

DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE WEDNESDAY OCTOBER 2: VECTORS (BASIC STUFF)

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

30 people took this 3-question quiz. The score distribution was as follows:

- Score of 1: 2 people
- Score of 2: 10 people
- Score of 3: 18 people

The mean score was about 2.5.

The question-wise answers and performance review were as follows:

- (1) Option (A): 22 people
- (2) Option (E): 24 people
- (3) Option (C): 30 people

2. SOLUTIONS

Many of you are familiar with vectors, either from Math 195 or some exposure to vectors in high school (or perhaps both). This quiz is to help gauge your level of understanding coming in. We will not get to start using the ideas in their full depth until a few weeks later.

For the benefit of those who haven't seen vectors at all, the definitions are briefly provided.

There are many ways of describing the vector in \mathbb{R}^n with coordinates a_1, a_2, \dots, a_n . You may have seen the vector described using angled braces as $\langle a_1, a_2, \dots, a_n \rangle$. In this linear algebra course, we will customarily write the vector as a *column* vector, i.e., the coordinates will be written in a vertical column. For instance, the vector $\langle 2, 3, 7 \rangle$ will be written as the column vector $\begin{bmatrix} 2 \\ 3 \\ 7 \end{bmatrix}$.

Two vectors in \mathbb{R}^n can be added with each other (note that both vectors need to be in the *same* \mathbb{R}^n in order to be added). The addition is coordinate-wise:

$$\begin{bmatrix} v_1 \\ v_2 \\ \cdot \\ \cdot \\ \cdot \\ v_n \end{bmatrix} + \begin{bmatrix} w_1 \\ w_2 \\ \cdot \\ \cdot \\ \cdot \\ w_n \end{bmatrix} = \begin{bmatrix} v_1 + w_1 \\ v_2 + w_2 \\ \cdot \\ \cdot \\ \cdot \\ v_n + w_n \end{bmatrix}$$

Also, given any real number λ (called a *scalar* to distinguish from a vector) and a vector $\vec{v} = \begin{bmatrix} v_1 \\ v_2 \\ \cdot \\ \cdot \\ \cdot \\ v_n \end{bmatrix}$, we

can define:

$$\lambda \vec{v} = \lambda \begin{bmatrix} v_1 \\ v_2 \\ \cdot \\ \cdot \\ v_n \end{bmatrix} := \begin{bmatrix} \lambda v_1 \\ \lambda v_2 \\ \cdot \\ \cdot \\ \lambda v_n \end{bmatrix}$$

We can identify the set of n -dimensional vectors with the set of points in \mathbb{R}^n . The vector $\vec{v} = \begin{bmatrix} v_1 \\ v_2 \\ \cdot \\ \cdot \\ v_n \end{bmatrix}$ in

this case corresponds to the point with coordinates (v_1, v_2, \dots, v_n) .

- (1) *Do not discuss this!* For a n -dimensional vector \vec{v} , the *set of scalar multiples* of \vec{v} is the set of vectors that can be expressed in the form $\lambda \vec{v}$, $\lambda \in \mathbb{R}$. Assume that \vec{v} is a nonzero vector. What can we say geometrically about the set of points in \mathbb{R}^n that correspond to the scalar multiples of \vec{v} ?
- (A) It is a straight line in \mathbb{R}^n that passes through the origin.
 (B) It is a straight line in \mathbb{R}^n . However, it need not pass through the origin.
 (C) It is a straight half-line in \mathbb{R}^n with the endpoint at the origin.
 (D) It is a straight half-line in \mathbb{R}^n , but the endpoint need not be at the origin.
 (E) It is a line segment in \mathbb{R}^n .

Answer: Option (A)

Explanation: This is clear from the way scalar multiplication of vectors is described visually. Note that if we restrict λ to the nonnegative reals (i.e., $\lambda \in (0, \infty)$), then we get a half-line (excluding the endpoint). However, the definition of scalar multiple allows for negative values of λ (which give the opposite half-line) and for $\lambda = 0$ (which gives the origin). We therefore get the whole line.

Performance review: 22 out of 30 people got this. 2 each chose (B), (C), (D), and (E).

- (2) *Do not discuss this!* Given two n -dimensional vectors \vec{v} and \vec{w} , the *set of linear combinations* of \vec{v} and \vec{w} is the set of all vectors that can be written in the form $\lambda \vec{v} + \mu \vec{w}$ where $\lambda, \mu \in \mathbb{R}$ (note that λ and μ can take arbitrary real values, and are allowed to be equal to each other). In other words, you can take scalar multiples, and you can then add these scalar multiples.

The set of linear combinations of \vec{v} and \vec{w} is sometimes also called the *span* of \vec{v} and \vec{w} .

What is the span of the vectors $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ in \mathbb{R}^2 ?

- (A) The zero vector only, because that is the only vector that can be expressed both as a multiple of $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and as a multiple of $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$.
 (B) The set of vectors that can be expressed as a scalar multiple either of $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ or of $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$.
 (C) The set of vectors that can be expressed as a scalar multiple of at least one of these three vectors: $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$, and $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$.
 (D) All vectors in the first quadrant of \mathbb{R}^2 , including the bounding half-lines. In other words, the set of vectors of the form $\begin{bmatrix} x \\ y \end{bmatrix}$ where $x \geq 0$ and $y \geq 0$.
 (E) All vectors in \mathbb{R}^2 .

Answer: Option (E)

Explanation: For a vector $\begin{bmatrix} x \\ y \end{bmatrix}$, we can write:

$$\begin{bmatrix} x \\ y \end{bmatrix} = x \begin{bmatrix} 1 \\ 0 \end{bmatrix} + y \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

Thus, we have expressed our arbitrary vector of \mathbb{R}^2 as a linear combination of the vectors $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$. In the notation given, $\lambda = x$ and $\mu = y$.

Note the key point that the λ and μ values can vary arbitrarily. They are not restricted to be nonnegative (Option (D)) and they are not required to have at least one of them zero, or to be equal to each other (Options (A)-(C)).

Performance review: 24 out of 30 people got this. 3 chose (C), 2 chose (D), 1 chose (A).

- (3) *Do not discuss this!:* Consider the transformation from \mathbb{R}^2 to \mathbb{R}^2 that interchanges the coordinates of a vector. Explicitly, the transformation is given as:

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto \begin{bmatrix} y \\ x \end{bmatrix}$$

Which of the following describes the transformation geometrically, with \mathbb{R}^2 viewed as the xy -plane?

- (A) It is a reflection about the x -axis in \mathbb{R}^2 , i.e., the axis for the first coordinate.
- (B) It is a reflection about the y -axis in \mathbb{R}^2 , i.e., the axis for the second coordinate.
- (C) It is a reflection about the line $y = x$ in \mathbb{R}^2 , i.e., the line of vectors where both coordinates are equal.
- (D) It is a reflection about the line $y = -x$ in \mathbb{R}^2 , i.e., the line of vectors where the coordinates are negatives of each other.

Answer: Option (C)

Explanation: You need to visualize this geometrically. The midpoint of (x_0, y_0) and (y_0, x_0) is $((x_0 + y_0)/2, (x_0 + y_0)/2)$, which lies on the $y = x$ line. Further, the line joining the points (x_0, y_0) and (y_0, x_0) has slope -1 , hence is perpendicular to the $y = x$ line. Therefore, the line $y = x$ is the perpendicular bisector of the line segment joining (x_0, y_0) and (y_0, x_0) . In other words, reflecting about the line $y = x$ sends (x_0, y_0) to the point (y_0, x_0) .

You've probably seen this in calculus already: remember that to obtain the graph of f^{-1} from the graph of f , you reflected about the $y = x$ line? That's because reflection about the $y = x$ line reverses the roles of x and y .

Performance review: All 30 people got this.

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY OCTOBER 4: LINEAR FUNCTIONS AND EQUATION-SOLVING (PART 1)

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

27 people took this 9-question quiz. The score distribution was as follows:

- Score of 1: 2 people
- Score of 3: 1 person
- Score of 4: 2 people
- Score of 5: 4 people
- Score of 6: 3 people
- Score of 7: 8 people
- Score of 8: 5 people
- Score of 9: 2 people

The mean score was about 6.1.

The question-wise answers and performance review are below:

- (1) Option (D): 20 people
- (2) Option (C): 16 people
- (3) Option (B): 21 people
- (4) Option (A): 23 people
- (5) Option (E): 9 people
- (6) Option (D): 19 people
- (7) Option (C): 26 people
- (8) Option (C): 11 people
- (9) Option (A): 20 people

2. SOLUTIONS

This quiz covers some basics involving linear functions and equation-solving (notes at **Linear functions: a primer** and **Equation-solving with a special focus on the linear case**). The quiz tests for the following:

- What it means to be (affine) linear, and in particular, the significance of the intercept as an additional parameter to track.
 - The distinction between behavior relative to the variables (the inputs) and behavior relative to the parameters.
 - Using the linear paradigm to study functional forms that are not themselves linear.
 - A small taste of dealing with measurement uncertainty to obtain upper and lower bounds (not covered in the notes, so this is where your famed ability to think out of the box should manifest).
 - Solving “triangular” systems of equations.
- (1) A function f of 3 variables x, y, z defined everywhere is (affine) linear in the variables. (The “affine” is to indicate that the intercept may be nonzero). Based on the above information and some input-output pairs for f , we would like to determine f uniquely. What is the minimum number of input-output pairs that we would need in order to achieve this?
- (A) 1
 - (B) 2
 - (C) 3

(D) 4

(E) 5

Answer: Option (D)

Explanation: The general expression for a linear function f of the variables x , y , and z is:

$$f(x, y, z) := ax + by + cz + d$$

There are four unknown parameters here (one coefficient for each variable, and one parameter d for the intercept). Our goal is to determine uniquely the values of the parameters. Since there are four parameters, we need four equations to determine them uniquely. Note that the equations themselves are linear. We should choose our four inputs in a manner that there are no linear dependencies between them (what this means will become clearer as we study more linear algebra).

Performance review: 20 out of 27 got this. 6 chose (C), 1 chose (E).

Historical note (last time; however, discussion was not permitted for this question last time): 7 out of 29 people got this. 20 chose (C), 2 chose (B).

- (2) Which of the following gives an example of a function F of three variables x, y, z whose third-order mixed partial derivative F_{xyz} is zero everywhere, but for which none of the second-order mixed partial derivatives F_{xy} , F_{xz} , F_{yz} is zero everywhere?

(A) $\sin(xy) - z^2$

(B) $\cos(x^2 + y^2) - \sin(y^2 + z^2)$

(C) $e^{xy} + (y - z)^2 + 3xz$

(D) $x^2 + y^2 + z^2$

(E) xyz

Answer: Option (C)

Explanation: Look for functions of the form:

$$F(x, y, z) = f(x, y) + g(y, z) + h(x, z)$$

where f , g , and h are all nonzero and none of them is additively separable.

The only option fitting this description is Option (C).

As for the other options:

- Option (A): Both F_{xz} and F_{yz} are zero. This is because neither x nor y interacts with z .
- Option (B): F_{xz} is zero, because there is no interaction between x and z .
- Option (D): F is completely additively separable in terms of x , y , and z , so all the second-order mixed partials F_{xy} , F_{xz} , and F_{yz} are zero.
- Option (E): $F_{xyz} = 1$.

Performance review: 16 out of 27 got this. 7 chose (E), 3 chose (B), 1 chose (A).

Historical note (last time): 20 out of 29 people got this. 8 chose (E), 1 chose (D).

- (3) Consider a function of the form $F(x, y) := Ca^xb^y$ where C, a, b are all positive reals that serve as parameters and x, y are restricted to the positive reals. We wish to study F using the paradigm of linear functions. What is the best way of doing this?

(A) Express $\ln(F(x, y))$ in terms of $\ln x$ and $\ln y$

(B) Express $\ln(F(x, y))$ in terms of x and y

(C) Express $F(x, y)$ in terms of $\ln x$ and $\ln y$

(D) Express $\ln(F(x, y))$ in terms of a^x and b^y

(E) Express $F(x, y)$ in terms of a^x and b^y

Answer: Option (B)

Explanation: We take logarithms to get:

$$\ln(F(x, y)) = \ln C + x \ln a + y \ln b$$

Note that a , b , and C are positive constants. Hence, $\ln C$, $\ln a$, and $\ln b$ are also constants. Thus, $\ln(F(x, y))$ is a linear function of x and y .

Performance review: 21 out of 27 got this. 4 chose (A), 2 chose (D).

Historical note (last time): 26 out of 29 got this. 3 chose (A).

- (4) Consider a function of the form $F(x, y) := Cx^a y^b$ where C, a, b are all positive reals that serve as parameters and x, y are restricted to the positive reals. We wish to study F using the paradigm of linear functions. What is the best way of doing this?
- (A) Express $\ln(F(x, y))$ in terms of $\ln x$ and $\ln y$
 - (B) Express $\ln(F(x, y))$ in terms of x and y
 - (C) Express $F(x, y)$ in terms of $\ln x$ and $\ln y$
 - (D) Express $\ln(F(x, y))$ in terms of x^a and y^b
 - (E) Express $F(x, y)$ in terms of x^a and y^b

Answer: Option (A)

Explanation: We take logarithms to get:

$$\ln(F(x, y)) = \ln C + a \ln x + b \ln y$$

Note that C is a positive constant, so $\ln C$ is a constant. $\ln(F(x, y))$ is a linear function of $\ln x$ and $\ln y$.

Performance review: 23 out of 27 got this. 1 each chose (B), (C), and (D), 1 left the question blank.

Historical note (last time): 28 out of 29 got this. 1 chose (B).

- (5) (**) *This is a hard question!* The population in the island of Andrognesia as a function of time is believed to be an exponential function. On January 1, 1984, the population was measured to be $3 * 10^5$ with a measurement error of up to 10^5 on either side, i.e., the population was measured to be between $2 * 10^5$ and $4 * 10^5$. On January 1, 1998, the population was measured to be $1.2 * 10^6$ with a measurement error of up to $4 * 10^5$ on either side, i.e., the population was measured to be between $8 * 10^5$ and $1.6 * 10^6$. If the population is an exponential function of time (i.e., the increment in population per year is a fixed proportion of the population that year), what is the **range of possible values** of the population measured on January 1, 2012? *Hint: Think of the umbral versus penumbral region for an eclipse*
- (A) Between $3.2 * 10^6$ and $6.4 * 10^6$
 - (B) Between $3.2 * 10^6$ and $1.28 * 10^7$
 - (C) Between $1.6 * 10^6$ and $3.2 * 10^6$
 - (D) Between $1.6 * 10^6$ and $6.4 * 10^6$
 - (E) Between $1.6 * 10^6$ and $1.28 * 10^7$

Answer: Option (E)

Explanation: Note first that $2012 - 1998 = 1998 - 1984 = 14$.

The key idea is that the lowest estimate occurs if the 1998 population was measured as low as possible *and* the rate of population growth estimated using the 1984 and 1998 populations is as low as possible. The lowest possible rate of growth we can measure occurs if we choose the highest possible 1984 value and the lowest possible 1998 value. Picking these, we obtain that the population estimate for 1984 is $4 * 10^5$ and the population estimate for 1998 is $8 * 10^5$. Since the multiplicative growth of the population depends on the time elapsed, the total population in 2012 will be the solution x to:

$$\frac{x}{8 * 10^5} = \frac{8 * 10^5}{4 * 10^5}$$

which solves to $x = 1.6 * 10^6$.

Similarly, the highest estimate will occur if we take the highest estimate possible for the 1998 population and the lowest estimate possibly for the 1984 population.

This relates to the idea of linear models as follows. Consider a plot of the logarithm of the population with respect to time. Since the growth is exponential, this should be a linear plot. If we knew the precise values of the populations in 1984 and 1998, we could fit a straight line through those and use that to determine the population in 2012. Uncertainty regarding the values of the population in 1984 and 1998, however, means that instead of having points in the graph, we have an interval (represented by a vertical line segment) for the time coordinate value of 1984 and another interval (represented by another vertical line segment) for the time coordinate value of 1998.

The upper end estimate is obtained by making a line through the lower end of the 1984 estimate range and the upper end of the 1998 estimate range. The lower end estimate is obtained by making a line through the upper end of the 1984 estimate range and the lower end of the 1998 estimate range.

Performance review: 9 out of 27 got this. 14 chose (A), 3 chose (D), 1 chose (C).

Historical note (last time): 6 out of 29 got this. 15 chose (A), 6 chose (C), 1 each chose (B) and (D).

- (6) Suppose, according to our model, a particular function $f(x, y)$ is of the form $f(x, y) = a_1 + a_2x + a_3y + a_4x^2y^2$ where a_1, a_2, a_3, a_4 are parameters. Our goal is to determine the values of the parameters a_1, a_2, a_3, a_4 . We do this by collecting a number of (input,output) pairs for the function f and then setting up equations in terms of the parameters using the (input,output) pairs. What can we say about the nature of f and the nature of the system of equations that we will need to solve? *Note that “nonlinear” as used here simply means that the expression is not guaranteed to be linear, though it may turn out to be linear in some cases. Similarly, “non-polynomial” means not guaranteed to be polynomial, though it may turn out to be polynomial in some cases.*
- (A) f is a linear function of x and y , hence we need to solve a linear system of equations to determine the parameters a_1, a_2, a_3, a_4 .
- (B) f is a nonlinear polynomial function of x and y , hence we need to solve a nonlinear polynomial system of equations to determine the parameters a_1, a_2, a_3, a_4 .
- (C) f is a linear function of x and y . However, we need to solve a nonlinear polynomial system of equations to determine the parameters a_1, a_2, a_3, a_4 .
- (D) f is a nonlinear polynomial function of x and y . However, we need to solve a linear system of equations to determine the parameters a_1, a_2, a_3, a_4 .
- (E) f is a nonlinear polynomial function of x and y . However, we need to solve a non-polynomial system of equations to determine the parameters a_1, a_2, a_3, a_4 .

Answer: Option (D)

Explanation: f is polynomial in x and y , and it is not linear because one of its terms is x^2y^2 (note that it may *happen* to be linear if $a_4 = 0$, but we do not know this in advance).

However, f is linear in the parameters, hence the system of equations that we get from input-output pairs is a linear system of equations.

Performance review: 19 out of 27 got this. 6 chose (B), 1 each chose (C) and (E).

Historical note (last time): 24 out of 29 got this. 4 chose (B), 1 chose (A).

- (7) Consider the system of equations:

$$\begin{aligned}x^2 - x &= 2 \\y^2 + xy &= x + 13\end{aligned}$$

What is the number of solutions to this system for real x and y ?

- (A) 0
(B) 2
(C) 4
(D) 6
(E) 8

Answer: Option (C)

Explanation: Solving the first equation, we get:

$$(x - 2)(x + 1) = 0$$

Thus, we have $x = 2$ or $x = -1$.

For each choice of x , we need to solve the second equation by plugging in that value of x .
For the choice $x = 2$, we have:

$$y^2 + 2y = 15$$

This simplifies to:

$$(y - 3)(y + 5) = 0$$

Thus, $y = 3$ or $y = -5$, so we get the solutions $x = 2, y = 3$ and $x = 2, y = -5$.
For the choice $x = -1$, we get:

$$y^2 - y = 12$$

This gives:

$$(y - 4)(y + 3) = 0$$

Thus, $y = 4$ or $y = -3$, so we get the solutions $x = -1, y = 4$ and $x = -1, y = -3$.
Overall, we have four solutions:

- $x = 2, y = 3$
- $x = 2, y = -5$
- $x = -1, y = 4$
- $x = -1, y = -3$

Performance review: 26 out of 27 got this. 1 chose (D).

Historical note (last time): 28 out of 29 got this. 1 chose (D).

- (8) Consider the system of equations:

$$\begin{aligned}x^2 - x &= 2 \\y^2 + xy &= x + 13 \\z^2 &= xy\end{aligned}$$

What is the number of solutions to this system for real x , y , and z ?

- (A) 0
- (B) 2
- (C) 4
- (D) 6
- (E) 8

Answer: Option (C)

Explanation: Recall from the preceding question that we have the following solutions to the first two equations:

- $x = 2, y = 3$
- $x = 2, y = -5$
- $x = -1, y = 4$
- $x = -1, y = -3$

For each of these, our goal is to find the corresponding z -values that work. Note that if xy is positive, there are two z -values. If $xy = 0$, there is a unique z -value, and if xy is negative, there are no z -values.

Of the four possible (x, y) -value pairs, only two give positive products. The other two give negative products. In both the positive product cases, we get 2 values of z , so we overall get four solutions as listed below:

- $x = 2, y = 3, z = \sqrt{6}$
- $x = 2, y = 3, z = -\sqrt{6}$
- $x = -1, y = -3, z = \sqrt{3}$
- $x = -1, y = -3, z = -\sqrt{3}$

Performance review: 11 out of 27 got this. 10 chose (B), 3 chose (D), 2 chose (E), 1 chose (A).

Historical note (last time): 25 out of 29 got this. 2 chose (D), 1 each chose (B) and (E).

(9) Consider the system of equations:

$$\begin{aligned}x^2 - x &= 2 \\y^2 + xy &= x + 13 \\z^2 &= x^2 - y^2\end{aligned}$$

What is the number of solutions to this system for real x , y , and z ?

- (A) 0
- (B) 2
- (C) 4
- (D) 6
- (E) 8

Answer: Option (A)

Explanation: The solutions to the first two equations are:

- $x = 2, y = 3$
- $x = 2, y = -5$
- $x = -1, y = 4$
- $x = -1, y = -3$

In all cases, $x^2 - y^2 < 0$. Thus, there is no z -value that works in any of the cases. Thus, there are no solutions to this system.

Performance review: 20 out of 27 got this. 5 chose (E), 1 each chose (C) and (D).

Historical note (last time): 22 out of 29 got this. 3 chose (D), 2 chose (C), 1 each chose (B) and (E).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE MONDAY OCTOBER 7: LINEAR FUNCTIONS AND EQUATION-SOLVING (PART 2)

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

28 people took this 5-question quiz. The score distribution was as follows:

- Score of 0: 3 people
- Score of 1: 3 people
- Score of 2: 9 people
- Score of 3: 5 people
- Score of 4: 6 people
- Score of 5: 2 people

The question-wise answers and performance summary are given below.

- (1) Option (C): 25 people
- (2) Option (D): 12 people
- (3) Option (E): 12 people
- (4) Option (C): 15 people
- (5) Option (D): 6 people

On comparison with last time: Last year, I had set the due date for the take-home quiz as Wednesday rather than Monday, so that might partly explain the better performance of students on the quiz last time. The main question where performance diverged considerable from last time was Q5, and this is the question where the sophistication of one extra class can go a long way.

2. SOLUTIONS

This quiz covers some basics involving linear functions and equation-solving (notes at **Linear functions: a primer** and **Equation-solving with a special focus on the linear case**). The quiz tests for the following:

- The distinction between behavior relative to the variables (the inputs) and behavior relative to the parameters.
 - Counting the number of parameters by creating the explicit general functional form from a verbal description (with a special focus on polynomial functional forms).
 - Figuring out how to “ask the right questions” with respect to input choices, so that the answers provide meaningful information. This builds towards the ideas of hypothesis testing, rank, and overdetermination that we will see in the future.
- (1) Suppose f is a polynomial function of x of degree at most a *known number* n . What is the minimum number of (input,output) pairs that we need in order to determine f uniquely? *Extra information: Somewhat surprisingly, in this case, we do not need to be judicious about our input choices. Any set of distinct inputs of the required number will do. This has something to do with the “Vandermonde matrix” and “Vandermonde determinant” and is also related to the Lagrange interpolation formula.*
- (A) $n - 1$
 - (B) n
 - (C) $n + 1$
 - (D) $2n$
 - (E) n^2

Answer: Option (C)

Explanation: There are $n + 1$ unknown coefficients, namely, the coefficients of x^i for $0 \leq i \leq n$. In general, a polynomial of degree n is of the form:

$$a_0 + a_1x + a_2x^2 + \cdots + a_nx^n$$

Every input-output pair gives rise to a linear equation in terms of the coefficients. With $n + 1$ input-output pairs, we get $n + 1$ equations in $n + 1$ variables where the “variables” here are the parameters). Due to linearity in the parameters, we actually get a system of $n + 1$ *linear* equations in the $n + 1$ variables. We would expect the system to have a unique solution if the inputs are chosen judiciously.

In fact, it turns out that the system always has a unique solution. The existence of a solution is given by the Lagrange interpolation formula, and its uniqueness follows from the fact that any polynomial of degree $\leq n$ that has $n + 1$ distinct roots must be the zero polynomial. This is also related to the idea of the Vandermonde determinant, but that is beyond the scope of the current discussion.

Update for after understanding coefficient matrices and ranks: Choosing the inputs judiciously basically amounts to choosing the inputs in a manner that the coefficient matrix (a $(n + 1) \times (n + 1)$ square matrix) has full rank $n + 1$. The statement above, that distinct inputs always work, is the statement that the coefficient matrix always has full rank $n + 1$ as long as the inputs are distinct. This type of matrix is called a Vandermonde matrix and there is a considerable theory related to these in algebra.

Performance review: 25 out of 28 got this. 3 chose (B).

Historical note (last time): 27 out of 28 got this. 1 chose (B).

- (2) f is a polynomial function of two variables x and y of total degree at most 2. In other words, for each monomial occurring in f , the total of the degrees of x and y in that monomial is at most 2. No other information is given about f . What is the minimum number of judiciously chosen (input,output) pairs we need in order to determine f uniquely?
- (A) 2
(B) 3
(C) 4
(D) 6
(E) 7

Answer: Option (D)

Explanation: First, a little bit to clarify the question. The total degree of a monomial $x^i y^j$ is $i + j$. For instance, the degree of $x^3 y^2$ is $3 + 2 = 5$.

The total degree of a polynomial that involves addition and scalar multiplication of monomials is the maximum of the degrees of the individual monomials. This is very similar to the situation with polynomials of one variable: the degree of the polynomial is the maximum of the degrees of the constituent monomials. Thus, for instance, $x - y + 17xy - x^2 y$ has total degree 3 because the constituent monomials have degree 1, 1, 2, and 3, and the maximum of these is 3.

In order to determine how many input-output pairs we would need, we need to count the number of parameters in the generic functional form. So, the first step is to figure out a *general* expression for all the possibilities we can have such polynomials of total degree at most 2. Since all such polynomials are obtained by (linearly) combining monomials of the sort, our first step is to list the monomials that are allowed.

The possible monomials involved in f are 1, x , y , x^2 , xy , and y^2 . There are six of them, hence there are 6 coefficients to be determined by setting up and solving the linear system. The generic form is:

$$a_1 + a_2x + a_3y + a_4x^2 + a_5xy + a_6y^2$$

Each (input,output) pair gives a *linear* equation in terms of a_1, a_2, \dots, a_6 . Six suitably chosen inputs will give six linear equations that we can then solve to uniquely determine the inputs. Note

that unlike the functions of one variable case, it is possible to choose the inputs “badly”, i.e., in a way that does not reveal information.

Update for after understanding coefficient matrices and ranks: Choosing the inputs judiciously basically amounts to choosing the inputs in a manner that the coefficient matrix (a 6×6 square matrix) has full rank 6. In this case, a bad choice of inputs could lead to a coefficient matrix of lower rank, leading to redundant equations. Unfortunately, it is beyond the current scope to describe the geometry of sets of inputs for which the system does not have full rank (though one example would be where all inputs have the same y -value), but anyway, “most” random choices of inputs will give systems with full rank.

Performance review: 12 out of 28 got this. 6 chose (B), 5 chose (C), 4 chose (A), 1 chose (E).

Historical note (last time): 15 out of 28 got this. 7 chose (B), 3 chose (A), 2 chose (C), 1 chose (E).

- (3) f is a polynomial function of two variables x and y of total degree at most 3. In other words, for each monomial occurring in f , the total of the degrees of x and y in that monomial is at most 3. No other information is given about f . What is the minimum number of judiciously chosen (input,output) pairs we need in order to determine f uniquely?

- (A) 3
 (B) 6
 (C) 8
 (D) 9
 (E) 10

Answer: Option (E)

Explanation: In addition to the 6 monomials for the preceding question, there are 4 monomials of degree three, namely x^3 , x^2y , xy^2 , and y^3 . Thus, there is a total of 10 monomials with unknown coefficients, so we need 10 data points to pin down the values of the coefficients. Explicitly, the general form looks like:

$$a_1 + a_2x + a_3y + a_4x^2 + a_5xy + a_6y^2 + a_7x^3 + a_8x^2y + a_9xy^2 + a_{10}y^3$$

Each (input,output) pair gives a *linear* equation in terms of a_1, a_2, \dots, a_{10} . We want 10 such equations, so we want 10 judiciously chosen input-output pairs.

Update for after understanding coefficient matrices and ranks: Choosing the inputs judiciously basically amounts to choosing the inputs in a manner that the coefficient matrix (a 10×10 square matrix) has full rank 10. In this case, a bad choice of inputs could lead to a coefficient matrix of lower rank, leading to redundant equations. Unfortunately, it is beyond the current scope to describe the geometry of sets of inputs for which the system does not have full rank (though one example would be where all inputs have the same y -value), but anyway, “most” random choices of inputs will give systems with full rank.

Performance review: 12 out of 28 got this. 7 chose (A), 5 chose (B), 2 each chose (C) and (D).

Historical note (last time): 14 out of 28 got this. 8 chose (A), 2 each chose (B), (C), and (D).

- (4) (*) *The perils of overfitting; see also Occam’s Razor:* Suppose we are trying to model a function that we expect to behave in a polynomial-like manner, though we don’t really have a good reason to believe this. Additionally, there is a possibility for measurement error in our observations. Our goal is to find the parameters so that we can both predict unmeasured values and do a qualitative analysis of the nature of the function and its derivatives and integrals.

We have a large number of observations (say, several thousands). We could attempt to “fit” the function using a polynomial of degree n for some fixed n using all those data points, and we will get a certain “best fit” that minimizes the deviation between the curve used for fitting and the function being fit. For instance, for $n = 1$, we are trying to find the best fit by a straight line function. For $n = 2$, we are trying to find the best fit by a polynomial of degree at most 2. We could try fitting using different values of n . Which of the following is true?

If you are interested in more on this, look up “overfitting”. A revealing quote is by mathematician and computer scientist John von Neumann: “With four parameters I can fit an elephant. And with five I can make him wiggle his trunk.” Another is by prediction guru Nate Silver: “The wide array of

statistical methods available to researchers enables them to be no less fanciful and no more scientific than a child finding animal patterns in clouds.”

- (A) Larger values of n give better fits, therefore the larger the value of n we use, the better.
- (B) Smaller values of n give better fits, therefore the smaller the value of n we use, the better.
- (C) Larger values of n give better fits, therefore the larger the value of n we use, the less impressive a good fit (i.e., low deviation between the polynomial and the actual set of observations) should be.
- (D) Smaller values of n give better fits, therefore the smaller the value of n we use, the less impressive a good fit (i.e., low deviation between the polynomial and the actual set of observations) should be.
- (E) The value of n we use for trying to get a good fit is irrelevant. A good fit is a good fit, regardless of the type of function used.

Answer: Option (C)

Explanation: Larger values of n mean more parameters, and we can use more parameters to get a better fit, generally speaking. However, that better fit may well be *fitting the noise in the measurements rather than the signal*. Even without measurement error, if we do not have *a priori* theoretical reasons to be sure of the model, it may just be fitting idiosyncracies of the observed values that do not extend to other values. For both these reasons, using too many parameters gives a misleading sense of complacency based on what appears like a good fit, but is a result of sheer chance.

For instance, *any* collection of n data points can be fitted (without need for error tolerance!) on a polynomial of degree at most $n - 1$. But what if there were actually measurement error? Then that polynomial of degree $n - 1$ would be a fake good fit. Imagine that the actual function is $f(x) = x$, but we are measuring values near 0 and the best fit for the measured values turns out to be $f(x) = x + x^3$. This may very well be a much better fit for values close to the origin because of biased measurement errors, but extrapolating it to a larger domain could go really awry.

Performance review: 15 out of 28 got this. 9 chose (E), 2 chose (B), 1 chose (A), and 1 did not attempt the question.

Historical note (last time): 14 out of 28 people got this. 9 chose (D), 4 chose (E), and 1 chose (A).

- (5) (*) F is an affine linear function of two variables x and y , i.e., it has the form $F(x, y) := ax + by + c$ with a , b , and c real numbers. We want to determine the values of the parameters a , b , and c by using input-output pairs. It is, however, costly to find input-output pairs. We have already found $F(1, 3)$ and $F(3, 7)$. We want to find F for one other pair of inputs to determine a , b , and c . Which of these will *not* be a good choice?
 - (A) $F(2, 2)$, i.e., the input $x = 2$, $y = 2$
 - (B) $F(2, 3)$, i.e., the input $x = 2$, $y = 3$
 - (C) $F(2, 4)$, i.e., the input $x = 2$, $y = 4$
 - (D) $F(2, 5)$, i.e., the input $x = 2$, $y = 5$
 - (E) $F(2, 6)$, i.e., the input $x = 2$, $y = 6$

Answer: Option (D)

Explanation: The point $(2, 5)$ is collinear with the points $(1, 3)$ and $(3, 7)$ (in fact, it is their midpoint) so the value at that point can be predicted based on the values at $(1, 3)$ and $(3, 7)$ as being the arithmetic mean between these values. Thus, it does not provide new information. Mathematically, if we use this point to get the third equation, that equation will be redundant with the existing equations. In equational form:

$$F(2, 5) = \frac{F(1, 3) + F(3, 7)}{2}$$

Purely arithmetic version of observation: We can obtain this observation computationally even without explicitly noting the observation about the midpoint. Explicitly, we have:

$$\begin{aligned}F(1, 3) &= a + 3b + c \\F(3, 7) &= 3a + 7b + c\end{aligned}$$

Adding, we get:

$$F(1, 3) + F(3, 7) = 4a + 10b + 2c = 2(2a + 5b + c) = 2F(2, 5)$$

Thus, we get:

$$F(2, 5) = \frac{F(1, 3) + F(3, 7)}{2}$$

Alternative geometric explanation: Think of the inputs as living in the xy -plane, and the output axis as the z -axis. The graph $z = F(x, y) = ax + by + c$ gives a plane in the three-dimensional space with coordinates x, y, z . We know that a plane is determined by knowing three non-collinear points on it. The points are of the form $(x, y, F(x, y))$ where x and y vary freely. The graph is a plane *because* F is a linear function. In general, the graph would be a surface.

The inputs $(1, 3)$, $(2, 5)$, and $(3, 7)$ being collinear, along with the fact that F is affine linear, tells us that the triples $(1, 3, F(1, 3))$, $(2, 5, F(2, 5))$, and $(3, 7, F(3, 7))$ are collinear in three-dimensional space. In fact, they lie on the line obtained by intersecting the plane that is the graph of F with the plane parallel to the z -axis whose intersection with the xy -plane passes through the points $(1, 3)$ and $(3, 7)$ (explicitly, this is the plane with equation $y = 2x + 1$). In this particular case, since $(2, 5)$ is the midpoint between the points $(1, 3)$ and $(3, 7)$, the point $(2, 5, F(2, 5))$ is the midpoint between the points $(1, 3, F(1, 3))$ and $(3, 7, F(3, 7))$.

Thus, if we use these three input-output pairs, then we get three *collinear* points in the plane we are trying to find, and we cannot determine the plane uniquely (any plane through the line joining the points works). If we chose an input (x_0, y_0) that was not collinear with the points $(1, 3)$ and $(3, 7)$, we would get a point $(x_0, y_0, F(x_0, y_0))$ that was not collinear with the points $(1, 3, F(1, 3))$ and $(3, 7, F(3, 7))$, and therefore, the plane would be determined uniquely.

Update for after understanding coefficient matrices and ranks: When trying to find the parameters a , b , and c , we need to set up a system of simultaneous linear equations. Each input-output pair gives one equation, and therefore, one row of the augmenting matrix. The augmenting column corresponds to the outputs, and the coefficient matrix part corresponds to the inputs. In other words, each input determines a row of the coefficient matrix.

For an input (x_i, y_i) , the corresponding equation is:

$$ax_i + by_i + c = F(x_i, y_i)$$

We are now viewing this as an equation in the variables a , b , and c . The row of the coefficient matrix when we set up the linear system in terms of the parameters is:

$$[x_i \quad y_i \quad 1]$$

For our system with three input-output pairs, the coefficient matrix becomes:

$$\begin{bmatrix} 1 & 3 & 1 \\ 3 & 7 & 1 \\ x & y & 1 \end{bmatrix}$$

where (x, y) is the third input that we choose. The input that would be bad to choose is the input for which the coefficient matrix has rank two rather than the expected rank of three. We can see that $x = 2, y = 5$ is the only such input. In other words, the matrix:

$$\begin{bmatrix} 1 & 3 & 1 \\ 3 & 7 & 1 \\ 2 & 5 & 1 \end{bmatrix}$$

has rank two. We can verify this by carrying out Gauss-Jordan elimination (the row reduction process) and obtaining that the last column is all zeros. On the other hand, the coefficient matrices for all the other choices of the third input have full rank three.

Performance review: 6 out of 28 got this. 10 chose (B), 7 chose (E), 3 chose (A), 2 chose (C).

Historical note (last time): 16 out of 28 got this. 9 chose (E), 2 chose (B), 1 chose (C).

**DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE FRIDAY OCTOBER 11:
GAUSS-JORDAN ELIMINATION (ORIGINALLY DUE WEDNESDAY OCTOBER 9,
BUT POSTPONED)**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

28 people took this 6-question quiz. The score distribution was as follows:

- Score of 3: 3 people
- Score of 4: 7 people
- Score of 5: 12 people
- Score of 6: 6 people

The mean score was 4.75.

The question-wise answers and performance review are as follows:

- (1) Option (B): 28 people (everybody!)
- (2) Option (B): 26 people
- (3) Option (A): 28 people (everybody!)
- (4) Option (B): 15 people
- (5) Option (C): 17 people
- (6) Option (D): 19 people

Note: This quiz was not administered the last time I taught the course, so there is no previous performance to compare against.

2. SOLUTIONS

PLEASE DO *NOT* DISCUSS ANY QUESTIONS

The quiz covers basics related to Gauss-Jordan elimination (notes titled **Gauss-Jordan elimination**, corresponding section in the book Section 1.2). Explicitly, the quiz covers:

- Setting up linear systems and interpreting the coefficient matrix in terms of the setup.
- Knowledge of the permissible rules for manipulating linear systems.
- Metacognition of the process of Gauss-Jordan elimination and its eventual result, the reduced row-echelon form, as well as the interpretation in terms of the solution set.

The questions are fairly easy questions if you understand the material. But it's important that you be able to answer them, otherwise what we study later will not make much sense.

- (1) *Do not discuss this!* The row operations that we can perform on the augmented matrix of a linear system include adding or subtracting another row. However, they do not include multiplying another row. In other words, suppose we start with:

$$\left[\begin{array}{cc|c} 1 & 2 & 5 \\ 2 & 7 & 6 \end{array} \right]$$

What we're not allowed to do is multiply row 2 by row 1 and obtain:

$$\left[\begin{array}{cc|c} 1 & 2 & 5 \\ 2 & 14 & 30 \end{array} \right]$$

What's the most compelling reason for our not being allowed to perform this operation?

- (A) The row operations arise from the corresponding operations on equations. For the “multiplication of rows” operation to be legitimate, it must correspond to multiplication of the corresponding equations, and multiplying equations is not a legitimate operation.
- (B) The row operations arise from the corresponding operations on equations. However, the “multiplication of rows” operation does not correspond to any legitimate operation on equations. Note that it does not correspond to multiplying the equations, because that is not how multiplication of linear polynomials work (in fact, if we multiplied the equations, we would end up with an equation that is not linear).

Answer: Option (B)

Explanation: In point of fact, both (A) and (B) are true in different ways, but (B) is more compelling. Let’s first understand why (B) is true. We will then turn to why (A) is (somewhat) true.

Every augmented matrix hides behind it a linear system. Let’s call the variables x_1 and x_2 . The original linear system is:

$$\begin{aligned}x_1 + 2x_2 &= 5 \\2x_1 + 7x_2 &= 6\end{aligned}$$

Now, if we multiply the two equations, we get:

$$(x_1 + 2x_2)(2x_1 + 7x_2) = 30$$

This simplifies to:

$$2x_1^2 + 11x_1x_2 + 14x_2^2 = 30$$

Note that this is quite different from the linear equation that is represented by the row obtained by multiplying the two rows. That equation is:

$$2x_1 + 14x_2 = 30$$

So (B) is the main reason why the operation makes no sense.

It’s worth noting that to a lesser extent, (A) is also an issue. Multiplying equations is legitimate: if two equations hold, their product also holds. However, replacing one of the equations by the product equation could lead to a potential loss of information. This loss of information occurs if the other equation being multiplied has both sides equal to zero (basically, multiplication by zero throws away information). Note that that problem does not occur with this linear system.

Performance review: All 28 got this.

- (2) *Do not discuss this!:* Consider a model where the functional form is linear in the parameters (though not necessarily in the inputs). We can use (input, output) pairs to set up a system of linear equations in the parameters. Given enough such equations, we can determine the values of the parameters.

What is the relation between the coefficient matrix and the parameters and (input, output) pairs?

- (A) The columns of the coefficient matrix correspond to the (input, output) pairs and the rows correspond to the parameters.
- (B) The rows of the coefficient matrix correspond to the (input, output) pairs and the columns correspond to the parameters.

Answer: Option (B)

Explanation: Each (input, output) pair gives an equation. Since the functional form is linear in the parameters, the equation is a linear equation. Each equation corresponds to a row of the augmented matrix (the input part affects the left side of the equation, and hence the coefficient matrix row, while the output part is the constant on the right side of the equation, i.e., the augmenting value).

The parameters are the variables that we are trying to solve for. These thus correspond to the columns.

Performance review: 26 out of 28 got this. 2 chose (A).

- (3) *Do not discuss this!*: Consider a model where the functional form is linear in the parameters (though not necessarily in the inputs). We can use (input, output) pairs to set up a system of linear equations in the parameters. Given enough such equations, we can determine the values of the parameters.

What is the relation between the inputs, the outputs, the coefficient matrix, and the augmenting column?

- (A) The inputs correspond to the coefficient matrix and the outputs correspond to the augmenting column. In other words, knowing the values of the inputs allows us to write down the coefficient matrix. Knowing the values of the outputs allows us to write down the augmenting column.
- (B) The outputs correspond to the coefficient matrix and the inputs correspond to the augmenting column. In other words, knowing the values of the outputs allows us to write down the coefficient matrix. Knowing the values of the inputs allows us to write down the augmenting column.

Answer: Option (A)

Explanation: Read the explanation for the preceding question.

Performance review: All 28 people got this.

- (4) *Do not discuss this!*: Consider the following rule to check for consistency using the augmented matrix: the system is inconsistent if and only if there is a zero row of the coefficient matrix with a nonzero value for that row in the augmenting column. In what sense does this rule work?

- (A) The rule can be applied to the augmented matrix directly in both the *if* and the *only if* direction.
- (B) The rule can be applied to the augmented matrix only in the *if* direction in general. In the *only if* direction, the rule can be applied to the augmented matrix *after* we have reduced the system to a situation where the coefficient matrix is in row-echelon form (note: it's not necessary to reach reduced row-echelon form).
- (C) The rule can be applied to the augmented matrix only in the *only if* direction in general. In the *if* direction, the rule can be applied to the augmented matrix *after* we have reduced the system to a situation where the coefficient matrix is in row-echelon form (note: it's not necessary to reach reduced row-echelon form).
- (D) The rule can be applied in either direction only *after* we have reduced the system to a situation where the coefficient matrix is in row-echelon form (note: it's not necessary to reach reduced row-echelon form).

Answer: Option (B)

Explanation: The rule obviously applies in the *if* direction: if there is any row with zeros in the coefficient matrix and a nonzero augmenting value, that equation has no solutions. Therefore, the system as a whole is inconsistent.

It does not work in the *only if* direction in general. This is because the inconsistency may be more subtle: it may arise due to equations being inconsistent *when viewed together*, rather than any individual equation being inconsistent. For instance, consider:

$$\begin{aligned}x + y &= 1 \\2x + 2y &= 3\end{aligned}$$

The augmented matrix is:

$$\left[\begin{array}{cc|c} 1 & 1 & 1 \\ 2 & 2 & 3 \end{array} \right]$$

Note that there is no zero row for the coefficient matrix. However, if we subtract twice the first row from the second row, we obtain:

$$\left[\begin{array}{cc|c} 1 & 1 & 1 \\ 0 & 0 & 1 \end{array} \right]$$

The coefficient matrix is now in rref, and we can now see that the system is inconsistent based on the second row (the coefficient matrix is all zeros, and the augmenting entry is nonzero).

Performance review: 15 out of 28 got this. 6 chose (D), 3 each chose (A) and (C), and 1 left the question blank.

- (5) *Do not discuss this!* Which of the following is *not* a possibility for the number of solutions to a system of simultaneous linear equations? Please see Options (D) and (E) before answering.
- (A) 0
 - (B) 1
 - (C) 2
 - (D) All of the above, i.e., none of them is a possibility
 - (E) None of the above, i.e., they are all possibilities

Answer: Option (C)

Explanation: 0 is a possibility that occurs when the system is inconsistent. 1 is a possibility that occurs when all the variables are leading variables and the system is consistent.

For there to be more than one solution, there must be a non-leading variable. This variable serves as a parameter that can take arbitrary real values. Thus, if there is more than one solution, there must be infinitely many solutions.

Performance review: 17 out of 28 got this. 9 chose (E), 1 each chose (B) and (D).

- (6) *Do not discuss this!* Which of the following describes the situation for a consistent system of simultaneous linear equations?
- (A) The leading variables are the parameters used to describe the general solution, and the number of leading variables equals the number of nonzero equations in the reduced row-echelon form (here nonzero equation makes an equation that does not have a zero row in the augmented matrix).
 - (B) The non-leading variables are the parameters used to describe the general solution, and the number of non-leading variables equals the number of nonzero equations in the reduced row-echelon form (here nonzero equation makes an equation that does not have a zero row in the augmented matrix).
 - (C) The leading variables are the parameters used to describe the general solution, and the number of leading variables equals the value (number of variables) - (number of nonzero equations in the reduced row-echelon form) (here nonzero equation makes an equation that does not have a zero row in the augmented matrix).
 - (D) The non-leading variables are the parameters used to describe the general solution, and the number of non-leading variables equals the value (number of variables) - (number of nonzero equations in the reduced row-echelon form) (here nonzero equation makes an equation that does not have a zero row in the augmented matrix).

Answer: Option (D)

Explanation: The non-leading variables serve as parameters. Also, the number of *leading variables* is the number of nonzero rows in the reduced row-echelon form. The number of non-leading variables is thus the total number of variables minus this quantity.

Note that we are given that the system is consistent, so we do not have to worry about the solution set being empty.

Performance review: 19 out of 28 got this. 3 each chose (A), (B) and (C).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY OCTOBER 11: LINEAR SYSTEMS

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

29 people took this 10-question quiz. The score distribution was as follows:

- Score of 2: 1 person
- Score of 3: 2 people
- Score of 4: 2 people
- Score of 5: 4 people
- Score of 6: 3 people
- Score of 7: 7 people
- Score of 8: 3 people
- Score of 9: 3 people
- Score of 10: 4 people

The mean score was 6.69.

The question-wise answers and performance review were as follows:

- (1) Option (B): 12 people
- (2) Option (A): 20 people
- (3) Option (B): 17 people
- (4) Option (C): 15 people
- (5) Option (D): 19 people
- (6) Option (D): 19 people
- (7) Option (B): 27 people
- (8) Option (A): 23 people
- (9) Option (C): 22 people
- (10) Option (B): 20 people

2. SOLUTIONS

The quiz questions here, although not hard *per se*, are conceptually demanding: answering them requires a clear understanding of multiple concepts and an ability to execute them conjunctively. Even if you feel that you've understood the material as presented in class, you will need to think through each question carefully. Some of the questions are related to similar homework problems (Homeworks 1 and 2), and they test a conceptual understanding of the solutions to these problems. You might want to view them in conjunction with the homework problems. Other questions sow the seeds of ideas we will explore later. The quiz should seem relatively easier when you review it later, assuming that you work hard on attempting the questions right now and read the solutions once they're put up.

- (1) (*) Rashid and Riena are trying to study a function f of two variables x and y . Rashid is convinced that the function is linear (i.e., it is of the form $f(x, y) := ax + by + c$) but has no idea what a , b , and c are. Riena thinks a linear model is completely out-of-place. Rashid is eager to find a , b , and c , whereas Riena is eager to disprove Rashid's linear model. Unfortunately, all they have is a black box that will output the value of the function for a given input pair (x, y) , and that black box can only be called three times. What should Rashid and Riena try for?
 - (A) Rashid and Riena would both like to provide three input pairs that are non-collinear as points in the xy -plane

- (B) Rashid would like to provide three input pairs that are non-collinear, while Riena would like to provide three input pairs that are collinear as points in the xy -plane.
- (C) Rashid and Riena would both like to provide three input pairs that are collinear as points in the xy -plane.
- (D) Rashid would like to provide three input pairs that are collinear, while Riena would like to provide three input pairs that are non-collinear as points in the xy -plane.
- (E) Both Rashid and Riena are indifferent regarding how the three input pairs are picked.

Answer: Option (B)

Explanation: Rashid wants to pick three independent constraints, i.e., pick three linear equations that are independent of each other. This would give him three independent equations in three variables, allowing him to uniquely determine the values of the parameters a , b , and c . Thus, he wishes to pick three non-collinear points. If he picked three collinear points, one of his equations would be redundant (deducible from the other two).

