0$. Then A_{ij} must have distinct columns. Consider deleting row i. The only way this could lead to non-distinct columns is if column j, after removing the ith entry, is equal to a column of A_{ij}, say column k. Then $a_{ik} = 0$, since otherwise column j and k would be equal in A.

This is not a contradiction, but suppose we pick the column j so that it has the least number of 1s. Then $a_{ik} = 0$ would mean that k is a column with fewer 1s, so we would get a contradiction. Therefore we can delete row i.

On the other hand, suppose that $\det(A) = 0$. Then there is a row i that is a linear combination of the other rows. We claim that we can remove this row. Suppose that deleting i leaves two columns j and k equal. Then $a_{ij} \neq a_{ik}$, because otherwise column j and k would be equal in A. But this contradicts row i being a linear combination of other rows, because taking a linear combination of indentical entries (from column j and k) should give identical results for a_{ij} and $= a_{ik}$. This finishes the proof.

Proof 2: (Antoine Imboden had this solution.) Suppose on the contrary that for each x_i there are two sets A_{j_i} and A_{k_i} such that $A_{j_i} \backslash \{x_i\} = A_{k_i} \backslash \{x_i\}$. Then the corresponding characteristics vectors satisfy

$$a_{j_i} - a_{k_i} = e_i \text{ or } -e_i,$$

where e_i is the vector with ith entry 1 and the other entries 0. So we have for $i = 1, \dots, n$ that

$$e_i \in \operatorname{span}\{a_j - a_k : 1 \leq j, k \leq n\} = \operatorname{span}\{a_j - a_{j+1} : 1 \leq j \leq n-1\} =: V.$$

But this is a contradiction because clearly $\dim(V) = n - 1$, while the e_i are n linearly independent vectors.

Proof 3: (Several students did something like this.) Suppose on the contrary that for each x there are two sets A_{j_x} and A_{k_x} such that $A_{j_x} \backslash \{x\} = A_{k_x} \backslash \{x\}$. We define a graph with vertices the A_j, and we connect A_{j_x} and A_{k_x} for each x. Then we have n vertices and n edges, so by a basic result from graph theory, there must be a cycle in this graph (an acyclic graph on n vertices is called a forest and can have at most $n - 1$

edges; if it also connected, it is called a tree and must have exactly $n - 1$ edges).
But this is a contradiction because this graph cannot have a cycle. Indeed, looking at the edge corresponding to x, we must have $x \in A_{j_x}$ and $x \notin A_{k_x}$ or vice versa. For every other edge, say between A_{j_y} and A_{k_y}, x is either in both sets or in neither. Such edges can never connect A_{j_x} to A_{k_x}.

Remark: Note that you have be a bit careful with this graph approach. If you define the graph to have vertices A_i and you put an edge whenever there is an x such that $A_{j_x}\backslash\{x\} = A_{k_x}\backslash\{x\}$, then the graph *will* have cycles.
You could try to fix this by taking a directed graph, directing each edge from the smaller set to the larger one. Then it is true that this graph has no *directed cycle*, but this does not imply that the number of edge is $\leq n - 1$.
You can still fix this approach by observing that in every cycle there must be an even number of edges corresponding to each x. So you can break each cycle by removing an edge, while still leaving the graph with an edge corresponding to each x. This gives an acyclic graph with $\geq n$ edges, and your contradiction. Proof 3 circumvents this by right away taking a graph with only one edge per x.

8. * *The examples that you found in problem 3 are not uniform. Show that there exists a tight example for the bound $|\mathcal{S}| \leq |X|$ for uniform 1-intersecting set systems in the case $|X| = 7$.*
Show that in general this is only possible if $|X| = q^2 + q + 1$ for some integer $q > 0$.
Here is such a system (writing $\{a, b, c\} = abc$):

$$\{123, 156, 147, 246, 257, 345, 367\}.$$

It is the only one up to permutations. This is a famous example that shows up all over discrete math, called the *Fano plane* (google it to see a picture or learn more).
It is an open research problem to determine for which q such an example really exists.

I'll give two proofs of the second (main) part. The first is by your fellow student Samuel Regamey, and unlike my own solution uses the incidence matrix (appropriately for this course).
First a little notation: We write k for the size of the $|S_i|$, and

$$r_x = \#\{S_i : x \in S_i\}$$

for the *representation number* of x.
We'll ignore trivial cases likes $|X| = 1$ or $k = 1$ below.

Proof 1: Let A be the incidence matrix, of size $|X| \times |\mathcal{S}|$, and consider the adjacency matrix $A^T A = J + (k - 1)I$. We count the number of 1s in $A^T A$ in two ways:

$$n^2 - n = \sum_{x \in X} r_x(r_x - 1).$$

Consider the S_i that contain x: they cannot intersect in any other points, so the size of their union is $\geq r_x \cdot (k - 1) + 1$, and it is also $\leq n$. This implies $r_x \leq n/(k - 1)$, which gives us

$$n^2 - n \leq \sum_{x \in X} \frac{n}{k - 1} \cdot (r_x - 1).$$

This we have for all x, so we also have

$$n^2 - n \leq \sum_{x\in X} \frac{n}{k-1} \cdot (\min(r_x) - 1) = \frac{n^2}{k-1} \cdot (\min(r_x) - 1),$$

so

$$1 - \frac{1}{n} \leq \frac{\min(r_x) - 1}{k-1}.$$

Since r_x and k are integers $< n$, this implies that $r_x \geq k$ for all x.
Now we count the number of 1s in A in two ways:

$$nk = \sum_{x\in X} r_x \geq \sum_{x\in X} k = nk.$$

This means that the inequality $r_x \geq k$ must always be tight, ie $r_x = k$ for all x. Going back to the equation for the number of 1s in $A^T A$ we get

$$n^2 - n = \sum_{x\in X} r_x(r_x - 1) = nk(k-1) \quad \Rightarrow \quad n = k^2 - k + 1.$$

Now put $k = q + 1$ to get the required form $q^2 + q + 1$.

Proof 2: Assume that we have a tight example with $|\mathcal{S}| = |X|$. We claim that then for all $x \neq y \in X$, there is a unique $S \in \mathcal{S}$ such that $x, y \in S$.
If we prove this claim, then the number of pairs in X equals the number of pairs within the subsets from $\mathcal{S}$:

$$\binom{|X|}{2} = |\mathcal{S}| \cdot \binom{k}{2} = |X| \cdot \binom{k}{2}.$$

So $|X| - 1 = k(k-1)$ and $|X| = k^2 - k + 1$; $k = q + 1$ gives the required form $q^2 + q + 1$.
Now we prove the claim. Consider x and a set $S \in \mathcal{S}$ such that $x \notin S$. Then every set S_i that does contain x must intersect S in some other element, call it y_i. These y_i are all distinct because these S_i are 1-intersecting and already intersect in x. Hence

$$r_x = \#\{S_i : x \in S_i\} = \# y_i \leq |S| = k.$$

On the other hand we have

$$\sum_{x\in X} r_x = \sum_{S\in\mathcal{S}} |S| = |\mathcal{S}| \cdot k = |X| \cdot k,$$

so the average of the r_x is k. This implies that all $r_x = k$.
Consider a pair $x \neq y \in X$, and take some $S \in \mathcal{S}$ such that $y \in S$ but $x \notin S$. Then each of the $r_x = k$ sets S_i that contain x must intersect S in some distinct y_i, and because $|S| = k$ we must have $y_i = y$ for some i. Therefore $x, y \in S_i$. No other set of $\mathcal{S}$ can contain x and y because then they would be 2-intersecting.