Remarks on collections, Hilary Term, 2012

Moderations Pure Math

1. (a) Even though many of you got this right, I did want to emphasize the importance of understanding the *definition* of an invertible map. In many different contexts, we are interested in a collection of mathematical objects and maps between them. For a map

$$f : A \longrightarrow B$$

to be invertible, it will always be the case that we are required to have a map

$$g : B \longrightarrow A,$$

for which $g \circ f$ and $f \circ g$ are both the identity map. When we are dealing with just sets, it is an easy theorem that a bijective map is an invertible map in this sense. However, when we deal with mathematical objects having more structure, the maps will not be arbitrary, but rather, will be required to preserve the structures. And then, bijectivity might not be enough to get the set-theoretic inverse to have the right property. You will see this when you study bits of topology. You will see, for example, continuous bijections

$$f : A \longrightarrow B$$

with the property that the set-theoretic inverse is not continuous. In this case, f is invertible as a map of sets, but not as a map of topological spaces.

This discussion might be somewhat confusing right now, but you should bear in mind that the existence of the map g as above is really the conceptually correct definition of invertibility. The general principle is that when the structures are 'purely algebraic' in some sense, then bijectivity is sufficient for invertibility. But as soon as you have spatial or analytic structures, matters get more complicated.

(b)(iv) A surprising number of people were confused by this. Given a map

$$f : A \longrightarrow A,$$

you should say in words what the condition

$$f^{-1}(f(X)) = X$$

means:

> The only things that get mapped to $f(X)$ are the elements of X;

or

> If something gets mapped to the image of X, then the thing must be in X.

Stated in this form, you see that any injective map will satisfy the condition. I will leave it to you to cook up injective maps that are not isomorphisms.
In this problem, you should be careful to note that the notation

$$f^{-1}$$

refers to *inverse image*. Its usage does not indicate that the map f is invertible. More precisely, for any $a \in A$, $f^{-1}(a)$ is the *set* of all things that get mapped to a. Somewhat unfortunately, the same notation is used for the inverse map when it exists. You should understand these concepts well enough to distinguish the two uses (which are loosely consistent) from the context.

2. (b) (ii) A common approach was to remove elements of S one by one. At some point in such a proof, one is tempted to say '... and so on.' While this is in principle correct, and certainly conveys the right idea, for a beginner in mathematics, it is not bad practice to try to avoid such mildly ambiguous arguments. One way to clean this up is to do the proof by induction on the size $|S|$ of S. The statement is clear with $B = S$ if $|S| = 1$. (Perhaps you should deal right at the beginning with the possibility that $V = 0$, the zero vector space, so that you can assume $|S| \geq 1$.) Now assume the statement true for $|S| = k \geq 1$ and let $|S| = k + 1$. Then it becomes clear that you can stop now right after removing one vector that's in the span of the rest. (Unless S is linearly independent, in which case we again have $B = S$.)
The other elegant approach is to use the notion of a *maximal linearly independent subset.*

3. Write a matrix A as a collection of column vectors:

$$A = (A_1, A_2, \ldots, A_n).$$

For a column vector $\mathbf{x} = (\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_n})^{\mathbf{T}}$, what does the product

$$A\mathbf{x}$$

give in terms of the columns? From the rules of matrix multiplication we get that

$$A\mathbf{x} = \mathbf{x_1}\mathbf{A_1} + \mathbf{x_2}\mathbf{A_2} + \cdots + \mathbf{x_n}\mathbf{A_n},$$

the linear combination of the columns of A with coefficients provided by the components of the vector $\mathbf{x}$. This elementary fact is so important you should never lose sight of it in linear algebra. We see therefore, that as we run over different vectors $\mathbf{x}$, the vector $A\mathbf{x}$ exactly runs over *the span of the columns of A*. This is sometimes called the *column space of A*. So when does

$$A\mathbf{x} = \mathbf{c}$$

have a solution? Exactly when $\mathbf{c}$ is in the column space of A. This is true whether or not A is invertible[1]. Recall that invertibility of (square) A is equivalent to the fact that the columns of A are linearly independent, and hence, form a basis of our space. In this case, the column space of A is the entire vector space $\mathbb{R}^n$, giving another explanation for why the equation is always solvable in that case. But bear in mind that for specific $\mathbf{c}$, the equation might be solvable even when A is not invertible. In the problem at hand, this will happen when $\alpha = 2$.
By the way, when does the equation have many solutions? You will see that this happens when $A\mathbf{x} = \mathbf{0}$ has a non-zero solution, say, $\mathbf{v}$. This is because for any solution $\mathbf{a}$ to

$$A\mathbf{x} = \mathbf{c},$$

we can always add anything in the line $\{\lambda\mathbf{v} \mid \lambda \in \mathbb{R}\}$, without changing the result of applying A:

$$A(\mathbf{a} + \lambda\mathbf{v}) = \mathbf{A}(\mathbf{a}).$$

Clearly, for such a $\mathbf{v}$ to exist, we need a non-trivial solution to

$$x_1A_1 + x_2A_2 + \cdots + x_nA_n = 0,$$

that is, we need the columns of A to be linearly dependent. You will have learned the determinant criterion for this as well:

$$\det(A) = 0.$$

[1]Obviously, this notion requires A to be a square matrix. Much of the previous discussion does not need A to be square.

If we put this all together, we are saying that

$$A\mathbf{x} = \mathbf{c}$$

has infinitely many solutions if $\mathbf{c}$ is in the column space of A *and* $\det(A) = 0$.
You should check the following as an exercise:
If the equation

$$A\mathbf{x} = \mathbf{c}$$

has two distinct solutions, then it has infinite many solutions.