Riena, on the other hand, wants to pick three points with some redundancy between them. If she picks three collinear points, then, assuming Rashid's model is correct, the values of the outputs at two of the three points determine the value of the output at the third. If it turns out that the actual value differs from the predicted value, this *falsifies* Rashid's model.

The reason Riena wants them to be collinear is thus the same as the reason Rashid wants them to be non-collinear. Rashid already thinks he knows the predicted result from the collinear case, so he is not interested in it. Riena wants to show Rashid that his prediction is wrong, so she cares about it. The non-collinear case would not be interesting to Riena because it offers no evidence against Rashid.

Note that if 4 or more inputs can be queried, then both Rashid's and Riena's concerns can be addressed. *Scarcity breeds conflict.*

More about this is discussed in the notes on **Hypothesis testing, rank, and overdetermination** that we intend to discuss next week.

Not clear to you?: Pick some function of two variables that is not linear, and put yourself in Riena's shoes. How would you convince Rashid that the function is not linear? What types of inputs would you choose? Similarly, put yourself in Rashid's shoes and try to determine what types of inputs will help you determine uniqueness.

Performance review: 12 out of 29 got this. 7 each chose (A) and (D), 2 chose (E), 1 chose (C).

Historical note (last time): 18 out of 29 got this. 6 chose (D), 4 chose (C), 1 chose (A).

- (2) (*) Let m and n be natural numbers with $m \geq 3$. We are given a bunch of numbers $x_1 < x_2 < \dots < x_m$ and another bunch of numbers y_1, y_2, \dots, y_m . We want to find a continuous function f on $[x_1, x_m]$, such that $f(x_i) = y_i$ for all $1 \leq i \leq m$, and such that the restriction of f to any interval of the form $[x_i, x_{i+1}]$ (for $1 \leq i \leq m - 1$) is a polynomial of degree $\leq n$. What is the smallest value of n for which we are guaranteed to be able to find such a function f ?
- (A) 1
 - (B) 2
 - (C) 3
 - (D) 4
 - (E) 5

Answer: Option (A)

Explanation: The brief explanation is that we can construct a piecewise linear function: a function whose graph comprises straight line segments joining (x_i, y_i) to (x_{i+1}, y_{i+1}) .

The more sophisticated explanation is that for each pair of adjacent points, we have to choose a functional form between the two points, and there are two input-output pairs constraining the functional form (namely, the two endpoints). We thus need a functional form with two parameters. The linear functional form works well for that purpose. If we were to count the *total* number of parameters, it would work out to $2(m - 1)$ (2 parameters for each of the $m - 1$ intervals (x_1, x_2) , (x_2, x_3) , \dots , (x_{m-1}, x_m)). The number of equations would also work out to $2(m - 1)$. The number of parameters equals the number of equations. However, the equality of the number of parameters and the number of equations in the abstract is not *completely sufficient* to argue that we can always

find a solution, because of concerns about redundant equations. In this case, it is sufficient, based on the explanation offered earlier. It is still a good working heuristic, though.

Note on subtle difference in goals from previous situations: The situation of this question, and the next three questions, is different in a subtle but important way from earlier situations. Broadly, consider two types of situations.

- (a) *Prediction/forecasting/extrapolation situation:* The situation where we are trying to model some pre-existing phenomenon and we have a model with a general functional form (involving not-yet-known parameters) that we believe applies to the phenomenon. We are trying to find the parameters using input-output pairs. In this case, we are trying to find information about a function that, in some sense, already exists. This means that our job isn't just to find *some function that fits the known data*, but to find *the right function*. In this situation, having more input-output pairs (and hence equations) than parameters (that become our new variables) is desirable.
- (b) *Engineering solution situation:* The situation where our goal is just to find some function that fits the existing data, rather than to find a particular function. In this case, having more variables than equations, resulting in the system being underdetermined (i.e., the solution being non-unique and having degrees of freedom) is a *good* thing.

The type of situation we are dealing with in these questions is type (b). We are interested in just finding some function satisfying certain constraints, not in finding a specific function that already exists.

Note that all situations that involve prediction and trend forecasting fall under (a) rather than (b). Situation (b) arises in cases where we are solving specific engineering problems to come up with one-time constructs that satisfy constraints (such as the roller-coaster ride example in the book). From the social science perspective, the more common situation is situation (a), where we are trying to predict or forecast for an existing model that we do not yet fully understand.

Not clear to you?: Make a picture with points marked for the values of the function at x_1, x_2, \dots, x_m . Now notice that we can make a graph that uses a straight line segment for each interval between adjacent x_i s.

Performance review: 20 out of 29 people got this. 8 chose (B), 1 chose (C).

Historical note (last time): 17 out of 29 got this. 10 chose (C), 2 chose (D).

- (3) (*) Let m and n be natural numbers with $m \geq 3$. We are given a bunch of numbers $x_1 < x_2 < \dots < x_m$ and another bunch of numbers y_1, y_2, \dots, y_m . We want to find a continuous function f on $[x_1, x_m]$, such that $f(x_i) = y_i$ for all $1 \leq i \leq m$, and such that the restriction of f to any interval of the form $[x_i, x_{i+1}]$ (for $1 \leq i \leq m - 1$) is a polynomial of degree $\leq n$. In addition, we want to make sure that f is differentiable on the open interval (x_1, x_m) . What is the smallest value of n for which we are guaranteed to be able to find such a function f ?
 - (A) 1
 - (B) 2
 - (C) 3
 - (D) 4
 - (E) 5

Answer: Option (B)

Explanation: We can take the first piece to be linear (simply join the points (x_1, y_1) and (x_2, y_2)). For each subsequent piece, we have three constraints. For the i^{th} piece, the constraints include the value at x_i , the value at x_{i+1} , and the derivative at x_i (which must equal the corresponding derivative from the preceding piece). In order to be able to fit all three constraints, we need to use $n = 2$ so as to have three parameters. It is also straightforward to check that the system we get this way is not redundant.

Here is the accounting regarding parameters and equations. The number of equations is $2(m - 1) + (m - 2)$. The $2(m - 1)$ equations arise from the values at the endpoints of each of the $m - 1$ intervals. The extra $m - 2$ equations arise from the equality of the formal expressions for the left hand derivative and right hand derivative at each of the $m - 2$ transition points x_2, x_3, \dots, x_{m-1} . The total number of equations is therefore $3m - 4$. Whatever value of n we choose, the total number

of parameters is $(n + 1)(m - 1)$ (we need $n + 1$ parameters for the piece definition on each of the $m - 1$ intervals $(x_1, x_2), (x_2, x_3), \dots, (x_{m-1}, x_m)$). We want to choose n such that the number of parameters is greater than or equal to the number of equations, and the smallest value that works is $n = 2$, in which case we get $(n + 1)(m - 1) = 3m - 3 \geq 3m - 4$. Note that we have one extra parameter over the number of equations, so we have a bit of “slack” here, at least in principle. Again, the crude comparison of parameters and equations is not conclusive because of the potential for redundancy, so we need an explanation of the sort offered in the preceding paragraph to be sure. Note that the slight excess of parameters over equations corresponds to the fact that for the piece definition in the very first interval (x_1, x_2) , we have some flexibility.

Note on subtle difference in goals from previous situations: See the note at the end of the answer to Question 2.

Not clear to you?: Make a picture similar to the preceding question, but notice now that from the second interval onward, you have a constraint on the initial slope.

Performance review: 17 out of 29 got this. 10 chose (C), 2 chose (A).

Historical note (last time): 25 out of 29 got this. 2 chose (C), 1 each chose (A) and (D).

- (4) (*) Let m and n be natural numbers with $m \geq 3$. We are given a bunch of numbers $x_1 < x_2 < \dots < x_m$ and another bunch of numbers y_1, y_2, \dots, y_m . We want to find a continuous function f on $[x_1, x_m]$, such that $f(x_i) = y_i$ for all $1 \leq i \leq m$, and such that the restriction of f to any interval of the form $[x_i, x_{i+1}]$ (for $1 \leq i \leq m - 1$) is a polynomial of degree $\leq n$. In addition, we want to make sure that f is differentiable on the open interval (x_1, x_m) . In addition, we are told the value of the right hand derivative of f at x_1 and the left hand derivative of f at x_m . What is the smallest value of n for which we are guaranteed to be able to find such a function f ?

- (A) 1
 (B) 2
 (C) 3
 (D) 4
 (E) 5

Answer: Option (C)

Explanation: This is similar to the cubic spline question on your advanced homework.

The key reason why this differs from the preceding question is that, here, we have a specification of the one-sided derivatives at the endpoint. The specification at x_1 constrains how we start, but that is still acceptable. We’ll need something quadratic to begin with, but we can keep choosing quadratics all the way, since at each piece we need to satisfy three conditions. The problem occurs in the final stage, i.e., the interval $[x_{m-1}, x_m]$. Here, we have *four* constraints: the values at the endpoints, and the derivatives at both endpoints. So, we need a functional form with four parameters. In other words, we need a cubic for the last piece.

It would be acceptable to use quadratics for all except the last piece, and a cubic for the last piece. The options available to us, however, require us to declare a single degree that works for all, so the smallest available is 3.

Using the accounting for parameters and equations similar to the preceding question, we obtain $3m - 2$ equations (the $3m - 4$ equations of the preceding question, plus 2 additional equations for the one-sided derivative conditions at the endpoints). We therefore need to choose n such that $(n + 1)(m - 1) \geq 3m - 2$. The value $n = 2$ *just* falls short: it would give $3m - 3$ parameters, and this corresponds to the fact that if we tried to fit a quadratic, we could run into trouble in the very last piece. The value $n = 3$, on the other hand, offers us considerable slack.

Note on subtle difference in goals from previous situations: See the note at the end of the answer to Question 2.

Performance review: 15 out of 29 got this. 7 chose (B), 3 chose (A), 2 chose (D), 1 chose (E), 1 left the question blank.

Historical note (last time): 11 out of 29 got this. 10 chose (B), 6 chose (D), 2 chose (A).

- (5) (*) Let k, m , and n be natural numbers with $m \geq 3$. We are given a bunch of numbers $x_1 < x_2 < \dots < x_m$ and another bunch of numbers y_1, y_2, \dots, y_m . We want to find a continuous function f on $[x_1, x_m]$, such that $f(x_i) = y_i$ for all $1 \leq i \leq m$, and such that the restriction of f to any interval of

the form $[x_i, x_{i+1}]$ (for $1 \leq i \leq m - 1$) is a polynomial of degree $\leq n$. In addition, we want to make sure that f is at least k times differentiable on the open interval (x_1, x_m) . What is the smallest value of n for which we are guaranteed to be able to find such a function f ?

- (A) $k - 2$
- (B) $k - 1$
- (C) k
- (D) $k + 1$
- (E) $k + 2$

Answer: Option (D)

Explanation: For the part $[x_1, x_2]$, we can simply fit a straight line. For each subsequent piece, we have $k + 1$ constraints: the left and right endpoint values, and all the derivative values up to the k^{th} derivative at the left endpoint. This gives a total of $k + 2$ constraints. A polynomial of degree $k + 1$ has the necessary number of parameters, namely $k + 2$. We can also use Taylor polynomials to see that the system can always be solved.

If we account for the total number of parameters and equations, the numbers we get are as follows. We have $2(m - 1) + k(m - 2)$ equations. The $2(m - 1)$ equations are from the values at the endpoints of the $m - 1$ intervals. The $k(m - 2)$ equations are from the equality of the first k derivatives of the piece functions at the transition points x_2, x_3, \dots, x_{m-1} . The total number of equations is $(k + 2)(m - 1) - k$. We have $(n + 1)(m - 1)$ parameters. Therefore, choosing $n = k + 1$ gives more parameters than equations. On the other hand, choosing $n = k$ gives $(k + 1)(m - 1)$ parameters, so that the number of parameters - the number of equations is $k - (m - 1)$. Since we are not given any concrete information about the sign relationship between k and $m - 1$, we cannot be sure that this will work.

Note on subtle difference in goals from previous situations: See the note at the end of the answer to Question 2.

Performance review: 19 out of 29 got this. 5 chose (B), 3 chose (E), 2 chose (C).

Historical note (last time): 12 out of 29 got this. 14 chose (B), 2 chose (C), 1 chose (A).

The next few questions are framed deterministically, though similar real-world applications would be probabilistic, with some square roots floating around. Unfortunately, we do not have the tools yet to deal with the probabilistic versions of the questions.

- (6) (*) A function f of one variable is known to be linear. We know that $f(1) = 1.5 \pm 0.5$ and $f(2) = 2.5 \pm 0.5$. Assume these are the full ranges, without any probability distribution known. Assuming nothing is known about how the measurement errors for f at different points are related, what can we say about $f(3)$?
- (A) $f(3) = 3.5$ (exactly)
 - (B) $f(3) = 3.5 \pm 0.5$
 - (C) $f(3) = 3.5 \pm 1$
 - (D) $f(3) = 3.5 \pm 1.5$
 - (E) $f(3) = 3.5 \pm 2.5$

Answer: Option (D)

Explanation: The highest value occurs if we take $f(1) = 1$ and $f(2) = 3$, giving $f(3) = 5$. The lowest value occurs if we take $f(1) = f(2) = 2$, giving $f(3) = 2$. Overall, $f(3)$ is between 2 and 5, so 3.5 ± 1.5 is a reasonable description.

Graphical interpretation: Make a picture of the coordinate xy -plane with a vertical line segment joining the points with coordinates $(1, 1)$ and $(1, 2)$ (depicting the range of possibilities for $f(1)$) and another line segment joining the points with coordinates $(2, 2)$ and $(2, 3)$ (depicting the range of possibilities for $f(2)$). The actual value of $f(1)$ could be anywhere on the first line segment and the actual value of $f(2)$ could be anywhere on the second line segment. The slowest growth case is the case where $f(1) = 2$ and $f(2) = 2$, so in fact we get a constant function $f(x) = 2$ in this case, and the graph of this is the line $y = 2$, and $f(3) = 2$ in this case.. The fastest growth case is the case $f(1) = 1$ and $f(2) = 3$, and we get the function $f(x) = 2x - 1$ in this case, and the graph is the line $y = 2x - 1$. We get $f(3) = 5$ in this case. Thus, $f(3)$ could be any value between 2 and 5,

or equivalently, the range of values for $f(3)$ is given by the line segment in the xy -plane joining the points $(3, 2)$ and $(3, 5)$.

The picture might remind you of eclipses. The region in the “shadow” so to speak is the penumbral region of the eclipse.

Performance review: 19 out of 29 got this. 6 chose (B), 3 chose (C), 1 chose (A).

Historical note (last time): 25 out of 29 got this. 3 chose (C), 1 chose (B).

- (7) (*) A function f of one variable is known to be linear. We know that $f(1) = 1.5 \pm 0.5$ and $f(2) = 2.5 \pm 0.5$. Assume these are the full ranges, without any probability distribution known. Assume also that the measurement error for f at all points is the same in magnitude and sign. What can we say about $f(3)$?
- (A) $f(3) = 3.5$ (exactly)
 - (B) $f(3) = 3.5 \pm 0.5$
 - (C) $f(3) = 3.5 \pm 1$
 - (D) $f(3) = 3.5 \pm 1.5$
 - (E) $f(3) = 3.5 \pm 2.5$

Answer: Option (B)

Explanation: $f(1) = 1.5 + m$ where m is measurement error with $|m| \leq 0.5$. Since all measurement errors are equal, $f(2) = 2.5 + m$. Thus, $f(3) = 3.5 + m$, with $|m| \leq 0.5$, giving the answer.

Graphical interpretation: The vertical line segments are the same as before: one joins $(1, 1)$ and $(1, 2)$, the other joins $(2, 2)$ and $(2, 3)$. However, since we now know that the errors are the same, the lower bounding line connects the lowest ends of both vertical line segments, and the upper bounding line connects the upper ends of both vertical line segments. In other words, we get a pair of parallel lines. The lower line passes through $(3, 3)$ and the upper line passes through $(3, 4)$. The range of possibilities for $f(3)$ is therefore the set of values between 3 and 4, i.e., it is 3.5 ± 0.5 .

Performance review: 27 out of 29 got this. 2 chose (C).

Historical note (last time): 26 out of 29 got this. 2 chose (D), 1 chose (C).

- (8) (*) Suppose f is a linear function on a bounded interval $[a, b]$ but our measurement of outputs for given inputs has some measurement error (with the range of measurement error the same regardless of the input, and no known correlation between the magnitude of measurement error at different points). Assume we can get the outputs for any two specified inputs we desire, and we will then fit a line through the (input,output) pairs to get the graph of f . How should we choose our inputs?
- (A) Choose the inputs as far as possible from each other, i.e., choose them as a and b .
 - (B) Choose the inputs to be as close to each other as possible, i.e., choose them to be nearby points but not equal to each other.
 - (C) It does not matter. Any choice of two distinct inputs is good enough.

Answer: Option (A)

Explanation: If the inputs are chosen close together, then even small errors in the input can cause large errors in the measurement of the slope. The error in slope is the signed difference of measurement errors divided by the distance between the inputs. The former is beyond our control, because we noted that the magnitude and sign of measurement error does not depend on where we choose the inputs. Thus, choosing the inputs as far as possible brings us a larger denominator, and therefore keeps the slope error at a minimum.

Performance review: 23 out of 29 got this. 4 chose (B), 2 chose (C).

Historical note (last time): 11 out of 29 got this. 15 chose (C), 3 chose (B).

- (9) (*) f is a function of one variable defined on an interval $[a, b]$. You are trying to find an explicit function that fits f well. You initially try a straight line fit that works at the points a and b . It turns out that this fit systematically overestimates f for points in between (i.e., the actual function f is below the linear function) with the maximum magnitude of discrepancy occurring at the midpoint $(a + b)/2$. Based on this information, what kind of fit should you try to look for?
- (A) Try to fit f using a logarithmic function
 - (B) Try to fit f using an exponential function
 - (C) Try to fit f using a quadratic function
 - (D) Try to fit f using a polynomial of degree at most 3

(E) Try to fit f using the reciprocal of a linear function

Answer: Option (C)

Explanation: Let L be the linear function obtained and let $g = f - L$. g is zero at both endpoints a and b , below zero in between, and has its minimum at the midpoint. These characteristics strongly suggest that g is a quadratic function with positive leading coefficient. Any quadratic function with positive leading coefficient that has zeros has its absolute minimum precisely at the midpoint between its zeros.

Since g is expected to be quadratic, $f = g + L$ is also expected to be quadratic.

Note that we could try fitting using a polynomial of degree at most 3. This, however, might run us into overfitting problems. As a general rule, we should try to fit using a function with as few parameters as possible and where the functional form is justified by broad theoretical considerations. If after fitting the quadratic, we discover some systemic errors that are best explained by a cubic type of discrepancy, we could then try a cubic.

Performance review: 22 out of 29 got this. 3 each chose (B) and (D), 1 chose (A).

Historical note (last time): 25 out of 29 got this. 3 chose (B), 1 chose (D).

(10) (*) Recall the Leontief input-output model. Recall that the GDP is defined as the total money value of all the *final* goods and services produced in the economy, which in this case means only those that go into meeting consumer demand, not interindustry demand (note that we are assuming away the existence of investment and government spending, which complicate the GDP calculation). Assuming that the unit prices of the goods are constant (a very unrealistic assumption given that price itself responds to supply and demand, but fortunately it does not affect the conclusion we draw here) what might be a way of increasing GDP while keeping the magnitude of output of each industry the same?

(A) Increase interindustry dependence, i.e., increase the amount needed from each industry that is necessary to produce a given amount in another industry.

(B) Reduce interindustry dependence, i.e., reduce the amount needed from each industry that is necessary to produce a given amount in another industry.

(C) Changes in interindustry dependence have no effect.

Answer: Option (B)

Explanation: The lower the interindustry dependence, the larger the share of the industries' output that can be used to meet consumer demand, i.e., contribute to GDP.

Performance review: 20 out of 29 got this. 6 chose (C), 3 chose (A).

Historical note (last time): 26 out of 29 got this. 3 chose (C).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE MONDAY OCTOBER 14: MATRIX COMPUTATIONS

MATH 196, SECTION 57 (VIPUL NAIK)

PLEASE DO *NOT* DISCUSS ANY QUESTIONS EXCEPT THE STARRED OR DOUBLE-STARRED QUESTIONS.

1. PERFORMANCE REVIEW

28 people took this 5-question quiz. The score distribution was as follows:

- Score of 1: 1 person
- Score of 2: 3 people
- Score of 3: 6 people
- Score of 4: 11 people
- Score of 5: 7 people

The question-wise answers and performance review are below:

- (1) Option (C): 27 people
- (2) Option (B): 25 people
- (3) Option (B): 13 people
- (4) Option (D): 18 people
- (5) Option (B): 21 people

2. SOLUTIONS

This quiz has a few questions on the mechanics of the computational execution of Gauss-Jordan elimination, and it has one question on setting up a linear system.

Suppose f is a function on the positive integers that takes positive integer values. Suppose n is a parameter related to the input size of an algorithm. We say that the running time of an algorithm (respectively, the space requirement of the algorithm) is:

- $O(f(n))$ if, for large enough n , it can be bounded from above by a positive constant times $f(n)$.
- $\Omega(f(n))$ if, for large enough n , it can be bounded from below by a positive constant times $f(n)$.
- $\Theta(f(n))$ if it is both $O(f(n))$ and $\Omega(f(n))$.

You can read more at:

http://en.wikipedia.org/wiki/Big_O_notation

- (1) (*) If you treat each arithmetic operation (addition, subtraction, multiplication, division) of numbers as taking constant time, and all entry rewrites and changes as again taking constant time per entry, what would be the best description of the worst-case running time of the algorithm to convert a $n \times n$ matrix to reduced row-echelon form? (Note that this complexity is termed *arithmetic complexity* and can be distinguished from the *bit complexity* of the algorithm, which could be considerably higher).
 - (A) $\Theta(n)$
 - (B) $\Theta(n^2)$
 - (C) $\Theta(n^3)$
 - (D) $\Theta(n^4)$
 - (E) $\Theta(n^5)$

Answer: Option (C)

Explanation: We have $\Theta(n^2)$ row operations, and the row operations all take $\Theta(n)$ time. Overall, we get $\Theta(n^3)$ as the arithmetic complexity.

More can be found in the lecture notes on Gauss-Jordan elimination. You can also learn more about the arithmetic complexity of Gaussian elimination by looking it up online.

Performance review: 27 out of 28 got this. 1 chose (B).

Historical note (last time): 24 out of 27 got this. 1 each chose (A), (B), and (E).

- (2) (*) If you treat each arithmetic operation (addition, subtraction, multiplication, division) of numbers as taking constant space, and all matrix entries as taking constant space, what would be the best description of the worst-case space requirement of the algorithm to convert a $n \times n$ matrix to reduced row-echelon form? Assume that space is reusable, i.e., it is possible to rewrite over existing space used.

- (A) $\Theta(n)$
- (B) $\Theta(n^2)$
- (C) $\Theta(n^3)$
- (D) $\Theta(n^4)$
- (E) $\Theta(n^5)$

Answer: Option (B)

Explanation: We can reuse the matrix space to keep rewriting over existing entries. We need some additional workspace for working memory, but this is quite small relative to the size of the matrix, and does not affect the order estimate.

More in the lecture notes on Gauss-Jordan elimination.

Performance review: 25 out of 28 got this. 2 chose (D), 1 chose (C).

Historical note (last time): 26 out of 27 got this. 1 chose (A).

- (3) (*) Suppose the coefficient matrix of a linear system with n variables and n equations is known in advance, and we can spend as much time processing it as we desire in advance (this time will not count towards the running time of the algorithm). In other words, we can use Gauss-Jordan elimination to row-reduce the coefficient matrix in advance. However, we do not have the output column with us in advance. What is the worst-case running time of the part of the algorithm that runs after the output column is known?

- (A) $\Theta(n)$
- (B) $\Theta(n^2)$
- (C) $\Theta(n^3)$
- (D) $\Theta(n^4)$
- (E) $\Theta(n^5)$

Answer: Option (B)

Explanation: If we store the sequence of row operations used to convert the coefficient matrix to reduced row-echelon form (there are $\Theta(n^2)$ such operations) we simply need to apply these operations to the output column, then read off the solutions. The arithmetic time complexity of this is $\Theta(n^2)$.

Performance review: 13 out of 28 got this. 11 chose (A), 3 chose (C), 1 chose (D).

Historical note (last time): 17 out of 27 got this. 5 chose (A), 4 chose (C), 1 chose (E).

- (4) Which of the following matrices does *not* have the identity matrix as its reduced row-echelon form?
(A)

$$\begin{bmatrix} 2 & 0 & 0 \\ 0 & 5 & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

(B)

$$\begin{bmatrix} 1 & 2 & 3 \\ 0 & 3 & 5 \\ 0 & 0 & 7 \end{bmatrix}$$

(C)

$$\begin{bmatrix} 4 & 0 & 0 \\ 3 & 1 & 0 \\ 0 & 5 & -6 \end{bmatrix}$$

(D)

$$\begin{bmatrix} 1 & 2 & -3 \\ 4 & -3 & -1 \\ -2 & 1 & 1 \end{bmatrix}$$

(E)

$$\begin{bmatrix} 1 & 2 & 3 \\ 2 & 4 & 7 \\ 3 & 7 & 11 \end{bmatrix}$$

Answer: Option (D)

Explanation: We can work to convert to rref to verify this, but one easy way of seeing that this matrix does not have full rank is to note that each row sum is 0, which indicates that $\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$ are both solutions to the system of linear equations with this as the coefficient matrix and outputs all zeros. In other words, the solution space is not zero-dimensional, and hence, the matrix does not have full rank.

As for the other options:

- Option (A): This is a diagonal matrix with all entries nonzero. Thus, its rref is the identity matrix. It is easy to see this: just divide each row by its diagonal entry.
- Option (B): This is upper-triangular. We can first convert the diagonal entries to 1 by dividing. Then, we can do row subtractions and clear out everything above the diagonal.
- Option (C): This is lower-triangular, and works for a reason similar to Option (B).
- Option (E): This actually needs to be worked out using row reduction.

Performance review: 18 out of 28 got this. 4 chose (E), 2 each chose (B) and (C), 2 wrote incorrect free-form responses.

Historical note (last time): 24 out of 27 got this. 1 each chose (B), (C), and (E).

- (5) A number of different consumer price indices have been constructed. All of them use the market prices for an existing collection of commodities (though not all of them use every commodity in the collection) and take a different “weighted” linear combination of those. For instance, one price index might be 3 times (the price per ton of wheat on the Chicago wheat market) + 4 times (the price of 1 gallon of unleaded gasoline at a particular gas station) + 17 times (the price of Burt’s chapstick). Another price index might use 30 times (the price of Transcend’s 32 GB flash drive) + 14 times (the price of 1 gallon of gasoline at a particular gas station).

What is a good way of modeling these?

- (A) The prices of the various goods are stored in a matrix, the different weightings used in various indices are stored in a vector, and the consumer price indices arise as the output vector of the matrix-vector product.
- (B) The different weightings used in various indices are stored in a matrix, the prices of the various goods are stored in a vector, and the consumer price indices arise as the output vector of the matrix-vector product.
- (C) The prices of the various goods are stored in a matrix, the consumer price indices are stored as a vector, and the weightings used in the indices arise as the output vector of the matrix-vector product.
- (D) The different weightings used in various indices are stored in a matrix, the consumer price indices are stored in a vector, and the prices of the various goods arise as the output vector of the matrix-vector product.
- (E) The consumer price indices are stored in a matrix, the prices of the various goods are stored in a vector, and the weightings used in the indices arise as the output vector of the matrix-vector product.

Answer: Option (B)

Explanation: Each row represents the weightings used for a particular index. The input column is the input vector of prices. The output column is the vector of the various index values.

Performance review: 21 out of 28 got this. 7 chose (A).

Historical note (last time): 22 out of 27 got this. 4 chose (A), 1 chose (C).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY OCTOBER 18: LINEAR SYSTEMS: RANK AND DIMENSION CONSIDERATIONS

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

27 people took this 8-question quiz. The score distribution was as follows:

- Score of 3: 1 person
- Score of 4: 4 people
- Score of 5: 7 people
- Score of 6: 6 people
- Score of 7: 8 people
- Score of 8: 1 person

The question-wise answers and performance review are as follows:

- (1) Option (E): 24 people
- (2) Option (B): 21 people
- (3) Option (E): 10 people
- (4) Option (E): 24 people
- (5) Option (B): 24 people
- (6) Option (B): 20 people
- (7) Option (A): 17 people
- (8) Option (C): 14 people

2. SOLUTIONS

The questions here consider a wide range of theoretical and practical settings where linear systems appear, and prompt you to think about the notion of rank and its relationship with whether we can uniquely acquire the information that we want. It relates approximately with the material in the **Linear systems and matrix algebra** notes (the corresponding section in the book is Section 1.3).

- (1) (*) Let m and n be positive integers. It turns out that *almost all* $m \times n$ matrices over the real numbers have a particular rank. What is that rank? (Unfortunately, it is beyond our current scope to define “almost all”).
 - (A) m (regardless of whether m or n is bigger)
 - (B) n (regardless of whether m or n is bigger)
 - (C) $(m + n)/2$
 - (D) $\max\{m, n\}$
 - (E) $\min\{m, n\}$

Answer: Option (E)

Explanation: In some sense, this is the only viable option, because we know that the rank is at *most* $\min\{m, n\}$. The intuitive reason why the theoretical maximum is usually attained is that “random things are most likely to be unrelated.”

Performance review: 24 out of 27 got this. 2 chose (A), 1 chose (D).

Historical note (last time): 23 out of 25 got this. 1 each chose (C) and (D).

- (2) (*) A container has a mix of two known gases that do not react with each other. The temperature and pressure of the container are known. Assume that $PV = nRT$. The volume of the container is also known, and so is the total mass of the gases in the container. Under what conditions can we

predict the amount (say, in the form of the number of moles) of each gas that is present from this information?

- (A) It is possible if both gases have the same molecular mass, because in that case, the coefficient matrix of the linear system has full rank 2.
- (B) It is possible if both gases have different molecular masses, because in that case, the coefficient matrix of the linear system has full rank 2.
- (C) It is possible if both gases have the same molecular mass, because in that case, the coefficient matrix of the linear system has rank 1.
- (D) It is possible if both gases have different molecular masses, because in that case, the coefficient matrix of the linear system has rank 1.
- (E) It is not possible to deduce the amount of each gas from the given information.

Answer: Option (B)

Explanation: Let n_1 and n_2 be the number of moles of each gas. Let m_1 and m_2 be their respective molecular masses. Then, we have two equations, where M is the total mass:

$$\begin{aligned}n_1 + n_2 &= PV/(RT) \\ n_1 m_1 + n_2 m_2 &= M\end{aligned}$$

The coefficient matrix for this is:

$$\begin{bmatrix} 1 & 1 \\ m_1 & m_2 \end{bmatrix}$$

This has rank one iff $m_1 = m_2$, and rank two otherwise.

If the coefficient matrix has rank one, then the second equation is redundant (note that it must be consistent since we are getting these numbers from an actual situation). In other words, we get a one-dimensional solution space, with a freely varying parameter. The nonnegativity of the number of moles of each gas does constrain the parameter to a closed bounded interval instead of all reals, but it still has infinitely many candidate values.

Performance review: 21 out of 27 got this. 4 chose (A), 2 chose (C).

Historical note (last time): 17 out of 25 got this. 4 chose (C), 3 chose (E), 1 chose (A).

- (3) (*) A container has a mix of three known gases with no reactions between the gases. The temperature and pressure of the container are known. Assume that $PV = nRT$. The volume of the container is also known, and so is the total mass of the gases in the container. Under what conditions can we predict the amount (say, in the form of the number of moles) of each gas that is present from this information?
- (A) It is possible if all three gases have the same molecular mass, because in that case, the coefficient matrix of the linear system has full rank 3.
 - (B) It is possible if all three gases have different molecular masses, because in that case, the coefficient matrix of the linear system has full rank 3.
 - (C) It is possible if all three gases have the same molecular mass, because in that case, the coefficient matrix of the linear system has rank 2.
 - (D) It is possible if all three gases have different molecular masses, because in that case, the coefficient matrix of the linear system has rank 2.
 - (E) It is not possible to deduce the amount of each gas from the given information.

Answer: Option (E)

Explanation: We have two equations in three variables. Explicitly, if the total mass is M , the number of moles of the three gases are n_1 , n_2 , and n_3 , and the molecular masses are m_1 , m_2 , and m_3 , then we have:

$$\begin{aligned}n_1 + n_2 + n_3 &= PV/(RT) \\ n_1m_1 + n_2m_2 + n_3m_3 &= M\end{aligned}$$

The coefficient matrix is:

$$\begin{bmatrix} 1 & 1 & 1 \\ m_1 & m_2 & m_3 \end{bmatrix}$$

This matrix can have rank either one or two. The rank is one if all three gases have the same molecular mass. The rank is two otherwise. Note that in either case, the rank is less than three, i.e., the system does not have full column rank. Thus, it will not be possible to solve the system and uniquely determine the values of n_1 , n_2 , and n_3 .

What if we include the nonnegativity constraints? Even in the presence of these constraints, a unique solution is not possible if there are nonzero amounts of all gases in the actual solution.

Performance review: 10 out of 27 got this. 10 chose (B), 4 chose (C), 3 chose (A).

Historical note (last time): 10 out of 25 got this. 13 chose (B), 1 each chose (A) and (C).

The branch of chemistry called quantitative analysis has historically used stoichiometric methods to determine the proportions of various chemicals present in a given mix. The idea is to use information about the amounts needed and produced in various reactions to estimate the quantities of chemicals present (the possible chemicals are first identified via “qualitative analysis” techniques). We generally find that these conditions give linear systems, and the coefficient matrices of these systems have (or can be written in a manner as to have) small integer entries.

- (4) (*) Consider a situation where we have a material that is a mix (in fixed proportion) of three known chemicals X , Y , and Z . Our goal is to find the amount of X , Y , and Z present. Suppose we want to set up a collection of experiments so that the coefficient matrix is diagonal, i.e., we are effectively solving a diagonal system of equations and can recover the quantities of each of X , Y , and Z . Which of the following is the best approach? Assume that we can measure, for each reagent, the amount of the reagent that gets used up for the reaction(s) to proceed to completion, but cannot isolate or separate the outputs from each other.
- (A) Choose a single reagent that reacts with all of X , Y , and Z .
 - (B) Choose a single reagent that reacts with only one of X , Y , and Z .
 - (C) Choose three separate reagents, each of which reacts with *all* of X , Y , and Z .
 - (D) Choose three separate reagents, each of which reacts only with X .
 - (E) Choose three separate reagents, one of which reacts only with X , one of which reacts only with Y , and one of which reacts only with Z .

Answer: Option (E)

Explanation: Consider a matrix with the columns indexed by X , Y , and Z and with the rows indexed by the reagents. The entry in each cell of the matrix is the amount of the row reagent needed to react with a unit amount of the column substance. It turns out that this is the coefficient matrix of the system of simultaneous linear equations that we construct.

In order to get a diagonal matrix, we need to have three reagents (so the number of rows equals the number of columns) and further, we need to choose them so that the off-diagonal entries are zero, i.e., our first reagent should react only with X , our second reagent should react only with Y , and our third reagent should react only with Z .

Performance review: 24 out of 27 got this. 2 chose (C), 1 chose (A).

Historical note (last time): 22 out of 25 got this. 1 each chose (A), (B), and (D).

- (5) (*) Suppose we are given an aqueous solution with two known dissolved substances. There are two different types of reactions. One is an acid-base reaction and the other is a redox reaction. For both reactions, we can use titrations (separately) to deduce the quantity of reagent needed. What type of system should we expect to get if only one of the solutes participates in the redox reaction but both participate in the acid-base reaction?
- (A) A diagonal system, i.e., the coefficient matrix is a diagonal matrix.

- (B) A triangular system, i.e., the coefficient matrix is a triangular matrix (whether it is upper or lower triangular depends on the order in which we write the rows).
- (C) A system of rank one, i.e., the coefficient matrix has rank one.

Answer: Option (B)

Explanation: The coefficient matrix has columns corresponding to the solutes, and rows corresponding to the reagents, with the matrix entries describing the amount of the row reagent consumed per unit of the solute. The row for the redox reagent has one entry zero. If we write that solute on the left, and that row as the second row, we get a matrix of the form:

$$\begin{bmatrix} * & * \\ 0 & * \end{bmatrix}$$

This is upper triangular.

Performance review: 24 out of 27 got this. 2 chose (C), 1 chose (A).

Historical note (last time): 20 out of 25 got this. 4 chose (A), 1 chose (C).

- (6) (*) Suppose we are given an aqueous solution with two known dissolved substances. There are two different types of reactions. One is an acid-base reaction and the other is a redox reaction. For both reactions, we can use titrations (separately) to deduce the quantity of reagent needed. Suppose we are given an aqueous solution with two known dissolved substances. Suppose both solutes participate in both reactions. What should we desire if we want to use the data from the two titrations to determine the amounts of each of the substances?
- (A) The proportions in which the two substances react should be the same for the two reactions.
- (B) The proportions in which the two substances react should differ for the two reactions.
- (C) It does not matter; we will be able to determine the amounts of each of the substances in both cases.
- (D) It does not matter; we will not be able to determine the amounts of each of the substances in either case.

Answer: Option (B)

Explanation: If they react in the same proportions, then the coefficient matrix will have rank one, i.e., the equations we get from the two reactions will convey the same algebraic information.

Performance review: 20 out of 27 got this. 4 chose (A), 2 chose (D), 1 left the question blank.

Historical note (last time): 12 out of 25 got this. 7 chose (C), 4 chose (D), 2 chose (A).

- (7) *Do not discuss this!:* A consumer price index is obtained from a “goods basket” by multiplying the price of each good in the basket by a fixed weight, and then adding up all the price \times weight products. The weights are kept fixed, but the prices vary from year to year. Thus, the consumer price index value itself fluctuates from year to year.

What is a good way of modeling this?

- (A) The prices of the various goods in various years are stored in a matrix, the weights used in the index are stored in a vector, and the consumer price index values arise as the output vector of the matrix-vector product.
- (B) The weights used in the index are stored in a matrix, the prices of the various goods in various years are stored in a vector, and the consumer price index values arise as the output vector of the matrix-vector product.
- (C) The prices of the various goods in various years are stored in a matrix, the consumer price index values are stored as a vector, and the weights used in the index arise as the output vector of the matrix-vector product.
- (D) The weights used in the index are stored in a matrix, the consumer price index values are stored in a vector, and the prices of the various goods in various years arise as the output vector of the matrix-vector product.
- (E) The consumer price index values are stored in a matrix, the prices of the various goods in various years are stored in a vector, and the weights used in the index arise as the output vector of the matrix-vector product.

Answer: Option (A)

Explanation: The weight vector is fixed. There are many price vectors, one for each year. It makes sense to store these as different rows of a matrix, that is then multiplied with the weight vector to give the vector of consumer price index values. The intuition is that the fixed vector is chosen as the vector and the many different vectors that need to be dotted (in the sense of “taking the dot product”) with the fixed vector are taken as rows of a matrix.

Performance review: 17 out of 27 got this. 10 chose (B).

Historical note (last time): 4 out of 25 got this. 20 chose (B), 1 chose (E).

- (8) Amelia wants to choose a healthy balanced diet. She has access to 30 different types of foods. There are 400 different nutrients that she wants a good amount of. Each of the foods that Amelia consumes offers a positive amount of each nutrient per unit foodstuff. Amelia is interested in meeting the daily value requirements for all nutrients. For some nutrients, her daily value requirements specify only a minimum. For some nutrients, both a minimum and a maximum are specified. Assume that the total amount of any nutrient can be obtained by adding up the amounts obtained from each of the foodstuffs Amelia consumes. Amelia wants to determine how much of each foodstuff she should consume. How should she model the situation?
- (A) The matrix with information on the nutritional contents of the foodstuffs is a 400×400 matrix, and the vector of amounts of each foodstuff consumed is a 400×1 column vector.
- (B) The matrix with information on the nutritional contents of the foodstuffs is a 30×30 matrix, and the vector of amounts of each foodstuff consumed is a 30×1 column vector.
- (C) The matrix with information on the nutritional contents of the foodstuffs is a 400×30 matrix, and the vector of amounts of each foodstuff consumed is a 30×1 column vector.
- (D) The matrix with information on the nutritional contents of the foodstuffs is a 30×400 matrix, and the vector of amounts of each foodstuff consumed is a 400×1 column vector.
- (E) The matrix with information on the nutritional contents of the foodstuffs is a 400×400 matrix, and the vector of amounts of each foodstuff consumed is a 30×1 column vector.

Answer: Option (C)

Explanation: The input food vector is 30×1 , the output nutrient vector is 400×1 , so the matrix must be 400×30 .

Performance review: 14 out of 27 got this. 12 chose (D), 1 left the question blank.

Historical note (last time): 20 out of 25 got this. 2 each chose (A) and (B), 1 chose (E).

DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE FRIDAY OCTOBER 18: LINEAR TRANSFORMATIONS

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

27 people took this 5-question quiz. The score distribution was as follows:

- Score of 1: 2 people
- Score of 2: 4 people
- Score of 3: 5 people
- Score of 4: 8 people
- Score of 5: 8 people

The mean score was 3.59 out of 5.

The question-wise answers and performance review were as follows:

- (1) Option (B): 24 people
- (2) Option (B): 22 people
- (3) Option (D): 19 people
- (4) Option (E): 19 people
- (5) Option (D): 13 people

2. SOLUTIONS

PLEASE DO *NOT* DISCUSS ANY QUESTIONS.

The quiz covers basics related to linear transformations (notes titled **Linear transformations**, corresponding section in the book Section 2.1). Explicitly, the quiz covers:

- Representation of a linear transformation using a matrix, and identifying the domain and co-domain in terms of the row and column counts of the matrix.
- Relationship between injectivity, surjectivity, rank, row count, and column count.
- Relationship between the entries of the matrix and the images of the standard basis vectors under the corresponding linear transformation.

The questions are fairly easy if you understand the material. But it's important that you be able to answer them, otherwise what we study later will not make much sense.

- (1) *Do not discuss this!*: Which of the following correctly describes a $m \times n$ matrix?
 - (A) There are m rows, and each row gives a vector with m coordinates. There are n columns, and each column gives a vector with n coordinates.
 - (B) There are m rows, and each row gives a vector with n coordinates. There are n columns, and each column gives a vector with m coordinates.
 - (C) There are n rows, and each row gives a vector with m coordinates. There are m columns, and each column gives a vector with n coordinates.
 - (D) There are n rows, and each row gives a vector with n coordinates. There are m columns, and each column gives a vector with m coordinates.

Answer: Option (B)

Explanation: This should be obvious by looking at the matrix. For instance, a 2×3 matrix is of the form:

$$\begin{bmatrix} a & b & c \\ d & e & f \end{bmatrix}$$

Performance review: 24 out of 27 got this. 3 chose (A).

- (2) *Do not discuss this!:* For a $p \times q$ matrix A , we can define a linear transformation T_A by $T_A(\vec{x}) := A\vec{x}$. What type of linear transformation is T_A ?
- (A) T_A is a linear transformation from \mathbb{R}^p to \mathbb{R}^q
 - (B) T_A is a linear transformation from \mathbb{R}^q to \mathbb{R}^p
 - (C) T_A is a linear transformation from $\mathbb{R}^{\max\{p,q\}}$ to $\mathbb{R}^{\min\{p,q\}}$
 - (D) T_A is a linear transformation from $\mathbb{R}^{\min\{p,q\}}$ to $\mathbb{R}^{\max\{p,q\}}$

Answer: Option (B)

Explanation: The way matrix-vector multiplication works, a $p \times q$ matrix multiplies with a $q \times 1$ vector to give a $p \times 1$ vector. Thus the linear transformation takes as input a q -dimensional vector and gives as output a p -dimensional vector, and is hence a transformation from \mathbb{R}^q to \mathbb{R}^p .

Performance review: 22 out of 27 got this. 3 chose (A), 2 chose (C).

- (3) *Do not discuss this!:* With the same notation as for the preceding question, which of the following is true?
- (A) If $p < q$, T_A must be injective
 - (B) If $p > q$, T_A must be injective
 - (C) If $p = q$, T_A must be injective
 - (D) If $p < q$, T_A cannot be injective
 - (E) If $p > q$, T_A cannot be injective

Answer: Option (D)

Explanation: For the linear transformation T_A to be injective, the matrix A needs to have full column rank q , because that is what it means for there to be no non-leading variables and for the solution to therefore be unique if it exists (see the lecture notes for more). In other words, we require that the matrix have rank q . However, we know that the rank of a matrix is at most equal to the minimum of the number of rows and the number of columns. Thus, if $p < q$, the matrix cannot have full column rank, and the linear transformation cannot be injective.

Intuitively, the linear transformation goes from \mathbb{R}^q to \mathbb{R}^p , so in order for it to be injective, the target space should be at least as big as the domain. Thus, $p < q$ is incompatible with injectivity.

Performance review: 19 out of 27 got this. 4 chose (E), 3 chose (C), 1 left the question blank.

- (4) *Do not discuss this!:* With the same notation as for the previous two questions, which of the following is true?
- (A) If $p < q$, T_A must be surjective
 - (B) If $p > q$, T_A must be surjective
 - (C) If $p = q$, T_A must be surjective
 - (D) If $p < q$, T_A cannot be surjective
 - (E) If $p > q$, T_A cannot be surjective

Answer: Option (E)

Explanation: For the linear transformation T_A to be surjective, the matrix A needs to have full row rank p , because we want the system to always be consistent and therefore we want that there should be no zero rows in the rref. We also know that the rank of the matrix is at most equal to the minimum of the number of rows and number of columns. Therefore, if $p > q$, A cannot have full row rank and T_A cannot be surjective.

Performance review: 19 out of 27 got this. 5 chose (D), 2 chose (C).

- (5) *Do not discuss this!:* With the same notation as for the last three questions, which of the following is true?
- (A) The rows of A are the images under T_A of the standard basis vectors of \mathbb{R}^p .
 - (B) The columns of A are the images under T_A of the standard basis vectors of \mathbb{R}^p .
 - (C) The rows of A are the images under T_A of the standard basis vectors of \mathbb{R}^q .
 - (D) The columns of A are the images under T_A of the standard basis vectors of \mathbb{R}^q .

Answer: Option (D)

Explanation: See the lecture notes for details. Note, however, that dimension considerations can get the answer here immediately. T_A is a map from \mathbb{R}^q to \mathbb{R}^p , so only Options (C) and (D) make sense from the domain perspective. Further, the rows of A are q -dimensional whereas the columns

are p -dimensional, so only the columns have the right dimension, thus making (D) the only legitimate option.

Performance review: 13 out of 27 got this. 11 chose (D), 2 chose (C), 1 chose (A).

DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE FRIDAY OCTOBER 25: MATRIX MULTIPLICATION (BASIC)

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

26 people took this 5-question quiz. The score distribution was as follows:

- Score of 0: 4 people
- Score of 1: 1 person
- Score of 2: 4 people
- Score of 3: 9 people
- Score of 4: 3 people
- Score of 5: 5 people

The mean score was 2.8.

The question-wise answers and performance review were as follows:

- (1) Option (D): 17 people
- (2) Option (B): 8 people
- (3) Option (A): 14 people
- (4) Option (D): 14 people
- (5) Option (B): 20 people

2. SOLUTIONS

PLEASE DO *NOT* DISCUSS ANY QUESTIONS

This quiz tests for basic comprehension of the setup for matrix multiplication. It corresponds to the material from Sections 1-6 (excluding Section 4) of the **Matrix multiplication and inversion** notes, and also to Section 2.3 of the book.

- (1) *Do not discuss this!* Suppose A and B are (not necessarily square) matrices. Then, which of the following describes correctly the relationship between the existence and value of the (alleged) matrix product AB and the existence and value of the (alleged) matrix product BA ?
 - (A) AB is defined if and only if BA is defined, and if so, they are equal.
 - (B) AB is defined if and only if BA is defined, but they need not be equal.
 - (C) If AB and BA are both defined, then $AB = BA$. However, it is possible for one of AB and BA to be defined and the other to not be defined.
 - (D) It is possible for only one of AB and BA to be defined. It is also possible for both AB and BA to be defined, but to not be equal to each other.

Answer: Option (D)

Explanation: Suppose A is a $m \times n$ matrix (i.e., it has m rows and n columns) and B is a $p \times q$ matrix (i.e., it has p rows and q columns), where $m, n, p,$ and q are positive integers. AB is defined if and only if $n = p$, i.e., the number of columns of A equals the number of rows of B . BA is defined if and only if $m = q$, i.e., the number of columns of B equals the number of rows of A . The conditions are independent of one another, so it is possible for only one of AB and BA to be defined.

Suppose now that AB and BA are both defined. Then, $n = p$ and $m = q$, so B is a $n \times m$ matrix. Thus, AB is a $m \times m$ matrix and BA is a $n \times n$ matrix, with m and n possibly different. Therefore, AB and BA do not even necessarily have the same dimensions, and therefore they definitely are not required to be equal.

Not clear to you?: Try picking actual numerical values of m , n , p , and q , write down actual example matrices, and see how the multiplication works.

Performance review: 17 out of 26 got this. 4 chose (B), 3 chose (C), 2 chose (A).

- (2) *Do not discuss this!:* Suppose A and B are matrices such that both AB and BA are defined. Which of the following correctly describes what we know about AB and BA ?
- (A) Both AB and BA are square matrices and have the same dimensions, i.e., in both AB and BA , the number of rows equals the number of columns, and further, the number of rows of AB equals the number of rows of BA .
 - (B) Both AB and BA are square matrices (the number of rows equals the number of columns) but they may not have the same dimensions: the number of rows in AB need not equal the number of rows in BA .
 - (C) AB and BA need not be square matrices but both must have the same dimensions: the number of rows in AB equals the number of rows in BA , and the number of columns in AB equals the number of columns in BA .
 - (D) AB and BA need not be square matrices and they need not have the same row count or the same column count, i.e., the number of rows in AB need not equal the number of rows in BA , and the number of columns in AB need not equal the number of columns in BA .

Answer: Option (B)

Explanation: See the second paragraph of the explanation for the preceding question. The conclusion there was that if AB and BA are both defined, then for A a $m \times n$ matrix, B is a $n \times m$ matrix, and thus AB is $m \times m$ and BA is $n \times n$. m need not be equal to n .

Not clear to you?: Just try picking positive integer values of m and n that are not equal, and try constructing a $m \times n$ matrix A and a $n \times m$ matrix B . Then, compute AB and BA and verify that they are square matrices of different dimensions.

Performance review: 8 out of 26 got this. 8 chose (D), 7 chose (A), 3 chose (C).

- (3) *Do not discuss this!:* Suppose A and B are matrices such that both AB and $A + B$ are defined. Which of the following correctly describes what we know about A and B ?
- (A) Both A and B are square matrices and have the same dimensions, i.e., in both A and B , the number of rows equals the number of columns, and further, the number of rows of A equals the number of rows of B .
 - (B) Both A and B are square matrices (the number of rows equals the number of columns) but they may not have the same dimensions: the number of rows in A need not equal the number of rows in B .
 - (C) A and B need not be square matrices but both must have the same dimensions: the number of rows in A equals the number of rows in B , and the number of columns in A equals the number of columns in B .
 - (D) A and B need not be square matrices and they need not have the same row count or the same column count, i.e., the number of rows in A need not equal the number of rows in B , and the number of columns in A need not equal the number of columns in B .

Answer: Option (A)

Explanation: Suppose A is a $m \times n$ matrix (i.e., it has m rows and n columns) and B is a $p \times q$ matrix (i.e., it has p rows and q columns). The condition that $A + B$ is defined tells us that $m = p$ (i.e., A and B have the same number of rows as each other) and that $n = q$ (i.e., A and B have the same number of columns as each other). Thus, both A and B are $m \times n$ matrices. In order for AB to make sense, we need n (the number of columns of A) to equal m (the number of rows of B). Thus, $m = n$, so that both A and B are $m \times m$ matrices.

Not clear to you?: Try writing an example that violates the conditions of Option (A) and see for yourself that you'll run into trouble either with computing $A + B$ or with computing AB .

Performance review: 14 out of 26 got this. 12 chose (C).

- (4) *Do not discuss this!:* Suppose A is a $p \times q$ matrix and B is a $q \times r$ matrix. The product matrix AB is a $p \times r$ matrix. Using the convention of matrices as linear transformations via their action by multiplication on column vectors, what is the appropriate interpretation of the matrix product in terms of composing linear transformations?

- (A) A corresponds to a linear transformation T_A from \mathbb{R}^p to \mathbb{R}^q , and B corresponds to a linear transformation T_B from \mathbb{R}^q to \mathbb{R}^r . The product AB therefore corresponds to a linear transformation from \mathbb{R}^p to \mathbb{R}^r that is the composite of the two linear transformations, with T_A applied first (to the domain) and then T_B (T_B being applied to the intermediate space obtained after applying T_A).
- (B) A corresponds to a linear transformation T_A from \mathbb{R}^p to \mathbb{R}^q , and B corresponds to a linear transformation T_B from \mathbb{R}^q to \mathbb{R}^r . The product AB therefore corresponds to a linear transformation from \mathbb{R}^p to \mathbb{R}^r that is the composite of the two linear transformations, with T_B applied first (to the domain) and then T_A (T_A being applied to the intermediate space obtained after applying T_B).
- (C) A corresponds to a linear transformation T_A from \mathbb{R}^q to \mathbb{R}^p , and B corresponds to a linear transformation T_B from \mathbb{R}^r to \mathbb{R}^q . The product AB therefore corresponds to a linear transformation from \mathbb{R}^r to \mathbb{R}^p that is the composite of the two linear transformations, with T_A applied first (to the domain) and then T_B (T_B being applied to the intermediate space obtained after applying T_A).
- (D) A corresponds to a linear transformation T_A from \mathbb{R}^q to \mathbb{R}^p , and B corresponds to a linear transformation T_B from \mathbb{R}^r to \mathbb{R}^q . The product AB therefore corresponds to a linear transformation from \mathbb{R}^r to \mathbb{R}^p that is the composite of the two linear transformations, with T_B applied first (to the domain) and then T_A (T_A being applied to the intermediate space obtained after applying T_B).

Answer: Option (D)

Explanation: Review the lecture notes regarding the interpretation of matrix multiplication as composition.

Performance review: 14 out of 26 got this. 9 chose (C), 2 chose (B), 1 chose (A).

- (5) *Do not discuss this!:* Suppose A , B , and C are matrices. Which of the following is true?
 - (A) If ABC is defined, then so are BCA and CAB .
 - (B) If ABC and BCA are both defined, then so is CAB . However, it is possible to have a situation where ABC is defined but BCA and CAB are not defined.
 - (C) It is possible to have a situation where ABC and BCA are both defined but CAB is not defined.

Answer: Option (B)

Explanation: Suppose A is a $m \times n$ matrix (i.e., it has m rows and n columns), B is a $p \times q$ matrix (i.e., it has p rows and q columns), and C is a $r \times s$ matrix (i.e., it has r rows and s columns). In order for ABC to be defined, we need $n = p$ (for AB to be defined) and $q = r$ (for BC to be defined). However, we do not need $s = m$, the condition that would allow us to multiply C with A . Therefore, we have no guarantee that the products BCA and CAB are defined.

If both ABC and BCA are defined, then we get the additional condition that $s = m$, and this allows us to define CAB .

Performance review: 20 out of 26 got this. 4 chose (A), 2 chose (C).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE MONDAY OCTOBER 28: MATRIX MULTIPLICATION AND INVERSION AS COMPUTATIONAL PROBLEMS

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

26 people took this 16-question quiz. The score distribution is as follows:

- Score of 3: 1 person
- Score of 6: 2 people
- Score of 7: 1 person
- Score of 8: 9 people
- Score of 9: 3 people
- Score of 10: 5 people
- Score of 11: 3 people
- Score of 12: 1 person
- Score of 15: 1 person

The question-wise answers and performance review are below:

- (1) Option (A): 25 people
- (2) Option (C): 23 people
- (3) Option (B): 12 people
- (4) Option (A): 15 people
- (5) Option (C): 15 people
- (6) Option (E): 20 people
- (7) Option (B): 9 people
- (8) Option (B): 11 people
- (9) Option (A): 8 people
- (10) Option (B): 5 people
- (11) Option (B): 21 people
- (12) Option (A): 24 people
- (13) Option (D): 10 people
- (14) Option (E): 1 person
- (15) Option (D): 18 people
- (16) Option (A): 14 people

2. SOLUTIONS

This quiz tests for a strong *conceptualization* (i.e., a metacognition) of the processes used for matrix multiplication and inversion. It is based on part of the **Matrix multiplication and inversion** notes and is related to Sections 2.3 and 2.4. It does not, however, test all aspects of that material.

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

- (1) How many arithmetic operations are needed for naive matrix multiplication of a $m \times n$ matrix and a $n \times p$ matrix?
 - (A) $O(mnp)$ additions and $O(mnp)$ multiplications
 - (B) $O(m + n + p)$ additions and $O(mnp)$ multiplications
 - (C) $O(mn)$ additions and $O(np)$ multiplications
 - (D) $O(mn + mp)$ additions and $O(mnp)$ multiplications
 - (E) $O(m + n + p)$ additions and $O(m + n + p)$ multiplications

Answer: Option (A)

Explanation: The actual number of multiplications is mnp and the actual number of additions is $m(n-1)p$, both of which are $O(mnp)$. In fact, it is $\Theta(mnp)$. (The O notation means an upper bound on order, the Θ notation means upper and lower bounds on order).

Performance review: 25 out of 26 got this. 1 person chose (D).

Historical note (last time): 21 out of 26 got this. 5 chose (D).

- (2) What is the arithmetic complexity (in terms of total number of arithmetic operations needed) for naive matrix multiplication of two generic $n \times n$ matrices?
- (A) $\Theta(n)$
 - (B) $\Theta(n^2)$
 - (C) $\Theta(n^3)$
 - (D) $\Theta(n^4)$
 - (E) $\Theta(n^5)$

Answer: Option (C)

Explanation: Based on Question 1, where m, n, p are all equal to n .

Performance review: 23 out of 26 got this. 3 people chose (D).

Historical note (last time): 23 out of 26 got this. 2 chose (B), 1 chose (D).

- (3) Which of the following is the tightest “obvious” lower bound on the possible arithmetic complexity of any generic algorithm for multiplying two $n \times n$ matrices? We use Ω to denote *at least that order*.
- (A) $\Omega(n)$
 - (B) $\Omega(n^2)$
 - (C) $\Omega(n^3)$
 - (D) $\Omega(n^4)$
 - (E) $\Omega(n^5)$

Answer: Option (B)

Explanation: The product matrix has n^2 entries, all of which could in principle be different and require at least some nonzero computation. In other words, just filling in the output matrix takes n^2 steps. Thus, the obvious lower bound for matrix multiplication is $\Omega(n^2)$.

Performance review: 12 out of 26 got this. 8 chose (A), 6 chose (C).

Historical note (last time): 12 out of 26 got this. 10 chose (C), 4 chose (A).

- (4) What is the minimum number of arithmetic operations needed to compute the product of two generic diagonal $n \times n$ matrices?
- (A) n
 - (B) $n + 1$
 - (C) $2n - 1$
 - (D) $2n$
 - (E) n^2

Answer: Option (A)

Explanation: All the off-diagonal entries are 0. The diagonal entries are obtained by entry-wise multiplication. Explicitly, if $AB = C$, then:

$$c_{ii} = a_{ii}b_{ii}$$

Thus, each of the n diagonal entries of the matrix requires one multiplication to compute, so the total number of operations necessary is n .

Performance review: 15 out of 26 got this. 8 chose (E), 2 chose (C), 1 chose (D).

Historical note (last time): 20 out of 26 got this. 4 chose (E), 2 chose (C).

- (5) What is the minimum number of arithmetic operations needed to compute the product of a generic $1 \times n$ matrix and a generic $n \times 1$ matrix?
- (A) n
 - (B) $n + 1$
 - (C) $2n - 1$
 - (D) $2n$
 - (E) n^2

Answer: Option (C)

Explanation: This is a dot product computation. It involves n multiplications and $n - 1$ additions, so a total of $2n - 1$ operations. Alternatively, recall that in general, for multiplying a $m \times n$ matrix and a $n \times p$ matrix, we require mnp multiplications and $m(n - 1)p$ additions. Here, $m = p = 1$.

Performance review: 15 out of 26 got this. 8 chose (A), 2 chose (B), 1 chose (E).

Historical note (last time): 16 out of 26 got this. 9 chose (A), 1 chose (D).

- (6) What is the minimum number of arithmetic operations needed to compute the product of a generic $n \times 1$ matrix and a generic $1 \times n$ matrix?
- (A) n
 - (B) $n + 1$
 - (C) $2n - 1$
 - (D) $2n$
 - (E) n^2

Answer: Option (E)

Explanation: This is the Hadamard product or outer product. The product is a $n \times n$ matrix and each entry is simply one product. Thus, there are n^2 multiplications and 0 additions, so a total of n^2 operations.

We can think of this in terms of the general setting of multiplication of a $m \times n$ matrix and a $n \times p$ matrix. The n in our current situation equals both the m and the p of the generic setup and the n of our generic setup equals 1. The number of multiplications is $mnp = n(1)n = n^2$ and the number of additions is $n(1 - 1)n = 0$.

Performance review: 20 out of 26 got this. 4 chose (A), 2 chose (C).

Historical note (last time): 14 out of 26 got this. 6 chose (A), 3 chose (C), 2 chose (D), 1 chose (B).

- (7) In order terms, what is the minimum number of arithmetic operations needed to compute the product of a generic $n \times n$ diagonal matrix and a generic $n \times n$ upper triangular matrix? The upper triangular matrix has zero entries below the diagonal. The entries on or above the diagonal may be nonzero (and generically, they will be nonzero).
- (A) $n(n - 1)/2$
 - (B) $n(n + 1)/2$
 - (C) $n(n - 1)$
 - (D) n^2
 - (E) $n(n + 1)$

Answer: Option (B)

Explanation: To compute the $(ij)^{th}$ entry of the product, we need to take the dot product of the i^{th} row of the diagonal matrix and the j^{th} column of the other matrix. There is only one nonzero term in the corresponding summation if $i \leq j$, and no nonzero term if $i > j$. So, we have to do 0 additions, and the number of multiplications is the number of entries in the upper triangular part.

So, we need to calculate the number of entries in the upper triangle including the diagonal. There are n positions on the diagonal. There are thus $n^2 - n$ off-diagonal entries, with half of them above the diagonal and half of them below the diagonal. Thus, there are $(n^2 - n)/2 = n(n - 1)/2$ entries above the diagonal, so a total of $n(n - 1)/2 + n = n(n + 1)/2$ entries on and above the diagonal.

Performance review: 9 out of 26 got this. 7 each chose (A) and (C), 2 chose (D), 1 left the question blank.

Historical note (last time): 19 out of 26 got this. 5 chose (D), 2 chose (B).

Adding n numbers to each other requires $n - 1$ addition operations. In a non-parallel setting, there is no way of improving this.

However, using the associativity of addition, we can write a faster parallelizable algorithm. A simple parallelization is to split the list being added into two sublists of length about $n/2$ each. Delegate the task of adding up within each sublist to different processors running in parallel. Then, add up the numbers obtained. This takes about half the time, with a little overhead (of collecting and adding up). This type of strategy is called a *divide and conquer* strategy. Using a divide and conquer strategy repeatedly, we can demonstrate that the parallelized arithmetic complexity of this approach is $\Theta(\log_2 n)$.

- (8) Suppose A is a $1 \times n$ matrix and B is a $n \times 1$ matrix. Assume an unlimited number of processors that all have free read access to both A and B , free write access to the product matrix, and a shared workspace where they can store intermediate results. What is the arithmetic complexity in this context (i.e., the parallelized arithmetic complexity) for computing AB ? What we mean here is: what is the smallest depth of a computational tree to compute AB ?

- (A) $\Theta(1)$
- (B) $\Theta(\log_2 n)$
- (C) $\Theta(n \log_2 n)$
- (D) $\Theta(n^2)$
- (E) $\Theta(n^2 \log_2 n)$

Answer: Option (B)

Explanation: Computing the dot product requires computing n individual products, and then adding them up. The computation of the individual products can be done in parallel, so that takes time $\Theta(1)$. The adding up involves the addition of n numbers, which takes then $\Theta(\log_2 n)$. The total time taken is thus $\Theta(\log_2 n)$.

Performance review: 11 out of 26 got this. 11 chose (C), 1 chose (D), 3 chose (A).

Historical note (last time): 22 out of 26 got this. 2 chose (C), 1 each chose (A) and (D).

- (9) Suppose A is a $n \times 1$ matrix and B is a $1 \times n$ matrix. Assume an unlimited number of processors that all have free read access to both A and B , free write access to the product matrix, and a shared workspace where they can store intermediate results. What is the arithmetic complexity in this context (i.e., the parallelized arithmetic complexity) for computing AB ? What we mean here is: what is the smallest depth of a computational tree to compute AB ?

- (A) $\Theta(1)$
- (B) $\Theta(\log_2 n)$
- (C) $\Theta(n \log_2 n)$
- (D) $\Theta(n^2)$
- (E) $\Theta(n^2 \log_2 n)$

Answer: Option (A)

Explanation: Since the number of columns in A = the number of rows in B is 1, we do not need to perform any additions. Rather, we need to do n^2 multiplications. All these multiplications can be performed in parallel, so the time taken is $\Theta(1)$.

Performance review: 8 out of 26 got this. 7 chose (B), 6 chose (D), 4 chose (C), 1 chose (E).

Historical note (last time): 10 out of 26 got this. 10 chose (C), 4 chose (B), 2 chose (D).

- (10) Suppose A and B are two $n \times n$ matrices. Assume an unlimited number of processors that all have free read access to both A and B , free write access to the product matrix, and a shared workspace where they can store intermediate results. What is the arithmetic complexity in this context (i.e., the parallelized arithmetic complexity) for computing AB ? What we mean here is: what is the smallest depth of a computational tree to compute AB ? Use naive matrix multiplication and speed it up using the parallelized processes discussed here.

- (A) $\Theta(1)$
- (B) $\Theta(\log_2 n)$
- (C) $\Theta(n \log_2 n)$
- (D) $\Theta(n^2)$
- (E) $\Theta(n^2 \log_2 n)$

Answer: Option (B)

Explanation: We need to calculate n^2 matrix entries in the product, but all the entries can be computed in parallel, and the time taken to compute a matrix entry is the time taken to perform a single dot product. As we saw earlier, this is $\Theta(\log_2 n)$.

We are given a $n \times n$ matrix A and we want to use *repeated squaring* to calculate powers of A . For instance, to calculate A^4 , we can simply calculate $(A^2)^2$, which requires two multiplications. To calculate A^5 , we calculate $(A^2)^2A$, which requires three multiplications. Assume that we can store any number of intermediate matrices, i.e., storage space is not a constraint.

Performance review: 5 out of 26 got this. 19 chose (E), 2 chose (C).

Historical note (last time): 13 out of 26 got this. 9 chose (E), 2 chose (D), 1 each chose (A) and (C).

- (11) What is the smallest number of matrix multiplications needed to calculate A^7 using repeated squaring?
- (A) 3
 (B) 4
 (C) 5
 (D) 6
 (E) 7

Answer: Option (B)

Explanation: We compute A^2 (first operation), then compute $(A^2)^2 = A^4$ (second operation), then multiply them to get A^6 (third operation), then multiply that by A to get A^7 (fourth operation).

Performance review: 21 out of 26 got this. 3 chose (A), 2 chose (C).

Historical note (last time): 21 out of 26 got this. 3 chose (C), 2 chose (A).

- (12) What is the smallest number of matrix multiplications needed to calculate A^8 using repeated squaring?
- (A) 3
 (B) 4
 (C) 5
 (D) 6
 (E) 7

Answer: Option (A)

Explanation: Three squarings will do the trick.

Performance review: 24 out of 26 got this. 2 chose (B).

Historical note (last time): 23 out of 26 got this. 3 chose (C).

- (13) What is the smallest number of matrix multiplications needed to calculate A^{21} using repeated squaring?
- (A) 3
 (B) 4
 (C) 5
 (D) 6
 (E) 7

Answer: Option (D)

Explanation: Square A four times to get to A^{16} . In the process, we have also found A^2 , A^4 , and A^8 . Multiply A^{16} by A^4 (fifth operation) and then multiply by A (sixth operation).

Performance review: 10 out of 26 got this. 15 chose (E), 1 left the question blank.

Historical note (last time): 12 out of 26 got this. 8 chose (E), 5 chose (C), 1 chose (A).

Suppose A is an *invertible* $n \times n$ matrix. It is possible to invert A using $\Theta(n^3)$ (worst-case) arithmetic operations via Gauss-Jordan elimination. We can thus add computation of the inverse to our toolkit when calculating powers. It is helpful even when calculating positive powers.

Count each matrix multiplication and each matrix inversion as one “matrix operation.”

- (14) What is the smallest positive r where we can achieve a saving on the total number of matrix operations to calculate A^r by also computing A^{-1} , rather than just using repeated squaring?
- (A) 3
 (B) 7
 (C) 15
 (D) 23
 (E) 31

Answer: Option (E)

Explanation: The case where we are likely to get the best saving is where $r = 2^s - 1$. In this case, doing the calculation without finding the inverse takes $2s - 2$ steps and doing the calculation using the inverse takes $s + 2$ steps: s to calculate A^{2^s} , 1 to calculate A^{-1} , and 1 more to multiply them and get A^{2^s-1} . For $s + 2 < 2s - 2$, we need $s \geq 5$, so the smallest number that works is $2^5 - 1 = 31$.

Performance review: 1 out of 26 got this. 16 chose (C), 4 chose (D), 3 chose (B), 2 chose (A), and 1 left the question blank.

Historical note (last time): 8 out of 26 got this. 9 chose (C), 5 chose (B), 2 chose (D), 1 chose (A).

- (15) *Strassen’s algorithm* is a *fast matrix multiplication* algorithm that can multiply two $n \times n$ matrices using $O(n^{\log_2 7})$ arithmetic operations. In practice, however, a lot of existing computer code for matrix multiplication, written long after Strassen’s algorithm was discovered, uses naive matrix multiplication. Which of the following reasons explain this? Please see Options (D) and (E) before answering.
- (A) Strassen’s algorithm becomes faster than naive matrix multiplication only for very large matrix sizes.
 (B) Strassen’s algorithm is more complicated to code.
 (C) Strassen’s algorithm is not as easily parallelizable as naive matrix multiplication.
 (D) All of the above.
 (E) None of the above.

Answer: Option (D)

Explanation: Strassen’s algorithm for the 2×2 case requires 7 multiplications and 14 additions. The number of additions is more. The saving on multiplication does dominate for large enough matrix sizes, but this does not occur immediately.

Strassen’s algorithm is definitely more complicated to code. If it were easy to code, it would also be easier to explain, and we would have discussed it in class.

Parallelization is harder because on the “conquer” part of the divide and conquer strategy. With the naive approach, each entry can be computed totally separately. With Strassen’s, there are many different types of computations that need to be done and then pieced together.

Performance review: 18 out of 26 got this. 5 chose (C), 3 chose (A).

Historical note (last time): 20 out of 26 got this. 3 chose (C), 2 chose (E), 1 chose (A).

There exist even faster algorithms for matrix multiplication than Strassen’s algorithm. The best known algorithm currently is the *Coppersmith-Winograd algorithm*, which can multiply two $n \times n$ matrices in time $O(n^{2.3727})$. However, the Coppersmith-Winograd algorithm is even more rarely implemented than Strassen’s for practical matrix multiplication code (according to some sources, Coppersmith-Winograd has *never* been implemented). The same reasons as those cited above for

the reluctance to use Strassen's algorithm apply. There are some additional obstacles to practical implementations of the Coppersmith-Winograd algorithm that make it even more difficult to use.

- (16) Suppose A and B are $n \times n$ matrices. What is the minimum number of matrix multiplications needed generically to compute the product $ABABABABA$?
- (A) 4
 - (B) 5
 - (C) 6
 - (D) 7
 - (E) 8

Answer: Option (A)

Explanation: We calculate AB (first), then square it to get $ABAB$ (second), then square again to get $ABABABAB$ (third), then multiply by A to get $ABABABABA$ (fourth).

Performance review: 14 out of 26 got this. 7 chose (B), 3 chose (E), 1 each chose (C) and (D).

Historical note (last time): 13 out of 26 got this. 8 chose (B), 2 chose (C), 2 chose (E), 1 chose (D).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY OCTOBER 30: LINEAR TRANSFORMATIONS AND FINITE STATE AUTOMATA

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

28 people took this 7-question quiz. The score distribution was as follows:

- Score of 0: 2 people
- Score of 1: 1 person
- Score of 2: 1 person
- Score of 3: 5 people
- Score of 4: 3 people
- Score of 5: 3 people
- Score of 6: 8 people
- Score of 7: 5 people

The mean score was 4.57.

The question-wise answers and performance review are below:

- (1) Option (E): 19 people
- (2) Option (B): 23 people
- (3) Option (A): 21 people
- (4) Option (C): 10 people
- (5) Option (D): 17 people
- (6) Option (B): 18 people
- (7) Option (C): 20 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS *ALL* QUESTIONS.

The purpose of this quiz is to explore in greater depth particular types of matrices, the corresponding linear transformations, and the relationship between operations on sets and similar operations on vector spaces. The material covered in the quiz will also prove to be a fertile source of *examples* and *counterexamples* for later content: in the future, when you are asked to come up with matrices that satisfy some very loosely stated conditions, the matrices of the type described here can be a place to begin your search.

Let n be a natural number greater than 1. Suppose $f : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ is a function satisfying $f(0) = 0$. Let T_f denote the linear transformation from \mathbb{R}^n to \mathbb{R}^n satisfying the following for all $i \in \{1, 2, \dots, n\}$:

$$T_f(\vec{e}_i) = \begin{cases} \vec{e}_{f(i)}, & f(i) \neq 0 \\ \vec{0}, & f(i) = 0 \end{cases}$$

Let M_f denote the matrix for the linear transformation T_f . M_f can be described explicitly as follows: the i^{th} column of M_f is $\vec{0}$ if $f(i) = 0$ and is $\vec{e}_{f(i)}$ if $f(i) \neq 0$.

Note that if $f, g : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ are functions with $f(0) = g(0) = 0$, then $M_{f \circ g} = M_f M_g$ and $T_{f \circ g} = T_f \circ T_g$.

We will also use the following terminology:

- A $n \times n$ matrix A is termed *idempotent* if $A^2 = A$.
- A $n \times n$ matrix A is termed *nilpotent* if there exists a positive integer r such that $A^r = 0$.

- A $n \times n$ matrix A is termed a *permutation matrix* if every row contains one 1 and all other entries 0, and every column contains one 1 and all other entries 0.

(1) What condition on a function $f : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ (satisfying $f(0) = 0$) is equivalent to requiring M_f to be idempotent?

- (A) $(f(x))^2 = x$ for all $x \in \{0, 1, 2, \dots, n\}$
- (B) $f(x^2) = x$ for all $x \in \{0, 1, 2, \dots, n\}$
- (C) $(f(x))^2 = f(x)$ for all $x \in \{0, 1, 2, \dots, n\}$
- (D) $f(f(x)) = x$ for all $x \in \{0, 1, 2, \dots, n\}$
- (E) $f(f(x)) = f(x)$ for all $x \in \{0, 1, 2, \dots, n\}$

Answer: Option (E)

Explanation: As noted above, $M_f^2 = M_f M_f = M_{f \circ f}$. In order to have $M_f^2 = M_f$, we need $M_{f \circ f} = M_f$. Since the matrix determines the function, we get $f \circ f = f$, so $f(f(x)) = f(x)$ for all $x \in \{0, 1, 2, \dots, n\}$.

Performance review: 19 out of 28 got this. 5 chose (C), 3 chose (D), 1 chose (B).

Historical note (last time): 12 out of 26 got this. 9 chose (C), 5 chose (D).

(2) What condition on a function $f : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ (satisfying $f(0) = 0$) is equivalent to requiring M_f to be nilpotent?

- (A) Composing f enough times with itself gives the identity function (i.e., the function that sends everything to itself).
- (B) Composing f enough times with itself gives the function that sends everything to 0.
- (C) Composing f enough times with itself gives the function that sends everything to 1.
- (D) Multiplying f enough times with itself gives the identity function (i.e., the function that sends everything to itself).
- (E) Multiplying f enough times with itself gives the function that sends everything to 0.

Answer: Option (B)

Explanation: The matrix M_f^r is the same as $M_{f \circ \dots \circ f}$ where f is composed with itself r times. Thus, $M_f^r = 0$ if and only if $M_{f \circ \dots \circ f} = 0$, which means that the r -fold composite of f is the function that sends everything to zero.

Performance review: 23 out of 28 got this. 3 chose (A), 2 chose (E).

Historical note (last time): 23 out of 26 got this. 2 chose (C), 1 chose (E).

(3) What condition on a function $f : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ (satisfying $f(0) = 0$) is equivalent to requiring M_f to be a permutation matrix?

- (A) Composing f enough times with itself gives the identity function (i.e., the function that sends everything to itself).
- (B) Composing f enough times with itself gives the function that sends everything to 0.
- (C) Composing f enough times with itself gives the function that sends everything to 1.
- (D) Multiplying f enough times with itself gives the identity function (i.e., the function that sends everything to itself).
- (E) Multiplying f enough times with itself gives the function that sends everything to 0.

Answer: Option (A)

Explanation: f itself is bijective, which means it cycles around the coordinates. Repeating f enough times should bring everything back where it was.

Performance review: 21 out of 28 got this. 7 chose (D).

Historical note (last time): 23 out of 26 got this. 2 chose (C), 1 chose (B).

(4) Consider a function $f : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ with the property that $f(0) = 0$ and, for each $i \in \{1, 2, \dots, n\}$, $f(i)$ is either i or 0. Note that the behavior may be different for different values of i (so some of them may go to themselves, and others may go to 0). What can we say M_f must be?

- (A) M_f must be the identity matrix.
- (B) M_f must be the zero matrix.
- (C) M_f must be an idempotent matrix.
- (D) M_f must be a nilpotent matrix.
- (E) M_f must be a permutation matrix.

Answer: Option (C)

Explanation: The function f given here satisfies the condition that $f(f(i)) = f(i)$ for all i . To see this, note that:

- if $f(i) = i$, then $f(f(i)) = f(i) = i$, i.e., we stay looped at i .
- if $f(i) = 0$, then $f(f(i)) = f(0) = 0$, i.e., once we reach 0, we stay there.

Performance review: 10 out of 28 got this. 9 chose (D), 7 chose (E), 2 chose (A).

Historical note (last time): 12 out of 26 got this. 6 chose (D), 5 chose (E), 3 chose (A).

- (5) Which of the following pairs of candidates for $f, g : \{0, 1, 2\} \rightarrow \{0, 1, 2\}$ satisfies the condition that $M_f M_g = 0$ but $M_g M_f \neq 0$?

- (A) $f(0) = 0, f(1) = 1, f(2) = 2$, whereas $g(0) = 0, g(1) = 2, g(2) = 1$
(B) $f(0) = 0, f(1) = 0, f(2) = 1$, whereas $g(0) = 0, g(1) = 2, g(2) = 0$
(C) $f(0) = 0, f(1) = 1, f(2) = 0$, whereas $g(0) = 0, g(1) = 0, g(2) = 2$
(D) $f(0) = 0, f(1) = 0, f(2) = 1$, whereas $g(0) = 0, g(1) = 1, g(2) = 0$
(E) $f(0) = 0, f(1) = 1, f(2) = 0$, whereas $g(0) = 0, g(1) = 0, g(2) = 1$

Answer: Option (D)

Explanation: For this option, $f \circ g$ sends everything to 0. $g \circ f$ sends 0 and 1 to 0 and sends 2 to 1.

Performance review: 17 out of 28 got this. 5 chose (E), 2 each chose (A), (B), and (C).

Historical note (last time): 20 out of 26 got this. 2 each chose (A), (B), and (E).

- (6) Which of the following pairs of candidates for $f, g : \{0, 1, 2\} \rightarrow \{0, 1, 2\}$ satisfies the condition that M_f and M_g are both nilpotent but $M_f M_g$ is not nilpotent?

- (A) $f(0) = 0, f(1) = 1, f(2) = 2$, whereas $g(0) = 0, g(1) = 2, g(2) = 1$
(B) $f(0) = 0, f(1) = 0, f(2) = 1$, whereas $g(0) = 0, g(1) = 2, g(2) = 0$
(C) $f(0) = 0, f(1) = 1, f(2) = 0$, whereas $g(0) = 0, g(1) = 0, g(2) = 2$
(D) $f(0) = 0, f(1) = 0, f(2) = 1$, whereas $g(0) = 0, g(1) = 1, g(2) = 0$
(E) $f(0) = 0, f(1) = 1, f(2) = 0$, whereas $g(0) = 0, g(1) = 0, g(2) = 1$

Answer: Option (B)

Explanation: $f \circ f$ is the function sending everything to 0. $g \circ g$ is also the function sending everything to 0. On the other hand, $f \circ g$ is the function that sends 1 to 1 and sends 0 and 2 to 0.

Performance review: 18 out of 28 got this. 3 each chose (A) and (E), 2 chose (D), 1 chose (C), and 1 left the question blank.

Historical note (last time): 22 out of 26 got this. 3 chose (A), 1 chose (C).

- (7) Which of the following pairs of candidates for $f, g : \{0, 1, 2\} \rightarrow \{0, 1, 2\}$ satisfies the condition that neither M_f and M_g is nilpotent but $M_f M_g$ is nilpotent?

- (A) $f(0) = 0, f(1) = 1, f(2) = 2$, whereas $g(0) = 0, g(1) = 2, g(2) = 1$
(B) $f(0) = 0, f(1) = 0, f(2) = 1$, whereas $g(0) = 0, g(1) = 2, g(2) = 0$
(C) $f(0) = 0, f(1) = 1, f(2) = 0$, whereas $g(0) = 0, g(1) = 0, g(2) = 2$
(D) $f(0) = 0, f(1) = 0, f(2) = 1$, whereas $g(0) = 0, g(1) = 1, g(2) = 0$
(E) $f(0) = 0, f(1) = 1, f(2) = 0$, whereas $g(0) = 0, g(1) = 0, g(2) = 1$

Answer: Option (C)

Explanation: M_f and M_g are both idempotent, since $f \circ f = f$ and $g \circ g = g$. However, $f \circ g$ is the function sending everything to 0, so $M_f M_g = 0$.

Performance review: 20 out of 28 got this. 3 chose (A), 2 chose (B), 1 each chose (D) and (E), 1 left the question blank.

Historical note (last time): 20 out of 26 got this. 3 chose (B), 1 each chose (A), (D), and (E).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY NOVEMBER 1: MATRIX MULTIPLICATION AND INVERSION: ABSTRACT BEHAVIOR PREDICTION

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

26 people took this 6-question quiz. The score distribution was as follows:

- Score of 0: 1 person
- Score of 1: 3 people
- Score of 2: 3 people
- Score of 3: 7 people
- Score of 4: 5 people
- Score of 5: 6 people
- Score of 6: 1 person

The mean score was 3.3.

The question-wise answers and performance review are as follows:

- (1) Option (B): 12 people
- (2) Option (C): 12 people
- (3) Option (E): 21 people
- (4) Option (E): 19 people
- (5) Option (A): 16 people
- (6) Option (E): 6 people

Note: Question 6 erroneously had “5 points” printed in front of it in the print copy handed out to students. That was because the question was copy-pasted to the quiz from a previous year’s test. We are not giving it additional weight.

Note on comparison with last time: Last time, I gave these questions (all except Question 6) *before* I gave the “linear transformations and finite state automata” questions. This might explain why the overall performance was somewhat worse last time compared to this time.

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS *ALL* QUESTIONS.

This quiz tests for *abstract behavior prediction* related to the structure of matrices defined based on the operations of matrix multiplication and inversion. It is based on part of the **Matrix multiplication and inversion** notes and is related to Sections 2.3 and 2.4. It does not, however, test all aspects of that material.

To understand this abstract behavior, we will consider *nilpotent*, *invertible*, and *idempotent* matrices.

- (1) Suppose A and B are $n \times n$ matrices such that B is invertible. Suppose r is a positive integer. What can we say that $(BAB^{-1})^r$ definitely equals?
 - (A) A^r
 - (B) BA^rB^{-1}
 - (C) $B^rA^rB^{-r}$
 - (D) B^rAB^{-r}
 - (E) BAB^{-1-r}

Answer: Option (B)

Explanation: Write:

$$BAB^{-1}BAB^{-1} \dots BAB^{-1}$$

Each B^{-1} and subsequent B multiply to the identity matrix, which disappears. So, we are left with:

$$BAA \dots AB^{-1}$$

where A appears r times. We thus get $BA^r B^{-1}$.

This is related to some deep facts about group structure. We will hint at these later in the course, but will not be able to appreciate the full depth of these.

Performance review: 12 out of 26 got this. 10 chose (A), 4 chose (C).

Historical note (last time): 8 out of 26 got this. 13 chose (A), 5 chose (C).

- (2) Suppose A and B are $n \times n$ matrices (n not too small) such that $(AB)^2 = 0$. What is the smallest r for which we can conclude that $(BA)^r$ is definitely 0?

- (A) 1
- (B) 2
- (C) 3
- (D) 4
- (E) 5

Answer: Option (C)

Explanation: We have:

$$(BA)^3 = BABABA = B(ABAB)A = B(AB)^2 A = B(0)A = 0$$

Thus, $(BA)^3$ is definitely 0. There are examples of matrices A and B such that $(AB)^2 = 0$ but $(BA)^2 \neq 0$. The smallest examples are 3×3 . For instance:

$$A = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

We obtain:

$$AB = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, BA = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$$

We have $(AB)^2 = 0$ but $(BA)^2 \neq 0$. However, $(BA)^3 = 0$.

How did we construct this example?: Building on the “linear transformations and finite state automata” framework, the idea is to choose $f, g : \{0, 1, 2, 3\} \rightarrow \{0, 1, 2, 3\}$ (with $f(0) = g(0) = 0$) such that $f \circ g$ is a function whose composite with itself sends everything to zero, whereas $g \circ f$ is a function whose composite with itself does not send everything to zero. The above matrices arise from one such example:

$$f(0) = 0, f(1) = 0, f(2) = 1, f(3) = 2, \quad g(0) = 0, g(1) = 1, g(2) = 2, g(3) = 0$$

The composites are:

$$(f \circ g)(0) = 0, (f \circ g)(1) = 0, (f \circ g)(2) = 1, (f \circ g)(3) = 0$$

and:

$$(g \circ f)(0) = 0, (g \circ f)(1) = 0, (g \circ f)(2) = 1, (g \circ f)(3) = 2$$

Notice that $f \circ g$ sends 2 to 1 and everything else to 0, hence its composite with itself sends everything to zero. On the other hand, the composite of $g \circ f$ with itself sends 3 to 1, and therefore does not send everything to zero.

In symbols, $A = M_f$, $B = M_g$, $AB = M_{f \circ g}$, and $BA = M_{g \circ f}$.

Performance review: 12 out of 26 got this. 12 chose (B), 1 each chose (A) and (E).

Historical note (last time): 6 out of 26 got this. 16 chose (B), 3 chose (A), 1 chose (E).

- (3) Suppose $n > 1$. A $n \times n$ matrix A is termed *nilpotent* if there exists a positive integer r such that A^r is the zero matrix. It turns out that if A is nilpotent, then $A^n = 0$. Which of the following describes correctly the relationship between being invertible and being nilpotent for $n \times n$ matrices?
- (A) A matrix is nilpotent if and only if it is invertible.
 - (B) Every nilpotent matrix is invertible, but not every invertible matrix is nilpotent.
 - (C) Every invertible matrix is nilpotent, but not every nilpotent matrix is invertible.
 - (D) An invertible matrix may or may not be nilpotent, and a nilpotent matrix may or may not be invertible.
 - (E) A matrix cannot be both nilpotent and invertible.

Answer: Option (E)

Explanation: Suppose $A^r = 0$ and A has inverse A^{-1} . We have $(A^{-1})^r A^r = I_n$ (where I_n denotes the $n \times n$ identity matrix), but it also equals $(A^{-1})^r 0 = 0$, so $I_n = 0$ is a contradiction.

Performance review: 21 out of 26 got this. 3 chose (B), 1 each chose (A) and (D).

Historical note (last time): 9 out of 26 got this. 9 chose (C), 6 chose (B), 2 chose (D).

- (4) Suppose A and B are $n \times n$ matrices. Which of the following is true? Please see Option (E) before answering.
- (A) AB is nilpotent if and only if A and B are both nilpotent.
 - (B) AB is nilpotent if and only if at least one of A and B is nilpotent.
 - (C) If both A and B are nilpotent, then AB is nilpotent, but AB being nilpotent does not imply that both A and B are nilpotent.
 - (D) If AB is nilpotent, then both A and B are nilpotent. However, both A and B being nilpotent does not imply that AB is nilpotent.
 - (E) None of the above.

Answer: Option (E)

Explanation: Here is an example where A and B are both nilpotent but AB is not:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, AB = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

Here is an example where neither A nor B is nilpotent but AB is the zero matrix, and therefore, is nilpotent:

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

How did we construct these examples: These examples fall out of the “linear transformations and finite state automata” framework. In fact, there are questions on that quiz that directly correspond to the construction of examples for this situation. Can you locate them?

Performance review: 19 out of 26 got this. 5 chose (C), 1 each chose (A) and (B).

Historical note (last time): 6 out of 26 got this. 9 chose (D), 5 each chose (B) and (C), 1 chose (A).

- (5) Suppose A and B are $n \times n$ matrices. Which of the following is true? Please see Option (E) before answering.
- (A) AB is invertible if and only if A and B are both invertible.
 - (B) AB is invertible if and only if at least one of A and B is invertible.
 - (C) If both A and B are invertible, then AB is invertible, but AB being invertible does not imply that both A and B are invertible.
 - (D) If AB is invertible, then both A and B are invertible. However, both A and B being invertible does not imply that AB is invertible.
 - (E) None of the above.

Answer: Option (A)

Explanation: If A and B are both invertible, then $(AB)^{-1} = B^{-1}A^{-1}$. If AB is invertible with inverse C , then $C(AB) = I_n$, so CA is an inverse for B , and $(AB)C = I_n$, so BC is an inverse for A . We are using the (somewhat nontrivial fact) that if a square matrix has a one-sided inverse, that inverse is actually a two-sided inverse.

Performance review: 16 out of 26 got this. 5 each chose (C) and (D).

Historical note (last time): 11 out of 26 got this. 8 chose (D), 3 each chose (C) and (E), 1 chose (B).

- (6) Suppose A and B are $n \times n$ matrices. Which of the following is true? We call a $n \times n$ matrix *idempotent* if it equals its own square. Please see Option (E) before answering.
- (A) AB is idempotent if and only if A and B are both idempotent.
 - (B) AB is idempotent if and only if at least one of A and B is idempotent.
 - (C) If both A and B are idempotent, then AB is idempotent, but AB being idempotent does not imply that both A and B are idempotent.
 - (D) If AB is idempotent, then both A and B are idempotent. However, both A and B being idempotent does not imply that AB is idempotent.
 - (E) None of the above.

Answer: Option (E)

Explanation: The following is an example where both A and B are idempotent but AB is not idempotent:

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

The product is:

$$AB = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

The following is an example where neither A nor B is idempotent but AB is idempotent:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

The product matrix is:

$$AB = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

Performance review: 6 out of 26 got this. 8 each chose (C) and (D), 3 chose (A), and 1 chose (B).

Historical note (last time: final): This question appeared in last year's final. 9 out of 30 got it correct at the time (note that students had significantly more exposure to the concepts by the time of the final). 12 chose (C), 4 each chose (A) and (D), 1 chose (B).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY NOVEMBER 6:
GEOMETRY OF LINEAR TRANSFORMATIONS (ABSTRACT)**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

27 people took this 14-question quiz. The score distribution was as follows:

- Score of 2: 1 person
- Score of 4: 1 person
- Score of 6: 3 people
- Score of 7: 2 people
- Score of 8: 4 people
- Score of 9: 4 people
- Score of 10: 5 people
- Score of 11: 4 people
- Score of 12: 2 people
- Score of 14: 1 person

The mean score was 8.8.

The question-wise answers and performance review are as follows:

- (1) Option (B): 18 people
- (2) Option (C): 22 people
- (3) Option (C): 17 people
- (4) Option (C): 18 people
- (5) Option (A): 10 people
- (6) Option (E): 22 people
- (7) Option (C): 25 people
- (8) Option (E): 16 people
- (9) Option (C): 13 people
- (10) Option (B): 17 people
- (11) Option (B): 12 people
- (12) Option (A): 20 people
- (13) Option (C): 15 people
- (14) Option (C): 13 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

This quiz tests for a deep abstract understanding of linear transformations and their geometry. It is related to Section 2.2 of the book and also to the **Geometry of linear transformations** lecture notes.

For the questions here, please use the following terminology.

Suppose n is a fixed natural number greater than 1. For ease of geometric visualization, you can take $n = 2$ for the discussion.

- A *linear automorphism* of \mathbb{R}^n is defined as a bijective linear transformation from \mathbb{R}^n to \mathbb{R}^n .
- An *affine linear automorphism* of \mathbb{R}^n is defined as a bijective function from \mathbb{R}^n to itself that preserves collinearity, i.e., it sends lines to lines. In addition, it preserves the ratios of lengths within each line. This can be included as part of the definition or deduced from the fact that collinearity is preserved for $n > 1$.

- A *self-isometry* of \mathbb{R}^n is defined as a bijective function from \mathbb{R}^n to itself that preserves Euclidean distance: for all pairs of points $\vec{x}, \vec{y} \in \mathbb{R}^n$, the Euclidean distance between \vec{x} and \vec{y} equals the Euclidean distance between $T(\vec{x})$ and $T(\vec{y})$.
- A *self-homothety* (or *similitude transformation* or *similarity transformation*) of \mathbb{R}^n is defined as a bijective function from \mathbb{R}^n to itself that multiplies all distances by a fixed number called the *factor of similitude* (dependent on the transformation): if the factor of similitude is λ , then for all pairs of points $\vec{x}, \vec{y} \in \mathbb{R}^n$, the distance between $T(\vec{x})$ and $T(\vec{y})$ equals λ times the distance between \vec{x} and \vec{y} .

- (1) What is the relationship between linear automorphisms and affine linear automorphisms of \mathbb{R}^n ?
- (A) Being a linear automorphism is precisely the same as being an affine linear automorphism.
 (B) Every linear automorphism is an affine linear automorphism, but not every affine linear automorphism is a linear automorphism.
 (C) Every affine linear automorphism is a linear automorphism, but not every linear automorphism is an affine linear automorphism.
 (D) A linear automorphism need not be affine linear, and an affine linear automorphism need not be linear.

Answer: Option (B)

Explanation: Linear automorphisms are precisely those affine linear automorphisms that send the origin to itself. In general, an affine linear automorphism can be expressed as a composite of a translation and a linear automorphism.

Performance review: 18 out of 27 got this. 8 chose (C), 1 chose (A).

Historical note (last time): 16 out of 27 got this. 10 chose (C), 1 chose (D).

- (2) What is the relationship between self-homotheties and self-isometries of \mathbb{R}^n ?
- (A) Being a self-homothety is precisely the same as being a self-isometry.
 (B) Every self-homothety is a self-isometry, but not every self-isometry is a self-homothety.
 (C) Every self-isometry is a self-homothety, but not every self-homothety is a self-isometry.
 (D) A self-homothety need not be a self-isometry, and a self-isometry need not be a self-homothety.

Answer: Option (C)

Explanation: We can define a self-isometry as a self-homothety where the factor of homothety (the factor of similitude) is 1.

Performance review: 22 out of 27 got this. 4 chose (B), 1 chose (D).

Historical note (last time): 21 out of 27 got this. 4 chose (D), 1 each chose (A) and (B).

- (3) What is the relationship between self-homotheties and affine linear automorphisms of \mathbb{R}^n ?
- (A) Being an affine linear automorphism is precisely the same as being a self-homothety.
 (B) Every affine linear automorphism is a self-homothety, but not every self-homothety is an affine linear automorphism.
 (C) Every self-homothety is an affine linear automorphism, but not every affine linear automorphism is a self-homothety.
 (D) An affine linear automorphism need not be a self-homothety, and a self-homothety need not be affine linear.

Answer: Option (C)

Explanation: A self-homothety must preserve ratios of distances. In particular, thanks to the triangle inequality, it must preserve collinearity with betweenness. In other words, if T is a self-homothety of \mathbb{R}^n and A, B, C are points in \mathbb{R}^n with B between A and C , then $T(A)$, $T(B)$, and $T(C)$ are also points in \mathbb{R}^n with $T(B)$ between $T(A)$ and $T(C)$, because of the triangle inequality.

There are affine linear automorphisms, such as shear automorphisms, that are not self-homotheties. Alternatively, we could also consider a dilation along one axis keeping the other axis fixed.

Performance review: 17 out of 27 got this. 6 chose (B), 4 chose (A).

Historical note (last time): 18 out of 27 got this. 6 chose (B), 2 chose (D), 1 chose (A).

- (4) What is the relationship between affine linear automorphisms and self-isometries of \mathbb{R}^n ?
- (A) Being an affine linear automorphism is precisely the same as being a self-isometry.
 (B) Every affine linear automorphism is a self-isometry, but not every self-isometry is an affine linear automorphism.

(C) Every self-isometry is an affine linear automorphism, but not every affine linear automorphism is a self-isometry.

(D) An affine linear automorphism need not be a self-isometry, and a self-isometry need not be affine linear.

Answer: Option (C)

Explanation: This can be deduced from the two preceding questions, and can also be seen directly.

Performance review: 18 out of 27 got this. 3 each chose (B) and (D), 2 chose (A), 1 left the question blank.

Historical note (last time): 22 out of 27 got this. 4 chose (D), 1 chose (A).

(5) There is a special kind of bijection from \mathbb{R}^n to \mathbb{R}^n called a *translation*. A translation with translation vector \vec{v} is defined as the bijection $\vec{x} \mapsto \vec{x} + \vec{v}$. A *nontrivial* translation is a translation whose translation vector is not the zero vector. Which of the following is an automorphism type that nontrivial translations are *not*? Please see Option (E) before answering.

(A) Linear automorphism

(B) Affine linear automorphism

(C) Self-isometry

(D) Self-homothety

(E) None of the above, i.e., nontrivial translations are of all these types

Answer: Option (A)

Explanation: Translations are self-isometries. Hence, they are also self-homotheties and affine linear automorphisms. They are not linear automorphisms because they do not preserve the origin.

Performance review: 10 out of 27 got this. 7 chose (E), 5 chose (C), 4 chose (D), 1 chose (B).

Historical note (last time): 12 out of 27 got this. 11 chose (E), 3 chose (C), and 1 chose (D).