4. (a) There was an answer that went ' $X \subset V$ is a subspace if it is itself a vector space. ' This statement is somwhat ambiguous, since it's not clear what's meant by 'is itself a vector space'. For example, the line

$$y = x + 1$$

in $\mathbb{R}^2$ can be quite naturally be considered a vector space as well, since it's a line, but it's certainly not a subspace of $\mathbb{R}^2$. So what was meant was that X is itself a vector space *when the operations are induced from those of* V. Of course, this reduces to checking that X is closed under those operations.

Part A, AC1

1. Much of this question can be answered without recourse to explicit computation. For example,

$$\mathrm{tr}(B^tA) = \mathrm{tr}((B^tA)^t) = \mathrm{tr}(A^t(B^t)^t) = \mathrm{tr}(A^tB),$$

where all we have used is that $\mathrm{tr}(C) = \mathrm{tr}(C^t)$ for any square matrix C. This last fact is obvious because the diagonal is unchanged after taking the transpose. Note that the equality

$$\mathrm{tr}(AB) = \mathrm{tr}(BA)$$

is also true, but requires some computation to prove.
Try re-examining the problem in this light.

2. Most people did this problem correctly. You might want to ask yourself the following questions:
(1) Is T diagonalizable? Why or why not?
(2) Find a matrix with the same characteristic polynomial for which the minimal polynomial is

$$(x+3)^2.$$

Find one for which it is $(x+3)$.

3. Here are some further questions:
True or false:

$$\mathbb{C}[x]/((x-1)^2) \simeq \mathbb{C} \times \mathbb{C},$$

as rings.
Compute

$$\mathbb{C}[x]/(x^2-1).$$

Part A, AC2

1, Ask yourself the following question and make sure you understand the answer: The dual basis E^* is defined by

$$e_i^*(e_j) = \delta_{ij},$$

that is, e_i^* is the linear map that sends e_i to 1 and all other basis vectors to zero. Why does this give a well-defined linear map? Let $V = \mathbb{R}^3$, for example. Does it make sense to speak of the linear map

$$\mathbb{R}^3 \longrightarrow \mathbb{R}$$

that takes $(1,0,0)^t$ to 1, $(0,1,1)^t$ to 0 and $(1,1,1)^t$ to 0? How does this situation differ from the definition of E^*?

On rather important point concerns the definition of the matrix M of a linear map T with respect to a basis $B = \{b_1, b_2, \ldots, b_n\}$. The usual convention is to say

$$M = (a_{ij})$$

where

$$Tb_j = \sum_i a_{ij} b_i. \quad (1)$$

Some people wrote

$$Tb_i = \sum_{ij} a_{ij} b_j \quad (2)$$

instead. This kind of distinction is somewhat painful to think about, but can be important in certain contexts. In any case, you have to settle upon the right convenyion if only for psychological comfort. The point of the first formula is that the image of any given basis vector under T determines a *column* of the matrix. That is to say, Tb_j determines the *j-th column of M*. That is what it means to sum over i in the formula. So to compute the matrix of a T, you compute Tb_1 and write it in terms of the original basis. Then the coefficients you get make up the first column of M. Now continue with Tb_2 and so on[2]. The main reason formula (1) is preferrable to (2) is that we are used to thinking of vectors in $\mathbb{R}^n$ as column vectors, and multiplying a matrix on the left. In this case, when we apply a matrix A to the j-th standard basis of $\mathbb{R}^n$, then we get exactly the j-th column of A. (Do it and make sure.) Finally, I suppose the reason for multiplying the matrix on the left is because we write $f(x)$ for the effect of applying a function f to the argument x. There have been attempts in the past to write $(x)f$ for the same thing, without much success. You can come up with some plausible arguments for a change.

2. There are so many things to say about this problem we will have to put it off until the tutorials. However, you should prove the following: A matrix A with entries in a field K is diagonalizable over K if and only if its minimal polynomial factorizes into distinct linear factors over K. For example, when $K = \mathbb{C}$, this is saying a matrix is diagonalizable if and only if its minimal polynomial has distinct roots.
Ask yourself the question: If you choose an $n \times n$ matrix with complex entries *randomly*, do you expect it to be diagonalizable?

3.
Take a polynomial $f(x) \in K[x]$ where K is any field. Prove: If f has degree 2 or 3, then it is irreducible if and only if it has no roots in K. That is, the only way a polynomial of degree 2 or 3 can be reducible is by having a root in K. When the polynomial has degree 4 or above, it can be reducible without having roots: consider

$$(x^2 + 1)(x^2 + 2) \in \mathbb{R}[x].$$

Equipped with this fact, it is easy to prove that $x^2 + 1 \in \mathbb{Z}_7[x]$ is irreducible.
Now try to produce an irreducible polynomial of degree 3 in $\mathbb{Z}_7[x]$ and relate it to a field with 343 elements.

[2]See, I'm allowed to use this phrase because I've already paid my dues giving clear-cut arguments when I was young.

By the way, a number of you seemed to think that if $f(x) \in \mathbb{Z}[x]$ is irreducible, then the same polynomial reduced mod 7 is irreducible. But consider something as simple as

$$x^2 - 7$$

which is clearly irreducible in $\mathbb{Z}[x]$. Reduced mod 7, it becomes x^2.
Finally, prove that

$$\mathbb{Z}[i] \simeq \mathbb{Z}[x]/(x^2 + 1).$$