(6) A collection of bijections from \mathbb{R}^n to itself is said to form a *group* if it satisfies all these three conditions:

- The composite of any two (possibly equal, possibly distinct) bijections in the collection is also in the collection.
- The identity bijection (i.e., the map sending every vector to itself) is in the collection.
- For every bijection in the collection, the inverse bijection is also in the collection.

For fixed n , which of the following collections of bijections from \mathbb{R}^n to itself does *not* form a group?

Please see Option (E) before answering.

(A) The collection of all linear automorphisms of \mathbb{R}^n

(B) The collection of all affine linear automorphisms of \mathbb{R}^n

(C) The collection of all self-isometries of \mathbb{R}^n

(D) The collection of all self-homotheties of \mathbb{R}^n

(E) None of the above, i.e., each of them is a group

Answer: Option (E)

Explanation: This is obvious from the definition and concept of symmetry.

Performance review: 22 out of 27 got this. 2 each chose (B) and (D), 1 chose (C).

Historical note (last time): 24 out of 27 got this. 2 chose (C), 1 chose (D).

For the remaining questions, we deal with the case $n = 2$.

We consider two special types of bijections from \mathbb{R}^2 to \mathbb{R}^2 : *rotations* (a rotation is specified by the center of rotation and the angle of rotation) and *reflections* (a reflection is specified by the line of reflection).

The identity map (i.e., the map sending every point to itself) is considered both a translation and a rotation. It is the translation by the zero vector. It can be viewed as a rotation about any point by the zero angle.

Note that for a rotation, the angle of rotation is determined uniquely up to additive multiples of 2π . The center of rotation is determined uniquely for all nontrivial rotations.

(7) What is the composite of two rotations centered at the same point in \mathbb{R}^2 ? Assume for simplicity that the composite is not the identity, i.e., the two rotations do not cancel each other. Note that the rotations must commute, so the order of operation does not matter.

(A) It must be a reflection about a line passing through that center point.

- (B) It must be a reflection about a line *not* passing through that center point.
- (C) It must be a rotation centered at the same point
- (D) It must be a rotation but it need not be centered at the same point.
- (E) It must be a translation.

Answer: Option (C)

Explanation: This is physically obvious.

Performance review: 25 out of 27 got this. 1 chose (D), 1 left the question blank.

Historical note (last time): 18 out of 27 got this. 5 chose (D), 3 chose (E), 1 chose (A).

- (8) What is the composite of two reflections about lines in \mathbb{R}^2 , if the two lines of reflection are known to be parallel but distinct? Although the two reflections do not commute, the *type* of their composite does not depend upon the order in which we compose them.

- (A) It must be a reflection about a third line which is parallel to both the lines and is equidistant from them.
- (B) It must be a reflection about a third line which is perpendicular to both the lines.
- (C) It must be a rotation about a point that is equidistant from both lines.
- (D) It must be a translation by a vector parallel to the lines about which we are reflecting.
- (E) It must be a translation by a vector perpendicular to the lines about which we are reflecting.

Answer: Option (E)

Explanation: The vector will in fact be twice the perpendicular difference vector between the lines.

You can think about it in terms of double mirrors.

Performance review: 16 out of 27 got this. 6 chose (A), 3 chose (D), 1 chose (B).

Historical note (last time): 20 out of 27 got this. 3 chose (D), 2 each chose (A) and (C).

- (9) What is the composite of two reflections about lines in \mathbb{R}^2 , if the two lines of reflection are distinct and intersect? Once again, the reflections do not in general commute, but the *type* of the composite does not depend on the order of composition.

- (A) It must be a reflection about a third line which passes through the point of intersection of the two lines of reflection.
- (B) It must be a reflection about a third line which does not pass through the point of intersection of the two lines of reflection.
- (C) It must be a rotation about the point of intersection.
- (D) It must be a translation by a vector that makes equal angles with both the lines.
- (E) It need not be a translation, rotation, or reflection.

Answer: Option (C)

Explanation: In fact, it will be a rotation with the angle of rotation equal to *twice* the angle between the lines. We can easily see this using basic Euclidean geometry. The simplest case to think of is the case of perpendicular lines of reflection. In this case, the composite is a rotation by π , also known as a half-turn.

Performance review: 13 out of 27 people got this. 8 chose (A), 2 each chose (D) and (E), 1 chose (B).

Historical note (last time): 18 out of 27 got this. 5 chose (A), 2 each chose (B) and (E).

- (10) What is the composite of a nontrivial rotation in \mathbb{R}^2 (i.e., the angle of rotation is not a multiple of 2π) and a nontrivial translation?

- (A) It must be a rotation with the same center of rotation but with a different angle of rotation.
- (B) It must be a rotation with the same angle of rotation but with a different center of rotation.
- (C) It must be a reflection about a line passing through the center of rotation.
- (D) It must be a reflection about a line *not* passing through the center of rotation.
- (E) It must be a translation.

Answer: Option (B)

Explanation: The angle of rotation must remain the same. We can see this by imagining some kind of physical figure that is made to undergo the rotation and then the translation. The rotation changes the direction of the figure by the angle of rotation. The translation then preserves the direction of the figure. The composite should also change the direction of the figure by the same

angle, hence must be a rotation by the same angle. The center of rotation may well be different, and in fact, must be different if the translation is nontrivial.

Performance review: 17 out of 27 got this. 3 each chose (A) and (E), 2 chose (D), 1 chose (C), 1 left the question blank.

Historical note (last time): 16 out of 27 got this. 7 chose (A), 2 chose (C), 1 each chose (D) and (E).

- (11) An affine linear automorphism of \mathbb{R}^2 is termed *area-preserving* if it preserves areas, i.e., the area of the image of any triangle under the automorphism is the same as the area of the original triangle. What is the relation between being a self-isometry and being an area-preserving affine linear automorphism of \mathbb{R}^2 ?
- (A) Being a self-isometry is precisely the same as being an area-preserving affine linear automorphism.
- (B) Every self-isometry is area-preserving, but not every area-preserving affine linear automorphism is a self-isometry.
- (C) Every area-preserving affine linear automorphism is a self-isometry, but not every self-isometry is area-preserving.
- (D) A self-isometry need not be an area-preserving affine linear automorphism, and an area-preserving affine linear automorphism need not be a self-isometry.

Answer: Option (B)

Explanation: If something preserves lengths, it sends triangles to congruent triangles, hence it also preserves areas.

There can be area-preserving affine linear automorphisms that are not self-isometries. Shear operations, such as the one with this matrix:

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

are examples. The point is that areas still get preserved. For instance, a triangle based on the x -axis gets bent out of shape, but its base and height remain the same, hence its area remains the same.

Performance review: 12 out of 27 got this. 9 chose (A), 4 chose (C), 1 chose (D), 1 left the question blank.

Historical note (last time): 15 out of 27 got this. 5 chose (A), 3 each chose (C) and (D), 1 left the question blank.

- (12) An affine linear automorphism of \mathbb{R}^2 is termed *orientation-preserving* if it preserves orientation, i.e., it does not interchange left with right. An affine linear automorphism of \mathbb{R}^2 is termed *orientation-reversing* if it reverses orientation, i.e., it interchanges the roles of left and right. Obviously, the composite of two orientation-preserving affine linear automorphisms is orientation-preserving. What can we say about the composite of two orientation-reversing affine linear automorphisms?
- (A) It must be orientation-preserving
- (B) It must be orientation-reversing
- (C) It may be orientation-preserving or orientation-reversing

Answer: Option (A)

Explanation: Reversing the orientation means switching the roles of left and right. Reversing the orientation a second time means switching the roles back. Thus, the composite of orientation-preserving.

Performance review: 20 out of 27 got this. 6 chose (C), 1 left the question blank.

Historical note (last time): 24 out of 27 got this. 2 chose (B), 1 chose (C).

- (13) The linear automorphism of \mathbb{R}^2 with matrix:

$$\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

is an example of a *shear automorphism*. Which of the following is this automorphism *not*? Please see options (D) and (E) before answering.

- (A) Area-preserving
- (B) Orientation-preserving
- (C) Self-isometry
- (D) None of the above, i.e., it is area-preserving, orientation-preserving, and a self-isometry
- (E) All of the above, i.e., it is not area-preserving, not orientation-preserving, and not a self-isometry of \mathbb{R}^2 .

Answer: Option (C)

Explanation: The vector $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ maps to the vector $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$. The original vector had length 1, but the new vector has length $\sqrt{2}$, so the map does not preserve length. It does preserve area and orientation.

Performance review: 15 out of 27 got this. 6 chose (E), 3 chose (D), 2 chose (A).

Historical note (last time): 16 out of 27 got this. 4 each chose (A) and (D), 3 chose (B).

- (14) Which of the following is guaranteed to send any triangle in \mathbb{R}^2 to a similar triangle? Please see Options (D) and (E) before answering.

- (A) Linear automorphism
- (B) Affine linear automorphism
- (C) Self-homothety
- (D) All of the above
- (E) None of the above

Answer: Option (C)

Explanation: Follows from the definition.

Performance review: 13 out of 27 got this. 10 chose (D), 2 chose (E), 1 chose (B), 1 left the question blank.

Historical note (last time): 20 out of 27 got this. 5 chose (D), 1 each chose (A) and (E).

**DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE WEDNESDAY NOVEMBER 6:
IMAGE AND KERNEL (BASIC)**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

28 people took this 5-question quiz. The score distribution was as follows:

- Score of 2: 2 people
- Score of 3: 7 people
- Score of 4: 9 people
- Score of 5: 10 people

The mean score was slightly under 4.

The question-wise answers and performance review are below:

- (1) Option (B): 26 people
- (2) Option (A): 22 people
- (3) Option (C): 17 people
- (4) Option (B): 19 people
- (5) Option (C): 27 people

2. SOLUTIONS

PLEASE DO NOT DISCUSS ANY QUESTIONS.

The questions here test for a very rudimentary understanding of the ideas covered in the lectures notes titled **Image and kernel of a linear transformation**. The corresponding section of the book is Section 3.1.

- (1) *Do not discuss this!* For a linear transformation $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$, the kernel of T is defined as the set of vectors $\vec{x} \in \mathbb{R}^m$ satisfying the condition that $T(\vec{x}) = \vec{0}$. Which of the following correctly describes the type of subset of \mathbb{R}^m that the kernel must be? Note that, as usual, we identify a set of vectors with the set of corresponding points.
 - (A) The kernel is a line segment in \mathbb{R}^m .
 - (B) The kernel is a linear subspace of \mathbb{R}^m , i.e., it passes through the origin and, for any two points in the kernel, the line joining them is completely inside the kernel.
 - (C) The kernel is an affine linear subspace of \mathbb{R}^m but it need not be linear, i.e., it is non-empty and the line joining any two points in it is also in it, but it need not contain the origin.
 - (D) The kernel is a curve in \mathbb{R}^m with a parametric description.

Answer: Option (B)

Explanation: See Section 4.3 of the lecture notes titled **Image and kernel of a linear transformation**.

Briefly: we can readily verify that $T(\vec{0}) = \vec{0}$, and we can verify that the kernel is a linear subspace based on our earlier definition (closed under addition and scalar multiplication). It's easy to see that this also coincides with our new definition of linear subspace.

Performance review: 26 out of 28 got this. 2 chose (C).

- (2) *Do not discuss this!* For a linear transformation $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$, the kernel of T is defined as the set of vectors $\vec{x} \in \mathbb{R}^m$ satisfying the condition that $T(\vec{x}) = \vec{0}$. Given a vector $\vec{y} \in \mathbb{R}^n$, the set of solutions to $T(\vec{x}) = \vec{y}$ is either empty, or it bears some relation with the kernel. What relation does it bear to the kernel if it is nonempty?

- (A) The solution set is an affine linear subspace of \mathbb{R}^m (see definition in Option (C) of Q1) that is a translate of the kernel, i.e., there is a vector \vec{v} such that the vectors in the solution set are precisely the vectors expressible as (\vec{v} plus a vector in the kernel).
- (B) The solution set coincides precisely with the kernel.
- (C) The solution set comprises a single point (i.e., a single vector) that is not in the kernel.

Answer: Option (A)

Explanation: See Section 5 of the lecture notes titled **Image and kernel of a linear transformation**.

Performance review: 22 out of 28 got this. 5 chose (B), 1 chose (C).

- (3) *Do not discuss this!:* Given a linear transformation $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$, recall that we say that T is *injective* if for every $\vec{y} \in \mathbb{R}^n$, there exists *at most one* $\vec{x} \in \mathbb{R}^m$ satisfying $T(\vec{x}) = \vec{y}$. Another way of formulating this is that if A is the $n \times m$ matrix for T , then the linear system $A\vec{x} = \vec{y}$ has at most one solution for \vec{x} for each fixed value of \vec{y} . We had earlier worked out that this condition is equivalent to full column rank (recall: all the variables need to be leading variables), which in this case means rank m .

What is the relationship between the injectivity of T and the kernel of T ?

- (A) T is injective if and only if the kernel of T is empty.
- (B) If T is injective, then the kernel of T is empty. However, the converse is not in general true.
- (C) T is injective if and only if the kernel of T comprises only the zero vector.
- (D) If T is injective, then the kernel of T comprises only the zero vector. However, the converse is not in general true.
- (E) If the kernel of T comprises only the zero vector, then T is injective. However, the converse is not in general true.

Answer: Option (C)

Explanation: See Sections 5 and 6 of the lecture notes titled **Image and kernel of a linear transformation**.

Performance review: 17 out of 28 got this. 8 chose (D), 3 chose (E).

- (4) *Do not discuss this!:* For a linear transformation $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$, the image of T is defined as the set of vectors $\vec{y} \in \mathbb{R}^n$ satisfying the condition that there exists a vector $\vec{x} \in \mathbb{R}^m$ satisfying $T(\vec{x}) = \vec{y}$. In other words, the image of T equals the range of T as a function. Which of the following correctly describes the type of subset of \mathbb{R}^n that the image must be? Note that, as usual, we identify a set of vectors with the set of corresponding points.

- (A) The image is a line segment in \mathbb{R}^n .
- (B) The image is a linear subspace of \mathbb{R}^n , i.e., it passes through the origin and, for any two points in the image, the line joining them is completely inside the image.
- (C) The image is an affine linear subspace of \mathbb{R}^n but it need not be linear, i.e., it is non-empty and the line joining any two points in it is also in it, but it need not contain the origin.
- (D) The image is a curve in \mathbb{R}^n with a parametric description.

Answer: Option (B)

Explanation: See Section 4.1 of the **Image and kernel of a linear transformation** lecture notes.

Performance review: 19 out of 28 got this. 7 chose (C), 2 chose (A).

- (5) *Do not discuss this!:* Given a linear transformation $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$, recall that we say that T is *surjective* if for every $\vec{y} \in \mathbb{R}^n$, there exists *at least one* $\vec{x} \in \mathbb{R}^m$ satisfying $T(\vec{x}) = \vec{y}$. Another way of formulating this is that if A is the $n \times m$ matrix for T , then the linear system $A\vec{x} = \vec{y}$ has at least one solution for \vec{x} for each fixed value of \vec{y} . We had earlier worked out that this condition is equivalent to full row rank (recall: we need all rows in the rref to be nonzero in order to avoid the potential for inconsistency), which in this case means rank n .

What is the relationship between the surjectivity of T and the image of T ?

- (A) T is surjective if and only if the image of T is empty.
- (B) T is surjective if and only if the image of T comprises only the zero vector.
- (C) T is surjective if and only if the image of T is all of \mathbb{R}^n .

Answer: Option (C)

Explanation: This is obvious from the definition.

Performance review: 27 out of 28 got this. 1 chose (B).

**DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE FRIDAY NOVEMBER 8
(DELAYED TO MONDAY NOVEMBER 11): IMAGE AND KERNEL
(COMPUTATIONAL)**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

26 people took this 3-question quiz. The score distribution was as follows:

- Score of 0: 2 people
- Score of 1: 5 people
- Score of 2: 8 people
- Score of 3: 11 people

The mean score was 2.08.

The question-wise answers and performance review were as follows:

- (1) Option (E): 20 people
- (2) Option (A): 15 people
- (3) Option (B): 19 people

Note: Performance on these questions was notably better than last year, even though last year these questions were in the 'please feel free to discuss' category. Seems like people are understanding the image and kernel better this year!

2. SOLUTIONS

PLEASE DO NOT DISCUSS ANY QUESTIONS.

- (1) *Do not discuss this!* Consider the linear transformation $\text{Avg} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ defined as:

$$\text{Avg} = \begin{bmatrix} x \\ y \end{bmatrix} \mapsto \begin{bmatrix} (x+y)/2 \\ (x+y)/2 \end{bmatrix}$$

What can we say about the kernel and image of Avg? Note that in our descriptions of the kernel and the image below, we use x to denote the first coordinate of the vector and y to denote the second coordinate of the vector.

Note: One way you can do that is to write the matrix for Avg, but in this particular situation, it's easiest to just do things directly.

- (A) The kernel is the zero subspace and the image is all of \mathbb{R}^2
- (B) The kernel is the line $y = x$ and the image is also the line $y = x$
- (C) The kernel is the line $y = x$ and the image is the line $y = -x$
- (D) The kernel is the line $y = -x$ and the image is also the line $y = -x$
- (E) The kernel is the line $y = -x$ and the image is the line $y = x$

Answer: Option (E)

Explanation: The kernel must satisfy that both coordinates of the output are zero, so $(x+y)/2 = 0$ and $(x+y)/2 = 0$. The solution is $y = -x$.

The image must satisfy that both coordinates are equal, so it is the line $y = x$.

Explicitly, the matrix in question is:

$$\begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix}$$

The ref for this is:

$$\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$$

The kernel is generated by the vector $\begin{bmatrix} -1 \\ 1 \end{bmatrix}$. The image is generated by $\begin{bmatrix} 1/2 \\ 1/2 \end{bmatrix}$.

Performance review: 20 out of 26 got this. 5 chose (D), 1 chose (B).

Historical note (last time): 10 out of 26 got this. 7 chose (B), 6 chose (D), 3 chose (C).

- (2) *Do not discuss this!:* Consider the *average of other two* linear transformation $\nu : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given as follows:

$$\nu = \begin{bmatrix} x \\ y \\ z \end{bmatrix} \mapsto \begin{bmatrix} (y+z)/2 \\ (z+x)/2 \\ (x+y)/2 \end{bmatrix}$$

What can we say about the kernel and image of ν ?

Note that in our descriptions of the kernel and the image below, we use x to denote the first coordinate of the vector, y to denote the second coordinate of the vector, and z to denote the third coordinate of the vector.

Note: This can both be reasoned directly (without any knowledge of linear algebra) or alternatively it can be done by writing the matrix of ν and computing its rank, image, and kernel.

- (A) The kernel is the zero subspace and the image is all of \mathbb{R}^3
- (B) The kernel is the line $x = y = z$ (one-dimensional) and the image is the plane $x + y + z = 0$ (two-dimensional)
- (C) The kernel is the plane $x + y + z = 0$ (two-dimensional) and the image is the line $x = y = z$ (one-dimensional)
- (D) The kernel is the plane $x = y = z$ (two-dimensional) and the image is the line $x + y + z = 0$ (one-dimensional)
- (E) The kernel is the line $x + y + z = 0$ (one-dimensional) and the image is the plane $x = y = z$ (two-dimensional)

Answer: Option (A)

Explanation: We can check that if all three outputs are zero, then $x = y = z = 0$. Alternatively, we can verify that the matrix:

$$\begin{bmatrix} 0 & 1/2 & 1/2 \\ 1/2 & 0 & 1/2 \\ 1/2 & 1/2 & 0 \end{bmatrix}$$

has full rank 3. Its inverse is the matrix:

$$\begin{bmatrix} -1 & 1 & 1 \\ 1 & -1 & 1 \\ 1 & 1 & -1 \end{bmatrix}$$

Intuitively, x is the sum of the second and third output minus the first output, and so on.

Performance review: 15 out of 26 got this. 6 chose (C), 5 chose (B).

Historical note (last time): 4 out of 26 got this. 11 chose (E), 7 chose (D), 2 each chose (B) and (C).

- (3) *Do not discuss this!:* Consider the *difference of other two* linear transformation $\mu : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by:

$$\mu = \begin{bmatrix} x \\ y \\ z \end{bmatrix} \mapsto \begin{bmatrix} y - z \\ z - x \\ x - y \end{bmatrix}$$

What can we say about the kernel and image of μ ?

Note that in our descriptions of the kernel and the image below, we use x to denote the first coordinate of the vector, y to denote the second coordinate of the vector, and z to denote the third coordinate of the vector.

Note: This can both be reasoned directly (without any knowledge of linear algebra) or alternatively it can be done by writing the matrix of μ and computing its rank, image, and kernel.

- (A) The kernel is the zero subspace and the image is all of \mathbb{R}^3
- (B) The kernel is the line $x = y = z$ (one-dimensional) and the image is the plane $x + y + z = 0$ (two-dimensional)
- (C) The kernel is the plane $x + y + z = 0$ (two-dimensional) and the image is the line $x = y = z$ (one-dimensional)
- (D) The kernel is the plane $x = y = z$ (two-dimensional) and the image is the line $x + y + z = 0$ (one-dimensional)
- (E) The kernel is the line $x + y + z = 0$ (one-dimensional) and the image is the plane $x = y = z$ (two-dimensional)

Answer: Option (B)

Explanation: The kernel must satisfy that all the three output coordinates are 0. This means that $y - z = 0$, $z - x = 0$, and $x - y = 0$, so we get $x = y = z$. The image must satisfy that the sum of the three coordinates is zero, so it lies in the plane $x + y + z = 0$. A little more effort can show that it is equal to the entire plane.

The matrix for the linear transformation μ is:

$$\begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}$$

We can row reduce this to get the rref:

$$\begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \\ 0 & 0 & 0 \end{bmatrix}$$

The rref shows that the linear transformation has rank two. The third variable is non-leading, so the kernel is generated by the vector:

$$\begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

Explicitly, the kernel is the line $x = y = z$.

For computing the image: the first two variables are leading variables, so the image of the linear transformation is the space spanned by the first two columns of the original matrix, i.e., it is the subspace spanned by the vectors:

$$\begin{bmatrix} 0 \\ -1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ -1 \end{bmatrix}$$

We want to express this as a plane. The appropriate approach to determining the equation of the plane is to take the cross product of the vectors, albeit that is a construct specific to three dimensions. Alternatively, we can solve another linear system. Doing this from scratch is somewhat beyond our current scope, but it is relatively easy to figure out the plane from the collection of options presented.

Performance review: 19 out of 26 got this. 4 chose (D), 2 chose (A), 1 chose (E).

Historical note (last time): 8 out of 26 got this. 15 chose (D), 2 chose (C), and 1 chose (E).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY NOVEMBER 8: LINEAR TRANSFORMATIONS: SEEDS FOR REAPING LATER

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 13-question quiz. The score distribution was as follows:

- Score of 4: 3 people
- Score of 5: 1 person
- Score of 6: 4 people
- Score of 7: 3 people
- Score of 8: 2 people
- Score of 9: 4 people
- Score of 10: 3 people
- Score of 11: 5 people

The question-wise answers and performance review are below:

- (1) Option (E): 21 people
- (2) Option (D): 15 people
- (3) Option (D): 14 people
- (4) Option (C): 18 people
- (5) Option (D): 23 people
- (6) Option (B): 21 people
- (7) Option (E): 14 people
- (8) Option (C): 5 people
- (9) Option (A): 16 people
- (10) Option (A): 15 people
- (11) Option (B): 17 people
- (12) Option (E): 4 people
- (13) Option (E): 16 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

In this quiz, we will sow the seeds of ideas that we will reap later. There are two broad classes of ideas that we touch upon here:

- Conjugation, similarity transformations, and products of matrices: This will be of relevance later when we discuss change of coordinates. We cover change of coordinates in more detail in Section 3.4 of the text.
 - Kernel and image for linear transformations arising from calculus, typically for infinite-dimensional spaces: This will be helpful in understanding linear transformations in an *abstract* sense, a topic that we cover in more detail in Chapter 4 of the text.
- (1) Suppose A and B are (possibly equal, possibly distinct) $n \times n$ matrices for some $n > 1$. Recall that the *trace* of a matrix is defined as the sum of its diagonal entries. Suppose $C = AB$ and $D = BA$. Which of the following is true?
 - (A) It must be the case that $C = D$
 - (B) The *set* of entry values in C is the same as the set of entry values in D , but they may appear in a different order.

- (C) C and D need not be equal, but the sum of all the matrix entries of C must equal the sum of all the matrix entries of D .
- (D) C and D need not be equal, but they have the same diagonal, i.e., every diagonal entry of C equals the corresponding diagonal entry of D .
- (E) C and D need not be equal and they need not even have the same diagonal. However, they must have the same trace, i.e., the sum of the diagonal entries of C equals the sum of the diagonal entries of D .

Answer: Option (E)

Explanation: We have the following formula for the i^{th} diagonal entry of the product:

$$c_{ii} = \sum_{j=1}^n a_{ij}b_{ji}$$

The sum of the diagonal entries of C is thus:

$$\sum_{i=1}^n c_{ii} = \sum_{i=1}^n \sum_{j=1}^n a_{ij}b_{ji}$$

The j^{th} diagonal entry of D is:

$$d_{jj} = \sum_{i=1}^n b_{ji}a_{ij}$$

The sum of the diagonal entries of D is thus:

$$\sum_{j=1}^n d_{jj} = \sum_{j=1}^n \sum_{i=1}^n b_{ji}a_{ij}$$

Thus, the sum of the diagonal entries of C is the same as the sum of the diagonal entries of D .

Performance review: 21 out of 25 got this. 2 chose (C), 1 each chose (B) and (D).

Historical note (last time): 23 out of 26 got this. 1 each chose (B), (C), and (D).

Suppose A is an invertible $n \times n$ matrix. The *conjugation operation* corresponding to A is the map that sends any $n \times n$ matrix X to AXA^{-1} . We can verify that the following hold for any two (possibly equal, possibly distinct) $n \times n$ matrices X and Y :

$$\begin{aligned} A(X+Y)A^{-1} &= AXA^{-1} + AYA^{-1} \\ A(XY)A^{-1} &= (AXA^{-1})(AYA^{-1}) \\ AX^rA^{-1} &= (AXA^{-1})^r \end{aligned}$$

The conceptual significance of this will (hopefully!) become clearer as we proceed.

- (2) Which of the following is guaranteed to be the same for X and AXA^{-1} ?
- (A) The sum of all entries
- (B) The sum of squares of all entries
- (C) The product of all entries
- (D) The sum of all diagonal entries (i.e., the trace)
- (E) The sum of squares of all diagonal entries

Answer: Option (D)

Explanation: We can write $X = A^{-1}(AX)$, whereas $AXA^{-1} = (AX)A^{-1}$. Thus, both X and AXA^{-1} are products of two matrices A^{-1} and AX but in opposite orders. Hence, by the preceding question, they have the same trace.

Performance review: 15 out of 25 got this. 3 each chose (A) and (C), 2 each chose (B) and (E).

Historical note (last time): 20 out of 26 got this. 5 chose (A), 1 chose (C).

- (3) A and X are $n \times n$ matrices, with A invertible. Which of the following is/are true? Please see Options (D) and (E) before answering, and select a single option that best reflects your view.
- (A) X is invertible if and only if AXA^{-1} is invertible.
 - (B) X is nilpotent if and only if AXA^{-1} is nilpotent.
 - (C) X is idempotent if and only if AXA^{-1} is idempotent.
 - (D) All of the above.
 - (E) None of the above.

Answer: Option (D)

Explanation: Essentially, the conjugation operation preserves all aspects of the multiplicative structure, hence it preserves the properties of being invertible, nilpotent, and idempotent.

Let us illustrate this with idempotent. We have that:

$$AX^2A^{-1} = (AXA^{-1})^2$$

If $X^2 = X$, we get:

$$AXA^{-1} = (AXA^{-1})^2$$

showing that AXA^{-1} is also idempotent. We can work backwards to show that the reverse implication also holds.

Performance review: 14 out of 25 got this. 4 chose (B), 3 chose (A), 2 chose (D), 1 chose (C), and 1 left the question blank.

Historical note (last time): 18 out of 26 got this. 3 each chose (A) and (B), 2 chose (C).

- (4) A and X are $n \times n$ matrices, with A invertible. Which of the following is equivalent to the condition that $AXA^{-1} = X$?
- (A) $A + X = X + A$
 - (B) $A - X = X - A$
 - (C) $AX = XA$
 - (D) $XA^{-1} = AX^{-1}$
 - (E) None of the above

Answer: Option (C)

Explanation: Start with:

$$AXA^{-1} = X$$

Multiply both sides of the equation on the *right* with the matrix A . We get:

$$AX = XA$$

Note that the algebraic manipulation is reversible: starting from $AX = XA$ and multiplying both sides by A^{-1} on the right gives the original relation.

Performance review: 18 out of 25 got this. 4 chose (E), 1 each chose (A), (B), and (D).

Historical note (last time): 18 out of 26 got this. 4 chose (D) 2 chose (A), 1 each chose (B) and (E).

Let's look at a computational application of matrix conjugation.

One computational application is power computation. Suppose we have a $n \times n$ matrix B and we need to compute B^r for a very large r . This requires $O(\log_2 r)$ multiplications, but note that each multiplication, if done naively, takes time $O(n^3)$ for a generic matrix. Suppose, however, that there exists a matrix A such that the matrix $C = ABA^{-1}$ is diagonal. If we can find A (and hence C) efficiently, then we can compute $C^r = (ABA^{-1})^r = AB^rA^{-1}$, and therefore $B^r = A^{-1}C^rA$. Note that each multiplication of diagonal matrices takes $O(n)$ multiplications, so this reduces the overall

arithmetic complexity from $O(n^3 \log_2 r)$ to $O(n \log_2 r)$. Note, however, that this is contingent on our being able to find the matrices A and C first. We will later see a method for finding A and C . Unfortunately, this method relies on finding the set of solutions to a polynomial equation of degree n , which requires operations that go beyond ordinary arithmetic operations of addition, subtraction, multiplication, and division. Even in the case $n = 2$, it requires solving a quadratic equation. We do have the formula for that.

- (5) Consider the following example of the above general setup with $n = 2$:

$$B = \begin{bmatrix} 1 & -1 \\ 0 & 2 \end{bmatrix}$$

We can choose:

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

The matrix $C = ABA^{-1}$ is a diagonal matrix. What diagonal matrix is it?

- (A) $\begin{bmatrix} 1 & 0 \\ 0 & -2 \end{bmatrix}$
 (B) $\begin{bmatrix} -1 & 0 \\ 0 & 2 \end{bmatrix}$
 (C) $\begin{bmatrix} 2 & 0 \\ 0 & -1 \end{bmatrix}$
 (D) $\begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$
 (E) $\begin{bmatrix} -1 & 0 \\ 0 & -2 \end{bmatrix}$

Answer: Option (D)

Explanation: We can just carry out the matrix multiplication. Note that:

$$A^{-1} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}$$

A *sanity check* is that the new matrix C should have the same trace as the original matrix B , because the trace is invariant under conjugation. Among the options, the only matrix with the correct trace (3) is Option (D).

Performance review: 23 out of 25 got this. 2 chose (B).

Historical note (last time): 23 out of 26 got this. 2 chose (B), 1 chose (C).

- (6) With A , B , and C as in the preceding question, what is the value of B^8 ? Use that $2^8 = 256$.

- (A) $\begin{bmatrix} 1 & -1 \\ 0 & 256 \end{bmatrix}$
 (B) $\begin{bmatrix} 1 & -255 \\ 0 & 256 \end{bmatrix}$
 (C) $\begin{bmatrix} 1 & 253 \\ 0 & 256 \end{bmatrix}$
 (D) $\begin{bmatrix} 1 & 253 \\ 254 & 256 \end{bmatrix}$
 (E) $\begin{bmatrix} 16 & -8 \\ 0 & 256 \end{bmatrix}$

Answer: Option (B)

Explanation: We calculate:

$$C^8 = \begin{bmatrix} 1 & 0 \\ 0 & 256 \end{bmatrix}$$

We now recover B^8 as $A^{-1}C^8A$.

Performance review: 21 out of 25 got this. 2 chose (A), 1 each chose (C) and (E).

Historical note (last time): 22 out of 26 got this. 3 chose (C), 1 chose (E).

- (7) Suppose $n > 1$. Let A be a $n \times n$ matrix such that the linear transformation corresponding to A is a self-isometry of \mathbb{R}^n , i.e., it preserves distances. Which of the following must necessarily be true? You can use the case $n = 2$ and the example of rotations to guide your thinking.

- (A) The trace of A (i.e., the sum of the diagonal entries of A) must be equal to 0
- (B) The trace of A (i.e., the sum of the diagonal entries of A) must be equal to 1
- (C) The sum of the entries in each column of A must be equal to 1
- (D) The sum of the absolute values of the entries in each column of A must be equal to 1
- (E) The sum of the squares of the entries in each column of A must be equal to 1

Answer: Option (E)

Explanation: The linear transformation corresponding to A preserves lengths of vectors. Thus, the images of the standard basis vectors are all unit vectors. Recall that the images of the standard basis vectors under the linear transformation corresponding to a matrix are the columns of that matrix. Therefore, each column of the matrix is a unit vector, i.e., the sum of the squares of the coordinates there is 1.

As an illustration, consider the case of a rotation matrix in two dimensions:

$$\begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

Performance review: 14 out of 25 got this. 6 chose (B), 3 chose (A), 2 chose (D).

Historical note (last time): 12 out of 26 got this. 4 each chose (A), (C), and (D). 1 chose (B). 1 left the question blank.

A *real vector space* (just called *vector space* for short) is a set V equipped with the following structures:

- A binary operation $+$ on V called addition that is commutative and associative.
- A special element $0 \in V$ that is an identity for addition.
- A scalar multiplication operation $\mathbb{R} \times V \rightarrow V$ denoted by concatenation such that:
 - $0\vec{v} = 0$ (the 0 on the right side being the vector 0) for all $\vec{v} \in V$.
 - $1\vec{v} = \vec{v}$ for all $\vec{v} \in V$.
 - $a(b\vec{v}) = (ab)\vec{v}$ for all $a, b \in \mathbb{R}$ and $\vec{v} \in V$.
 - $a(\vec{v} + \vec{w}) = a\vec{v} + a\vec{w}$ for all $a \in \mathbb{R}$ and $\vec{v}, \vec{w} \in V$.
 - $(a + b)\vec{v} = a\vec{v} + b\vec{v}$ for all $a, b \in \mathbb{R}$, $\vec{v} \in V$.

A *subspace* of a vector space is defined as a nonempty subset that is closed under addition and scalar multiplication. In particular, any subspace must contain the zero vector. A subspace of a vector space can be viewed as being a vector space in its own right.

Suppose V and W are vector spaces. A function $T : V \rightarrow W$ is termed a *linear transformation* if T preserves addition and scalar multiplication, i.e., we have the following two conditions:

- $T(\vec{v}_1 + \vec{v}_2) = T(\vec{v}_1) + T(\vec{v}_2)$ for all vectors $\vec{v}_1, \vec{v}_2 \in V$.
- $T(a\vec{v}) = aT(\vec{v})$ for all $a \in \mathbb{R}$, $\vec{v} \in V$.

The *kernel* of a linear transformation T is defined as the set of all vectors \vec{v} such that $T(\vec{v})$ is the zero vector. The *image* of a linear transformation T is defined as its range as a set map.

Denote by $C(\mathbb{R})$ (or alternatively by $C^0(\mathbb{R})$) the vector space of all continuous functions from \mathbb{R} to \mathbb{R} , with pointwise addition and scalar multiplication. Note that the elements of this vector space, which we would ordinarily call “vectors”, are now *functions*.

For k a positive integer, denote by $C^k(\mathbb{R})$ the subspace of $C(\mathbb{R})$ comprising those continuous functions that are at least k times *continuously* differentiable. Note that $C^{k+1}(\mathbb{R})$ is a subspace of $C^k(\mathbb{R})$, so we have a descending chain of subspaces:

$$C(\mathbb{R}) = C^0(\mathbb{R}) \supseteq C^1(\mathbb{R}) \supseteq C^2(\mathbb{R}) \supseteq \dots$$

The intersection of these spaces is the vector space $C^\infty(\mathbb{R})$, defined as the subspace of $C(\mathbb{R})$ comprising those functions that are *infinitely* differentiable.

- (8) We can think of differentiation as a linear transformation. Of the following options, which is the broadest way of viewing differentiation as a linear transformation? By “broadest” we mean “with the largest domain that makes sense among the given options.”

- (A) From $C^\infty(\mathbb{R})$ to $C^\infty(\mathbb{R})$
- (B) From $C^0(\mathbb{R})$ to $C^1(\mathbb{R})$
- (C) From $C^1(\mathbb{R})$ to $C^0(\mathbb{R})$
- (D) From $C^1(\mathbb{R})$ to $C^2(\mathbb{R})$
- (E) From $C^2(\mathbb{R})$ to $C^1(\mathbb{R})$

Answer: Option (C)

Explanation: $C^1(\mathbb{R})$ is the space of continuously differentiable functions, and differentiating any continuously differentiable function gives rise to a continuous function. We cannot go as broad as $C^0(\mathbb{R})$ because not all functions in $C^0(\mathbb{R})$ are differentiable. For instance, the absolute value function is not differentiable at 0.

Note that everything in $C^0(\mathbb{R})$ does get hit, because every continuous function is the derivative of its antiderivative.

Performance review: 5 out of 25 got this. 8 chose (A), 6 chose (B), 5 chose (E), 1 chose (D).

Historical note (last time): 18 out of 26 got this. 4 chose (A), 2 each chose (B) and (E).

- (9) Under the differentiation linear transformation, what is the image of $C^k(\mathbb{R})$ for a positive integer k ?

- (A) $C^{k-1}(\mathbb{R})$
- (B) $C^k(\mathbb{R})$
- (C) $C^{k+1}(\mathbb{R})$
- (D) $C^1(\mathbb{R})$
- (E) $C^\infty(\mathbb{R})$

Answer: Option (A)

Explanation: $C^k(\mathbb{R})$ is the space of functions that are at least k times continuously differentiable. Thus, the derivative of any function here is at least $(k - 1)$ times continuously differentiable. Thus, the image is in $C^{k-1}(\mathbb{R})$. The image is the whole of $C^{k-1}(\mathbb{R})$ because any function in $C^{k-1}(\mathbb{R})$ can be integrated to get a function in $C^k(\mathbb{R})$.

Performance review: 16 out of 25 got this. 9 chose (C).

Historical note (last time): 20 out of 26 got this. 2 each chose (C) and (D). 1 each chose (B) and (E).

- (10) What is the kernel of differentiation?

- (A) The vector space of all constant functions
- (B) The vector space of all linear functions (i.e., functions of the form $x \mapsto mx + c$ with $m, c \in \mathbb{R}$)
- (C) The vector space of all polynomial functions
- (D) $C^\infty(\mathbb{R})$
- (E) $C^1(\mathbb{R})$

Answer: Option (A)

Explanation: The derivative of a function on \mathbb{R} is zero if and only if the function is a constant function.

Performance review: 15 out of 25 got this. 7 chose (E), 2 chose (B), 1 chose (D).

Historical note (last time): 11 out of 26 got this. 8 chose (B), 4 chose (C), 2 chose (E), and 1 left the question blank.

- (11) Suppose k is a positive integer greater than 2. Consider the operation of “differentiating k times.” This is a linear transformation that can be defined as the k -fold composite of differentiation with

itself. Viewed most generally, this is a linear transformation from $C^k(\mathbb{R})$ to $C(\mathbb{R})$. What is the kernel of this linear transformation?

- (A) The set of all constant functions
- (B) The set of all polynomial functions of degree at most $k - 1$
- (C) The set of all polynomial functions of degree at most k
- (D) The set of all polynomial functions of degree at most $k + 1$
- (E) The set of all polynomial functions

Answer: Option (B)

Explanation: Each differentiation reduces the degree of the polynomial by 1, unless we are already at a constant, in which case we differentiate to 0. So, if the degree of the polynomial is at most $k - 1$, differentiating k times gives 0. Conversely, if differentiating k times gives zero, then repeated integration gives a generic polynomial of degree at most $k - 1$.

Performance review: 17 out of 25 got this. 4 chose (D), 3 chose (C), 1 chose (E).

Historical note (last time): 17 out of 26 got this. 3 each chose (A) and (D), 2 chose (C), and 1 left the question blank.

- (12) Suppose k is a positive integer greater than 2. Consider the set P_k of all polynomial functions of degree at most k . This set is a vector subspace of $C(\mathbb{R})$. Of the following subspaces of $C(\mathbb{R})$, which is the *smallest* subspace of which P_k is a subspace?

- (A) $C^1(\mathbb{R})$
- (B) $C^{k-1}(\mathbb{R})$
- (C) $C^k(\mathbb{R})$
- (D) $C^{k+1}(\mathbb{R})$
- (E) $C^\infty(\mathbb{R})$

Answer: Option (E)

Explanation: All polynomials are infinitely differentiable, so they are all in $C^\infty(\mathbb{R})$, the smallest of the spaces listed.

Performance review: 4 out of 25 got this. 11 chose (D), 7 chose (C), 2 chose (B), 1 left the question blank.

Historical note (last time): 7 out of 26 got this. 11 chose (D), 6 chose (C), 1 chose (B), and 1 left the question blank.

Two more definitions of use. A *linear functional* on a vector space V is a linear transformation from V to \mathbb{R} , where \mathbb{R} is viewed as a one-dimensional vector space over itself in the obvious way.

We define $C([0, 1])$ as the set of all continuous functions from $[0, 1]$ to \mathbb{R} with pointwise addition and scalar multiplication.

- (13) Which of the following is *not* a linear functional on $C([0, 1])$?

- (A) $f \mapsto f(0)$
- (B) $f \mapsto f(1)$
- (C) $f \mapsto \int_0^1 f(x) dx$
- (D) $f \mapsto \int_0^1 f(x^2) dx$
- (E) $f \mapsto \int_0^1 (f(x))^2 dx$

Answer: Option (E)

Explanation: This is easy to see from the description. The key point here is that if we square *after* evaluation, then that is not linear. Squaring prior to evaluation is fine. Explicitly, the point is that:

$$(f + g)(x^2) = f(x^2) + g(x^2)$$

But the following is *not* true generally:

$$((f + g)(x))^2 = (f(x))^2 + (g(x))^2$$

In fact, the left side simplifies to $(f(x))^2 + (g(x))^2 + 2f(x)g(x)$.

Performance review: 16 out of 25 got this. 6 chose (D), 1 each chose (A), (B), and (C).

Historical note (last time): 15 out of 26 got this. 8 chose (D), 1 each chose (B) and (C), and 1 left the question blank.

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE MONDAY NOVEMBER 11: USING LINEAR SYSTEMS FOR MEASUREMENT

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

27 people took this 4-question quiz. The score distribution was as follows:

- Score of 0: 2 people
- Score of 2: 9 people
- Score of 3: 9 people
- Score of 4: 7 people

The mean score was about 2.7.

The question-wise answers and performance review are as follows:

- (1) Option (E): 21 people
- (2) Option (D): 18 people
- (3) Option (B): 17 people
- (4) Option (D): 17 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS *ALL* QUESTIONS.

The purpose of this quiz is two-fold. First, many of the ideas you saw early on in the course (in the second and third week) may be on the verge of fading out. Drawing from the best research on *spaced repetition* (see for instance http://en.wikipedia.org/wiki/Spaced_repetition) it's high time we tried recalling some of that stuff. But with a twist, because we can now use some concepts from later topics we've seen to refine our past understanding.

The second purpose is to prepare you for what we hope to eventually get to: a deep and rich understanding of linear algebra as it's *used*: in linear regressions, computing correlations, and more fancy applications like factor analysis and principal component analysis. The third question, in particular, relates to the central idea behind linear regression (specifically, ordinary least squares regression). The questions also relate, albeit not very directly, to the broad ideas behind factor analysis and principal component analysis.

For the questions here, assume two dimensions of a person's general cognitive ability: verbal and mathematical. Denote by g_v the person's general verbal ability and by g_m the person's general mathematical ability.

Various ability tests can be devised that aim to test for the person's abilities. However, it is not possible to construct a test that *solely* measures g_v or *solely* measure g_m . Different tests measure g_v and g_m to different extents. For instance, an ordinary numerical computation test might measure mostly g_m . On the other hand, a test similar to the quizzes in this course might measure both g_v and g_m a fair amount, given how much you have to read to answer the quiz questions.

Of course, the score on a given test could depend on a lot of factors other than general abilities. Some of them could be systematic: a student with poor mathematical abilities in general may have "trained for the test." As an example, using a calculus test to test for general mathematical ability might mean that people who have happened to take calculus do a lot better than people who haven't, but have similar general mathematical ability. Some are more ephemeral, such as students guessing answers, mood fluctuations, and other context-specific factors that affect scores.

For simplicity, we will assume that there are no systemic factors other than general verbal and general mathematical ability that the test is measuring. For even greater simplicity, assume that the (expected) test score is linear in g_v and g_m with zero intercept, i.e., the score is of the form $w_v g_v + w_m g_m$ where w_v and w_m are real numbers that serve as *weights*. Note that this assumes that there is no *interaction* between the verbal and mathematical skills in determining the scores.

The assumption may or may not be realistic. For instance, a question (such as those on this quiz!) that requires a lot of reading *and* strong math skills would probably have an expected score formula that is *multiplicative* in g_v and g_m : having zero or near-zero verbal ability means you will be unable to do the question, even if your mathematical ability is awesome. Similarly, having zero or near-zero mathematical ability means you will be unable to do the question, even if your verbal ability is awesome. Multiplicatively separable functions are better suited to capture this sort of dependence. However, even if the test has questions of this sort, we can take logs on test scores and make them additively separable, so the additive model may still work well.

The assumption of *linearity* goes further, but this too might be realistic.

My goal is to use one or more tests in order to determine the true values of a student's g_v and g_m . Another formulation is that my goal is to determine the vector:

$$\vec{g} = \begin{bmatrix} g_v \\ g_m \end{bmatrix}$$

- (1) I administer two tests to a student. The student's score s_1 on the first test is $2g_v + 3g_m$ while the score s_2 on the second test is $3g_v + 5g_m$. How do I recover g_v and g_m from s_1 and s_2 ?
- (A) $g_v = 2s_1 + 3s_2$, $g_m = 3s_1 + 5s_2$
 (B) $g_v = 2s_1 - 3s_2$, $g_m = 3s_1 - 5s_2$
 (C) $g_v = 5s_1 + 3s_2$, $g_m = 3s_1 + 2s_2$
 (D) $g_v = 5s_1 - 3s_2$, $g_m = 3s_1 - 2s_2$
 (E) $g_v = 5s_1 - 3s_2$, $g_m = -3s_1 + 2s_2$

Answer: Option (E)

Explanation: We have the following:

$$\begin{bmatrix} s_1 \\ s_2 \end{bmatrix} = \begin{bmatrix} 2 & 3 \\ 3 & 5 \end{bmatrix} \begin{bmatrix} g_v \\ g_m \end{bmatrix}$$

We thus have:

$$\begin{bmatrix} g_v \\ g_m \end{bmatrix} = \begin{bmatrix} 2 & 3 \\ 3 & 5 \end{bmatrix}^{-1} \begin{bmatrix} s_1 \\ s_2 \end{bmatrix}$$

Recall that, for a general 2×2 matrix, the inverse is given by:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}^{-1} = \frac{1}{ad - bc} \begin{bmatrix} d & -b \\ -c & a \end{bmatrix}$$

In this case, the determinant $ad - bc$ equals $(2)(5) - (3)(3) = 1$, so that the inverse is:

$$\begin{bmatrix} 2 & 3 \\ 3 & 5 \end{bmatrix}^{-1} = \begin{bmatrix} 5 & -3 \\ -3 & 2 \end{bmatrix}$$

Plugging in, we get:

$$\begin{bmatrix} g_v \\ g_m \end{bmatrix} = \begin{bmatrix} 5 & -3 \\ -3 & 2 \end{bmatrix} \begin{bmatrix} s_1 \\ s_2 \end{bmatrix}$$

This simplifies to:

$$\begin{aligned} g_v &= 5s_1 - 3s_2 \\ g_m &= -3s_1 + 2s_2 \end{aligned}$$

Performance review: 21 out of 27 got this. 3 chose (A), 2 chose (D), 1 chose (C).

Historical note (last time): 23 out of 25 people got this. 1 each chose (B) and (C).

- (2) In order to combat the problem of uncertainty about my model, I decide to administer more than two tests. I administer a total of n tests. The score on the i^{th} test is $s_i = w_{i,v}g_v + w_{i,m}g_m$. The score vector \vec{s} has coordinates $s_i, 1 \leq i \leq n$.

If there is no measurement error and the student's actual score in each test equals the student's expected score, then we have a system of n simultaneous linear equations in 2 variables.

Let W be the matrix:

$$\begin{bmatrix} w_{1,v} & w_{1,m} \\ w_{2,v} & w_{2,m} \\ \cdot & \cdot \\ \cdot & \cdot \\ \cdot & \cdot \\ w_{n,v} & w_{n,m} \end{bmatrix}$$

Assume that all entries of W are positive, i.e., each test tests to a nonzero extent for both verbal and mathematical ability.

Our goal is to "solve for" \vec{g} the following vector equation:

$$W\vec{g} = \vec{s}$$

What is the necessary and sufficient condition on W so that the equation has at most one solution for \vec{g} for each \vec{s} ? If \vec{s} arises from an actual \vec{g} , i.e., it is a true score vector, then note that there will be a solution.

- (A) All the ratios $w_{i,v} : w_{i,m}$ are the same.
(B) All the ratios $w_{i,v} : w_{i,m}$ are different.
(C) At least two of the ratios $w_{i,v} : w_{i,m}$ are the same.
(D) At least two of the ratios $w_{i,v} : w_{i,m}$ are different.

Answer: Option (D)

Explanation: Recall that for a unique solution, the coefficient matrix needs to have full column rank, which in this case means rank 2. If all the ratios $w_{i,v} : w_{i,m}$ are the same, then the rank is 1. On the other hand, if two of the ratios differ, then those two rows alone give rank 2, and the remaining rows cannot affect the rank any more.

Performance review: 18 out of 27 got this. 6 chose (B), 2 chose (A), 1 chose (C).

Historical note (last time): 14 out of 25 people got this. 4 each chose (A) and (C). 3 chose (B).

- (3) Use notation as in the previous question. Suppose that there is some measurement error, so that instead of getting the true score vector \vec{s} , I have a somewhat distorted score vector \vec{t} . How do I go about recovering my "best guess" for \vec{s} from \vec{t} ?

- (A) Find the closest vector to \vec{t} in the kernel of the linear transformation corresponding to W .
(B) Find the closest vector to \vec{t} in the image of the linear transformation corresponding to W .
(C) Find the farthest vector from \vec{t} in the kernel of the linear transformation corresponding to W .
(D) Find the farthest vector from \vec{t} in the image of the linear transformation corresponding to W .

Answer: Option (B)

Explanation: Note that when talking of the kernel and image of W , we want to solve $W\vec{g} = \vec{s}$. Instead of \vec{s} , we have a distorted vector \vec{t} . Our best guess for \vec{s} is the vector closest to \vec{t} for which the equation can be solved, i.e., the vector closest to \vec{t} in the image of the linear transformation corresponding to W .

Performance review: 17 out of 27 got this. 5 chose (A), 4 chose (C), 1 chose (D).

Historical note (last time): 14 out of 25 people got this. 8 chose (C), 3 chose (A).

- (4) Suppose I want to introduce a *new* test that tests for both verbal and mathematical ability with expected score of the form $w_v g_v + w_m g_m$, but the values w_v and w_m are currently unknown. My strategy is as follows. I find two students. I administer a bunch of tests with *known* w_v and w_m values to those students. I use those tests to find the g_v and g_m values for both students. Then, I

administer the new test to both students and try to determine the values of w_v and w_m . Assume no measurement error.

Of course, I want the matrix W of the *known* tests to satisfy the condition of Question 2. What additional criteria would I wish of the two students I use for this in order to correctly determine g_v and g_m ? Note that it will not be possible to be sure of this in advance, but one can still pick student pairs who are more likely to satisfy the criterion and thus avoid waste of effort.

- (A) The students should have the same $g_v + g_m$ value.
- (B) The students should have different $g_v + g_m$ values.
- (C) The students should have the same $g_v : g_m$ ratio.
- (D) The students should have different $g_v : g_m$ ratios.

Answer: Option (D)

Explanation: The reasoning is similar to that used for Question 2.

Let $g_{v,1}$ and $g_{m,1}$ be the g_v and g_m values for the first student. Let $g_{v,2}$ and $g_{m,2}$ be the g_v and g_m values for the second student. We want to find the weights w_v and w_m for the test. Explicitly, this means solving the following equation for unknowns w_v and w_m :

$$\begin{bmatrix} g_{v,1} & g_{m,1} \\ g_{v,2} & g_{m,2} \end{bmatrix} \begin{bmatrix} w_v \\ w_m \end{bmatrix} = \begin{bmatrix} \text{score of first student} \\ \text{score of second student} \end{bmatrix}$$

In order to have a unique solution, we need the coefficient matrix to have full column rank 2, i.e., we need the students' $g_v : g_m$ ratios to differ.

Performance review: 17 out of 27 got this. 9 chose (C), 1 chose (B).

Historical note (last time): 14 out of 25 got this. 9 chose (C), 1 each chose (A) and (B).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY NOVEMBER 13:
MATRIX MULTIPLICATION: ROWS, COLUMNS, ORTHOGONALITY, AND OTHER
MISCELLANEA**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 13-question quiz. The score distribution was as follows:

- Score of 4: 2 people
- Score of 5: 1 person
- Score of 6: 1 person
- Score of 7: 3 people
- Score of 8: 3 people
- Score of 9: 5 people
- Score of 10: 4 people
- Score of 11: 4 people
- Score of 12: 2 people

The mean score was 8.68.

The question-wise answers and performance review are as follows:

- (1) Option (D): 24 people
- (2) Option (C): 22 people
- (3) Option (D): 22 people
- (4) Option (E): 19 people
- (5) Option (A): 17 people
- (6) Option (B): 23 people
- (7) Option (E): 14 people
- (8) Option (E): 20 people
- (9) Option (C): 7 people
- (10) Option (A): 12 people
- (11) Option (B): 7 people
- (12) Option (E): 14 people
- (13) Option (B): 16 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

The purpose of this quiz is two-fold. First, many of the ideas related to matrix multiplication are at the stage where a bit of review will help prevent their fading out. Drawing from the best research on *spaced repetition* (see for instance http://en.wikipedia.org/wiki/Spaced_repetition) we will try to recall some of the stuff. But with a twist, because we consider it from a somewhat different angle.

Second, the new angle will also turn out to be useful for later material.

For Questions 1-5: Given a n -dimensional vector $\langle a_1, a_2, \dots, a_n \rangle \in \mathbb{R}^n$, the vector can be interpreted as a $n \times 1$ matrix (a column vector). This is the default interpretation. But there are also two other interpretations: as a $1 \times n$ matrix (a row vector) and as a diagonal $n \times n$ matrix.

Also note that for Questions 1-5, all the three ways of representing vectors coincide with each other for $n = 1$, so the questions are uninteresting for $n = 1$ because all answer options are equivalent. You may

therefore assume that $n > 1$ for these questions, though obviously the correct answers are correct for $n = 1$ as well.

- (1) Suppose I want to add two vectors $\vec{a} = \langle a_1, a_2, \dots, a_n \rangle$ and $\vec{b} = \langle b_1, b_2, \dots, b_n \rangle$ to obtain the output vector $\langle a_1 + b_1, a_2 + b_2, \dots, a_n + b_n \rangle$ using matrix addition. What format (row vector, column vector, or diagonal matrix) should I use? Please see Option (D) before answering and select the option that best describes your view.
- (A) Represent both \vec{a} and \vec{b} as row vectors and interpret the sum as a row vector.
 (B) Represent both \vec{a} and \vec{b} as column vectors and interpret the sum as a column vector.
 (C) Represent both \vec{a} and \vec{b} as diagonal matrices and interpret the sum as a diagonal matrix.
 (D) We can use any of the above.

Answer: Option (D)

Explanation: Regardless of whether we represent the vectors as row vectors, column vectors, or diagonal matrices, matrix addition is executed coordinate-wise, as desired.

Performance review: 24 out of 25 people got this. 1 chose (B).

- (2) Suppose I want to perform coordinate-wise multiplication on two vectors. Explicitly, I have two vectors $\vec{a} = \langle a_1, a_2, \dots, a_n \rangle$ and $\vec{b} = \langle b_1, b_2, \dots, b_n \rangle$ and I want to obtain the output vector $\langle a_1 b_1, a_2 b_2, \dots, a_n b_n \rangle$ using matrix multiplication (with the matrix for \vec{a} written on the left and the matrix for \vec{b} written on the right). What format (row vector, column vector, or diagonal matrix) should I use? Please see Option (D) before answering and select the option that best describes your view.
- (A) Represent both \vec{a} and \vec{b} as row vectors and interpret the matrix product as a row vector.
 (B) Represent both \vec{a} and \vec{b} as column vectors and interpret the matrix product as a column vector.
 (C) Represent both \vec{a} and \vec{b} as diagonal matrices and interpret the matrix product as a diagonal matrix.
 (D) We can use any of the above.

Answer: Option (C)

Explanation: The multiplication of diagonal matrices is coordinate-wise along the diagonal. For instance, when $n = 2$, we get:

$$\begin{bmatrix} a_1 & 0 \\ 0 & a_2 \end{bmatrix} \begin{bmatrix} b_1 & 0 \\ 0 & b_2 \end{bmatrix} = \begin{bmatrix} a_1 b_1 & 0 \\ 0 & a_2 b_2 \end{bmatrix}$$

In general:

$$\begin{bmatrix} a_1 & 0 & \dots & 0 \\ 0 & a_2 & \dots & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & \dots & 0 & a_n \end{bmatrix} \begin{bmatrix} b_1 & 0 & \dots & 0 \\ 0 & b_2 & \dots & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & \dots & 0 & b_n \end{bmatrix} = \begin{bmatrix} a_1 b_1 & 0 & \dots & 0 \\ 0 & a_2 b_2 & \dots & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & \dots & 0 & a_n b_n \end{bmatrix}$$

Note that for Options (A) and (B), it does not even make sense to try computing the product.

Option (A): Here, \vec{a} and \vec{b} are both represented by n -dimensional row vectors, i.e., $1 \times n$ matrices, so we cannot multiply them because the number of columns of the first matrix does not equal the number of rows of the second matrix.

Option (B): Here, \vec{a} and \vec{b} are both represented by n -dimensional column vectors, i.e., $n \times 1$ matrices, so we cannot multiply them because the number of columns of the first matrix does not equal the number of rows of the second matrix.

Performance review: 22 out of 25 got this. 3 chose (D).

- (3) Suppose I am given two vectors $\vec{a} = \langle a_1, a_2, \dots, a_n \rangle$ and $\vec{b} = \langle b_1, b_2, \dots, b_n \rangle$ and I want to obtain a 1×1 matrix with entry $\sum_{i=1}^n a_i b_i$ using matrix multiplication (with the matrix for \vec{a} written on the left and the matrix for \vec{b} written on the right). What format (row vector, column vector, or diagonal matrix) should I use?
- (A) Represent both \vec{a} and \vec{b} as row vectors.
 (B) Represent both \vec{a} and \vec{b} as column vectors.
 (C) Represent both \vec{a} and \vec{b} as diagonal matrices.

(D) Represent \vec{a} as a row vector and \vec{b} as a column vector.

(E) Represent \vec{a} as a column vector and \vec{b} as a row vector.

Answer: Option (D)

Explanation: The dimensions match: \vec{a} is represented by a $1 \times n$ matrix and \vec{b} is represented by a $n \times 1$ matrix. The product definition also matches. Note that this particular form of product is the *dot product* of the vectors.

Note that Option (C) would give a diagonal matrix with the products $a_i b_i$ along the diagonal, but would not add them up. Options (A) and (B) do not make sense, for the same reason as discussed in the answer to Question 2. Option (E) would give a product that is a $n \times n$ matrix whose $(ij)^{th}$ entry is the product $a_i b_j$. The matrix described by Option (E) is termed the *Hadamard product*.

Performance review: 22 out of 25 got this. 2 chose (E), 1 left the question blank.

(4) Suppose I am given three vectors $\vec{a} = \langle a_1, a_2, \dots, a_n \rangle$, $\vec{b} = \langle b_1, b_2, \dots, b_n \rangle$, and $\vec{c} = \langle c_1, c_2, \dots, c_n \rangle$. I want to obtain a 1×1 matrix with entry $\sum_{i=1}^n (a_i b_i c_i)$ using matrix multiplication (with the matrix for \vec{a} written on the left, the matrix for \vec{b} written in the middle, and the matrix for \vec{c} written on the right). What format should I use?

(A) \vec{a} as a row vector, \vec{b} as a column vector, \vec{c} as a diagonal matrix.

(B) \vec{a} as a column vector, \vec{b} as a row vector, \vec{c} as a diagonal matrix.

(C) \vec{a} as a diagonal matrix, \vec{b} as a row vector, \vec{c} as a column vector.

(D) \vec{a} as a column vector, \vec{b} as a diagonal matrix, \vec{c} as a row vector.

(E) \vec{a} as a row vector, \vec{b} as a diagonal matrix, \vec{c} as a column vector.

Answer: Option (E)

Explanation: First, note that the dimensions match. \vec{a} is represented by a $1 \times n$ matrix, \vec{b} by a diagonal $n \times n$ matrix, and \vec{c} by a $n \times 1$ matrix. The multiplication makes sense, and the product is a 1×1 matrix.

Second, note that the matrix we get is the correct one. The product of the diagonal matrix for \vec{b} and the column vector for \vec{c} gives a column vector whose i^{th} coordinate is $b_i c_i$. The product of the row vector for \vec{a} with this column vector gives $\sum_{i=1}^n (a_i b_i c_i)$ as desired.

Here is how it looks in the case $n = 2$:

$$\begin{bmatrix} a_1 & a_2 \end{bmatrix} \begin{bmatrix} b_1 & 0 \\ 0 & b_2 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}$$

If we begin our simplification by multiplying the second and third matrix, we obtain:

$$\begin{bmatrix} a_1 & a_2 \end{bmatrix} \begin{bmatrix} b_1 c_1 \\ b_2 c_2 \end{bmatrix}$$

We now do the remaining multiplication and obtain the desired 1×1 matrix:

$$[a_1 b_1 c_1 + a_2 b_2 c_2]$$

Alternatively, we could simplify the original product by multiplying the first two matrices to begin with, and obtain:

$$\begin{bmatrix} a_1 b_1 & a_2 b_2 \end{bmatrix} \begin{bmatrix} c_1 \\ c_2 \end{bmatrix}$$

Multiplying them out gives the desired result. *Note:* It's not surprising that both ways of simplifying the product of three matrices give the same result. That follows from associativity of matrix multiplication. We did it both ways for a sanity check.

Performance review: 19 out of 25 got this. 2 each chose (A) and (C). 1 each chose (B) and (D).

The next few questions rely on the concept of orthogonality (*orthogonal* is a synonym for *perpendicular* or *at right angles*). We say that two vectors (of the same dimension) are orthogonal if their

dot product is zero. By this definition, the zero vector of a given dimension is orthogonal to every vector of that dimension. Note that it does not make sense to talk of orthogonality for vectors with different dimensions, i.e., with different numbers of coordinates.

- (5) Suppose A is a $n \times m$ matrix. We can think of solving the system $A\vec{x} = \vec{0}$ (where \vec{x} is a $m \times 1$ column vector of unknowns) as trying to find all the vectors orthogonal to all the vectors in a given set of vectors. What set of vectors is that?
- (A) The set of row vectors of A , i.e., the rows of A , viewed as m -dimensional vectors.
 (B) The set of column vectors of A , i.e., the columns of A , viewed as n -dimensional vectors.

Answer: Option (A)

Explanation: $A\vec{x}$ is a $n \times 1$ matrix (i.e., a column vector with n coordinates) and its i^{th} entry is the dot product of the i^{th} row of A and the vector \vec{x} , both of which are m -dimensional vectors. This is zero if and only if the i^{th} row of A and the vector \vec{x} are orthogonal to each other. In order to have $A\vec{x} = \vec{0}$, we need \vec{x} to be orthogonal to all the rows of A .

Performance review: 17 out of 25 got this. 8 chose (B).

- (6) Suppose A is a $p \times q$ matrix and B is a $q \times r$ matrix where p , q , and r are positive integers. The matrix product AB is a $p \times r$ matrix. What orthogonality condition corresponds to the condition that the matrix product AB is a zero matrix (i.e., all its entries are zero)?
- (A) Every row of A is orthogonal to every row of B .
 (B) Every row of A is orthogonal to every column of B .
 (C) Every column of A is orthogonal to every row of B .
 (D) Every column of A is orthogonal to every column of B .

Answer: Option (B)

Explanation: The ik^{th} entry of AB can be viewed as the dot product of the i^{th} row of A and the k^{th} column of B . In particular, this entry is zero if and only if the i^{th} row of A is orthogonal to the k^{th} column of B . In order for AB to be the zero matrix, we need this to hold for *every* row of A and *every* column of B , giving the answer option.

Performance review: 23 out of 25 got this. 2 chose (C).

- (7) Suppose A is an invertible $n \times n$ square matrix. Which of the following correctly characterizes the $n \times n$ matrix A^{-1} using orthogonality? Recall that AA^{-1} and $A^{-1}A$ are both equal to the $n \times n$ identity matrix.
- (A) For every i in $\{1, 2, \dots, n\}$, the i^{th} row of A is orthogonal to the i^{th} row of A^{-1} . The dot product of the i^{th} row of A and the j^{th} row of A^{-1} for distinct i, j in $\{1, 2, \dots, n\}$ equals 1.
 (B) For every i in $\{1, 2, \dots, n\}$, the i^{th} column of A is orthogonal to the i^{th} column of A^{-1} . The dot product of the i^{th} row of A and the j^{th} column of A^{-1} for distinct i, j in $\{1, 2, \dots, n\}$ equals 1.
 (C) For every distinct i, j in $\{1, 2, \dots, n\}$, the i^{th} row of A is orthogonal to the j^{th} row of A^{-1} . The dot product of the i^{th} row of A with the i^{th} row of A^{-1} equals 1.
 (D) For every i in $\{1, 2, \dots, n\}$, the i^{th} row of A is orthogonal to the i^{th} column of A^{-1} . The dot product of the i^{th} row of A and the j^{th} column of A^{-1} for distinct i, j in $\{1, 2, \dots, n\}$ equals 1.
 (E) For every distinct i, j in $\{1, 2, \dots, n\}$, the i^{th} row of A is orthogonal to the j^{th} column of A^{-1} . The dot product of the i^{th} row of A and the i^{th} column of A^{-1} equals 1.

Answer: Option (E)

Explanation: The product AA^{-1} is the identity matrix. The entries of this matrix can be described as follows:

- For distinct i, j in $\{1, 2, \dots, n\}$, the $(ij)^{\text{th}}$ entry of AA^{-1} is 0. This translates to saying that the i^{th} row of A is orthogonal to the j^{th} column of A^{-1} .
- For any i in $\{1, 2, \dots, n\}$, the $(ii)^{\text{th}}$ entry of AA^{-1} is 1. This translates to saying that the dot product of the i^{th} row of A and the i^{th} column of A^{-1} is equal to 1.

These correspond to Option (E).

Note that it is *also* true that $A^{-1}A$ is the identity matrix. We can use this to obtain an alternative characterization of A^{-1} . This condition will use the rows of A^{-1} and the columns of A . Explicitly:

- For distinct i, j in $\{1, 2, \dots, n\}$, the $(ij)^{\text{th}}$ entry of $A^{-1}A$ is 0. This translates to saying that the i^{th} row of A^{-1} is orthogonal to the j^{th} column of A .

- For any i in $\{1, 2, \dots, n\}$, the $(ii)^{th}$ entry of $A^{-1}A$ is 1. This translates to saying that the dot product of the i^{th} row of A^{-1} and the i^{th} column of A is equal to 1.

Performance review: 14 out of 25 got this. 6 chose (C), 3 chose (D), 1 each chose (A) and (B).

The remaining questions review your skills at abstract behavior prediction.

- (8) Suppose n is a positive integer greater than 1. Which of the following is always true for two invertible $n \times n$ matrices A and B ?
- (A) $A + B$ is invertible, and $(A + B)^{-1} = A^{-1} + B^{-1}$
 (B) $A + B$ is invertible, and $(A + B)^{-1} = B^{-1} + A^{-1}$
 (C) $A + B$ is invertible, though neither of the formulas of the preceding two options is correct
 (D) AB is invertible, and $(AB)^{-1} = A^{-1}B^{-1}$
 (E) AB is invertible, and $(AB)^{-1} = B^{-1}A^{-1}$

Answer: Option (E)

Explanation: A and B being invertible does not imply that $A + B$ is invertible. For instance, A may be the identity matrix and B may be its negative.

On the other hand, AB is invertible and $(AB)^{-1} = B^{-1}A^{-1}$. This is because when we invert (i.e., go backward) we must do so in reverse.

Performance review: 20 out of 25 got this. 3 chose (D), 1 each chose (A) and (B).

Historical note (last time, appeared in a midterm): 21 out of 30 people got this. 5 chose (C), 4 chose (D).

- (9) Suppose n is a positive integer greater than 1. For a nilpotent $n \times n$ matrix C , define the *nilpotency* of C as the smallest positive integer r such that $C^r = 0$. Note that the nilpotency is not defined for a non-nilpotent matrix. Given two $n \times n$ matrices A and B , what is the relation between the nilpotencies of AB and BA ?
- (A) AB is nilpotent if and only if BA is nilpotent, and if so, their nilpotencies must be equal.
 (B) AB is nilpotent if and only if BA is nilpotent, and if so, their nilpotencies must differ by 1.
 (C) AB is nilpotent if and only if BA is nilpotent, and if so, their nilpotencies must either be equal or differ by 1.
 (D) It is possible for AB to be nilpotent and BA to be non-nilpotent; however, if both are nilpotent, then their nilpotencies must be equal.
 (E) It is possible for AB to be nilpotent and BA to be non-nilpotent, however if both are nilpotent, then their nilpotencies must differ by 1.

Answer: Option (C)

Explanation: We have seen examples where the nilpotencies are not equal. For instance:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

$AB = 0$ but BA is not zero.

On the other hand, we also have examples where the nilpotencies are equal. For instance:

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

Note, however, that $(AB)^r = 0$ implies $(BA)^{r+1} = 0$ and $(BA)^s = 0$ implies $(AB)^{s+1} = 0$. Thus, the nilpotencies can differ by at most one, since each nilpotency is bounded by 1 more than the other.

Performance review: 7 out of 25 got this. 9 chose (D), 4 chose (E), 2 chose (A), 1 chose (B), and 2 left the question blank.

Historical note (last time, appeared in a midterm): 5 out of 30 got this. 12 chose (E), 11 chose (D), and 2 chose (B).

- (10) What is the smallest n for which there exist examples of invertible $n \times n$ matrices A and B such that $A \neq B$ but $A^2 = B^2$?
- (A) 1

- (B) 2
- (C) 3
- (D) 4
- (E) This is not possible for any n .

Answer: Option (A)

Explanation: We can take $A = [-1]$ and $B = [1]$.

Performance review: 12 out of 25 got this. 10 chose (B), 2 chose (E), 1 chose (C).

- (11) What is the smallest n for which there exist examples of invertible $n \times n$ matrices A and B such that $A \neq B$ but $A^3 = B^3$?

- (A) 1
- (B) 2
- (C) 3
- (D) 4
- (E) This is not possible for any n .

Answer: Option (B)

Explanation: Note first that $n = 1$ does not work, because if two real numbers have the same cube, they must be equal (this is because cubing is a one-one function).

However, $n = 2$ works. Explicitly, we can take A as a rotation matrix by $2\pi/3$ and B as the identity matrix. Both A^3 and B^3 equal the identity matrix. Explicitly:

$$A = \begin{bmatrix} -1/2 & -\sqrt{3}/2 \\ \sqrt{3}/2 & -1/2 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

If you prefer dealing only with matrices with integer entries, consider the following matrix. This is harder to think of, however:

$$A = \begin{bmatrix} -1 & -1 \\ 1 & 0 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Performance review: 7 out of 25 got this. 13 chose (C), 4 chose (E), 1 chose (D).

- (12) What is the smallest n for which there exist examples of invertible $n \times n$ matrices A and B such that $A \neq B$ but $A^2 = B^2$ and $A^3 = B^3$?

- (A) 1
- (B) 2
- (C) 3
- (D) 4
- (E) This is not possible for any n .

Answer: Option (E)

Explanation: Suppose $A^2 = B^2$ with A and B both invertible. Then, $A^{-2} = B^{-2}$ as well. We are also given $A^3 = B^3$. Multiplying the two equations, we get $A = B$. Thus, the specification required is not possible.

Performance review: 14 out of 25 got this. 7 chose (B), 3 chose (C), 1 chose (D).

- (13) What is the smallest n for which there exist examples of (not necessarily invertible) $n \times n$ matrices A and B such that $A \neq B$ but $A^2 = B^2$ and $A^3 = B^3$?

- (A) 1
- (B) 2
- (C) 3
- (D) 4
- (E) This is not possible for any n .

Answer: Option (B)

Explanation: Note that the condition $A^3 = B^3$ would imply $A = B$ if $n = 1$. Therefore, the smallest possible case is $n = 2$. We will furnish an example in this case. In our example, we take:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

Note how and why this example works: $A^2 = B^2 = 0$, so all higher powers of A and of B are equal to 0.

Performance review: 16 out of 25 got this. 5 chose (E), 3 chose (C), 1 chose (A).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY NOVEMBER 15: IMAGE AND KERNEL

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

27 people took this 18-question quiz. The score distribution was as follows:

- Score of 2: 3 people
- Score of 5: 3 people
- Score of 6: 1 person
- Score of 7: 4 people
- Score of 8: 2 people
- Score of 9: 2 people
- Score of 10: 3 people
- Score of 12: 2 people
- Score of 14: 3 people
- Score of 15: 2 people
- Score of 16: 1 person
- Score of 18: 1 person

The mean score was 9.22.

The question-wise answers and performance review are below:

- (1) Option (E): 18 people
- (2) Option (D): 22 people
- (3) Option (E): 18 people
- (4) Option (D): 9 people
- (5) Option (E): 16 people
- (6) Option (D): 23 people
- (7) Option (C): 9 people
- (8) Option (E): 13 people
- (9) Option (D): 12 people
- (10) Option (E): 18 people
- (11) Option (C): 12 people
- (12) Option (B): 8 people
- (13) Option (E): 7 people
- (14) Option (C): 11 people
- (15) Option (D): 12 people
- (16) Option (B): 16 people
- (17) Option (D): 15 people
- (18) Option (C): 10 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS *ALL* QUESTIONS.

The purpose of this quiz is to review in greater depth the ideas behind image and kernel. The goal of the first seven questions is to review the ideas of injectivity, surjectivity, and bijectivity in the context of arbitrary functions between sets. The purpose is two-fold: (i) to give a functions-based approach to justifying, intuitively and formally, facts about the effect of matrix multiplication on rank, and (ii) to hint at ways in which linear transformations behave better than other types of functions.

The corresponding lecture notes are titled **Image and kernel of a linear transformation** and the corresponding section of the text is Section 3.1.

Just as a reminder, a function $f : A \rightarrow B$ between sets A and B is said to be:

- *injective* if for every $b \in B$, there is *at most* one value of a such that $f(a) = b$. In other words, if we denote by $f^{-1}(b)$ the set $\{a \in A \mid f(a) = b\}$, then $|f^{-1}(b)| \leq 1$ for all $b \in B$ (here $|f^{-1}(b)|$ denotes the size of the set $f^{-1}(b)$).
- *surjective* if for every $b \in B$, there is *at least* one value of a such that $f(a) = b$. In other words, if we denote by $f^{-1}(b)$ the set $\{a \in A \mid f(a) = b\}$, then $|f^{-1}(b)| \geq 1$ for all $b \in B$.
- *bijective* if for every $b \in B$, there is *exactly* one value of a such that $f(a) = b$. In other words, if we denote by $f^{-1}(b)$ the set $\{a \in A \mid f(a) = b\}$, then $|f^{-1}(b)| = 1$ for all $b \in B$.

(1) Suppose $g : A \rightarrow B$ and $f : B \rightarrow C$ are functions. The composite $f \circ g$ is a function from A to C . What can we say the relationship between the injectivity of $f \circ g$, the injectivity of f , and the injectivity of g ?

- (A) $f \circ g$ is injective if and only if f and g are both injective.
- (B) If f and g are both injective, then $f \circ g$ is injective. However, $f \circ g$ being injective does not imply anything about the injectivity of either f or g .
- (C) If f and g are both injective, then $f \circ g$ is injective. If $f \circ g$ is injective, then at least one of f and g is injective, but we cannot conclusively say for any specific one of the two that it must be injective.
- (D) If f and g are both injective, then $f \circ g$ is injective. If $f \circ g$ is injective, then f is injective, but we do not have enough information to deduce whether g is injective.
- (E) If f and g are both injective, then $f \circ g$ is injective. If $f \circ g$ is injective, then g is injective, but we do not have enough information to deduce whether f is injective.

Answer: Option (E)

Explanation: See the lecture notes for more details (note that the roles of f and g are reversed in the lecture notes). The hard part is the direction from $f \circ g$ to g . To see this, note that if g has a collision $g(a_1) = g(a_2)$ (i.e., is non-injective) then that collision continues for $f \circ g$, i.e, we still have $f(g(a_1)) = f(g(a_2))$.

Performance review: 18 out of 27 got this. 6 chose (D), 1 each chose (A), (B), and (C).

Historical note (last time): 24 out of 26 got this. 1 each chose (A) and (D).

(2) Suppose $g : A \rightarrow B$ and $f : B \rightarrow C$ are functions. The composite $f \circ g$ is a function from A to C . What can we say the relationship between the surjectivity of $f \circ g$, the surjectivity of f , and the surjectivity of g ?

- (A) $f \circ g$ is surjective if and only if f and g are both surjective.
- (B) If f and g are both surjective, then $f \circ g$ is surjective. However, $f \circ g$ being surjective does not imply anything about the surjectivity of either f or g .
- (C) If f and g are both surjective, then $f \circ g$ is surjective. If $f \circ g$ is surjective, then at least one of f and g is surjective, but we cannot conclusively say for any specific one of the two that it must be surjective.
- (D) If f and g are both surjective, then $f \circ g$ is surjective. If $f \circ g$ is surjective, then f is surjective, but we do not have enough information to deduce whether g is surjective.
- (E) If f and g are both surjective, then $f \circ g$ is surjective. If $f \circ g$ is surjective, then g is surjective, but we do not have enough information to deduce whether f is surjective.

Answer: Option (D)

Explanation: See the lecture notes for more details (note that the roles of f and g are reversed in the lecture notes).

Performance review: 22 out of 27 got this. 3 chose (A), 1 each chose (B) and (E).

Historical note (last time): 23 out of 26 got this. 2 chose (E), 1 chose (A).

(3) Suppose $g : A \rightarrow B$ and $f : B \rightarrow C$ are functions. The composite $f \circ g$ is a function from A to C . Suppose $f \circ g$ is bijective. What can we say about f and g individually?

- (A) Both f and g must be bijective.
- (B) Both f and g must be injective, but neither of them need be surjective.
- (C) Both f and g must be surjective, but neither of them need be injective.

- (D) f must be injective but need not be surjective. g must be surjective but need not be injective.
 (E) f must be surjective but need not be injective. g must be injective but need not be surjective.

Answer: Option (E)

Explanation: The composite $f \circ g$ is bijective, so it is both injective and surjective. The results of the previous two questions now take care of things.

Performance review: 18 out of 27 got this. 4 chose (A), 2 each chose (C) and (D), 1 left the question blank.

Historical note (last time): 23 out of 26 got this. 2 chose (D), 1 chose (A).

- (4) $g : A \rightarrow B$ and $f : B \rightarrow C$ are functions. The composite $f \circ g$ is a function from A to C . Suppose both f and g are surjective. Further, suppose that for every $b \in B$, $g^{-1}(b)$ has size m (for a fixed positive integer m) and for every $c \in C$, $f^{-1}(c)$ has size n (for a fixed positive integer n). Then, what can we say about the sizes of the fibers (i.e., the inverse images of points in C) under the composite $f \circ g$?
- (A) The size is $\min\{m, n\}$
 (B) The size is $\max\{m, n\}$
 (C) The size is $m + n$
 (D) The size is mn
 (E) The size is m^n

Answer: Option (D)

Explanation: For any $c \in C$, $(f \circ g)^{-1}(c) = g^{-1}(f^{-1}(c))$ is the union, for all $b \in f^{-1}(c)$, of the sets $g^{-1}(b)$. Each of these sets has size m , and there is a total of n such sets, so we get a total of mn elements. *Remember, as you learned in kindergarten, that multiplication is repeated addition.*

Performance review: 9 out of 27 got this. 7 chose (A), 6 chose (B), 3 chose (E), 1 chose (C), 1 left the question blank.

Historical note (last time): 21 out of 26 got this. 4 chose (A), 1 chose (B).

- (5) **PLEASE READ THIS VERY CAREFULLY AND CONSIDER A WIDE VARIETY OF POLYNOMIAL EXAMPLES:** Suppose f is a polynomial function of degree $n > 2$ from \mathbb{R} to \mathbb{R} . What can we say about the fibers of f , i.e., the sets of the form $f^{-1}(x)$, $x \in \mathbb{R}$?

Hint: At the one extreme, consider a polynomial of the form x^n . Consider the sizes of the fibers $f^{-1}(0)$ and $f^{-1}(x)$ for a positive value of x (the fiber size for the latter will depend on whether n is even or odd). Alternatively, consider a polynomial of the form $(x - 1)(x - 2) \dots (x - n)$. Consider the size of the fiber $f^{-1}(0)$.

- (A) Every fiber has size n .
 (B) The minimum of the sizes of fibers is exactly n , but every fiber need not have size n .
 (C) The maximum of the sizes of fibers is exactly n , but every fiber need not have size n .
 (D) The minimum of the sizes of fibers is at least n , but need not be exactly n .
 (E) The maximum of the sizes of fibers is at most n , but need not be exactly n .

Answer: Option (E)

Explanation: Suppose we are trying to calculate the size of the fiber $f^{-1}(x)$ for a particular value of x . This is equivalent to solving the equation $f(t) = x$ in the variable t . This is a polynomial equation of degree n , so it has at most n roots. Thus, the size of each fiber is at most n . Thus, the maximum of the sizes of the fibers is at most n .

For $n = 2$, the maximum of the fiber sizes is always 2. However, for $n \geq 3$, there are examples where the maximum of the sizes of the fibers is less than n . Specifically, consider the example of x^n . The maximum of the fiber sizes here is 1 if n is odd and 2 if n is even. In both cases, it is less than n for $n \geq 3$.

Performance review: 16 out of 27 got this. 6 chose (C), 3 chose (B), 1 chose (A), 1 left the question blank.

Historical note (last time): 1 out of 26 got this (!!!). 16 chose (C), 5 chose (B), 3 chose (D), 1 chose (A).

- (6) Suppose f is a continuous injective function from \mathbb{R} to \mathbb{R} . What can we say about the nature of f ?
- (A) f must be an increasing function on all of \mathbb{R} .
 (B) f must be a decreasing function on all of \mathbb{R} .

- (C) f must be a constant function on all of \mathbb{R} .
- (D) f must be either an increasing function on all of \mathbb{R} or a decreasing function on all of \mathbb{R} , but the information presented is insufficient to decide which case occurs.
- (E) f must be either an increasing function or a decreasing function or a constant function on all of \mathbb{R} , but the information presented is insufficient for deciding anything stronger.

Answer: Option (D)

Explanation: This is easy to see pictorially, though a rigorous proof would invoke the intermediate value theorem and the extreme value theorem.

Performance review: 23 out of 27 got this. 4 chose (E).

Historical note (last time): 21 out of 26 got this. 5 chose (E).

- (7) **PLEASE READ THIS CAREFULLY, MAKE CASES, AND CHECK YOUR REASONING:** Suppose f , g , and h are continuous bijective functions from \mathbb{R} to \mathbb{R} . What can we say about the functions $f + g$, $f + h$, and $g + h$?

Hint: Based on the preceding question, you know something about the nature of f , g , and h individually as functions, but there is some degree of ambiguity in your knowledge. Make cases based on the possibilities and see what you can deduce in the best and worst case.

- (A) They are all continuous bijective functions from \mathbb{R} to \mathbb{R} .
- (B) At least two of them are continuous bijective functions from \mathbb{R} to \mathbb{R} . However, we cannot say more.
- (C) At least one of them is a continuous bijective function from \mathbb{R} to \mathbb{R} . However, we cannot say more.
- (D) Either all three sums are continuous bijective functions from \mathbb{R} to \mathbb{R} , or none is.
- (E) It is possible that none of the sums is a continuous bijective functions from \mathbb{R} to \mathbb{R} ; it is also possible that one, two, or all the sums are continuous bijective functions from \mathbb{R} to \mathbb{R} .

Answer: Option (C)

Explanation: Since f , g , and h are all continuous bijective functions $\mathbb{R} \rightarrow \mathbb{R}$, each one of them is either increasing or decreasing. Further, the functions that are increasing must have a limit of $-\infty$ at $-\infty$ and a limit of ∞ at ∞ , whereas the functions that are decreasing must have a limit of ∞ at $-\infty$ and a limit of $-\infty$ at ∞ . Thus:

- A sum of two continuous increasing surjective functions is also a continuous increasing surjective function, and hence is bijective: To see this, use the fact that the limit of the sum is the sum of the limits to deduce that for the sum, the limit at $-\infty$ is $-\infty$ and the limit at ∞ is ∞ , so that the function must be surjective.
- A sum of two continuous decreasing surjective functions is also a continuous decreasing surjective function, hence is bijective. To see this, use the fact that the limit of the sum is the sum of the limits to deduce that for the sum, the limit at $-\infty$ is ∞ and the limit at ∞ is $-\infty$, so that the function must be surjective.

We consider various cases:

- If all three functions are increasing, so are all the pairwise sums, and hence, all the sums $f + g$, $f + h$, and $g + h$ are bijective.
- If all three functions are decreasing, so are all the pairwise sums, and hence, all the sums $f + g$, $f + h$, and $g + h$ are decreasing.
- If two of the functions are increasing and the third function is decreasing, then we know for certain that the sum of the two increasing functions is bijective. But the sum of either of the increasing functions with the decreasing function may be increasing, decreasing, or neither. For instance, if $f(x) = g(x) = x$ and $h(x) = -x$, then $f + h$ and $g + h$ are both the zero function, which is neither increasing nor decreasing, and hence not one-to-one.
- If two of the functions are decreasing and the third function is increasing, then we know for certain that the sum of the two decreasing functions is bijective. We cannot say anything for sure about the other two sums, for the same reasons as in the previous case (specifically, we can just use the negatives of the functions in the preceding example).

It's clear from all these that (C) is the right option.

Performance review: 9 out of 27 got this. 10 chose (E), 4 chose (B), 3 chose (A), 1 chose (D).

Historical note (last time): 1 out of 26 got this. 15 chose (E), 8 chose (A), 2 chose (D).

The questions that follow tripped up students quite a bit last time, so I urge you to proceed with caution. You can do each of these questions in either of two ways:

- Using abstract, general reasoning.
- Constructing concrete examples.

While the former approach is one you should eventually be able to embrace without trepidation, feel free to rely on the latter approach for now. For this, consider matrices describing the linear transformations and use matrix multiplication to compute the composite where needed. Compute the kernel, image, and rank using the methods known to you. Take matrices such as those arising from finite state automata (as described in the “linear transformations and finite state automata” quiz) or their generalizations to rectangular matrices.

For instance, you might try taking a matrix such as $\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}$. This describes a linear

transformation $\mathbb{R}^5 \rightarrow \mathbb{R}^4$ and has rank three. The dimension of the kernel (inside \mathbb{R}^5) is 2 (explicitly, the kernel is precisely the set of vectors in \mathbb{R}^5 whose first three coordinates are zero) and the dimension of the image (inside \mathbb{R}^4) is 3 (explicitly, the image is precisely the set of vectors in \mathbb{R}^4 whose fourth coordinate is 0).

- (8) *This is the analogue for linear transformations of Question 1:* Suppose m, n, p are positive integers. Suppose A is a $m \times n$ matrix and B is a $n \times p$ matrix. The product AB is a $m \times p$ matrix. Denote by T_A , T_B , and T_{AB} respectively the linear transformations corresponding to A , B , and AB . We have $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m$, $T_B : \mathbb{R}^p \rightarrow \mathbb{R}^n$, and $T_{AB} : \mathbb{R}^p \rightarrow \mathbb{R}^m$. Note that $T_{AB} = T_A \circ T_B$.

Recall that a matrix has full column rank if and only if the corresponding linear transformation is injective.

Which of the following describes correctly the relationship between A having full column rank (i.e., rank n), B having full column rank (i.e., rank p), and AB having full column rank (i.e., rank p)?

- (A) AB has full column rank (i.e., rank p) if and only if A and B both have full column rank (ranks n and p respectively).
- (B) If A and B both have full column rank, then AB has full column rank. However, AB having full column rank does not imply anything (separately or jointly) regarding whether A or B has full column rank.
- (C) If A and B both have full column rank, then AB has full column rank. If AB has full column rank, then at least one of A and B has full column rank, but we cannot definitively say for any particular one of A and B that it must have full column rank.
- (D) If A and B both have full column rank, then AB has full column rank. AB having full column rank implies that A has full column rank, but it does not tell us for sure that B has full column rank.
- (E) If A and B both have full column rank, then AB has full column rank. AB having full column rank implies that B has full column rank, but it does not tell us for sure that A has full column rank.

Answer: Option (E)

Explanation: This is a special case of Question 1 (note that the letters do not match). Essentially, T_A plays the role of f and T_B plays the role of g .

An alternative way of thinking of this is that the rank of a product is less than or equal to the rank of each of the matrices being multiplied. If the rank of AB is p (full column rank), then that means that the rank of B is at least p . Since the number of columns of B equals p , this forces B to have full column rank p .

Note that A need not have full column rank. For instance, consider:

$$A = \begin{bmatrix} 1 & 0 \end{bmatrix}, B = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

Performance review: 13 out of 27 got this. 6 chose (D), 3 each chose (B) and (C), 2 chose (A).

- (9) *This is the analogue for linear transformations of Question 2:* Suppose m, n, p are positive integers. Suppose A is a $m \times n$ matrix and B is a $n \times p$ matrix. The product AB is a $m \times p$ matrix. Denote by T_A, T_B , and T_{AB} respectively the linear transformations corresponding to A, B , and AB . We have $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m, T_B : \mathbb{R}^p \rightarrow \mathbb{R}^n$, and $T_{AB} : \mathbb{R}^p \rightarrow \mathbb{R}^m$. Note that $T_{AB} = T_A \circ T_B$.

Recall that a matrix has full row rank if and only if the corresponding linear transformation is surjective.

Which of the following describes correctly the relationship between A having full row rank (i.e., rank m), B having full row rank (i.e., rank n), and AB having full row rank (i.e., rank m)?

- (A) AB has full row rank if and only if A and B both have full row rank.
 (B) If A and B both have full row rank, then AB has full row rank. However, AB having full row rank does not imply anything (separately or jointly) regarding whether A or B has full row rank.
 (C) If A and B both have full row rank, then AB has full row rank. If AB has full row rank, then at least one of A and B has full row rank, but we cannot definitively say for any particular one of A and B that it must have full row rank.
 (D) If A and B both have full row rank, then AB has full row rank. AB having full row rank implies that A has full row rank, but it does not tell us for sure that B has full row rank.
 (E) If A and B both have full row rank, then AB has full row rank. AB having full row rank implies that B has full row rank, but it does not tell us for sure that A has full row rank.

Answer: Option (D)

Explanation: This follows from Question 2. We can also think of it in terms of the rank of a product being less than or equal to the ranks of the individual matrices. This forces the rank of A to be at least m , and therefore exactly m .

Also, for an example of a situation where B does not have full row rank, we can use the same example as in the preceding question.

Performance review: 12 out of 27 got this. 6 chose (E), 4 chose (A), 3 chose (B), 2 chose (C).

- (10) *This is the analogue for linear transformations of Question 3:* Suppose m and n are positive integers. Suppose A is a $m \times n$ matrix and B is a $n \times m$ matrix. The product AB is a $m \times m$ matrix. The corresponding linear transformations are $T_A : \mathbb{R}^n \rightarrow \mathbb{R}^m, T_B : \mathbb{R}^m \rightarrow \mathbb{R}^n$, and $T_{AB} : \mathbb{R}^m \rightarrow \mathbb{R}^m$.

Suppose the square matrix AB has full rank m . What can we deduce about the ranks of A and B ?

- (A) Both A and B have full row rank, *and* both A and B have full column rank.
 (B) Both A and B have full column rank, but neither of them need have full row rank.
 (C) Both A and B have full row rank, but neither of them need have full column rank.
 (D) A must have full column rank but need not have full row rank. B must have full row rank but need not have full column rank.
 (E) A must have full row rank but need not have full column rank. B must have full column rank but need not have full row rank.

Answer: Option (E)

Explanation: Follows from Question 3. We can use the same example as for the preceding two questions.

Also note that in this case, we must have $m \leq n$, and therefore, the rank of both A and B , since it's $\leq \min\{m, n\}$ but also $\geq m$, must equal m . This means full row rank in the case of A , and full column rank in the case of B .

Performance review: 18 out of 27 got this. 5 chose (D), 2 chose (A), 1 each chose (B) and (C).

For the coming questions, we will denote vector spaces by letters such as U, V , and W . You can, however, consider them to be finite-dimensional vector spaces of the form \mathbb{R}^n . However, you

should take care not to use a letter for the dimension of a vector space if the letter is already in use elsewhere in the question. Also, you should take care to use different letters for the dimensions of different vector spaces, unless it is given to you that the vector spaces have the same dimension. The results also hold for infinite-dimensional vector spaces, but you can work on all the problems assuming you are working in the finite-dimensional setting.

- (11) *This is an analogue for linear transformations of Question 4:* Suppose $T_1 : U \rightarrow V$ and $T_2 : V \rightarrow W$ are linear transformations. The composite $T_2 \circ T_1$ is also a linear transformation, this time from U to W . Suppose the kernel of T_1 has dimension m and the kernel of T_2 has dimension n . Suppose both T_1 and T_2 are surjective. What can you say about the dimension of the kernel of $T_2 \circ T_1$?

Please note this carefully: Although this question is analogous to Question 4, the correct answer options differ for the two questions. Here is an intuitive explanation for the relationship between the questions. Question 4 asked about the *sizes* of the fibers. This question asks about the dimensions of the kernels. The fibers do correspond to the kernels. But the relationship between dimension and size is of a *logarithmic nature*. What we mean is that the dimension can be thought of as the logarithm of the size. This isn't literally true, because the size is infinite. But metaphorically, it makes sense, because, for instance, the dimension of \mathbb{R}^p is the exponent p , and that comports with the laws of logarithms (similar to how the $\log_2(2^p) = p$).

- (A) The dimension is $\min\{m, n\}$.
- (B) The dimension is $\max\{m, n\}$.
- (C) The dimension is $m + n$.
- (D) The dimension is mn .
- (E) The dimension is m^n .

Answer: Option (C)

Explanation: See the lecture notes for more.

Performance review: 12 out of 27 got this. 5 chose (D), 4 chose (A), 3 chose (B), 2 chose (E), 1 left the question blank.

Historical note (last time): 3 out of 26 got this. 13 chose (B), 7 chose (A), 3 chose (D).

- (12) Suppose $T_1 : U \rightarrow V$ and $T_2 : V \rightarrow W$ are linear transformations. The composite $T_2 \circ T_1$ is also a linear transformation, this time from U to W . Suppose the kernel of T_1 has dimension m and the kernel of T_2 has dimension n . However, unlike the preceding question, we are not given any information about the surjectivity of either T_1 or T_2 . The answer to the preceding question gives an (inclusive) *upper* bound on the dimension of the kernel of $T_2 \circ T_1$. Which of the following is the best *lower* bound we can manage in general?

- (A) $|m - n|$
- (B) m
- (C) n
- (D) $m + n$

Answer: Option (B)

Explanation: The kernel of $T_2 \circ T_1$ contains the kernel of T_1 , so m is a lower bound on the dimension.

Performance review: 8 out of 27 got this. 11 chose (A), 4 chose (C), 3 chose (D), 1 left the question blank.

Historical note (last time): 11 out of 26 got this. 10 chose (A), 4 chose (C), 1 chose (D).

- (13) Suppose $T_1, T_2 : U \rightarrow V$ are linear transformations. Which of the following is true? Please see Options (D) and (E) before answering and select the single option that best reflects your view.

- (A) If both T_1 and T_2 are injective, then $T_1 + T_2$ is injective.
- (B) If both T_1 and T_2 are surjective, then $T_1 + T_2$ is surjective.
- (C) If both T_1 and T_2 are bijective, then $T_1 + T_2$ is bijective.
- (D) All of the above
- (E) None of the above

Answer: Option (E)

Explanation: We can get counterexamples in one dimension: consider the situation where T_1 has matrix $[1]$ and T_2 has matrix $[-1]$. Then, both T_1 and T_2 are bijective (hence also injective and surjective) but the sum $T_1 + T_2$, which is the zero map, is neither injective nor surjective.

Performance review: 7 out of 27 got this. 13 chose (D), 4 chose (B), 2 chose (C), 1 left the question blank.

Historical note (last time): 6 out of 26 got this. 19 chose (D), 1 chose (C).

- (14) Suppose $T_1, T_2 : U \rightarrow V$ are linear transformations. Which of the following best describes the relation between the kernels of T_1 , T_2 , and $T_1 + T_2$?
- (A) The kernel of $T_1 + T_2$ equals the intersection of the kernel of T_1 and the kernel of T_2 .
- (B) The kernel of $T_1 + T_2$ is contained inside the intersection of the kernel of T_1 and the kernel of T_2 , but need not be equal to the intersection.
- (C) The kernel of $T_1 + T_2$ contains the intersection of the kernel of T_1 and the kernel of T_2 , but need not be equal to the intersection.
- (D) The kernel of $T_1 + T_2$ is contained inside the sum of the kernel of T_1 and the kernel of T_2 , but need not be equal to the sum.
- (E) The kernel of $T_1 + T_2$ contains the sum of the kernel of T_1 and the kernel of T_2 , but need not be equal to the sum.

Answer: Option (C)

Explanation: We will show this in steps:

- Suppose \vec{u} is in the intersection of the kernel of T_1 and the kernel of T_2 . Then, \vec{u} is in the kernel of $T_1 + T_2$: This is easy to see: $(T_1 + T_2)(\vec{u}) = T_1(\vec{u}) + T_2(\vec{u}) = 0 + 0 = 0$.
- It is possible to have a situation where the kernel of $T_1 + T_2$ does not contain the sum of the kernels of T_1 and T_2 . For instance, consider the case that T_1 has matrix $[0]$ and T_2 has matrix $[1]$.
- It is possible to have a situation where the kernel of $T_1 + T_2$ is not even contained inside the sum of the kernels of T_1 and T_2 , let alone the intersection: Consider T_1 to be the linear transformation with matrix $[1]$ and T_2 to be the linear transformation with matrix $[-1]$. Both have zero kernels, so the sum of the kernels is also zero. But the sum $T_1 + T_2$ is a linear transformation with matrix $[0]$, so its kernel is all of \mathbb{R} .

Performance review: 11 out of 27 got this. 6 chose (A), 4 each chose (B) and (D), 1 chose (E), 1 left the question blank.

Historical note (last time): 3 out of 26 got this. 10 chose (B), 6 each chose (A) and (D), 1 chose (E).

- (15) Suppose $T_1, T_2 : U \rightarrow V$ are linear transformations. Which of the following best describes the relation between the images of T_1 , T_2 , and $T_1 + T_2$?
- (A) The image of $T_1 + T_2$ equals the intersection of the image of T_1 and the image of T_2 .
- (B) The image of $T_1 + T_2$ is contained inside the intersection of the image of T_1 and the image of T_2 , but need not be equal to the intersection.
- (C) The image of $T_1 + T_2$ contains the intersection of the image of T_1 and the image of T_2 , but need not be equal to the intersection.
- (D) The image of $T_1 + T_2$ is contained inside the sum of the image of T_1 and the image of T_2 , but need not be equal to the sum.
- (E) The image of $T_1 + T_2$ contains the sum of the image of T_1 and the image of T_2 , but need not be equal to the sum.

Answer: Option (D)

Explanation: We show the following:

- Any vector in the image of $T_1 + T_2$ can be expressed as the sum of a vector in the image of T_1 and a vector in the image of T_2 : Suppose \vec{v} is in the image of $T_1 + T_2$. Thus, $\vec{v} = (T_1 + T_2)(\vec{u})$ which simplifies to $T_1(\vec{u}) + T_2(\vec{u})$, thus it is in the sum of the images.
- The image of $T_1 + T_2$ need not contain the intersection of the images of T_1 and T_2 : We can again use the example of linear transformations with matrices $[1]$ and $[-1]$ respectively.
- The image of $T_1 + T_2$ need not be contained in the intersection of the images of T_1 and T_2 : We can again use the example of linear transformations with matrices $[1]$ and $[0]$.

Performance review: 12 out of 27 got this. 8 chose (B), 4 chose (C), 2 chose (E), 1 left the question blank.

Historical note (last time): 11 out of 26 got this. 11 chose (B), 2 chose (C), 1 each chose (A) and (E).

- (16) Suppose T is a linear transformation from a vector space V to itself. Note that V may be an infinite-dimensional space, such as $C^\infty(\mathbb{R})$ (with T being differentiation), but for convenience, you can imagine V to be finite-dimensional (we will not reference the dimension of V in this question, however). Suppose the kernel of T has dimension n . What can you say from this information about the dimension of the kernel of T^r for a positive integer r ?
- (A) It is at least n and at most $n + r$.
(B) It is at least n and at most nr .
(C) It is at least $n + r$ and at most nr .
(D) It is at least $n + r$ and at most n^r .

Answer: Option (B)

Explanation: We know the kernel of a composite contains the kernel of the very first operation, so the dimension is at least n . But it could be bigger. Recall that the dimensions of the kernels could at worst add up, so the worst case scenario (which occurs if each time the kernel is also contained in the image) is that the total dimension is nr .

Performance review: 16 out of 27 got this. 9 chose (A), 1 chose (C), 1 left the question blank.

Historical note (last time): 13 out of 26 got this. 7 chose (A), 5 chose (C), 1 chose (D).

The next few questions deal with the relationship between the rows and columns of the matrix on the one hand, and the image and kernel of the linear transformation on the other hand.

- (17) Suppose A is a $n \times m$ matrix and $T_A : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the corresponding linear transformation. Which of the following correctly describes the relationship between the rows and columns of A and the image and kernel of T_A ?
- (A) The kernel of T_A is precisely the subspace of \mathbb{R}^m spanned by the rows of A . The image of T_A is precisely the subspace of \mathbb{R}^n spanned by the columns of A .
(B) The kernel of T_A is precisely the subspace of \mathbb{R}^m spanned by the columns of A . The image of T_A is precisely the subspace of \mathbb{R}^n spanned by the rows of A .
(C) The kernel of T_A is precisely the subspace of \mathbb{R}^m comprising the vectors that are *orthogonal* to the rows of A . The image of T_A is precisely the subspace of \mathbb{R}^n comprising the vectors that are *orthogonal* to the columns of A .
(D) The kernel of T_A is precisely the subspace of \mathbb{R}^m comprising the vectors that are *orthogonal* to the rows of A . The image of T_A is the subspace of \mathbb{R}^n spanned by the columns of A .
(E) The kernel of T_A is precisely the subspace of \mathbb{R}^m spanned by the rows of A . The image of T_A is precisely the subspace of \mathbb{R}^n comprising the vectors that are *orthogonal* to the columns of A .

Answer: Option (D)

Explanation: The kernel is the set of vectors $\vec{x} \in \mathbb{R}^m$ such that $A\vec{x} = \vec{0}$. The entries of $A\vec{x}$ are the dot products of the rows of A with \vec{x} . Therefore, a particular entry is zero if and only if the corresponding dot product is zero, i.e., \vec{x} is orthogonal to the corresponding row of A . Thus, all the entries of the output vector are zero iff \vec{x} is orthogonal to all rows of A (there were also similar questions in the preceding quiz).

The part of the statement about the image is a standard fact about the image that you have seen in lecture.

Performance review: 15 out of 27 got this. 5 chose (B), 4 chose (C), 2 chose (A), 1 left the question blank.

- (18) Suppose A and B are $n \times m$ matrices, $T_A : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the linear transformation corresponding to A , and $T_B : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is the linear transformation corresponding to B . Which of the following correctly describes the relation between the rows, columns, image and kernel? Please see Option (E) before answering.

- (A) If B can be obtained from A by a sequence of row interchange operations, then T_A and T_B have the same kernel as each other and also the same image as each other.
- (B) If B can be obtained from A by a sequence of column interchange operations, then T_A and T_B have the same kernel as each other and also the same image as each other.
- (C) If B can be obtained from A by a sequence of row interchange operations, then T_A and T_B have the same kernel as each other. If B can be obtained from A by a sequence of column interchange operations, then T_A and T_B have the same image as each other.
- (D) If B can be obtained from A by a sequence of row interchange operations, then T_A and T_B have the same image as each other. If B can be obtained from A by a sequence of column interchange operations, then T_A and T_B have the same kernel as each other.
- (E) None of the above.

Answer: Option (C)

Explanation: Interchanging the rows is a legitimate operation for computing the reduced row-echelon form, and the process does not affect the solution set for the system of linear equations, aka the kernel.

Also, the columns of the matrix are a spanning set for the image, so interchanging the columns should not affect the image.

However, interchanging the rows can alter the image. For instance:

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}$$

have the same set of rows, but different images.

Similarly, interchanging the columns can alter the kernel. For instance:

$$A = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$$

have the same set of columns (interchanged) but different kernels.

Performance review: 10 out of 27 got this. 6 each chose (D) and (E), 3 chose (B), 1 chose (A), 1 left the question blank.

**DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: ORIGINALLY DUE FRIDAY
NOVEMBER 15, DELAYED TO WEDNESDAY NOVEMBER 20: LINEAR
DEPENDENCE, BASES, AND SUBSPACES**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 3-question quiz. The score distribution was as follows:

- Score of 0: 1 person
- Score of 1: 8 people
- Score of 2: 13 people
- Score of 3: 3 people

The question-wise answers and performance review were as follows:

- (1) Option (E): 21 people
- (2) Option (D): 14 people
- (3) Option (C): 8 people

2. SOLUTIONS

PLEASE DO NOT DISCUSS ANY QUESTIONS

The purpose of this quiz is to review some basic ideas from part of the lecture notes titled **Linear dependence, bases, and subspaces**. The corresponding sections of the book are Sections 3.2 and 3.3.

- (1) *Do not discuss this!*: Suppose S is a finite nonempty set of vectors in \mathbb{R}^n , and T is a nonempty subset of S . What can we say about S and T ?
 - (A) S is linearly dependent if and only if T is linearly dependent. S is linearly independent if and only if T is linearly independent.
 - (B) If S is linearly dependent, then T is linearly dependent. If S is linearly independent, then T is linearly independent. However, we cannot deduce anything about the linear dependence or independence of S from the linear dependence or independence of T .
 - (C) If T is linearly dependent, then S is linearly dependent. If T is linearly independent, then S is linearly independent. However, we cannot deduce anything about the linear dependence or independence of T from the linear dependence or independence of S .
 - (D) If S is linearly dependent, then T is linearly dependent. If T is linearly independent, then S is linearly independent. We cannot make either of the two other deductions.
 - (E) If T is linearly dependent, then S is linearly dependent. If S is linearly independent, then T is linearly independent. We cannot make either of the other two deductions.

Answer: Option (E)

Explanation: Any linear relation between the vectors in T is also a linear relation between the vectors in S , because all the vectors in T are vectors in S .

The correct statement in the reverse direction is actually a logically equivalent statement, namely, the *contrapositive*. All it's saying is that if S is *not* linearly dependent, then T couldn't have been linearly dependent either, because if T had been linearly dependent, then S would have been too. Explicitly, given any implication $P \implies Q$, the implication $(\text{not } Q) \implies (\text{not } P)$, called the *contrapositive*, also holds. This is just a special case.

For more details, see the lecture notes or textbook.

Performance review: 21 out of 25 got this. 3 chose (B), 1 chose (D).

- (2) *Do not discuss this!* Suppose S is a finite set of vectors in \mathbb{R}^n . Consider the three statements: (i) S is linearly independent, (ii) S does not contain the zero vector, (iii) S does not contain any two vectors that are scalar multiples of one another. Which of the following options best describes the relationship between these statements?
- (A) (i) is equivalent to (ii), and both imply (iii), but the reverse implication does not hold.
 (B) (i) is equivalent to (iii), and both imply (ii), but the reverse implication does not hold.
 (C) (i) is equivalent to (ii) and (iii) combined.
 (D) (i) implies both (ii) and (iii), but (ii) and (iii), even if combined, do not imply (i).

Answer: Option (D)

Explanation: If either (ii) or (iii) is violated, we can obtain a linear relation within S , making S linearly dependent. Thus, the negation of (ii) or (iii) implies the negation of (i). Therefore, the contrapositive holds: (i) implies both (ii) and (iii).

However, (ii) and (iii), even if combined, do not imply (i). This is because there can be linear relations between the vectors that involve more than two vectors at a time. For instance, consider the set of vectors $\vec{v}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\vec{v}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, and $\vec{v}_3 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$. These three vectors are linearly dependent, because $\vec{v}_1 + \vec{v}_2 - \vec{v}_3 = \vec{0}$. But none of the vectors equals zero, and no two vectors are scalar multiples of each other.

Performance review: 14 out of 25 got this. 6 chose (C), 5 chose (B).

- (3) *Do not discuss this!* Suppose V is a linear subspace of \mathbb{R}^n for some n , and W is a linear subspace of V . Assume also that $W \neq V$, i.e., W is a *proper* subspace of V . Which of the following correctly describes the relationship between bases of V and bases of W ?
- (A) Given a basis of V , we can find a subset of that basis that is a basis of W . Also, given a basis of W , we can find a set containing that basis that is a basis of V .
 (B) Given a basis of V , we can find a subset of that basis that is a basis of W . However, given a basis of W , we may not necessarily be able to find a set containing that basis that is a basis of V .
 (C) Given a basis of V , we may not necessarily be able to find a subset of that basis that is a basis of W . However, given a basis of W , we can find a set containing that basis that is a basis of V .

Answer: Option (C)

Explanation: For a counterexample for the part about getting a basis for W from a basis for V , consider the case that $V = \mathbb{R}^2$ has the standard basis (\vec{e}_1 and \vec{e}_2) and W is the subspace spanned by $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$. Clearly, there is no subset of the standard basis of V that forms a basis for W . In fact, none of the standard basis vectors of V is in W .

The idea behind showing the other direction is: start with a basis of W . Then, keep adding vectors one by one, making sure that each new vector being added is not in the span of the vectors so far. We will eventually obtain a basis of V .

Performance review: 8 out of 25 got this. 10 chose (A), 7 chose (B).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY NOVEMBER 20:
IMAGE AND KERNEL: APPLICATIONS TO CALCULUS**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 7-question quiz. The score distribution was as follows:

- Score of 1: 1 person
- Score of 2: 3 people
- Score of 3: 5 people
- Score of 4: 3 people
- Score of 5: 9 people
- Score of 6: 3 people
- Score of 7: 1 person

The question-wise answers and performance review were as follows:

- (1) Option (B): 4 people
- (2) Option (B): 8 people
- (3) Option (C): 18 people
- (4) Option (B): 19 people
- (5) Option (C): 14 people
- (6) Option (A): 18 people
- (7) Option (A): 23 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

The goal of this quiz is to use the setting of calculus to practice our skill of understanding linear transformations, specifically their injectivity, surjectivity, bijectivity, kernel and image. It builds on the November 8 quiz, but goes further. Please refer back to the November 8 quiz for the definitions of vector space, subspace, and linear transformation.

Please read these questions *very* carefully. For the first few questions, the interpretation of the question in the language of calculus is provided. Please refer to that if the linear algebra-based description is unclear.

- (1) Let $\mathbb{R}[x]$ denote the vector space of all polynomials in one variable with real coefficients, with the usual addition and scalar multiplication of polynomials. There is an obvious linear transformation from $\mathbb{R}[x]$ to $C^\infty(\mathbb{R})$ that sends any polynomial to the function it describes, e.g., the polynomial $x^2 + 1$ gets sent to the function $x \mapsto x^2 + 1$. What can you say about this map $\mathbb{R}[x] \rightarrow C^\infty(\mathbb{R})$?

Please note: We are *not* talking here about whether the polynomial functions themselves are injective or surjective as functions from \mathbb{R} to \mathbb{R} . Rather, we are talking about whether the mapping from *the set of polynomials* (which itself is a vector space over the reals) to *the set of infinitely differentiable functions* (which itself is another vector space).

- (A) The map is neither injective nor surjective, i.e., different polynomials may define the same function, and not every infinitely differentiable function can be expressed using a polynomial.
- (B) The map is injective but not surjective, i.e., different polynomials always define different functions, and not every infinitely differentiable function can be expressed using a polynomial.
- (C) The map is surjective but not injective, i.e., different polynomials may define the same function, and every infinitely differentiable function can be expressed using a polynomial.

(D) The map is bijective, i.e., different polynomials always define different functions, and every infinitely differentiable function can be expressed using a polynomial.

Answer: Option (B)

Explanation: The map is linear, so to prove injectivity, it suffices to show that the kernel is zero. In other words, it suffices to show that if $p(x) \in \mathbb{R}[x]$ is in the kernel, then p is the zero polynomial.

Suppose p is in the kernel. Then, this means that $p(x) = 0$ for all $x \in \mathbb{R}$. This means that *every* real number is a root of p . But a nonzero polynomial can have only finitely many roots (the number of roots is at most equal to the degree of the polynomial), so this forces p to be the zero polynomial.

The map is not surjective because there exist lots of infinitely differentiable functions that are not polynomials. Examples include exponential and trigonometric functions.

Performance review: 4 out of 25 got this. 12 chose (C), 7 chose (D), 2 chose (A).

Historical note (last time): 3 out of 26 got this. 10 chose (D), 8 chose (C), and 5 chose (A).

- (2) Denote by $\mathbb{R}[[x]]$ the vector space of all *formal power series* in one variable with real coefficients, with coefficient-wise addition and scalar multiplication. Explicitly, an element $a \in \mathbb{R}[[x]]$ is of the form:

$$a = \sum_{i=0}^{\infty} a_i x^i = a_0 + a_1 x + a_2 x^2 + \dots$$

where $a_i \in \mathbb{R}$ for $i \in \mathbb{N}_0$. Addition is coefficient-wise, i.e., if:

$$a = \sum_{i=0}^{\infty} a_i x^i, b = \sum_{i=0}^{\infty} b_i x^i$$

Then we have:

$$a + b = \sum_{i=0}^{\infty} (a_i + b_i) x^i$$

and for any real number λ , we have:

$$\lambda a = \sum_{i=0}^{\infty} (\lambda a_i) x^i$$

Note that a formal power series may have any radius of convergence. The radius of convergence could range from being 0 (which means that the formal power series converges only at the point $\{0\}$) to being ∞ (which means that the formal power series converges on all of \mathbb{R}). In other words, a formal power series need not define an actual function on \mathbb{R} .

Aside: If you remember sequences and series from single-variable calculus, you will recall that the radius of convergence is the reciprocal of the exponential growth rate of coefficients. In particular, if the coefficients *grow superexponentially*, the radius of convergence is zero. On the other hand, if the coefficients *decay superexponentially*, the radius of convergence is ∞ . If the coefficients have exponential growth, the radius of convergence is less than 1. If the coefficients have exponential decay, the radius of convergence is greater than 1. Finally, if the coefficients grow or decay subexponentially, the radius of convergence is 1.

Note that $\mathbb{R}[x]$ can be viewed as a subspace of $\mathbb{R}[[x]]$ by thinking of each polynomial as a formal power series where there are only finitely many nonzero coefficients.

Let Ω be the subset of $\mathbb{R}[[x]]$ comprising those formal power series that converge globally, i.e., the radius of convergence is ∞ . Note that Ω is a *subspace* of $\mathbb{R}[[x]]$.

What is the relation between $\mathbb{R}[x]$ and Ω ?

Note that by *proper* subspace we mean a subspace that is not equal to the whole space.

(A) $\mathbb{R}[x] = \Omega$, i.e., a power series is globally convergent if and only if it is a polynomial (i.e., it has only finitely many nonzero coefficients).

(B) $\mathbb{R}[x]$ is a proper subspace of Ω , i.e., every polynomial is a globally convergent power series, but there exist globally convergent power series that are not polynomials.

- (C) Ω is a proper subspace of $\mathbb{R}[x]$, i.e., every globally convergent power series is a polynomial, but there are polynomials that are not globally convergent power series.
- (D) $\mathbb{R}[x]$ and Ω are incomparable, i.e., there exist polynomials that are not globally convergent power series and there exist globally convergent power series that are not polynomials.

Answer: Option (B)

Explanation: The power series given by a polynomial converges, because it is a finite sum.

However, there do exist globally convergent power series that are not polynomials. Specifically, any power series where the coefficients decay super-exponentially. An example is the power series of the exponential function.

Performance review: 8 out of 25 got this. 8 chose (C), 5 chose (D), 4 chose (A).

Historical note (last time): 12 out of 26 got this. 11 chose (C), 2 chose (A), 1 chose (D).

- (3) The *Taylor series operator* can be viewed as a linear transformation from $C^\infty(\mathbb{R})$ to $\mathbb{R}[[x]]$. This operator sends any infinitely differentiable function to its Taylor series centered at 0. Explicitly, the operator is:

$$f \mapsto \sum_{k=0}^{\infty} \frac{f^{(k)}(0)}{k!} x^k$$

What can we say about the kernel of this linear transformation?

- (A) The kernel is the set of functions f satisfying $f(0) = 0$
- (B) The kernel is the set of functions f satisfying $f'(0) = 0$
- (C) The kernel is the set of functions f such that f and all its derivatives take the value 0 at 0.
- (D) The kernel is the set of polynomial functions.
- (E) The kernel is the set of functions that have globally convergent power series.

Answer: Option (C)

Explanation: This is direct from the definition.

Performance review: 18 out of 25 got this. 4 chose (B), 2 chose (E), 1 chose (A).

Historical note (last time): 5 out of 26 got this. 10 chose (B), 9 chose (A), 1 each chose (D) and (E).

- (4) Which of the following is the best explanation for why we put the $+C$ when performing indefinite integration?
- (A) The kernel of differentiation is a zero-dimensional space (namely, the zero function only), hence the fibers (inverse images or pre-images) for differentiation are all zero-dimensional spaces, i.e., single functions.
- (B) The kernel of differentiation is a one-dimensional space (namely, the vector space of constant functions), hence the fibers (inverse images or pre-images) for differentiation are all one-dimensional spaces, i.e., lines that are translates of the space of constant functions.
- (C) The image of differentiation is a zero-dimensional space (namely, the zero function only), hence the fibers (inverse images or pre-images) for differentiation are all zero-dimensional spaces, i.e., single functions.
- (D) The image of differentiation is a one-dimensional space (namely, the vector space of constant functions), hence the fibers (inverse images or pre-images) for differentiation are all one-dimensional spaces, i.e., lines that are translates of the space of constant functions.

Answer: Option (B)

Explanation: This is obvious once you think about it.

Performance review: 19 out of 25 got this. 4 chose (D), 2 chose (A).

Historical note (last time): 11 out of 26 got this. 7 chose (C), 4 chose (D), 3 chose (A), 1 chose (E).

- (5) When finding all functions f on \mathbb{R} such that $f''(x) = g(x)$ for some known continuous function g on \mathbb{R} , we get a general description of the form $G(x) + C_1x + C_2$ where C_1, C_2 , are arbitrary real numbers. Which of the following is the best explanation for this?
- (A) The kernel of the operation of differentiating twice is precisely the set of constant functions.

- (B) The kernel of the operation of differentiating twice is precisely the set of nonconstant linear functions.
- (C) The kernel of the operation of differentiating twice is the union of the set of constant functions and the set of nonconstant linear functions.
- (D) The image of the operation of differentiating twice is precisely the set of constant functions.
- (E) The image of the operation of differentiating twice is precisely the set of nonconstant linear functions.

Answer: Option (C)

Explanation: Note that the kernel is the set of functions of the form $x \mapsto C_1x + C_2$ where $C_1, C_2 \in \mathbb{R}$. Note that C_1 and C_2 are allowed to be equal to zero. In the case that $C_1 = 0$, we get constant functions, and in the case $C_1 \neq 0$, we get nonconstant linear functions. The kernel includes both types.

Performance review: 14 out of 25 got this. 6 chose (A), 2 each chose (B) and (D), 1 chose (E).

Historical note (last time): 5 out of 26 got this. 11 chose (A), 6 chose (D), 3 chose (B), 1 chose (E).

- (6) Consider a second-order homogeneous linear differential equation of the form:

$$y'' + p_1(x)y' + p_2(x)y = 0$$

where x is the independent variable, y is the dependent variable, and p_1 and p_2 are known functions. We are trying to find global solutions, i.e., functions defined on all of \mathbb{R} . One way of thinking of this is to consider the linear transformation L that sends a function y of x to $L(y) = y'' + p_1(x)y' + p_2(x)y$, a new function of x . Which of the following best describes what we are trying to do?

- (A) L is a linear transformation $C^2(\mathbb{R}) \rightarrow C(\mathbb{R})$, and the solution space we are interested in is the kernel of L .
- (B) L is a linear transformation $C^2(\mathbb{R}) \rightarrow C(\mathbb{R})$, and the solution space we are interested in is the image of L .
- (C) L is a linear transformation $C(\mathbb{R}) \rightarrow C^2(\mathbb{R})$, and the solution space we are interested in is the kernel of L .
- (D) L is a linear transformation $C(\mathbb{R}) \rightarrow C^2(\mathbb{R})$, and the solution space we are interested in is the image of L .

Answer: Option (A)

Explanation: L makes sense for functions in $C^2(\mathbb{R})$ in so far as it requires differentiating twice. It does not make sense for other functions. The image of L could (a priori) land anywhere in $C(\mathbb{R})$. We are interested in the kernel of L , i.e., the functions whose image is 0, because the left side of the differential equation is precisely computing $L(y)$.

Performance review: 18 out of 25 got this. 4 chose (B), 2 chose (C), 1 chose (D).

Historical note (last time): 5 out of 26 got this. 9 chose (B), 8 chose (C), 4 chose (D).

- (7) Consider a second-order non-homogeneous linear differential equation of the form:

$$y'' + p_1(x)y' + p_2(x)y = q(x)$$

where x is the independent variable, y is the dependent variable, and p_1 , p_2 , and q are known functions. We are trying to find global solutions, i.e., functions defined on all of \mathbb{R} . One way of thinking of this is to consider the linear transformation L that sends a function y of x to $L(y) = y'' + p_1(x)y' + p_2(x)y$, a new function of x . Which of the following best describes what we are trying to do?

- (A) We are trying to find the inverse image under L of $q(x)$, and we know this is a translate of the solution space of the corresponding homogeneous linear differential equation (the one from the preceding question).
- (B) We are trying to find the image under L of $p_1(x)$, and we know this is a translate of the solution space of the corresponding homogeneous linear differential equation (the one from the preceding question).

Answer: Option (A)

Explanation: This follows from the fibers being translates of the kernel.

Performance review: 23 out of 25 got this. 2 chose (B).

TAKE-HOME CLASS QUIZ SOLUTIONS: DUE FRIDAY NOVEMBER 22: LINEAR DYNAMICAL SYSTEMS

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

24 people took this 12-question quiz. The score distribution was as follows:

- Score of 3: 3 people
- Score of 4: 5 people
- Score of 5: 6 people
- Score of 6: 4 people
- Score of 7: 3 people
- Score of 10: 3 people

The mean score was about 5.58.

The question-wise answers and performance review were as follows:

- (1) Option (A): 21 people
- (2) Option (D): 11 people
- (3) Option (E): 8 people
- (4) Option (E): 10 people
- (5) Option (A): 11 people
- (6) Option (B): 13 people
- (7) Option (C): 14 people
- (8) Option (B): 15 people
- (9) Option (A): 9 people
- (10) Option (C): 7 people
- (11) Option (A): 8 people
- (12) Option (A): 7 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS *ALL* QUESTIONS.

This quiz covers a topic that we will not be able to get to formally in the course due to time constraints. The corresponding section of the book is Section 7.1, and there is more relevant material discussed in the later sections of Chapter 7. However, you do not need to read those sections in order to attempt this quiz. Also, simply mastering the computational techniques in those sections of the book will not help you much with the quiz questions.

The questions here consider a linear dynamical system. Consider a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$. Let A be the matrix of T , so that A is a $n \times n$ matrix. For any positive integer r , the matrix A^r is the matrix for the linear transformation T^r (note here that T^r refers to the r -fold *composite* of T). The goal is to determine, starting off with an arbitrary vector $\vec{x} \in \mathbb{R}^n$, how the following sequence behaves:

$$\vec{x}, T(\vec{x}), T^2(\vec{x}), T^3(\vec{x}), \dots$$

More explicitly, each term of the sequence is obtained by applying T to the preceding term. In other words, the sequence is:

$$\vec{x}, T(\vec{x}), T(T(\vec{x})), T(T(T(\vec{x}))), \dots$$

- (1) What is the necessary and sufficient condition on A such that for *every* choice of $\vec{x} \in \mathbb{R}^n$, the sequence described above eventually reaches, and stays at, the zero vector? Note that if it reaches the zero vector, it must do so in at most n steps. Please see Option (E) before answering.
- (A) A is a nilpotent matrix.
 - (B) A is an idempotent matrix.
 - (C) A is an invertible matrix.
 - (D) A is a non-invertible matrix.
 - (E) None of the above.

Answer: Option (A)

Explanation: First, note that if A is a nilpotent matrix, there exists a positive integer r such that A^r is the zero matrix, which is equivalent to requiring that T^r be the zero linear transformation. In particular, this means that $T^r(\vec{x})$ is the zero vector for all initial vectors \vec{x} , so every sequence eventually reaches the zero vector.

Other direction: From the fact stated, if the sequence reaches the zero vector, it must do so in n steps, so that this means that if the condition holds for every $\vec{x} \in \mathbb{R}^n$, then T^n sends every vector to the zero vector. In particular, this means that A^n is the zero matrix, so A is nilpotent.

Performance review: 21 out of 24 got this. 3 chose (E).

Historical note (last time): 21 out of 24 got this. 2 chose (E), 1 chose (B).

- (2) What is the necessary and sufficient condition on A such that there *exists* a nonzero vector $\vec{x} \in \mathbb{R}^n$ for which the sequence described above eventually reaches, and stays at, the zero vector? Note that if it reaches the zero vector, it must do so in at most n steps. Please see Option (E) before answering.
- (A) A is a nilpotent matrix.
 - (B) A is an idempotent matrix.
 - (C) A is an invertible matrix.
 - (D) A is a non-invertible matrix.
 - (E) None of the above.

Answer: Option (D)

Explanation: Suppose \vec{x} is a nonzero vector and r is the smallest positive integer such that the vector $A^r \vec{x}$ is the zero vector. Then, $A(A^{r-1} \vec{x}) = 0$ but $A^{r-1} \vec{x} \neq 0$, so that $A^{r-1} \vec{x}$ is a nonzero vector in the kernel of T . Thus, T has a nonzero kernel, so it must be non-invertible, hence A must be a non-invertible matrix.

Conversely, if A is non-invertible, there is a nonzero vector, say \vec{x} , in the kernel of A . This vector can be used.

Performance review: 11 out of 24 got this. 7 chose (A), 5 chose (E), 1 chose (B).

Historical note (last time): 5 out of 24 got this. 15 chose (A), 3 chose (C), and 1 chose (E).

- (3) What is the necessary and sufficient condition on A such that for *every* choice of $\vec{x} \in \mathbb{R}^n$, the sequence described above returns to \vec{x} after a finite and positive number of steps? Please see Option (E) before answering.
- (A) A is a nilpotent matrix.
 - (B) A is an idempotent matrix.
 - (C) A is an invertible matrix.
 - (D) A is a non-invertible matrix.
 - (E) None of the above.

Answer: Option (E)

Explanation: It is definitely a *necessary* condition that A be invertible, otherwise there would be a nonzero vector in its kernel for which the sequence could never return to the vector. However, this is not sufficient. Consider the case where $n = 1$ and the matrix A is $[2]$. This is invertible, but applying it repeatedly to a nonzero vector can never get us back to the original vector.

Performance review: 8 out of 24 got this. 16 chose (B).

Historical note (last time): 5 out of 24 got this. 16 chose (B), 2 chose (C), 1 chose (A).

- (4) What is the necessary and sufficient condition on A such that there *exists* a nonzero vector $\vec{x} \in \mathbb{R}^n$ for which the sequence described above returns to \vec{x} after a finite and positive number of steps? Please see Option (E) before answering.

- (A) A is a nilpotent matrix.
- (B) A is an idempotent matrix.
- (C) A is an invertible matrix.
- (D) A is a non-invertible matrix.
- (E) None of the above.

Answer: Option (E)

Explanation: Similar to the preceding question, except now that A does not even need to be invertible.

Performance review: 10 out of 24 got this. 7 chose (C), 6 chose (B), 1 chose (A).

Historical note (last time): 3 out of 24 got this. 11 chose (A), 9 chose (B), 1 chose (C).

- (5) Suppose $n = 2$ and T is a rotation by an angle that is a rational multiple of π . What can we say about the range of the sequence

$$\vec{x}, T(\vec{x}), T^2(\vec{x}), T^3(\vec{x}), \dots$$

starting from a nonzero vector \vec{x} ?

- (A) The range is finite, i.e., there are only finitely many distinct vectors in the sequence.
- (B) The range is infinite and forms a dense subset of the circle centered at the origin and with radius equal to the length of the vector \vec{x} . However, it is not the entire circle.
- (C) The range is infinite and is the entire circle centered at the origin and with radius equal to the length of the vector \vec{x} .
- (D) The range is infinite and forms a dense subset of the line of the vector \vec{x} (excluding the origin), but is not the entire line (excluding the origin).
- (E) The range is infinite and is the entire line of the vector \vec{x} , excluding the origin.

Answer: Option (A)

Explanation: If the angle of rotation for T is θ , then the angle of rotation for T^r is $r\theta$. This is because angles of rotation add up when we compose the rotations.

Since the angle of rotation θ is a rational multiple of π , it is of the form $p\pi/q$ where p and q are integers with $q \neq 0$. Then, $2q\theta = 2\pi p$. This implies that T^{2q} is rotation by an integer multiple of 2π , and hence, is the identity transformation. In particular, this means that for any nonzero vector \vec{x} , the sequence $\vec{x}, T(\vec{x}), T^2(\vec{x}), \dots$ returns to \vec{x} at $T^{2q}(\vec{x})$. Beyond that point, it will just cycle the same set of vectors.

Performance review: 11 out of 24 got this. 9 chose (C), 3 chose (B), 1 chose (D).

Historical note (last time): 4 out of 24 got this. 12 chose (B), 5 chose (C), 2 chose (E), 1 chose (D).

- (6) Suppose $n = 2$ and T is a rotation by an angle that is an irrational multiple of π . What can we say about the range of the sequence

$$\vec{x}, T(\vec{x}), T^2(\vec{x}), T^3(\vec{x}), \dots$$

starting from a nonzero vector \vec{x} ?

- (A) The range is finite, i.e., there are only finitely many distinct vectors in the sequence.
- (B) The range is infinite and forms a dense subset of the circle centered at the origin and with radius equal to the length of the vector \vec{x} . However, it is not the entire circle.
- (C) The range is infinite and is the entire circle centered at the origin and with radius equal to the length of the vector \vec{x} .
- (D) The range is infinite and forms a dense subset of the line of the vector \vec{x} (excluding the origin), but is not the entire line (excluding the origin).
- (E) The range is infinite and is the entire line of the vector \vec{x} , excluding the origin.

Answer: Option (B)

Explanation: Since T is rotation by an *irrational* multiple θ of π , there is no positive integer multiple of θ that is also an integer multiple of 2π . Thus, T^r is not the identity map for any positive integer r , so there is no cycling back, so we do get an infinite set of vectors. All of them have the same length as \vec{x} , because rotations preserve length. Thus, they lie on the circle centered at the origin with length equal to the length of the vector \vec{x} . In fact, we can show that they form a dense

subset of the circle, i.e., they come arbitrarily close to every point of the circle. Note, however, that we do not get all points on the circle. For instance, $-\vec{x}$ is not in the range, because achieving it would require a nonzero rational (in fact, odd integer) multiple of π , and no integer multiple of an irrational multiple of π is of that form.

Performance review: 13 out of 24 got this. 8 chose (C), 2 chose (D), 1 chose (A).

Historical note (last time): 9 out of 24 got this. 4 chose (A), 7 chose (C), 3 chose (D), 1 chose (E).

We return to generic n now.

- (7) A nonzero vector \vec{x} is termed an *eigenvector* for a linear transformation $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ with *eigenvalue* a real number $\lambda \in \mathbb{R}$ if $T(\vec{x}) = \lambda\vec{x}$. Note that λ is allowed to be 0. We sometimes conflate the roles of T and its matrix A , so that we call \vec{x} an eigenvector for A and λ an eigenvalue for A .

If \vec{x} is an eigenvector of T (or equivalently, of A) with eigenvalue λ , which of the following is true? We denote by I_n the identity transformation from \mathbb{R}^n to itself.

- (A) \vec{x} must be in the kernel of the linear transformation $T + \lambda I_n$
- (B) \vec{x} must be in the image of the linear transformation $T + \lambda I_n$
- (C) \vec{x} must be in the kernel of the linear transformation $T - \lambda I_n$
- (D) \vec{x} must be in the image of the linear transformation $T - \lambda I_n$
- (E) \vec{x} must be in the kernel of the linear transformation λT

Answer: Option (C)

Explanation: We are trying to find the vectors \vec{x} such that $T(\vec{x}) = \lambda\vec{x}$. This can be rewritten as $T\vec{x} = \lambda I_n \vec{x}$ and hence as $(T - \lambda I_n)\vec{x} = 0$, so that \vec{x} is in the kernel of $T - \lambda I_n$.

Note that the other options are false for the following reasons:

- Option (A): Consider the case that $T = I_n$, \vec{x} is any nonzero vector, and $\lambda = 1$. Then, $T + \lambda I_n = 2I_n$, the kernel of which is zero, so \vec{x} is not in the kernel. In fact, Option (A) holds true if and only if $\lambda = 0$.
- Option (B): Consider the case that $T = 0$, \vec{x} is any nonzero vector, and $\lambda = 0$. Then, $T + \lambda I_n = 0$, the image of which is zero, so \vec{x} is not in the image. In fact, Option (B) holds true if $\lambda \neq 0$. If $\lambda \neq 0$, then $\vec{x} = (T + \lambda I_n)(\vec{x}/(2\lambda))$. For $\lambda = 0$, Option (B) may or may not hold. The example of T being the zero linear transformation gives a situation where it does not hold. Here is an example where it does hold. Consider:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

Then, \vec{e}_1 is an eigenvector for T with eigenvalue 0, and it is in the image of $T + \lambda I_n = T$, since $T(\vec{e}_2) = \vec{e}_1$.

- Option (D): Consider the case that $T = I_n$, \vec{x} is any nonzero vector, and $\lambda = 1$. Then, $T - \lambda I_n = 0$, the image of which is zero, so \vec{x} is not in the image. Option (D) is often true and often false, but its truth or falsehood does not have any direct relation with whether λ is zero or nonzero.
- Option (E): Consider the case that $T = I_n$, \vec{x} is any nonzero vector, and $\lambda = 1$. Then, $\lambda T = I_n$, the kernel of which is zero, so \vec{x} is not in the kernel. In fact, Option (E) holds true if and only if $\lambda = 0$.

Performance review: 14 out of 24 got this. 4 chose (E), 3 chose (B), 2 chose (A), 1 chose (D).

Historical note (last time): 17 out of 24 got this. 4 chose (A), 2 chose (B), 1 chose (E).

- (8) As above, let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear transformation with matrix A . Use the terminology of eigenvector and eigenvalue from the preceding question. Which of the following is a characterization of the situation that A is a diagonal matrix?

- (A) Every nonzero vector in \mathbb{R}^n is an eigenvector for T .
- (B) Every standard basis vector in \mathbb{R}^n is an eigenvector for T .

- (C) Every vector with at least one zero coordinate in \mathbb{R}^n is an eigenvector for T .
 (D) T has a unique eigenvector (up to scalar multiples, i.e., all eigenvectors of T are scalar multiples of each other).
 (E) T has no eigenvector.

Answer: Option (B)

Explanation: Suppose the matrix of T has c_i as the i^{th} diagonal entry. Then, $T(\vec{e}_i)$ is the i^{th} column of the diagonal matrix, which has 0s everywhere except in the diagonal entry, which is c_i . Thus, $T(\vec{e}_i) = c_i\vec{e}_i$, so that all standard basis vectors are eigenvectors.

Conversely, if every standard basis vector is an eigenvector, then for each i , $T(\vec{e}_i) = c_i\vec{e}_i$ for some c_i , which forces the i^{th} column to have the value c_i in the diagonal position and the value 0 elsewhere. This gives a diagonal matrix.

Performance review: 15 out of 24 got this. 5 chose (D), 2 chose (A), 1 each chose (C) and (E).

Historical note (last time): 13 out of 24 got this. 5 chose (E), 3 chose (D), 2 chose (C), 1 chose (A).

- (9) As above, let $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear transformation with matrix A . Use the terminology of eigenvector and eigenvalue from the preceding question. Which of the following is a characterization of the situation that A is a scalar matrix (i.e., a diagonal matrix with all diagonal entries equal)?
 (A) Every nonzero vector in \mathbb{R}^n is an eigenvector for T .
 (B) Every standard basis vector in \mathbb{R}^n is an eigenvector for T .
 (C) Every vector with at least one zero coordinate in \mathbb{R}^n is an eigenvector for T .
 (D) T has a unique eigenvector (up to scalar multiples, i.e., all eigenvectors of T are scalar multiples of each other).
 (E) T has no eigenvector.

Answer: Option (A)

Explanation: If the matrix for T is scalar with scalar value λ , then $T(\vec{x}) = \lambda\vec{x}$ for all vectors $\vec{x} \in \mathbb{R}^n$. Thus, all nonzero vectors are eigenvectors.

To establish the converse, we need to show that if every nonzero vector of T is an eigenvector, then all of them have the *same* eigenvalue. Note that two vectors in the same line have the same eigenvalue, so the statement is trivial for $n = 1$.

For $n \geq 2$, we already know that the matrix of T is diagonal on account of the preceding question. We want to show that all diagonal entries are equal to each other. Consider the i^{th} and j^{th} diagonal entries. Let's say these are c_i and c_j respectively. Then, $T(\vec{e}_i + \vec{e}_j) = T(\vec{e}_i) + T(\vec{e}_j)$. This simplifies to $c_i\vec{e}_i + c_j\vec{e}_j$. For this to be a multiple of $\vec{e}_i + \vec{e}_j$, we need that $c_i = c_j$. Since this is true for every pair of i and j , we get that all the diagonal entries are equal and that the matrix is a scalar matrix.

Performance review: 9 out of 24 got this. 8 chose (B), 7 chose (D).

Historical note (last time): 8 out of 24 got this. 11 chose (B), 2 each chose (C) and (E), 1 chose (D).

- (10) Suppose A is a strictly upper-triangular $n \times n$ matrix, i.e., all entries of A that are on or below the main diagonal are zero. T is the linear transformation corresponding to A . It will turn out that the only eigenvalue for T is 0. What can we say about the eigenvectors for T for this eigenvalue?
 (A) All nonzero vectors in \mathbb{R}^n are eigenvectors for T with eigenvalue 0.
 (B) All standard basis vectors in \mathbb{R}^n are eigenvectors for T with eigenvalue 0.
 (C) The vector \vec{e}_1 is an eigenvector for T with eigenvalue 0. The information presented is not sufficient to determine whether any of the other standard basis vectors is an eigenvector.
 (D) The vector \vec{e}_n is an eigenvector for T with eigenvalue 0. The information presented is not sufficient to determine whether any of the other standard basis vectors is an eigenvector.
 (E) At least one of the standard basis vectors is an eigenvector for T with eigenvalue 0. However, the information presented is not sufficient to say definitively for any particular standard basis vector that it is an eigenvector.

Answer: Option (C)

Explanation: The first column is the zero column, so the vector \vec{e}_1 gets mapped to zero. As for the remaining standard basis vectors, we do not have enough information to know whether they are

sent to zero. This is because there are entries above the diagonal in those columns, and these are allowed to be nonzero.

For instance, consider:

$$\begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

This sends \vec{e}_1 to the zero vector and sends \vec{e}_2 to \vec{e}_1 .

Performance review: 7 out of 24 got this. 5 each chose (A) and (D), 3 each chose (B) and (E), 1 left the question blank.

Historical note (last time): 5 out of 24 got this. 9 chose (A), 6 chose (E), 2 each chose (B) and (D).

- (11) Suppose A is a strictly upper-triangular $n \times n$ matrix, i.e., all entries of A that are on or below the main diagonal are zero. T is the linear transformation corresponding to A . Which of the following is A guaranteed to be? Please see Options (D) and (E) before answering.

- (A) Nilpotent
- (B) Idempotent
- (C) Invertible
- (D) All of the above
- (E) None of the above

Answer: Option (A)

Explanation: The image of T is contained in the span of the vectors $\vec{e}_1, \vec{e}_2, \dots, \vec{e}_{n-1}$. The image of T^2 is in the span of $\vec{e}_1, \vec{e}_2, \dots, \vec{e}_{n-2}$. Each time we apply T , we lose the last basis vector. Thus, T^n is the zero transformation, so A^n is the zero matrix, so A is nilpotent.

For instance, consider:

$$\begin{bmatrix} 0 & 1 & 2 \\ 0 & 0 & 3 \\ 0 & 0 & 0 \end{bmatrix}$$

The image of T is spanned by the column vectors. The first column is the zero column, so it is spanned by the other two column vectors:

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 2 \\ 3 \\ 0 \end{bmatrix}$$

This is the same as the span of \vec{e}_1 and \vec{e}_2 .

The image of T^2 is thus the image of the span of these two vectors. This is the span of the first two columns of A . The first column is zero, so the image of T^2 is simply the span of the second column vector:

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$$

This is the span of \vec{e}_1 . The image of this under T again is the zero space. Thus, the image of T^3 is zero, so $A^3 = 0$.

Performance review: 8 out of 24 got this. 10 chose (E), 5 chose (C), 1 chose (D).

Historical note (last time): 7 out of 24 got this. 10 chose (C), 5 chose (E), 1 chose (D). 1 left the question blank.

- (12) Consider the case $n = 2$ and let $T : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a rotation by an angle that is *not* an integer multiple of π . What can we say about the set of eigenvectors and eigenvalues for T ?

- (A) T has no eigenvectors
- (B) T has one eigenvector (up to scalar multiples) with eigenvalue 1
- (C) T has one eigenvector (up to scalar multiples) and the eigenvalue depends on the angle of rotation

- (D) T has two linearly independent eigenvectors (so that the set of all eigenvectors is obtained as the set of scalar multiples of either one of these vectors) with the same eigenvalue
- (E) T has two linearly independent eigenvectors (so that the set of all eigenvectors is obtained as the set of scalar multiples of either one of these vectors) with distinct eigenvalues

Answer: Option (A)

Explanation: Every nonzero vector gets rotated, so no nonzero vector goes to a scalar multiple of itself. Note that the case that the angle of rotation is an integer multiple of π differs: in that case, the transformation is either the identity or the negative identity and hence is scalar, so all nonzero vectors are eigenvectors for it.

Performance review: 7 out of 24 got this. 6 chose (C), 4 each chose (B) and (D), 3 chose (E).

Historical note (last time): 6 out of 24 got this. 10 chose (E), 4 chose (C), 2 chose (D), 1 chose (B).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE MONDAY NOVEMBER 25:
STOCHASTIC MATRICES**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

22 people took this 10-question quiz. The score distribution was as follows:

- Score of 4: 2 people
- Score of 5: 2 people
- Score of 6: 5 people
- Score of 7: 4 people
- Score of 8: 5 people
- Score of 9: 1 person
- Score of 10: 3 people

The mean score was about 7.05.

The question-wise answers and performance review were as follows:

- (1) Option (C): 11 people
- (2) Option (C): 13 people
- (3) Option (C): 18 people
- (4) Option (E): 20 people
- (5) Option (B): 21 people
- (6) Option (D): 11 people
- (7) Option (A): 16 people
- (8) Option (C): 19 people
- (9) Option (B): 15 people
- (10) Option (B): 11 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

This quiz can be viewed as a continuation of the quiz on linear dynamical systems. The book defines column-stochastic matrices using the jargon “transition matrix” on Page 53 (Definition 2.1.4) and uses them throughout the text when describing (a simplified version of) Google’s PageRank algorithm. The quiz questions are self-contained and do not require you to read the book, but you may benefit from skimming through the book’s discussion of PageRank to complement these questions. Note that the 4th Edition does not include the discussion of transition matrices and PageRank.

In this quiz, we discuss the dynamics of a very special type of linear transformation. A $n \times n$ matrix A is termed a *row-stochastic matrix* if all its entries are in the interval $[0, 1]$ and all the row sums are equal to 1. A $n \times n$ matrix is termed a *column-stochastic matrix* if all its entries are in the interval $[0, 1]$ and all the column sums are equal to 1. A $n \times n$ matrix A is termed a *doubly stochastic matrix* if it is both row-stochastic and column-stochastic, i.e., all the entries are in the interval $[0, 1]$, all the row sums are equal to 1, and all the column sums are equal to 1.

- (1) Suppose A and B are two $n \times n$ row-stochastic matrices. Which of the following is *guaranteed* to be row-stochastic? Please see Options (D) and (E) before answering.
 - (A) $A + B$
 - (B) $A - B$
 - (C) AB

- (D) All of the above
 (E) None of the above

Answer: Option (C)

Explanation: Suppose we are trying to compute the $(ik)^{th}$ entry of AB . This is the sum:

$$\sum_{j=1}^n a_{ij}b_{jk}$$

We now want to sum up all such entries in the i^{th} row of AB . Thus, the sum is:

$$\sum_{k=1}^n \sum_{j=1}^n a_{ij}b_{jk}$$

The sum can be rearranged as:

$$\sum_{j=1}^n \left(a_{ij} \sum_{k=1}^n b_{jk} \right)$$

Each of the inner sums is 1, on account of being a row sum of B . Thus, the sum simplifies to:

$$\sum_{j=1}^n a_{ij}$$

This is 1, on account of being a row sum of A .

Thus, every row sum of AB is 1. Further, because of the way we define matrix multiplication, all the entries of AB are nonnegative. Combined with the condition on sums, we get that all entries are in $[0, 1]$ with all row sums 1. Thus, the matrix AB is row-stochastic.

As for Options (A) and (B), note that for Option (A), the row sums will become 2 and for Option (B), the row sums will become 0.

Performance review: 11 out of 22 got this. 10 chose (E), 1 chose (A).

Historical note (last time): 14 out of 24 got this. 9 chose (E), 1 chose (D).

- (2) Suppose A and B are two $n \times n$ column-stochastic matrices. Which of the following is *guaranteed* to be column-stochastic? Please see Options (D) and (E) before answering.

- (A) $A + B$
 (B) $A - B$
 (C) AB
 (D) All of the above
 (E) None of the above

Answer: Option (C)

Explanation: Suppose we are trying to compute the $(ik)^{th}$ entry of AB . This is the sum:

$$\sum_{j=1}^n a_{ij}b_{jk}$$

We now want to sum up all such entries in the k^{th} column of AB . Thus, the sum is:

$$\sum_{i=1}^n \sum_{j=1}^n a_{ij}b_{jk}$$

The sum can be rearranged as:

$$\sum_{j=1}^n \left(b_{jk} \sum_{i=1}^n a_{ij} \right)$$

Each of the inner sums is 1, on account of being a row sum of A . Thus, the sum simplifies to:

$$\sum_{j=1}^n b_{jk}$$

This is 1, on account of being a column sum of B .

Thus, every column sum of AB is 1. Further, because of the way we define matrix multiplication, all the entries of AB are nonnegative. Combined with the condition on sums, we get that all entries are in $[0, 1]$ with all row sums 1. Thus, the matrix AB is column-stochastic.

As for Options (A) and (B), note that for Option (A), the column sums will become 2 and for Option (B), the column sums will become 0.

Performance review: 13 out of 22 got this. 8 chose (E), 1 chose (A).

Historical note (last time): 12 out of 24 got this. 11 chose (E), 1 chose (A).

- (3) Suppose A and B are two $n \times n$ doubly stochastic matrices. Which of the following is *guaranteed* to be doubly stochastic? Please see Options (D) and (E) before answering.
- (A) $A + B$
 (B) $A - B$
 (C) AB
 (D) All of the above
 (E) None of the above

Answer: Option (C)

Explanation: This follows by combining the two preceding questions.

Performance review: 18 out of 22 got this. 3 chose (E), 1 chose (A).

Historical note (last time): 11 out of 24 got this. 8 chose (D), 4 chose (E), 1 chose (A).

We now consider the case $n = 2$. In this case, the doubly stochastic matrices have the form:

$$\begin{bmatrix} a & 1-a \\ 1-a & a \end{bmatrix}$$

where $a \in [0, 1]$. Denote this matrix by $D(a)$ for short.

- (4) Suppose $a, b \in [0, 1]$ (they are allowed to be equal). The product $D(a)D(b)$ equals $D(c)$ for some $c \in [0, 1]$. What is that value of c ?
- (A) $a + b$
 (B) ab
 (C) $2ab + a + b$
 (D) $(1-a)(1-b)$
 (E) $1 - a - b + 2ab$

Answer: Option (E)

Explanation: We carry out the multiplication:

$$\begin{bmatrix} a & 1-a \\ 1-a & a \end{bmatrix} \begin{bmatrix} b & 1-b \\ 1-b & b \end{bmatrix} = \begin{bmatrix} ab + (1-a)(1-b) & a(1-b) + (1-a)b \\ a(1-b) + (1-a)b & ab + (1-a)(1-b) \end{bmatrix} = D(ab + (1-a)(1-b)) = D(1-a-b+2ab)$$

Performance review: 20 out of 22 got this. 1 each chose (A) and (D).

Historical note (last time): 19 out of 24 got this. 2 chose (B), 1 each chose (A), (C), and (D).

- (5) For what value(s) of a is the matrix $D(a)$ non-invertible? Note that when judging invertibility, we do not insist that the inverse matrix also be doubly stochastic.
- (A) $a = 0$ only
 (B) $a = 1/2$ only
 (C) $a = 1$ only
 (D) $0 < a < 1$ (i.e., $D(a)$ is invertible only at $a = 0$ and $a = 1$)
 (E) $a \neq 1/2$

Answer: Option (B)

Explanation: For the matrix to be non-invertible, we need the rows to be scalar multiples of each other. Since both row sums are 1, this can happen only if the rows are identical, which happens iff $a = 1/2$. Equivalently, we can note that invertibility requires a nonzero determinant, and the determinant is $a^2 - (1 - a)^2 = 2a - 1$, which is 0 iff $a = 1/2$.

Performance review: 21 out of 22 got this. 1 chose (A).

Historical note (last time): 15 out of 24 got this. 4 chose (A), 2 each chose (C) and (D), 1 chose (E).

- (6) For what value(s) of a is it true that the matrix $D(a)$ does not have an inverse that is a doubly stochastic matrix? In other words, either $D(a)$ should be non-invertible or it should be invertible but the inverse is not a doubly stochastic matrix.

(A) $a = 0$ only

(B) $a = 1/2$ only

(C) $a = 1$ only

(D) $0 < a < 1$ (i.e., $D(a)$ has an inverse that is also doubly stochastic only if $a = 0$ or $a = 1$)

(E) $a \neq 1/2$

Answer: Option (D)

Explanation: The inverse of $D(a)$, for $a \neq 1/2$, is:

$$\begin{bmatrix} a/(2a-1) & (a-1)/(2a-1) \\ (a-1)/(2a-1) & a/(2a-1) \end{bmatrix}$$

For this to be doubly stochastic, we need that $0 \leq a/(2a-1) \leq 1$. We make cases:

- $a = 0$: In this case, $D(a)$ is its own inverse.
- $a \neq 0$ (excluding $a = 1/2$): In this case, since $0 \leq a/(2a-1)$, we obtain that $2a - 1 > 0$. So, starting with $a/(2a-1) \leq 1$ we get $a \leq 2a - 1$, which simplifies to $1 \leq a$, forcing $a = 1$ (since we are constrained by $0 \leq a \leq 1$). The choice $a = 1$ works (in the sense that $D(1)$ is its own inverse).

The upshot is that the only cases where the inverse is doubly stochastic are the cases $a = 0$ and $a = 1$. Otherwise, the inverse either does not exist (case $a = 1/2$) or exists but is not doubly stochastic.

Performance review: 11 out of 22 got this. 9 chose (B), 2 chose (C).

Historical note (last time): 9 out of 24 got this. 9 chose (A), 2 each chose (B), (C), and (E).

For the next few questions, denote by T_a the linear transformation whose matrix is $D(a)$. For any vector $\vec{x} \in \mathbb{R}^2$, we can consider the sequence:

$$\vec{x}, T_a(\vec{x}), T_a^2(\vec{x}), \dots$$

Note that if we were to start with a vector $\vec{x} \in \mathbb{R}^2$ with both coordinates equal, it would be invariant under T_a .

Thus, for the questions below, assume that we start with a nonzero vector $\vec{x} \in \mathbb{R}^2$ for which the two coordinates are not equal to each other.

- (7) For what value of a is it the case that $\lim_{r \rightarrow \infty} T_a^r(\vec{x})$ does *not* exist?

(A) $a = 0$ only

(B) $a = 1/2$ only

(C) $a = 1$ only

(D) $0 < a < 1$

(E) $a \neq 1/2$

Answer: Option (A)

Explanation: In this case, applying T_a interchanges the coordinates. A second application interchanges them back. The sequence thus cycles between the vector \vec{x} and the vector obtained by interchanging its coordinates.

The failure of the remaining options will become clear from the answers to the rest of the questions.

Performance review: 16 out of 22 got this. 2 chose (B), 1 each chose (C), (D), and (E). 1 left the question blank.

Historical note (last time): 7 out of 24 got this. 11 chose (B), 4 chose (C), 1 each chose (D) and (E).

- (8) For what value of a is it the case that the sequence

$$\vec{x}, T_a(\vec{x}), T_a^2(\vec{x}), \dots$$

is a constant sequence?

- (A) $a = 0$ only
- (B) $a = 1/2$ only
- (C) $a = 1$ only
- (D) $0 < a < 1$
- (E) $a \neq 1/2$

Answer: Option (C)

Explanation: $a = 1$ gives the identity transformation. Any other choice of a replaces the first coordinate by a combination of the two coordinates that cannot be equal to it unless the two coordinates were equal to begin with. Explicitly:

$$D(a) \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} ax_1 + (1-a)x_2 \\ (1-a)x_1 + ax_2 \end{bmatrix}$$

For the output to equal the input, we need that:

$$\begin{aligned} x_1 &= ax_1 + (1-a)x_2 \\ x_2 &= (1-a)x_1 + ax_2 \end{aligned}$$

Solving the first equation alone gives $x_1 = x_2$ or $a = 1$. By assumption, $x_1 \neq x_2$, so we get $a = 1$. Note that the second equation yields a similar conclusion.

Performance review: 19 out of 22 got this. 2 chose (B), 1 chose (A).

Historical note (last time): 18 out of 24 got this. 3 chose (D), 2 chose (E), 1 chose (B).

- (9) For what value of a is it the case that the sequence

$$\vec{x}, T_a(\vec{x}), T_a^2(\vec{x}), \dots$$

is not a constant sequence but becomes constant from $T_a(\vec{x})$ onward?

- (A) $a = 0$ only
- (B) $a = 1/2$ only
- (C) $a = 1$ only
- (D) $0 < a < 1$
- (E) $a \neq 1/2$

Answer: Option (B)

Explanation: If the sequence is not constant, then $a \neq 1$. However, it becomes constant from $T_a(\vec{x})$ onward. Thus, $T_a(\vec{x})$ has both coordinates equal by the previous question. Hence, we get that:

$$ax_1 + (1-a)x_2 = (1-a)x_1 + ax_2$$

This simplifies to:

$$(2a-1)(x_1-x_2) = 0$$

Thus, either $a = 1/2$ or $x_1 = x_2$. Since $x_1 \neq x_2$ by assumption, we get $a = 1/2$.

Performance review: 15 out of 22 got this. 3 chose (A), 2 each chose (C) and (D).

Historical note (last time): 8 out of 24 got this. 10 chose (A), 4 chose (C), 1 each chose (D) and (E).

- (10) For a other than 0, $1/2$, or 1, what is the limit $\lim_{r \rightarrow \infty} (D(a))^r$? Here, when we talk of taking the limit of a sequence of matrices, we are taking the limit entry-wise.

- (A) The matrix $D(0)$
- (B) The matrix $D(1/2)$
- (C) The matrix $D(1)$
- (D) The matrix $D(a)$
- (E) The matrix $D(1 - a)$

Answer: Option (B)

Explanation: Intuitively, what's happening is that we start off with the point:

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

We now make the coordinates “come closer to each other” by converting to:

$$\begin{bmatrix} ax_1 + (1 - a)x_2 \\ (1 - a)x_1 + ax_2 \end{bmatrix}$$

We then iterate. Each time, the coordinates are coming closer to each other, but note also that the sum of the coordinates remains fixed. Thus, we hope to eventually converge to the vector with both coordinates $(x_1 + x_2)/2$. This is $D(1/2) \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$.

Here is a more formal demonstration: By one of the previous questions, we can verify that if $D(a)D(b) = D(c)$, then $(2a - 1)(2b - 1) = 2c - 1$. An immediate corollary is that if $(D(a))^r = D(u)$, then $(2a - 1)^r = 2u - 1$. As a result, if $\lim_{r \rightarrow \infty} (D(a))^r = D(u)$, then $2u - 1 = \lim_{r \rightarrow \infty} (2a - 1)^r = 0$, forcing $u = 1/2$.

The geometric intuition for why we look at $2a - 1$ is because that is the ratio of the signed difference between the output coordinates to the signed difference between the input coordinates. It describes an contraction factor, and the contraction factors multiply when we compose.

Performance review: 11 out of 22 got this. 5 chose (A), 4 chose (E), 1 each chose (C) and (D).

Historical note (last time): 3 out of 24 got this. 10 chose (D), 7 chose (E), 2 each chose (A) and (C).

**TAKE-HOME CLASS QUIZ: DUE MONDAY NOVEMBER 25: SUBSPACE, BASIS,
DIMENSION, AND ABSTRACT SPACES: APPLICATIONS TO CALCULUS**

MATH 196, SECTION 57 (VIPUL NAIK)

Your name (print clearly in capital letters): _____

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

This quiz builds on the November 8 and November 20 quizzes that apply ideas we are learning about linear transformations to the calculus setting. The November 8 quiz went over some basic ideas related to differentiation as a linear transformation. The November 20 quiz explored the ideas in greater depth. We now look at questions that apply the ideas of basis, dimension, and subspace to the calculus setting.

We begin by recalling some notation and facts we already saw in earlier quizzes. Denote by $C(\mathbb{R})$ (or alternatively by $C^0(\mathbb{R})$) the vector space of all continuous functions from \mathbb{R} to \mathbb{R} , with pointwise addition and scalar multiplication. Note that the elements of this vector space, which we would ordinarily call “vectors”, are now *functions*.

For k a positive integer, denote by $C^k(\mathbb{R})$ the subspace of $C(\mathbb{R})$ comprising those continuous functions that are at least k times *continuously* differentiable. Note that $C^{k+1}(\mathbb{R})$ is a subspace of $C^k(\mathbb{R})$, so we have a descending chain of subspaces:

$$C(\mathbb{R}) = C^0(\mathbb{R}) \supseteq C^1(\mathbb{R}) \supseteq C^2(\mathbb{R}) \supseteq \dots$$

The intersection of these spaces is the vector space $C^\infty(\mathbb{R})$, defined as the subspace of $C(\mathbb{R})$ comprising those functions that are *infinitely* differentiable.

We had also noted that:

- The kernel of differentiation is the vector space of constant functions.
- The kernel of k times differentiating is the vector space of polynomials of degree at most $k - 1$.
- The fiber of any function for differentiation is a translate of the space of constant functions. That’s what explains the $+C$ when you perform indefinite integration.

Note: For finite-dimensional spaces, a linear transformation T from a vector space to itself is injective if and only if it is surjective. This follows from dimension and rank considerations: T is injective if and only if its kernel is zero, which happens if and only if the matrix has full column rank, which happens if and only if the matrix has full row rank (because the matrix is a square matrix), which happens if and only if T is surjective. The rank-nullity theorem provides an equivalent explanation. We had also seen that if $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is injective, then $m \leq n$, and if $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is surjective, then $m \geq n$. In particular, we cannot have a surjective map from a proper subspace to the whole space.

With infinite-dimensional spaces, however, we can have funny phenomena. Examples of these phenomena are strewn across the quizzes.

- We can have a map from an infinite-dimensional vector space to itself that is injective but not surjective.
- We can have a map from an infinite-dimensional vector space to itself that is surjective but not injective.
- We can have a surjective map from a proper subspace to the whole space (for instance, differentiation $C^1(\mathbb{R}) \rightarrow C(\mathbb{R})$ is surjective, even though $C^1(\mathbb{R})$ is a proper subspace of $C(\mathbb{R})$).
- We can have an injective map from a space to a proper subspace.

Note that we will use the terms *subspace* and *vector subspace* synonymously with *linear subspace* in this quiz.

- (1) Suppose V is a vector subspace of the vector space $C^\infty(\mathbb{R})$. We know that differentiation is linear. How is that information computationally useful?
- (A) It tells us that knowing how to differentiate all functions in any spanning set for V tells us how to differentiate any function in V (assuming we know how to express any function in V as a linear combination of the functions in the spanning set).
- (B) It tells us that knowing how to differentiate all functions in any linearly independent set in V tells us how to differentiate any function in V .

Your answer: _____

- (2) Suppose V is a vector subspace of the vector space $C^\infty(\mathbb{R})$. We know that differentiation is linear. How is that information computationally useful?
- (A) It tells us that knowing the antiderivatives of all functions in any spanning set for V tells us the antiderivative of every function in V (assuming we know how to express any function in V as a linear combination of the functions in the spanning set).
- (B) It tells us that knowing the antiderivatives of all functions in any linearly independent set in V tells us the antiderivative of every function in V .

Your answer: _____

We now consider two related vector spaces. $\mathbb{R}[x]$ is defined as the vector space of polynomials with real coefficients in the single variable x , with the usual addition and scalar multiplication. There is a natural injective homomorphism from $\mathbb{R}[x]$ to $C^\infty(\mathbb{R})$ that sends any polynomial to the same polynomial viewed as a function.

$\mathbb{R}(x)$ is defined as the vector space of all rational functions where the numerator and denominator are both polynomials with the denominator nonzero, up to equivalence (i.e., two rational functions $p_1(x)/q_1(x)$ and $p_2(x)/q_2(x)$ are equivalent if $p_1(x)q_2(x) = q_1(x)p_2(x)$). Addition and scalar multiplication are defined the usual way. Note that there is a natural injective homomorphism from $\mathbb{R}[x]$ to $\mathbb{R}(x)$ that sends any polynomial $p(x)$ to the rational function $p(x)/1$.

Also note that $\mathbb{R}(x)$ does not map to $C^\infty(\mathbb{R})$, for the reason that a rational function, viewed *qua* function, is not necessarily defined everywhere. Specifically, if written in simplified form, it is not defined at the set of roots of its denominator.

Note that both $\mathbb{R}[x]$ and $\mathbb{R}(x)$ are infinite-dimensional vector spaces, i.e., they do not have finite spanning sets.

- (3) Which of the following is *not* a basis for $\mathbb{R}[x]$? Please see Option (E) before answering.
- (A) $1, x, x^2, x^3, \dots$
- (B) $1, x, x(x-1), x(x-1)(x-2), x(x-1)(x-2)(x-3), \dots$
- (C) $1, x+1, x^2+x+1, x^3+x^2+x+1, \dots$
- (D) $1, x, x^2-x, x^3-x^2, x^4-x^3, \dots$
- (E) None of the above, i.e., each of them is a basis.

Your answer: _____

Let's now revisit the topic of *partial fractions* as a tool for integrating rational functions. The idea behind partial fractions is to consider an integration problem with respect to a variable x with integrand of the following form:

$$\frac{a_0 + a_1x + a_2x^2 + \dots + a_{n-1}x^{n-1}}{p(x)}$$

where p is a polynomial of degree n . For convenience, we may take p to be a monic polynomial, i.e., a polynomial with leading coefficient 1. For p fixed, the set of all rational functions of the form above forms a vector subspace of dimension n inside $\mathbb{R}(x)$. A natural choice of basis for this subspace is:

$$\frac{1}{p(x)}, \frac{x}{p(x)}, \dots, \frac{x^{n-1}}{p(x)}$$

The goal of partial fraction theory is to provide an *alternate basis* for this space of functions with the property that those basis elements are particularly easy to integrate (recurring to one of our earlier questions). Let's illustrate one special case: the case that p has n distinct real roots $\alpha_1, \alpha_2, \dots, \alpha_n$. The alternate basis in this case is:

$$\frac{1}{x - \alpha_1}, \frac{1}{x - \alpha_2}, \dots, \frac{1}{x - \alpha_n}$$

The explicit goal is to rewrite a partial fraction:

$$\frac{a_0 + a_1x + a_2x^2 + \dots + a_{n-1}x^{n-1}}{p(x)}$$

in terms of the basis above. If we denote the numerator as $r(x)$, we want to write:

$$\frac{r(x)}{p(x)} = \frac{c_1}{x - \alpha_1} + \frac{c_2}{x - \alpha_2} + \dots + \frac{c_n}{x - \alpha_n}$$

The explicit formula is:

$$c_i = \frac{r(\alpha_i)}{\prod_{j \neq i} (\alpha_i - \alpha_j)}$$

Once we rewrite the original rational function as a linear combination of the new basis vectors, we can integrate it easily because we know the antiderivatives of each of the basis vectors. The antiderivative is thus:

$$\left(\sum_{i=1}^n \frac{r(\alpha_i)}{\prod_{j \neq i} (\alpha_i - \alpha_j)} \ln |x - \alpha_i| \right) + C$$

where the obligatory $+C$ is put for the usual reasons.

Note that this process only handles rational functions that are proper fractions, i.e., the degree of the numerator must be less than that of the denominator.

We now consider cases where p is a polynomial of a different type.

- (4) Suppose p is a monic polynomial of degree n that is a product of pairwise distinct irreducible factors that are all either monic linear or monic quadratic. Call the roots for the linear polynomials $\alpha_1, \alpha_2, \dots, \alpha_s$ and call the monic quadratic factors q_1, q_2, \dots, q_t . Which of the following sets forms a basis for the vector space that we are interested in, namely all rational functions of the form $r(x)/p(x)$ where the degree of r is less than n ? Please see Option (E) before answering.
- (A) All rational functions of the form $1/(x - \alpha_i), 1 \leq i \leq s$ together with all rational functions of the form $1/q_j(x), 1 \leq j \leq t$
- (B) All rational functions of the form $1/(x - \alpha_i), 1 \leq i \leq s$ together with all rational functions of the form $q'_j(x)/q_j(x), 1 \leq j \leq t$
- (C) All rational functions of the form $1/q_j(x), 1 \leq j \leq t$ together with all rational functions of the form $q'_j(x)/q_j(x), 1 \leq j \leq t$
- (D) All rational functions of the form $1/(x - \alpha_i), 1 \leq i \leq s$ together with all rational functions of the form $1/q_j(x), 1 \leq j \leq t$ and all rational functions of the form $q'_j(x)/q_j(x), 1 \leq j \leq t$
- (E) None of the above

Your answer: _____

- (5) Suppose $p(x) = (x - \alpha)^n$. Which of the following sets forms a basis for the vector space that we are interested in, namely all rational functions of the form $r(x)/p(x)$ where the degree of r is less than n ? Please see Options (D) and (E) before answering.
- (A) The single function $1/(x - \alpha)$
- (B) The single function $1/(x - \alpha)^n$
- (C) All the functions $1/(x - \alpha), 1/(x - \alpha)^2, \dots, 1/(x - \alpha)^n$
- (D) Any of the above works

(E) None of the above works

Your answer: _____

We now recall our earlier discussion of the solution process for first-order linear differential equations. Consider a first-order linear differential equation with independent variable x and dependent variable y , with the equation having the form:

$$y' + p(x)y = q(x)$$

where $p, q \in C^\infty(\mathbb{R})$.

We solve this equation as follows. Let H be an antiderivative of p , so that $H'(x) = p(x)$.

$$\frac{d}{dx} (ye^{H(x)}) = q(x)e^{H(x)}$$

This gives:

$$ye^{H(x)} = \int q(x)e^{H(x)} dx$$

So:

$$y = e^{-H(x)} \int q(x)e^{H(x)} dx$$

The indefinite integration gives a $+C$, so overall, we get:

$$y = Ce^{-H(x)} + \text{particular solution}$$

It's now time to understand this in terms of linear algebra.

Define a linear transformation $L : C^\infty(\mathbb{R}) \rightarrow C^\infty(\mathbb{R})$ as:

$$f(x) \mapsto f'(x) + p(x)f(x)$$

(6) The kernel of L is one-dimensional. Which of the following functions spans the kernel?

- (A) $p(x)$
- (B) $q(x)$
- (C) $H(x)$
- (D) $e^{H(x)}$
- (E) $e^{-H(x)}$

Your answer: _____

(7) I would like to argue that L is *surjective* as a linear transformation from $C^\infty(\mathbb{R})$ to $C^\infty(\mathbb{R})$. Why is that true?

- (A) The kernel of L is zero-dimensional.
- (B) The image of L is zero-dimensional.
- (C) The kernel of L is one-dimensional.
- (D) The image of L is one-dimensional.
- (E) For any q , we have a formula above that describes a solution function that maps to q .

Your answer: _____

Let n be a nonnegative integer. Denote by P_n the vector space of all polynomials in one variable x that have degree $\leq n$. P_n is a subspace of $\mathbb{R}[x]$, which in turn can be viewed as a subspace of $C^\infty(\mathbb{R})$ through the natural injective map. For convenience and completeness, define P_{-1} to be the zero subspace.

Differentiation defines a linear transformation from $C^\infty(\mathbb{R})$ to itself.

(8) What are the kernel and image of the restriction of differentiation to P_n ? The result should be valid for all positive integers n .

- (A) The kernel and image are both P_n
- (B) The kernel is the zero subspace and the image is P_n
- (C) The kernel is P_n and the image is the zero subspace
- (D) The kernel is P_{n-1} and the image is P_0 (the subspace of constant functions)
- (E) The kernel is P_0 and the image is P_{n-1}

Your answer: _____

- (9) What are the kernel and image of the restriction of differentiation to all of $\mathbb{R}[x]$?

- (A) The kernel and image are both $\mathbb{R}[x]$
- (B) The kernel is the zero subspace and the image is $\mathbb{R}[x]$
- (C) The kernel is $\mathbb{R}[x]$ and the image is the zero subspace
- (D) The kernel is $\mathbb{R}[x]$ and the image is P_0 (the subspace of constant functions)
- (E) The kernel is P_0 and the image is $\mathbb{R}[x]$

Your answer: _____

- (10) We can use differentiation to define a linear transformation from $\mathbb{R}(x)$ to $\mathbb{R}(x)$, where we differentiate a rational function using the quotient rule for differentiation and the known rules for differentiating polynomials. What can we say about this linear transformation?

- (A) The differentiation linear transformation is bijective from $\mathbb{R}(x)$ to $\mathbb{R}(x)$, i.e., every rational function is the derivative of a unique rational function.
- (B) The differentiation linear transformation is injective but not surjective from $\mathbb{R}(x)$ to $\mathbb{R}(x)$, i.e., every rational function is the derivative of *at most one* rational function, but there do exist rational functions that are not expressible as the derivative of any rational function.
- (C) The differentiation linear transformation is surjective but not injective from $\mathbb{R}(x)$ to $\mathbb{R}(x)$, i.e., every rational function is the derivative of *at least one* rational function, but there do exist rational functions that occur as derivatives of more than one rational function.
- (D) The differentiation linear transformation is neither injective nor surjective from $\mathbb{R}(x)$ to $\mathbb{R}(x)$.

Your answer: _____

- (11) Denote by $\mathbb{R}[[x]]$ the vector space of all formal power series with real coefficients in one variable, i.e., series of the form:

$$\sum_{i=0}^{\infty} a_i x^i$$

Formal differentiation defines a linear transformation from $\mathbb{R}[[x]]$ to itself. What can we say about this linear transformation?

- (A) The formal differentiation linear transformation is bijective from $\mathbb{R}[[x]]$ to $\mathbb{R}[[x]]$.
- (B) The formal differentiation linear transformation is injective but not surjective from $\mathbb{R}[[x]]$ to $\mathbb{R}[[x]]$.
- (C) The formal differentiation linear transformation is surjective but not injective from $\mathbb{R}[[x]]$ to $\mathbb{R}[[x]]$.
- (D) The formal differentiation linear transformation is neither injective nor surjective from $\mathbb{R}[[x]]$ to $\mathbb{R}[[x]]$.

Your answer: _____

- (12) Consider the following two linear transformations $T_1, T_2 : \mathbb{R}[x] \rightarrow \mathbb{R}[x]$: T_1 is differentiation, and T_2 is multiplication by x . Which of the following is true?

- (A) Both T_1 and T_2 are injective, but neither is surjective.
- (B) Both T_1 and T_2 are surjective, but neither is injective.
- (C) T_1 is injective but not surjective. T_2 is surjective but not injective.
- (D) T_1 is surjective but not injective. T_2 is injective but not surjective.
- (E) Neither T_1 nor T_2 is injective. Neither T_1 nor T_2 is surjective.

Your answer: _____

(13) Consider the linear transformations T_1 and T_2 of the preceding question. What can we say regarding whether T_1 and T_2 commute?

(A) T_1 and T_2 commute.

(B) T_1 and T_2 do not commute.

Your answer: _____

**DIAGNOSTIC IN-CLASS QUIZ SOLUTIONS: DUE MONDAY NOVEMBER 25:
SUBSPACE, BASIS, AND DIMENSION**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

23 people took this 3-question quiz. The score distribution was as follows:

- Score of 0: 2 people
- Score of 1: 8 people
- Score of 2: 6 people
- Score of 3: 7 people

The mean score was 1.78.

The question-wise answers and performance review were as follows:

- (1) Option (A): 15 people
- (2) Option (E): 11 people
- (3) Option (C): 15 people

2. SOLUTIONS

PLEASE DO NOT DISCUSS ANY QUESTIONS.

This quiz covers material related to the **Linear dependence, bases and subspaces** notes corresponding to Sections 3.2 and 3.3 of the text.

Keep in mind the following facts. Suppose $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear transformation. Suppose A is the matrix for T , so that $T(\vec{x}) = A\vec{x}$ for all $\vec{x} \in \mathbb{R}^m$. Then, A is a $n \times m$ matrix. Further, the following are true:

- The dimension of the image of T equals the rank of A .
 - The dimension of the kernel of T , called the *nullity* of A , is m minus the rank of A .
- (1) *Do not discuss this!*: Suppose $T : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a linear transformation. What is the best we can say about the dimension of the image of T ?
 - (A) It is at least 0 and at most $\min\{m, n\}$. However, we cannot be more specific based on the given information.
 - (B) It is at least 0 and at most $\max\{m, n\}$. However, we cannot be more specific based on the given information.
 - (C) It is at least $\min\{m, n\}$ and at most $\max\{m, n\}$. However, we cannot be more specific based on the given information.
 - (D) It is at least $\min\{m, n\}$ and at most $m + n$. However, we cannot be more specific based on the given information.
 - (E) It is at least $\max\{m, n\}$ and at most $m + n$. However, we cannot be more specific based on the given information.

Answer: Option (A)

Explanation: The dimension of the image equals the rank of the matrix for T , which is a $n \times m$ matrix, hence is at most $\min\{m, n\}$. It is at least 0 for obvious reasons. It is easy to see that there exist examples for each possible dimension ranging from 0 to $\min\{m, n\}$.

Performance review: 15 out of 23 people got this. 5 chose (E), 2 chose (B), 1 chose (C).

Historical note (last time): 15 out of 26 got this. 5 chose (C), 3 each chose (B) and (D).

- (2) *Do not discuss this!*: Suppose $T_1, T_2 : \mathbb{R}^m \rightarrow \mathbb{R}^n$ are linear transformations. Suppose the images of T_1 and T_2 have dimensions d_1 and d_2 respectively. What can we say about the dimension of the image of $T_1 + T_2$? Assume that both m and n are larger than $d_1 + d_2$.

- (A) It is precisely $|d_2 - d_1|$.
- (B) It is precisely $\min\{d_1, d_2\}$.
- (C) It is precisely $\max\{d_1, d_2\}$.
- (D) It is precisely $d_1 + d_2$.
- (E) Based on the information, it could be any integer r with $|d_2 - d_1| \leq r \leq d_1 + d_2$.

Answer: Option (E)

Explanation: The image of $T_1 + T_2$ is contained in the sum of the images of T_1 and T_2 , hence its dimension is at most the dimension of the sum of the images of T_1 and T_2 . The dimension of the sum of subspaces is at most equal to the sum of the dimensions (because we can take the union of the spanning sets). Thus, the dimension of the image of $T_1 + T_2$ is at most equal to $d_1 + d_2$.

For the lower bound of $|d_2 - d_1|$, note that the dimension of the image of T_1 is at most the sum of the dimensions of the images of $T_1 + T_2$ and of T_2 and also that the dimension of the image of T_2 is at most the sum of the dimensions of the images of $T_1 + T_2$ and of T_1 . This gives lower bounds of $d_1 - d_2$ and $d_2 - d_1$ respectively on the dimension of the image of $T_1 + T_2$. The maximum of these is $|d_2 - d_1|$, which is the lower bound.

It is easy to construct examples of diagonal matrices with 0, 1 and -1 as the diagonal entries to realize any r with $|d_2 - d_1| \leq r \leq d_1 + d_2$.

Performance review: 11 out of 23 people got this. 7 chose (C), 3 chose (B), 2 chose (D).

Historical note (last time): 12 out of 26 got this. 8 chose (B), 2 each chose (A) and (C), 1 chose (D).

- (3) *Do not discuss this!* Suppose V_1 and V_2 are subspaces of \mathbb{R}^n . We define the sum $V_1 + V_2$ as the subset of \mathbb{R}^n comprising all vectors that can be expressed as a sum of a vector in V_1 and a vector in V_2 . Define $V_1 \cup V_2$ as the set-theoretic union of V_1 and V_2 , i.e., the set of all vectors that are either in V_1 or in V_2 . What can we say about these?
- (A) $V_1 \cup V_2 = V_1 + V_2$ and it is a subspace of \mathbb{R}^n .
 - (B) $V_1 \cup V_2$ is contained in $V_1 + V_2$ and both are subspaces of \mathbb{R}^n .
 - (C) $V_1 \cup V_2$ is contained in $V_1 + V_2$, and $V_1 + V_2$ is a subspace of \mathbb{R}^n . $V_1 \cup V_2$ is generally not a subspace of \mathbb{R}^n (though it might be in special cases).
 - (D) $V_1 \cup V_2$ contains $V_1 + V_2$, and both are subspaces of \mathbb{R}^n .
 - (E) $V_1 \cup V_2$ contains $V_1 + V_2$, and $V_1 \cup V_2$ is a subspace of \mathbb{R}^n . $V_1 + V_2$ is generally not a subspace of \mathbb{R}^n (though it might be in special cases).

Answer: Option (C)

Explanation: $V_1 + V_2$ is defined as the set of all vectors that can be expressed as the sum of a vector in V_1 and a vector in V_2 . In particular, it contains both V_1 and V_2 . The reason it contains V_1 is that any vector in V_1 can be written as itself plus the *zero vector* of V_2 . The reason it contains V_2 is that any vector in V_2 can be written as the *zero vector* of V_1 plus that vector itself.

Thus, $V_1 \cup V_2$ is contained in $V_1 + V_2$. They need not be equal. For instance, consider the case that V_1 is the span of \vec{e}_1 and V_2 is the span of \vec{e}_2 inside \mathbb{R}^n . The union of these is the set of vectors that are on either of the axes. It is a union of two perpendicular lines. The sum on the other hand is the plane spanned by the first two coordinates. The sum includes linear combinations where both spaces contribute nontrivially.

This also hints at why $V_1 \cup V_2$ does not need to be a subspace: it contains both subspaces, but not the vectors that are obtained by combining nonzero vectors from both subspaces. In fact, $V_1 \cup V_2$ is a subspace if and only if it equals $V_1 + V_2$, and this happens if and only if either $V_1 \subseteq V_2$ or $V_2 \subseteq V_1$.

Performance review: 15 out of 23 got this. 5 chose (D), 2 chose (B), 1 chose (E).

Historical note (last time): 8 out of 26 got this. 10 chose (B), 4 chose (E), 2 chose (D), 1 chose (A).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY NOVEMBER 27:
SIMILARITY OF LINEAR TRANSFORMATIONS**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

23 people took this 12-question quiz. The score distribution was as follows:

- Score of 4: 3 people
- Score of 5: 3 people
- Score of 6: 6 people
- Score of 7: 5 people
- Score of 8: 3 people
- Score of 9: 3 people

The mean score was about 6.48.

The question-wise answers and performance review are as follows:

- (1) Option (D): 21 people
- (2) Option (A): 4 people
- (3) Option (E): 3 people
- (4) Option (E): 17 people
- (5) Option (A): 19 people
- (6) Option (A): 22 people
- (7) Option (A): 11 people
- (8) Option (B): 9 people
- (9) Option (B): 5 people
- (10) Option (A): 6 people
- (11) Option (B): 18 people
- (12) Option (D): 14 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

This quiz corresponds to material discussed in the lecture notes titled **Coordinates**. It also corresponds to Section 3.4 of the text.

Recall that $n \times n$ matrices A and B are termed *similar* if there exists an invertible $n \times n$ matrix S such that $A = SBS^{-1}$. The relation of matrices being similar is an *equivalence relation* (please refer to the notes for an explanation of the terminology).

For these questions, assume $n > 1$, because a lot of phenomena are much simpler in the case $n = 1$ and you may be misled if you look only at that case. In other words, just because an equality is true for 1×1 matrices, do not assume it is always true. On the other hand, if you can find *counterexamples* to a statement for 1×1 matrices, you can probably use that to construct counterexamples for all sizes of matrices by using scalar matrices.

- (1) Which of the following can we say about two (possibly equal, possibly distinct) similar $n \times n$ matrices A and B ? Please see Options (D) and (E) before answering.
 - (A) A is invertible if and only if B is invertible.
 - (B) A is nilpotent if and only if B is nilpotent.
 - (C) A is idempotent if and only if B is idempotent.
 - (D) All of the above.

(E) None of the above.

Answer: Option (D)

Explanation: Suppose A is similar to B . Then, there exists an invertible matrix S such that $A = SBS^{-1}$, or equivalently, $B = S^{-1}AS$.

- Option (A): If A is invertible, then so is B , and $B^{-1} = S^{-1}A^{-1}S$. Conversely, if B is invertible, so is A , and $A^{-1} = SB^{-1}S^{-1}$.
- Option (B): We know that for any positive integer r , $A^r = SB^rS^{-1}$ and $B^r = S^{-1}A^rS$. If B is nilpotent, then there exists a positive integer r such that $B^r = 0$, so $A^r = SB^rS^{-1} = S(0)S^{-1} = 0$, so $A^r = 0$. Conversely, if A is nilpotent, then there exists a positive integer r such that $A^r = 0$, so $B^r = S^{-1}A^rS = S^{-1}(0)S = 0$.
- Option (C): We know that $SB^2S^{-1} = A^2$ and $S^{-1}A^2S = B^2$, so $A^2 = A$ if and only if $B^2 = B$.

Performance review: 21 out of 23 people got this. 1 each chose (A) and (E).

Historical note (last time): 17 out of 19 got this. 1 each chose (A) and (B).

(2) Which of the following can we say about two (possibly equal, possibly distinct) similar $n \times n$ matrices A and B ? Please see Options (D) and (E) before answering.

- (A) A is scalar if and only if B is scalar.
(B) A is diagonal if and only if B is diagonal.
(C) A is upper triangular if and only if B is upper triangular.
(D) All of the above.
(E) None of the above.

Answer: Option (A)

Explanation: If A is scalar, then it commutes with every matrix. In particular, A commutes with S , so $B = S^{-1}AS = AS^{-1}S = A$. Thus, $A = B$, and so B is also scalar. Similarly, if B is scalar, then it commutes with S , so $A = SBS^{-1} = BSS^{-1} = B$, so A is also scalar. In other words, A is scalar if and only if B is scalar, and in the event this happens, they are both equal.

The other options fail for reasons described below:

- Option (B): It is possible for A to be diagonal and B to not be diagonal. The idea is to use a matrix S that sends the standard basis vectors to vectors that are not standard basis vectors.
- Option (C): It is possible for A to be upper triangular and B to not be. For instance, consider:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$$

Performance review: 4 out of 23 people got this. 10 chose (D), 7 chose (E), 1 each chose (B) and (C).

Historical note (last time): 11 out of 19 got this. 4 chose (D), 2 each chose (B) and (E).

(3) Suppose A_1, A_2, B_1, B_2 are $n \times n$ matrices such that A_1 is similar to B_1 and A_2 is similar to B_2 . Which of the following is *definitely* true? Please see Options (D) and (E) before answering.

- (A) $A_1 + A_2$ is similar to $B_1 + B_2$.
(B) $A_1 - A_2$ is similar to $B_1 - B_2$.
(C) A_1A_2 is similar to B_1B_2 .
(D) All of the above.
(E) None of the above.

Answer: Option (E)

Explanation: The key problem in each case is that the matrix we use for the similarity between A_1 and B_1 is not the same as the matrix we use for the similarity between A_2 and B_2 . If the matrix were the same, then all the conclusions stated above would hold.

For instance, consider the case that:

$$A_1 = A_2 = B_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B_2 = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Note that A_1 and B_1 are similar on account of being equal. A_2 and B_2 are similar using the (self-inverse) similarity matrix:

$$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Consider the sums:

$$A_1 + A_2 = \begin{bmatrix} 2 & 0 \\ 0 & 0 \end{bmatrix}, B_1 + B_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

These are not similar, because the latter is the identity matrix and hence is not similar to anything else.

Consider the differences:

$$A_1 - A_2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B_1 - B_2 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

These are not similar, because the former is the zero matrix, which is not similar to any other matrix.

Finally, consider the products:

$$A_1 A_2 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B_1 B_2 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

The latter is the zero matrix, hence is not similar to any other matrix.

Performance review: 3 out of 23 people got this. 11 chose (C), 9 chose (D).

Historical note (last time): 12 out of 19 got this. 3 each chose (C) and (D), 1 chose (A).

- (4) Suppose A_1, A_2, B_1, B_2 are $n \times n$ matrices such that A_1 is similar to B_1 and A_2 is similar to B_2 . Which of the following is *definitely* true? Please see Options (D) and (E) before answering.
- (A) $A_1 + B_1$ is similar to $A_2 + B_2$.
 - (B) $A_1 - B_1$ is similar to $A_2 - B_2$.
 - (C) $A_1 B_1$ is similar to $A_2 B_2$.
 - (D) All of the above.
 - (E) None of the above.

Answer: Option (E)

Explanation: There isn't even an *a priori* reason why any of the options should be true, unlike for the previous question where at least *a priori* the options are plausible. For Options (A) and (C), the following 1×1 counterexample works: $A_1 = B_1 = [1]$, $A_2 = B_2 = [2]$. For Option (B), we cannot use 1×1 counterexamples, because in the 1×1 case, we'd have $A_1 - B_1 = A_2 - B_2 = [0]$. We can, however, use 2×2 counterexamples. Explicitly, consider the example:

$$A_1 = A_2 = B_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$$

Here, we have:

$$A_1 - B_1 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, A_2 - B_2 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

Clearly, $A_1 - B_1$ is *not* similar to $A_2 - B_2$.

Performance review: 17 out of 23 people got this. 3 each chose (C) and (D).

- (5) Suppose A and B are both $n \times n$ matrices (but they are not given to be similar). Which of the following holds?
- (A) A is similar to B if and only if $-A$ is similar to $-B$.
 - (B) If A is similar to B , then $-A$ is similar to $-B$. However, $-A$ being similar to $-B$ does not imply that A is similar to B .
 - (C) If $-A$ is similar to $-B$, then A is similar to B . However, A being similar to B does not imply that $-A$ is similar to $-B$.
 - (D) A being similar to B does not imply that $-A$ is similar to $-B$. Also, $-A$ being similar to $-B$ does not imply that A is similar to B .

Answer: Option (A)

Explanation: We have that for any invertible matrix S , $S(-B)S^{-1} = -(SBS^{-1})$. In other words, if $A = SBS^{-1}$, then $-A = S(-B)S^{-1}$. Conversely, if $-A = S(-B)S^{-1}$, then $A = SBS^{-1}$. Thus, A is similar to B if and only if $-A$ is similar to $-B$, and the matrix used for similarity is the same in both cases.

Note that invertibility of A or B is not necessary for this question, and the inclusion of the adjective “invertible” in the original print version of the quiz was based on an erroneous copy-paste. However, the question is correct even with the “invertible” assumption.

Performance review: 19 out of 23 people got this. 2 chose (D), 1 chose (B), 1 left the question blank.

- (6) Suppose A and B are both $n \times n$ matrices (but they are not given to be similar). Which of the following holds?

(A) A is similar to B if and only if $2A$ is similar to $2B$.

(B) If A is similar to B , then $2A$ is similar to $2B$. However, $2A$ being similar to $2B$ does not imply that A is similar to B .

(C) If $2A$ is similar to $2B$, then A is similar to B . However, A being similar to B does not imply that $2A$ is similar to $2B$.

(D) A being similar to B does not imply that $2A$ is similar to $2B$. Also, $2A$ being similar to $2B$ does not imply that A is similar to B .

Answer: Option (A)

Explanation: We have that for any invertible matrix S , $S(2B)S^{-1} = 2(SBS^{-1})$. In other words, if $A = SBS^{-1}$, then $2A = S(2B)S^{-1}$. Conversely, if $2A = S(2B)S^{-1}$, then $A = SBS^{-1}$. Thus, A is similar to B if and only if $2A$ is similar to $2B$, and the matrix used for similarity is the same in both cases.

Performance review: 22 out of 23 people got this. 1 chose (B).

- (7) Suppose A and B are both invertible $n \times n$ matrices (but they are not given to be similar). Which of the following holds?

(A) A is similar to B if and only if A^{-1} is similar to B^{-1} .

(B) If A is similar to B , then A^{-1} is similar to B^{-1} . However, A^{-1} being similar to B^{-1} does not imply that A is similar to B .

(C) If A^{-1} is similar to B^{-1} , then A is similar to B . However, A being similar to B does not imply that A^{-1} is similar to B^{-1} .

(D) A being similar to B does not imply that A^{-1} is similar to B^{-1} . Also, A^{-1} being similar to B^{-1} does not imply that A is similar to B .

Answer: Option (A)

Explanation: Note that once we show one direction, the other direction follows, because the inverse operation is self-inverse: the inverse of the inverse is the inverse. This automatically narrows the space of possibilities to two: Option (A) and Option (D). To demonstrate that the correct answer is Option (A), we will show the forward implication: if A is similar to B , then A^{-1} is similar to B^{-1} .

Suppose A is similar to B . Then, there exists an invertible $n \times n$ matrix S such that $A = SBS^{-1}$. Then, $A^{-1} = (SBS^{-1})^{-1} = (S^{-1})^{-1}B^{-1}S^{-1} = SB^{-1}S^{-1}$ (note that the order of multiplication reverses when we take the inverse). Thus, A^{-1} and B^{-1} are also similar.

Performance review: 11 out of 23 people got this. 5 each chose (B) and (D), 1 each chose (C) and (E).

Historical note (last time): 14 out of 19 got this. 2 chose (B), 1 each chose (C), (D), and (E).

- (8) Suppose A and B are both $n \times n$ matrices (but they are not given to be similar). Which of the following holds?

(A) A is similar to B if and only if A^2 is similar to B^2 .

(B) If A is similar to B , then A^2 is similar to B^2 . However, A^2 being similar to B^2 does not imply that A is similar to B .

(C) If A^2 is similar to B^2 , then A is similar to B . However, A being similar to B does not imply that A^2 is similar to B^2 .

- (D) A being similar to B does not imply that A^2 is similar to B^2 . Also, A^2 being similar to B^2 does not imply that A is similar to B .

Answer: Option (B)

Explanation: If A is similar to B , that means there exists a $n \times n$ invertible matrix S such that $A = SBS^{-1}$. Thus, $A^2 = (SBS^{-1})^2 = SB^2S^{-1}$, so that A^2 and B^2 are both similar.

However, A^2 being similar to B^2 does not imply that A is similar to B . For instance, consider:

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}$$

Note that both A^2 and B^2 are the zero matrix, so A^2 and B^2 are similar. However, A is not similar to B . In fact, B , being the zero matrix, is the only matrix in its similarity class, for obvious reasons.

Performance review: 9 out of 23 people got this. 10 chose (A), 4 chose (D).

Historical note (last time): 10 out of 19 got this. 5 chose (D), 2 chose (A), 1 each chose (C) and (E).

- (9) Suppose A and B are $n \times n$ matrices (but they are not given to be similar and they are not given to be invertible). We say that A and B are *quasi-similar* (not a standard term!) if there exist $n \times n$ matrices C and D such that $A = CD$ and $B = DC$. What can we say is the relation between being similar and being quasi-similar?

(A) A and B are similar if and only if they are quasi-similar.

(B) If A and B are similar, they are quasi-similar. However, the converse is not necessarily true: A and B may be quasi-similar but not similar.

(C) If A and B are quasi-similar, they are similar. However, the converse is not necessarily true: A and B may be similar but not quasi-similar.

(D) Neither implies the other. A and B may be similar but not quasi-similar. Also, A and B may be quasi-similar but not similar.

Answer: Option (B)

Explanation: If A and B are similar, there exists an invertible $n \times n$ matrix S such that $A = SBS^{-1}$. In that case, we can set $C = SB$ and $D = S^{-1}$ to obtain that $A = CD$ and $B = DC$.

The converse is not necessarily true. To establish a counter-example, it suffices to construct matrices C and D such that $CD = 0$ but DC is nonzero. If we label CD as A and DC as B , we have constructed quasi-similar matrices that are not similar. Here are the examples:

$$C = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}, D = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

The products are:

$$A = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

These are quasi-similar but not similar.

Performance review: 5 out of 23 people got this. 15 chose (D), 2 chose (C), 1 chose (E).

Historical note (last time): 10 out of 19 got this. 7 chose (D), 2 chose (C).

- (10) With the notion of quasi-similar as defined in the preceding question, what can we say about the relation between being similar and being quasi-similar for $n \times n$ matrices A and B that are both given to be *invertible*?

(A) A and B are similar if and only if they are quasi-similar.

(B) If A and B are similar, they are quasi-similar. However, the converse is not necessarily true: A and B may be quasi-similar but not similar.

(C) If A and B are quasi-similar, they are similar. However, the converse is not necessarily true: A and B may be similar but not quasi-similar.

(D) Neither implies the other. A and B may be similar but not quasi-similar. Also, A and B may be quasi-similar but not similar.

Answer: Option (A)

Explanation: We already proved that similar implies quasi-similar. We want to prove the reverse implication under the assumption that A and B are invertible. So, suppose $A = CD$ and $B = DC$ with A and B both invertible.

First, note that C is invertible. In fact, $C(DA^{-1})$ is the identity matrix.

Now, note that:

$$A = CD = CDCC^{-1} = C(DC)C^{-1} = CBC^{-1}$$

Thus, A and B are similar.

Performance review: 6 out of 23 people got this. 6 each chose (B) and (D), 5 chose (C).

Historical note (last time): 9 out of 19 got this. 7 chose (B), 2 chose (D), 1 chose (C).

- (11) Suppose A and B are two $n \times n$ matrices. Which of the following best describes the relation between similarity and having the same rank?

- (A) A and B are similar if and only if they have the same rank.
- (B) If A and B are similar, then they have the same rank. However, it is possible for A and B to have the same rank but not be similar.
- (C) If A and B have the same rank, then they are similar. However, it is possible for A and B to be similar but not have the same rank.
- (D) A and B may be similar but have different ranks. Also, A and B may have the same rank but not be similar.

Answer: Option (B)

Explanation: Note that similar matrices represent the same linear transformation in different coordinates. In particular, this means that geometrically, the kernel and image remain the same, but they get re-labeled. Thus, the matrices must have the same rank.

Explicitly, if $A = SBS^{-1}$, then the image of A is the image of SBS^{-1} . Start with \mathbb{R}^n . Its image under SBS^{-1} can be computed by taking successive images under the linear transformations corresponding to S^{-1} , then B , then S . The first transformation, given by S^{-1} , is bijective from \mathbb{R}^n to \mathbb{R}^n on account of S being invertible. We then do B on the image. Since the image of S^{-1} is all of \mathbb{R}^n , the image of BS^{-1} is the same as the image of B . Then again, S is bijective. Therefore it has zero kernel. Thus, its restriction to the image of BS^{-1} sends that subspace of \mathbb{R}^n to an equal-dimensional subspace of \mathbb{R}^n . The upshot is that the image of $SBS^{-1} = A$ has the same dimension as the image of B . Thus, A and B have the same rank.

However, it is possible for matrices having the same rank to not be similar. For instance, *any* two invertible $n \times n$ matrices have the same rank, namely n . However, they need not be similar. In fact, we can take two different scalar matrices with different scalar values, such as $[1]$ and $[2]$. Or, we could take these two matrices:

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Performance review: 18 out of 23 people got this. 5 chose (A).

Historical note (last time): 15 out of 19 got this. 3 chose (D), 1 chose (A).

- (12) Suppose A and B are two $n \times n$ matrices. Which of the following best describes the relation between quasi-similarity and having the same rank?

- (A) A and B are quasi-similar if and only if they have the same rank.
- (B) If A and B are quasi-similar, then they have the same rank. However, it is possible for A and B to have the same rank but not be quasi-similar.
- (C) If A and B have the same rank, then they are quasi-similar. However, it is possible for A and B to be quasi-similar but not have the same rank.
- (D) A and B may be quasi-similar but have different ranks. Also, A and B may have the same rank but not be quasi-similar.

Answer: Option (D)

Explanation: For an example of quasi-similar matrices that have different ranks, consider the example provided earlier of the zero matrix being quasi-similar to a nonzero matrix. Explicitly:

$$A = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

For an example of matrices that have the same rank that are not quasi-similar, consider:

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

Both A and B have rank one. However, they are not quasi-similar. This can be seen in either of two ways:

- Given two quasi-similar matrices, one is nilpotent if and only if the other is, and their nilpotencies differ by at most one. However, in the example above, A is idempotent and not nilpotent, while B is nilpotent. The reason is roughly that if $(CD)^r = 0$, then $(DC)^{r+1} = 0$, and conversely, if $(DC)^s = 0$, then $(CD)^{s+1} = 0$.
- Any two quasi-similar matrices have the same trace (as mentioned below). However, A has trace 1 while B has trace 0.

Performance review: 14 out of 23 people got this. 5 chose (B), 3 chose (A), 1 chose (C).

Historical note (last time): 7 out of 19 got this. 7 chose (B), 3 chose (A), 2 chose (C).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE MONDAY DECEMBER 2:
SIMILARITY OF LINEAR TRANSFORMATIONS (APPLIED)**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 16-question quiz. The score distribution was as follows:

- Score of 4: 3 people
- Score of 5: 4 people
- Score of 6: 3 people
- Score of 8: 4 people
- Score of 9: 4 people
- Score of 10: 4 people
- Score of 12: 3 people

The mean score was 7.76.

The question-wise answers and performance review were as follows:

- (1) Option (A): 15 people
- (2) Option (B): 16 people
- (3) Option (A): 15 people
- (4) Option (B): 20 people
- (5) Option (D): 10 people
- (6) Option (B): 12 people
- (7) Option (A): 3 people
- (8) Option (B): 19 people
- (9) Option (B): 15 people
- (10) Option (B): 6 people
- (11) Option (C): 10 people
- (12) Option (E): 8 people
- (13) Option (D): 6 people
- (14) Option (C): 11 people
- (15) Option (D): 16 people
- (16) Option (A): 12 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

This quiz corresponds to material discussed in the lecture notes titled *Coordinates*. It also corresponds to Section 3.4 of the text.

Recall that $n \times n$ matrices A and B are termed *similar* if there exists an invertible $n \times n$ matrix S such that $A = SBS^{-1}$. The relation of matrices being similar is an *equivalence relation*.

Recall that $n \times n$ matrices A and B are termed *quasi-similar* if there exist $n \times n$ matrices C and D such that $A = CD$ and $B = DC$. Recall that similar matrices are always quasi-similar, but quasi-similar matrices need not be similar. However, for *invertible* matrices, similarity and quasi-similarity are equivalent.

Also, note that if A and B are quasi-similar matrices, then A and B have the same trace. However, the converse is not true: it is possible to have two matrices A and B that have the same trace but are not quasi-similar.

For these questions, assume $n > 1$, because a lot of phenomena are much simpler in the case $n = 1$ and you may be misled if you look only at that case.

Note also that the trace of a square matrix is defined as the sum of its diagonal entries. The *determinant* of a 2×2 matrix, denoted \det , is defined as follows:

$$\det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = ad - bc$$

The following are some important facts about the determinant:

- The determinant of a 2×2 diagonal matrix is the product of the diagonal entries.
- The determinant of a 2×2 matrix is zero if and only if the matrix is non-invertible.
- The determinant of the product of two 2×2 matrices is the product of the determinants.
- The determinant of the inverse of an invertible 2×2 matrix is the reciprocal of the determinant.
- If A and B are similar 2×2 matrices, they have the same determinant.
- If A and B are quasi-similar 2×2 matrices, they have the same determinant.
- If the determinant of A is positive, then the linear transformation given by A is an orientation-preserving linear automorphism of \mathbb{R}^2 .
- If the determinant of A is negative, then the linear transformation given by A is an orientation-reversing linear automorphism of \mathbb{R}^2 .

(1) Suppose A and B are both $n \times n$ matrices (but they are not given to be similar). Denote by I_n the $n \times n$ identity matrix. Which of the following holds?

- (A) A is similar to B if and only if $A - I_n$ is similar to $B - I_n$.
- (B) If A is similar to B , then $A - I_n$ is similar to $B - I_n$. However, $A - I_n$ being similar to $B - I_n$ does not imply that A is similar to B .
- (C) If $A - I_n$ is similar to $B - I_n$, then A is similar to B . However, A being similar to B does not imply that $A - I_n$ is similar to $B - I_n$.
- (D) A being similar to B does not imply that $A - I_n$ is similar to $B - I_n$. Also, $A - I_n$ being similar to $B - I_n$ does not imply that A is similar to B .

Answer: Option (A)

Explanation: We have that for any invertible matrix S , $S(B - I_n)S^{-1} = (SBS^{-1} - SI_nS^{-1} = SBS^{-1} - I_n$. In other words, if $A = SBS^{-1}$, then $A - I_n = S(B - I_n)S^{-1}$. Conversely, if $A - I_n = S(B - I_n)S^{-1}$, then $A = SBS^{-1}$. Thus, A is similar to B if and only if $A - I_n$ is similar to $B - I_n$, and the matrix used for similarity is the same in both cases.

Performance review: 15 out of 25 got this. 7 chose (D), 2 chose (C), 1 chose (B).

Suppose f is a polynomial of degree r in one variable with real coefficients. For a $n \times n$ matrix X , we denote by $f(X)$ we mean the matrix we get by applying the polynomial to f , where constant terms are interpreted as scalar matrices. For instance, if $f(x) = x^2 + 3x + 5$, then $f(X) = X^2 + 3X + 5I_n$.

(2) Suppose A and B are both $n \times n$ matrices (but they are not given to be similar). Suppose f is a polynomial of degree r in one variable, where $r \geq 2$. Which of the following holds?

- (A) A is similar to B if and only if $f(A)$ is similar to $f(B)$.
- (B) If A is similar to B , then $f(A)$ is similar to $f(B)$. However, $f(A)$ being similar to $f(B)$ does not imply that A is similar to B .
- (C) If $f(A)$ is similar to $f(B)$, then A is similar to B . However, A being similar to B does not imply that $f(A)$ is similar to $f(B)$.
- (D) A being similar to B does not imply that $f(A)$ is similar to $f(B)$. Also, $f(A)$ being similar to $f(B)$ does not imply that A is similar to B .

Answer: Option (B)

Explanation: For the forward direction, note that $A = SBS^{-1}$ implies that $f(A) = Sf(B)S^{-1}$. For the breakdown of the reverse direction, see the explanation for Q8 of the November 27 quiz. This covers the $f(x) = x^2$ case. Similar examples can be constructed for other polynomials.

Performance review: 16 out of 25 got this. 4 chose (D), 3 chose (A), 2 chose (C).

- (3) Suppose A and B are both $n \times n$ matrices (but they are not given to be similar). Suppose f is a polynomial of degree r in one variable, where $r = 1$. Which of the following holds?
- (A) A is similar to B if and only if $f(A)$ is similar to $f(B)$.
 - (B) If A is similar to B , then $f(A)$ is similar to $f(B)$. However, $f(A)$ being similar to $f(B)$ does not imply that A is similar to B .
 - (C) If $f(A)$ is similar to $f(B)$, then A is similar to B . However, A being similar to B does not imply that $f(A)$ is similar to $f(B)$.
 - (D) A being similar to B does not imply that $f(A)$ is similar to $f(B)$. Also, $f(A)$ being similar to $f(B)$ does not imply that A is similar to B .

Answer: Option (A)

Explanation: Degree one polynomials differ from higher degree polynomials in that we can recover the original matrix from a degree one polynomial of it. Thus, we can deduce that $A = SBS^{-1}$ if and only if $f(A) = Sf(B)S^{-1}$. Special cases of this were covered in Question 1 of this quiz and Questions 5 and 6 of the November 27 quiz.

Performance review: 15 out of 25 got this. 4 chose (D), 3 each chose (B) and (C).

Suppose p and q are real numbers (possibly equal, possibly distinct). The diagonal matrices:

$$A = \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix}$$

and

$$B = \begin{bmatrix} q & 0 \\ 0 & p \end{bmatrix}$$

are similar. Explicitly, the two matrices are similar under the change-of-basis transformation that interchanges the coordinates, i.e., if we set:

$$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

then:

$$S = S^{-1}$$

and we have:

$$B = S^{-1}AS$$

Moreover, the only diagonal matrices similar to A are A and B (in the special case that $p = q$, we get $A = B$ is a scalar matrix, so A is the only diagonal matrix similar to A).

- (4) What is the necessary and sufficient condition on p and q such that the matrix $A = \begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix}$ is similar to $-A$?
- (A) $p = q$
 - (B) $p = -q$
 - (C) $p = 1/q$
 - (D) $p = -1/q$
 - (E) $p + q = 1$

Answer: Option (B)

Explanation: We have:

$$-A = \begin{bmatrix} -p & 0 \\ 0 & -q \end{bmatrix}$$

Now, for A to be similar to $-A$, we have one of these two conditions:

$$-A = A \text{ or } -A = \begin{bmatrix} q & 0 \\ 0 & p \end{bmatrix}$$

The first case ($-A = A$) gives us that $p = q = 0$, so that A is the zero matrix. The second case gives us that $-p = q$ and $-q = p$. Both of these are equivalent to $p = -q$. We now notice that the first case $p = q = 0$ is subsumed within the second case, so that $p = -q$ describes the necessary and sufficient condition.

Performance review: 20 out of 25 got this. 2 each chose (A) and (D), 1 chose (C).

- (5) Which of the following is a necessary and sufficient condition on p and q so that the matrix $A =$

$\begin{bmatrix} p & 0 \\ 0 & q \end{bmatrix}$ is invertible and similar to $-A^{-1}$?

(A) $p = q$

(B) $p = -q$

(C) $p = 1/q$

(D) $p = -1/q$

(E) $p + q = 1$

Answer: Option (D)

Explanation: We have:

$$-A^{-1} = \begin{bmatrix} -1/p & 0 \\ 0 & -1/q \end{bmatrix}$$

For this to be similar to A , we must have:

$$-A^{-1} = A \text{ or } -A^{-1} = \begin{bmatrix} q & 0 \\ 0 & p \end{bmatrix}$$

The first case gives $-1/p = p$ and $-1/q = q$, solving to $p^2 = -1$ and $q^2 = -1$, which is not possible. Thus, the first case is ruled out.

This brings us to the second case. In this case, $-1/p = q$ and $-1/q = p$. Both of these are equivalent to $p = -1/q$. This is the correct answer.

Performance review: 10 out of 25 got this. 9 chose (B), 5 chose (C), 1 chose (E).

Consider the matrix:

$$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

used above. We have $S = S^{-1}$. For a general matrix:

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

we have:

$$S^{-1}AS = \begin{bmatrix} d & c \\ b & a \end{bmatrix}$$

In other words, it swaps the rows *and* swaps the columns. This observation may be useful for some of the following questions.

- (6) For an angle θ with $-\pi \leq \theta \leq \pi$, the rotation matrix for θ is given as:

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

Note that $R(-\pi) = R(\pi)$, but other than that equality, all the $R(\theta)$ s are distinct.

- Which of these describes the relation between the rotation matrices for different values of θ ?
- (A) All the rotation matrices $R(\theta)$, $-\pi < \theta \leq \pi$, are similar to each other.
 - (B) The rotation matrix $R(\theta)$ is similar to itself and to the rotation matrix $R(-\theta)$. However, it is not in general similar to any other rotation matrix.
 - (C) No two different rotation matrices are similar.
 - (D) The rotation matrix $R(\theta)$ is similar to itself and to the rotation matrix $R(\pi - \theta)$ (or $R(-\pi - \theta)$, depending on which of the two angles lies within the specified range). However, it is not in general similar to any other rotation matrix.

Answer: Option (B)

Explanation: The matrix we can use for the similarity transformation is the following self-inverse matrix:

$$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$$

This matrix essentially performs a reflection about the x -axis, and the net effect is to negate the angle of rotation. Explicitly:

$$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}$$

Performance review: 12 out of 25 got this. 9 chose (D), 4 chose (A).

Historical note (last time): 9 out of 19 got this. 4 each chose (A) and (D), 2 chose (C).

- (7) Consider the linear automorphisms of \mathbb{R}^2 that are given as *reflections* about lines in \mathbb{R}^2 through the origin. (Note that we need the line of reflection to pass through the origin for the automorphism to be *linear* rather than merely being *affine linear*). Which of these describes the relation between reflection matrices for different possible lines of reflection through the origin?
- (A) All the reflection matrices are similar to each other.
 - (B) No two reflection matrices for different lines of reflection are similar.
 - (C) The reflection matrices for two different lines of reflection are similar if and only if the lines of reflection are perpendicular.
 - (D) The reflection matrices for two different lines of reflection are similar if and only if the lines of reflection make an angle that is a rational multiple of π .

Answer: Option (A)

Explanation: The rotation matrix that rotates one line to the other can be used as the matrix for similarity.

Performance review: 3 out of 25 got this. 12 chose (C), 6 chose (B), 4 chose (D).

Historical note (last time): 2 out of 19 people got this. 12 chose (C), 4 chose (D), 1 chose (B).

- (8) Suppose m and n are positive integers with $m < n$. Denote by P_m the “orthogonal projection onto the first m coordinates” linear transformation from \mathbb{R}^n to \mathbb{R}^n , defined as follows. This takes as input a n -dimensional vector, sends each of the first m coordinates to itself, and sends the remaining coordinates to zero. What is the trace of the matrix of P_m ?
- (A) 1
 - (B) m
 - (C) n
 - (D) $n - m$
 - (E) $m - n$

Answer: Option (B)

Explanation: The matrix is diagonal with the first m diagonal entries equal to 1 and the remaining $n - m$ diagonal entries equal to 0. The trace is thus m .

Performance review: 19 out of 25 got this. 3 each chose (D) and (E).

Historical note (last time): 9 out of 19 got this. 5 chose (D), 4 chose (E), 1 chose (C).

- (9) It is a fact that if A, B are $n \times n$ matrices that describe orthogonal projections onto (possibly different) m -dimensional subspaces of \mathbb{R}^n , then A and B are similar. What can we say must be the trace of an orthogonal projection onto any m -dimensional subspace of \mathbb{R}^n ?

- (A) 1
- (B) m
- (C) n
- (D) $n - m$
- (E) $m - n$

Answer: Option (B)

Explanation: This follows from the preceding question.

Performance review: 15 out of 25 got this. 4 each chose (A) and (D), 2 chose (C).

Historical note (last time): 4 out of 19 got this. 6 chose (C), 5 chose (D), 4 chose (E).

- (10) Suppose A , B and C are $n \times n$ matrices. Which of the following matrices is *not* guaranteed (based on the given information) to have the same trace as the product ABC ? Please see (and read carefully) Options (D) and (E) before answering.

- (A) BCA
- (B) CBA
- (C) CAB

(D) None the above, i.e., they are all guaranteed to have the same trace as ABC .

(E) All of the above, i.e., none of them is guaranteed to have the same trace as ABC .

Answer: Option (B)

Explanation: First, let's note that the other two options don't work:

- Option (A): $ABC = A(BC)$ whereas $BCA = (BC)A$, so ABC and BCA are quasi-similar matrices. Therefore, they have the same trace.
- Option (C): $ABC = (AB)C$ and $CAB = C(AB)$, so ABC and CAB are quasi-similar matrices. Therefore, they have the same trace.

Consider the example:

$$A = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}, B = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

The product ABC is the zero matrix, because AB is the zero matrix. Thus, ABC has trace zero. The product CBA , on the other hand, is the matrix:

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$$

Performance review: 6 out of 25 got this. 11 chose (D), 7 chose (E), 1 chose (C).

Historical note (last time): 4 out of 19 got this. 8 chose (D), 5 chose (E), 1 each chose (A) and (C).

- (11) Which of the following gives a pair of matrices A and B that have the same trace as each other *and* the same determinant as each other, but that are *not* similar to each other?

- (A) $A = \begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}, B = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}$
- (B) $A = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, B = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$
- (C) $A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$
- (D) $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$
- (E) $A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}$

Answer: Option (C)

Explanation: The identity matrix is not similar to any non-identity matrix, because it is scalar, so conjugating it by anything leaves it as it is.

All other pairs of matrices are in fact similar:

- Option (A): We can use the coordinate interchange matrix $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.

- Option (B): We can use the conjugating matrix $\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$.
- Option (D) We can use the coordinate interchange matrix $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.
- Option (E): Use $\begin{bmatrix} 1 & 0 \\ 0 & 2 \end{bmatrix}$.

Performance review: 10 out of 25 got this. 7 chose (B), 5 chose (E), 2 chose (D), 1 left the question blank.

Historical note (last time): 6 out of 19 got this. 5 chose (B), 3 each chose (D) and (E), and 2 chose (A).

- (12) Suppose A and B are 2×2 matrices. Which of the following correctly describes the relation between $\det A$, $\det B$, and $\det(A + B)$? Please see Option (E) before answering.
- (A) $\det(A + B) = \det A + \det B$
 (B) $\det(A + B) \leq \det A + \det B$, but equality need not necessarily hold.
 (C) $\det(A + B) \geq \det A + \det B$, but equality need not necessarily hold.
 (D) $|\det(A + B)| \leq |\det A| + |\det B|$, but equality need not necessarily hold.
 (E) None of the above.

Answer: Option (E)

Explanation: The determinant does not interact in any meaningful manner with addition. In fact, for any (possibly equal, possibly distinct) real numbers p , q , and r , we can construct matrices A and B such that $\det A = p$, $\det B = q$, $\det(A + B) = r$. We can in fact just use 2×2 matrices to achieve this. The general description is tricky, so let's just give counter-examples for each part:

- Options (A), (B) and (D): Consider $A = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$, $B = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$. Note that $\det A = \det B = 0$, but $\det(A + B) = 1$.
- Option (C): Consider $A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $B = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$. Then, $\det A = \det B = 1$, but $\det(A + B) = 0$.

Performance review: 8 out of 25 got this. 7 chose (C), 5 chose (B), 3 chose (D), 2 chose (A).

Historical note (last time): 4 out of 19 got this. 5 chose (B), 4 chose (C), 3 each chose (A) and (D).

Let n be a natural number greater than 1. Suppose $f : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ is a function satisfying $f(0) = 0$. Let T_f denote the linear transformation from \mathbb{R}^n to \mathbb{R}^n satisfying the following for all $i \in \{1, 2, \dots, n\}$:

$$T_f(\vec{e}_i) = \begin{cases} \vec{e}_{f(i)}, & f(i) \neq 0 \\ \vec{0}, & f(i) = 0 \end{cases}$$

Let M_f denote the matrix for the linear transformation T_f . M_f can be described explicitly as follows: the i^{th} column of M_f is $\vec{0}$ if $f(i) = 0$ and is $\vec{e}_{f(i)}$ if $f(i) \neq 0$.

Note that if $f, g : \{0, 1, 2, \dots, n\} \rightarrow \{0, 1, 2, \dots, n\}$ are functions with $f(0) = g(0) = 0$, then $M_{f \circ g} = M_f M_g$ and $T_{f \circ g} = T_f \circ T_g$.

For the following questions, the discussion prior to Question 3 might be helpful. Note, however, that while that discussion gives one possible candidate for the matrix S of the similarity transformation, it is not the only possible candidate. For some but not all of the following questions, in the case that two matrices are similar, the matrix S described there works. In the case that they are not similar, the lack of similarity can be inferred from the traces not being equal, or from the determinants not being equal.

- (13) $n = 2$ for this question. For the following three functions f , g , and h , consider the corresponding matrices M_f, M_g, M_h . Either two of them are similar and the third is not similar to either (in which

case, select the matrix that is not similar to the other two), or all three are similar (if so, select Option (D)), or no two are similar (if so, select Option (E)).

- (A) $f(0) = 0, f(1) = 1, f(2) = 0$
- (B) $g(0) = 0, g(1) = 0, g(2) = 2$
- (C) $h(0) = 0, h(1) = 1, h(2) = 1$
- (D) All the above give similar matrices.
- (E) No two of the corresponding matrices are similar.

Answer: Option (D)

Explanation: All three matrices are similar. Here is what the matrices look like:

- Option (A): $M_f = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$
- Option (B): $M_g = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$
- Option (C): $M_h = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$

Clearly, they all have trace 1, rank 1, and determinant 0. Thus, we cannot *prima facie* rule out the possibility of their being similar. But to actually confirm that they are similar, it would help to demonstrate a matrix that accomplishes the similarity transformation.

The similarity of Options (A) and (B) is relatively easy: the two options are related in that they have interchanged roles of the first and second vector relative to each other. Thus, the following matrix works well for accomplishing similarity:

$$S = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

Explicitly, $M_f = SM_gS^{-1}$. Note that $S^{-1} = S$.

The similarity between M_f and M_h is trickier. If we set:

$$X = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$$

Then:

$$X^{-1} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}$$

Then, $M_f = XM_hX^{-1}$.

Now that we have shown the similarity of M_f with M_g and the similarity of M_g with M_h , the fact that similarity is an equivalence relation tells us that all three matrices are similar.

Performance review: 6 out of 25 got this. 10 chose (C), 4 chose (E), 3 chose (B), 2 chose (A).

Historical note (last time): 4 out of 19 got this. 10 chose (C), 3 chose (A), 1 each chose (A) and (E).

- (14) $n = 2$ for this question. For the following three functions f , g , and h , consider the corresponding matrices M_f, M_g, M_h . Either two of them are similar and the third is not similar to either (in which case, select the matrix that is not similar to the other two), or all three are similar (if so, select Option (D)), or no two are similar (if so, select Option (E)).

- (A) $f(0) = 0, f(1) = 0, f(2) = 1$
- (B) $g(0) = 0, g(1) = 2, g(2) = 0$
- (C) $h(0) = 0, h(1) = 2, h(2) = 1$
- (D) All the above give similar matrices.
- (E) No two of the corresponding matrices are similar.

Answer: Option (C)

Explanation: The three matrices are:

- Option (A): $M_f = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$

- Option (B): $M_g = \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix}$
- Option (C): $M_h = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$

M_f and M_g are similar. In fact, we can use the matrix M_h itself:

$$M_f = M_h M_g M_h^{-1}$$

M_h is not similar to either M_f or M_g . We can see this, for instance, by noting that M_h has full rank, but neither M_f nor M_g do. Alternatively, note that M_h has determinant -1 , unlike both M_f and M_g , that have determinant 0 .

Performance review: 11 out of 25 got this. 5 chose (D), 4 chose (E), 3 chose (B), 2 chose (A).

Historical note (last time): 8 out of 19 got this. 4 each chose (B) and (D), 3 chose (E).

- (15) $n = 3$ for this question. For the following three functions f , g , and h , consider the corresponding matrices M_f, M_g, M_h . Either two of them are similar and the third is not similar to either (in which case, select the matrix that is not similar to the other two), or all three are similar (if so, select Option (D)), or no two are similar (if so, select Option (E)).
- (A) $f(0) = 0, f(1) = 2, f(2) = 1, f(3) = 3$
 (B) $g(0) = 0, g(1) = 1, g(2) = 3, g(3) = 2$
 (C) $h(0) = 0, h(1) = 3, h(2) = 2, h(3) = 1$
 (D) All the above give similar matrices.
 (E) No two of the corresponding matrices are similar.

Answer: Option (D)

Explanation: In each case, the matrix interchanges two coordinates and leaves the third coordinate as is. Which two coordinates get interchanged just depends on how we label, and therefore, all the matrices are similar. The explicit descriptions are below:

- Option (A): $M_f = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$
- Option (B): $M_g = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$
- Option (C): $M_h = \begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$

Performance review: 16 out of 25 got this. 4 chose (E), 3 chose (C), 2 chose (B).

Historical note (last time): 8 out of 19 got this. 4 chose (B), 3 chose (E), 2 each chose (A) and (C).

- (16) $n = 3$ for this question. For the following three functions f , g , and h , consider the corresponding matrices M_f, M_g, M_h . Either two of them are similar and the third is not similar to either (in which case, select the matrix that is not similar to the other two), or all three are similar (if so, select Option (D)), or no two are similar (if so, select Option (E)).
- (A) $f(0) = 0, f(1) = 1, f(2) = 2, f(3) = 3$
 (B) $g(0) = 0, g(1) = 2, g(2) = 3, g(3) = 1$
 (C) $h(0) = 0, h(1) = 3, h(2) = 1, h(3) = 2$
 (D) All the above give similar matrices.
 (E) No two of the corresponding matrices are similar.

Answer: Option (A)

Explanation: Let's first write out the matrices:

- Option (A): $M_f = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

- Option (B): $M_g = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$
- Option (C): $M_h = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$

M_f is the identity matrix, and therefore, cannot be similar to anything else. M_g and M_h both describe matrices that cycle the three coordinates, albeit in opposite orders. A re-labeling can change M_g to M_h . Explicitly, if:

$$S = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Then $SM_gS^{-1} = M_h$.

Performance review: 12 out of 25 got this. 6 chose (D), 4 chose (E), 2 chose (C), 1 chose (B).

Historical note (last time): 7 out of 19 got this. 3 each chose (B), (C), (D), and (E).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY DECEMBER 4:
ORDINARY LEAST SQUARES REGRESSION**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 8-question quiz. The score distribution was as follows:

- Score of 3: 1 person
- Score of 4: 2 people
- Score of 5: 1 person
- score of 6: 8 people
- Score of 7: 8 people
- Score of 8: 5 people

The mean score was 6.4.

The question-wise answers and performance review were as follows:

- (1) Option (C): 23 people
- (2) Option (C): 24 people
- (3) Option (C): 22 people
- (4) Option (D): 20 people
- (5) Option (A): 13 people
- (6) Option (D): 20 people
- (7) Option (A): 18 people
- (8) Option (A): 20 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

- (1) Assume no measurement error. Consider the situation where we have a function f of the form $f(x) = a_0 + a_1x$ with unknown values of the parameters a_0 and a_1 . We collect n distinct input-output pairs, i.e., we collect n distinct inputs and compute the outputs for them. The coefficient matrix for the system is a $n \times 2$ matrix (the rows correspond to the input values, and the columns correspond to the unknown parameters). What is the rank of this matrix?
 - (A) It is always 2
 - (B) It is always n
 - (C) It is always $\min\{2, n\}$
 - (D) It is always $\max\{2, n\}$

Answer: Option (C)

Explanation: Note that the rank is at most $\min\{2, n\}$ because the rank is at most the minimum of the number of rows and the number of columns. To see that the rank is exactly that value, consider the cases $n = 1$ and $n = 2$. In the case $n = 1$, we have a single nonzero row $[1 \ x_1]$ so the rank is exactly one. In the case $n = 2$, we have a coefficient matrix as follows:

$$\begin{bmatrix} 1 & x_1 \\ 1 & x_2 \end{bmatrix}$$

Note that, since $x_1 \neq x_2$, the second row is *not* a scalar multiple of the first. Thus, the matrix has rank two. The case $n \geq 2$ follows, because we have already achieved the maximum possible rank of two using the first two rows.

Performance review: 23 out of 25 people got this. 2 chose (A).

- (2) Assume no measurement error. Consider the situation where we have a function f of the form $f(x) = a_0 + a_1x + a_2x^2 + \cdots + a_mx^m$ with unknown values of the parameters a_0, a_1, \dots, a_m . We collect n distinct input-output pairs, i.e., we collect n distinct inputs and compute the outputs for them. The coefficient matrix for the system is a $n \times (m + 1)$ matrix (the rows correspond to the input values, and the columns correspond to the unknown parameters). What is the rank of this matrix?
- (A) It is always $m + 1$
 (B) It is always n
 (C) It is always $\min\{m + 1, n\}$
 (D) It is always $\max\{m + 1, n\}$

Answer: Option (C)

Explanation: The rank can be at *most* $\min\{m + 1, n\}$ because it is at most the minimum of the number of rows and the number of columns. See Section 3.1 of the lecture notes on “hypothesis testing, rank, and overdetermination” but note that the roles of m and n are interchanged there relative to this question.

Performance review: 24 out of 25 people got this. 1 chose (D).

- (3) Assume no measurement error. Consider the situation where we have a function f of the form $f(x, y) = a_0 + a_1x + a_2y$ with unknown values of the parameters $a_0, a_1,$ and a_2 . We collect n distinct input-output pairs, i.e., we collect n distinct inputs (here an input specification involves specifying both the x -value and the y -value) and compute the outputs for them. The coefficient matrix for the system is a $n \times 3$ matrix (the rows correspond to the input values, and the columns correspond to the unknown parameters). What is the rank of this matrix?

- (A) It is always $\min\{3, n\}$
 (B) It is always $\max\{3, n\}$
 (C) For $n = 1$, it is 1. For $n \geq 2$, it is 2 if the input points are all collinear in the xy -plane. Otherwise, it is 3.
 (D) For $n = 1$, it is 1. For $n \geq 2$, it is 3 if the input points are all collinear in the xy -plane. Otherwise, it is 2.

Answer: Option (C)

Explanation: For $n \geq 2$, the first two rows are not scalar multiples of each other. Explicitly, the first two rows look like:

$$\begin{bmatrix} 1 & x_1 & y_1 \\ 1 & x_2 & y_2 \end{bmatrix}$$

In the case that all the points are collinear, later rows can be obtained as linear combinations of the first two rows. Equivalently, knowing the value of f at (x_1, y_1) and at (x_2, y_2) allows us to predict the value of f at all other points on the line joining these in \mathbb{R}^2 , and therefore other points on the line do not reveal new information. Thus, the rank of the coefficient matrix is 2.

See the discussion in Section 2.1 of the lecture notes on hypothesis testing, rank and overdetermination. Also see the answer to Question 1 of the Friday October 11 take-home quiz on linear systems, and the answer to Question 5 of the Monday October 7 take-home class quiz on linear functions and equation-solving (part 2).

Performance review: 22 out of 25 people got this. 2 chose (A), 1 chose (D).

- (4) Which of the following is closest to correct in the setting where we use a linear system to find the parameters using input-output pairs given a functional form that is linear in the parameters? Assume for simplicity that we are dealing with a functional form $y = f(x)$ with one input and one output, but possibly multiple parameters in the general description.
- (A) The solutions to the linear system that we set up correspond to possibilities for the inputs to the function, and geometrically correspond to choices of points x for the graph $y = f(x)$.
 (B) The solutions to the linear system that we set up correspond to possibilities for the inputs to the function, and geometrically correspond to different possible choices for the line or curve that is the graph $y = f(x)$.

- (C) The solutions to the linear system that we set up correspond to possibilities for the parameters, and geometrically correspond to choices of points x for the graph $y = f(x)$.
- (D) The solutions to the linear system that we set up correspond to possibilities for the parameters, and geometrically correspond to different possible choices for the line or curve that is the graph $y = f(x)$.

Answer: Option (D)

Explanation: We are trying to find the line or curve. That's the whole goal of regression. And the way we do this is by choosing a general functional form and then using regression to estimate the parameters in that functional form.

Performance review: 20 out of 25 people got this. 5 chose (C).

- (5) Continuing with the notation and setup of the preceding question, consider the coefficient matrix of the linear system. This matrix defines a linear transformation from the vector space of possible parameter values to the vector space of the outputs of the function. What is the image of this linear transformation?
 - (A) The image is the set of possible output values for which the linear system is consistent, i.e., we can find *at least one* function f of the required functional form that fits all the input-output pairs with *no measurement error*.
 - (B) The image is the set of possible output values for which the linear system has *at most one solution*, i.e., the set of output values for which we can find *at most one* function f of the required functional form that fits all the input-output pairs with *no measurement error*.

Answer: Option (A)

Explanation: This follows from the definition of image.

Performance review: 13 out of 25 people got this. 12 chose (B).

- (6) Consider the case of polynomial regression for a polynomial function of one variable, allowing for measurement error. We believe that a function has the form of a polynomial. We can tentatively choose a degree m for the polynomial we are trying to fit, and a value n for the number of distinct inputs for which we compute the corresponding outputs to obtain input-output pairs (i.e., data points). We will get a $n \times (m + 1)$ coefficient matrix. Which of the following correctly describes what we should try for?
 - (A) We should choose n and $m + 1$ to be exactly equal, so that we get a unique polynomial.
 - (B) We should choose n to be greater than $m + 1$, so that the system is guaranteed to be consistent and we can find the polynomial.
 - (C) We should choose n to be less than $m + 1$, so that the system is guaranteed to be consistent and we can find the polynomial.
 - (D) We should choose n to be greater than $m + 1$, so that the system is *not* guaranteed to be consistent, but we do have a unique solution after we project the output vector to a vector for which the system is consistent.
 - (E) We should choose n to be less than $m + 1$, so that the system is *not* guaranteed to be consistent, but we do have a unique solution after we project the output vector to a vector for which the system is consistent.

Answer: Option (D)

Explanation: If we chose $n > m + 1$, the coefficient matrix has full column rank $m + 1$, and therefore the linear transformation is injective, i.e., we have at most one solution. However, it does not have full row rank, so the linear transformation is not surjective, i.e., we do not necessarily have a solution for every output vector.

Both aspects are features for us. The existence of at most one solution means that we can find the parameters uniquely after we project to the closest vector in the image. The fact that the system does not have full row rank means that there is a potential for inconsistency, and in the presence of measurement error, the system will probably be inconsistent with the measured output vector. This is good because the potential for inconsistency allows us to test the validity of the model better. Also, in the case of measurement error, the more the number of inputs that we use, the better our estimate of the function is likely to be.

Performance review: 20 out of 25 people got this. 3 chose (E), 1 each chose (B) and (C).

- (7) Consider the general situation of linear regression. Denote by X the coefficient matrix for the linear system (also called the design matrix). Denote by $\vec{\beta}$ the parameter vector that we are trying to solve for. Denote by \vec{y} an observed output vector. The idea in ordinary least squares regression is to choose a suitable vector $\vec{\epsilon}$ such that the linear system $X\vec{\beta} = \vec{y} - \vec{\epsilon}$ can be solved for $\vec{\beta}$. Among the many possibilities that we can choose for $\vec{\epsilon}$, what criterion do we use to select the appropriate choice? Recall that the *length* of a vector is the square root of the sum of squares of its coordinates.
- (A) We choose $\vec{\epsilon}$ to have the minimum length possible subject to the constraint that $X\vec{\beta} = \vec{y} - \vec{\epsilon}$ has a solution.
- (B) We choose $\vec{\epsilon}$ such that the system $X\vec{\beta} = \vec{y} - \vec{\epsilon}$ can be solved and such that the solution vector $\vec{\beta}$ has the minimum possible length (among all such choices of $\vec{\epsilon}$).
- (C) We choose $\vec{\epsilon}$ such that the system $X\vec{\beta} = \vec{y} - \vec{\epsilon}$ can be solved and such that the difference vector $\vec{y} - \vec{\epsilon}$ has the minimum possible length (among all such choices of $\vec{\epsilon}$).

Answer: Option (A)

Explanation: We want to deviate as little as possible from the measured output. This is the idea behind using the orthogonal projection.

Performance review: 18 out of 25 people got this. 6 chose (C), 1 chose (B).

- (8) We have data for the logarithm of annual per capita GDP for a country for the last 100 years. We want to see if this fits a polynomial model. The idea is to try to first fit a polynomial of degree 0 (i.e., per capita GDP remains constant), then fit a polynomial of degree ≤ 1 (i.e., per capita GDP grows or decays exponentially), then fit a polynomial of degree ≤ 2 (i.e., per capita GDP grows or decays as the exponential of a quadratic function), and so on.

What happens to the length of the error vector $\vec{\epsilon}$ as we increase the degree of the polynomial that we are trying to fit?

- (A) The error vector $\vec{\epsilon}$ keeps getting smaller and smaller in length, with a near-certainty that it keeps *strictly* decreasing in length at each step, until the error vector becomes $\vec{0}$ (which we expect will happen when we get to the stage of trying to fit the function using a polynomial of degree 99).
- (B) The error vector $\vec{\epsilon}$ keeps getting larger and larger in length, with a near-certainty that it keeps *strictly* increasing in length at each step, until the error vector becomes \vec{y} (which we expect will happen when we get to the stage of trying to fit the function using a polynomial of degree 99).

Answer: Option (A)

Explanation: The space that we are projecting on keeps getting bigger and bigger, so the distance from the vector \vec{y} to the space keeps getting smaller and smaller. Note that it's highly unlikely that there would be *no* improvement possible by introducing a new parameter, so it's likely that the improvement would be strict at every stage until the time that we reach a polynomial of degree 99, where we can obtain a perfect fit because we have only 100 data points.

Performance review: 20 out of 25 people got this. 5 chose (B).

**TAKE-HOME CLASS QUIZ SOLUTIONS: DUE WEDNESDAY DECEMBER 4:
MATRIX TRANSPOSE: PRELIMINARIES**

MATH 196, SECTION 57 (VIPUL NAIK)

1. PERFORMANCE REVIEW

25 people took this 7-question quiz. The score distribution was as follows:

- Score of 4: 4 people
- Score of 5: 5 people
- Score of 6: 4 people
- Score of 7: 12 people

The question-wise answers and performance review were as follows:

- (1) Option (A): 25 people (everybody)
- (2) Option (D): 25 people (everybody)
- (3) Option (A): 23 people
- (4) Option (C): 22 people
- (5) Option (A): 18 people
- (6) Option (D): 18 people
- (7) Option (A): 18 people

2. SOLUTIONS

PLEASE FEEL FREE TO DISCUSS ALL QUESTIONS.

The following questions are related to material from parts of Chapter 5 that we are glossing over. You do not need to read that chapter, because we are using a very limited part of it in a very limited fashion and we've included all relevant definitions in the quiz. However, if you want to understand some of the constructs in more detail, please do read the chapter.

Note: Due to limited class time, I'm making this a take-home class quiz, but in an ideal world, this would have been a diagnostic in-class quiz.

For a $n \times m$ matrix A , denote by A^T (spoken as *A-transposed* and called the *transpose of A*) the $m \times n$ matrix whose $(ij)^{th}$ entry is defined as the $(ji)^{th}$ entry of A . In other words, the roles of rows and columns are interchanged when we transition from A to A^T . The T should not be interpreted as an exponent letter. Note that whereas A describes a linear transformation from \mathbb{R}^m to \mathbb{R}^n , A^T describes a linear transformation from \mathbb{R}^n to \mathbb{R}^m . Note, however, that although the domain and co-domain for A and A^T are interchanged with each other, A and A^T are not in general inverses of each other.

- (1) Suppose A is a $n \times m$ matrix and A^T is the transpose of A . Under what conditions does the sum $A + A^T$ make sense (i.e., exist as a matrix)?
 - (A) $A + A^T$ makes sense if and only if $m = n$.
 - (B) $A + A^T$ makes sense if and only if $m < n$.
 - (C) $A + A^T$ makes sense if and only if $m > n$.
 - (D) $A + A^T$ makes sense regardless of whether $m = n$, $m < n$, or $m > n$.

Answer: Option (A)

Explanation: A is a $n \times m$ matrix whereas A^T is a $m \times n$ matrix. We know that $A + A^T$ makes sense if and only if A and A^T have the same number of rows and also have the same number of columns. Both of these conditions are equivalent to requiring that $m = n$.

Performance review: All 25 got this correct.

- (2) Suppose A is a $n \times m$ matrix and A^T is the transpose of A . Under what conditions does the product AA^T make sense (i.e., exist as a matrix)?
- (A) AA^T makes sense if and only if $m = n$.
 (B) AA^T makes sense if and only if $m < n$.
 (C) AA^T makes sense if and only if $m > n$.
 (D) AA^T makes sense regardless of whether $m = n$, $m < n$, or $m > n$.

Answer: Option (D)

Explanation: A is a $n \times m$ matrix and A^T is a $m \times n$ matrix. The product AA^T makes sense because the number of columns in A equals the number of rows in A^T , so the product AA^T is a $n \times n$ matrix.

Performance review: All 25 got this correct.

- (3) Suppose A is a $n \times m$ matrix and A^T is the transpose of A . Under what conditions do both AA^T and $A^T A$ exist *and* have the same number of rows as each other and the same number of columns as each other (note that they still need not be equal)?
- (A) This happens if and only if $m = n$.
 (B) This happens if and only if $m < n$.
 (C) This happens if and only if $m > n$.
 (D) This happens always, regardless of whether $m = n$, $m < n$, or $m > n$.

Answer: Option (A)

Explanation: The product AA^T is a $n \times n$ matrix (see the explanation for Question 2). For similar reasons, the product $A^T A$ is a $m \times m$ matrix. These have the same size if and only if $m = n$.

Performance review: 23 out of 25 got this. 2 chose (D).

- (4) Suppose A is a $n \times n$ matrix such that $A^T = A^{-1}$. We describe this condition by saying that A is an *orthogonal* $n \times n$ matrix. Which of the following is a correct characterization of a matrix being orthogonal? Please see Option (C) before answering, and select the option that best reflects your view.
- (A) Every row vector of A is a unit vector, and any two distinct rows of A are orthogonal.
 (B) Every column vector of A is a unit vector, and any two distinct columns of A are orthogonal.
 (C) Both of the above work, i.e., they are equivalent to each other and to the condition that $A^T = A^{-1}$.

Answer: Option (C)

Explanation: The condition that $A^T = A^{-1}$ can be interpreted in two equivalent ways: $AA^T = I_n$ and $A^T A = I_n$.

With the $AA^T = I_n$ interpretation, we see that the dot product of each row of A with the corresponding column of A^T is 1, and the dot product of each row of A with a different column of A^T is 0. Since the “corresponding column of A^T ” agrees with the original row of A , we obtain that the dot product of each row of A with itself is 1 (i.e., each row of A is a unit vector) and the dot product of any two distinct rows of A is 0, i.e., any two distinct rows of A are orthogonal.

With the $A^T A = I_n$ interpretation, we obtain the analogous result for columns, because we are now dealing with dot products between rows of A^T and columns of A .

Performance review: 22 out of 25 got this. 2 chose (B), 1 chose (A).

A square matrix A is termed *symmetric* if $A = A^T$ and *skew-symmetric* if $A = -A^T$.

The following facts are true and can be easily verified:

- Suppose A and B are matrices such that $A + B$ makes sense. Then, $(A + B)^T = A^T + B^T$.
- Suppose A and B are matrices such that AB makes sense. Then, $(AB)^T = B^T A^T$. Note that the order of multiplication flips over. The rule is similar to the rule for inverses, even though the transpose is *not* the same as the inverse.
- For any matrix A , $(A^T)^T = A$.

- (5) Suppose A is a matrix. What can we say that the nature of the matrices $A + A^T$ and AA^T ?
- (A) $A + A^T$ is symmetric if it makes sense. AA^T is symmetric if it makes sense.

- (B) $A + A^T$ is symmetric if it makes sense. AA^T is skew-symmetric if it makes sense.
 (C) $A + A^T$ is skew-symmetric if it makes sense. AA^T is symmetric if it makes sense.
 (D) $A + A^T$ is skew-symmetric if it makes sense. AA^T is skew-symmetric if it makes sense.

Answer: Option (A)

Explanation: For the sum: $(A + A^T)^T = A^T + (A^T)^T = A^T + A = A + A^T$. We use that $(A^T)^T = A$ while simplifying.

For the product: $(AA^T)^T = (A^T)^T A^T$ (we use the rule for transpose of a product). This simplifies to AA^T using the fact that $(A^T)^T = A$.

Performance review: 18 out of 25 got this. 5 chose (B), 2 chose (C).

- (6) Suppose n is a positive integer. Consider the vector space $\mathbb{R}^{n \times n}$ of $n \times n$ matrices. The subset comprising symmetric matrices is a linear subspace and the subset comprising skew-symmetric matrices is also a linear subspace. The subset comprising diagonal matrices is also a linear subspace. Which of the following best describes the containment relation between the subspaces of diagonal matrices, symmetric matrices, and skew-symmetric matrices?

- (A) The subspace comprising all diagonal matrices is contained in the subspace comprising all skew-symmetric matrices, which in turn is contained in the subspace comprising all symmetric matrices.
 (B) The subspace comprising all diagonal matrices is contained both in the subspace comprising all skew-symmetric matrices and in the subspace comprising all symmetric matrices. However, neither of the two latter subspaces is contained in the other.
 (C) The subspace comprising all diagonal matrices is contained in the subspace comprising all skew-symmetric matrices, but neither of these subspaces is contained in the subspace comprising all symmetric matrices.
 (D) The subspace comprising all diagonal matrices is contained in the subspace comprising all symmetric matrices, but neither of these subspaces is contained in the subspace comprising all skew-symmetric matrices.
 (E) None of the three subspaces is fully contained in any of the others.

Answer: Option (D)

Explanation: The subspace of diagonal matrices is contained in the subspace of skew-symmetric matrices, because any diagonal matrix is symmetric. This can easily be seen by looking at such a matrix. In the 2×2 case, a diagonal matrix is a matrix of the form:

$$\begin{bmatrix} a_{11} & 0 \\ 0 & a_{22} \end{bmatrix}$$

Clearly, such a matrix equals its own transpose.

On the other hand, the *only* symmetric matrix that is also skew-symmetric is the zero matrix. Explicitly, if A is a matrix satisfying both the conditions $A = A^T$ and $A + A^T = 0$, then we get $2A = 0$, hence $A = 0$. In particular, this also means that the only *diagonal* matrix that is skew-symmetric is the zero matrix. Thus, the space of diagonal matrices is not contained in the space of skew-symmetric matrices.

Performance review: 18 out of 25 got this. 2 each chose (B), (C), and (E), 1 chose (A).

- (7) Suppose n is a positive integer. Consider the vector space $\mathbb{R}^{n \times n}$ of $n \times n$ matrices. This vector space has dimension $n \times n = n^2$. What are the respective dimensions of the subspaces comprising symmetric and skew-symmetric matrices? *Hint:* Try the case $n = 1$ and then the case $n = 2$. In both cases, try to write down an explicit basis for each of the subspaces. You might want to revisit the preceding question in light of your improved understanding after solving this question.

- (A) The subspace comprising all symmetric matrices has dimension $n(n + 1)/2$ and the subspace comprising all skew-symmetric matrices has dimension $n(n - 1)/2$.
 (B) The subspace comprising all symmetric matrices has dimension $n(n - 1)/2$ and the subspace comprising all skew-symmetric matrices has dimension $n(n + 1)/2$.
 (C) Both subspaces have dimension $n^2/2$.

Answer: Option (A)

Explanation: For a symmetric matrix A , we can freely choose all the diagonal entries (this gives n dimensions) and for each pair $i \neq j$ of distinct elements of $\{1, 2, \dots, n\}$, we can choose the entry $a_{ij} = a_{ji}$. There are $(n^2 - n)/2$ dimensions arising from the latter choices. The total number of choices we have is therefore $n + (n^2 - n)/2 = n(n + 1)/2$.

In the skew-symmetric case, the diagonal entries all need to be zero. For each pair $i \neq j$ of distinct elements of $\{1, 2, \dots, n\}$, we can choose the entry $a_{ij} = -a_{ji}$. There are $(n^2 - n)/2 = n(n - 1)/2$ such choices.

In the 1×1 case, all matrices are symmetric and only the zero matrix is skew-symmetric. This agrees with our counts: the dimension of the subspace of symmetric matrices is $n(n + 1)/2 = 1(1 + 1)/2 = 1$ (as it should be, because that's the dimension of the whole space) and the dimension of the subspace of skew-symmetric matrices is $n(n - 1)/2 = 1(1 - 1)/2 = 0$ (as it should be, because that's the dimension of the zero subspace).

In the 2×2 case, the following matrices form a basis for the subspace of symmetric matrices (the basis has size $2(2 + 1)/2 = 3$):

$$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

The subspace of skew-symmetric matrices has dimension $2(2 - 1)/2 = 1$ and has a basis given by the following matrix:

$$\begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

In the 3×3 case, the following matrices form a basis for the subspace of symmetric matrices (there are $3(3 + 1)/2 = 6$ of them):

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

The following matrices form a basis for the subspace of skew-symmetric matrices (there are $3(3 - 1)/2 = 3$ of them):

$$\begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}$$

Performance review: 18 out of 25 got this. 5 chose (B), 2 chose (C).