

The following content is provided under a Creative Commons license. Your support will help MIT OpenCourseWare continue to offer high-quality educational resources for free. To make a donation or view additional materials from hundreds of MIT courses, visit MIT OpenCourseWare at ocw.mit.edu.

PROFESSOR:

Hope everyone had a great Veteran's Day break yesterday, spending it, of course, as I check my test audience, spending it mostly doing your [INAUDIBLE] research or PSETs. But at least one person got to watch TV. So at least one person got to have a real break, and that's something truly amazing and special.

So now we're going to talk about SVMs. They're pretty much the hardest thing in 6.034. However, in recent years a few shortcuts have popped up that will sometimes allow you to solve the question, depending on what they're asking for, without solving some vast ugly set of equations with a vast ugly number of unknowns. So I'm going to show that to you guys.

I'm also going to try to explain all the alphabet soup that's in [INAUDIBLE] and what all the letters stand for because it took me a few times going through SVMs. It took me a few times going through SVMs to actually find out for sure what all those letters stood for. And if you guys figure it out first try, that's going to be great, and you guys will be just fine.

So let's take a look at the problem that's perhaps most optimized for using some of the shortcuts to solving it and not putting up all the equations. Then I will-- not because I'm sadistic, but because I'm being nice, I will force you with me to solve some of the things they didn't ask for us to solve so that you can see that we can't get away with everything without doing some of the harder stuff.

And of course, definitely ask questions as always, but this time even more so. You guys, well, if you were looking around, you saw nobody in this entire lecture hall raised their hand that they are already set and ready and know SVMs. So if you have a question, maybe everybody else does.

So let's go. We'll start right here. As always, pretend that I can draw, and that therefore all the pluses and minuses are only on integer coordinates. So we are asked in this problem to circle the support vectors. Draw the edges of the street and then the dotted line in the middle that separates them, the separator, as a dashed line. And then to give w and b .

So what are w and b ? Well, there's a few important equations in SVMs that we really hope-- and I'm going to tell you we're lucky in this because we don't have to-- but we really hope that we don't have to use because they provide a huge number of variables.

So one of those crucial equations is that for a plus support vector, w vector dot x plus, the plus support vector, plus b equals 1. w dot x minus plus b equals minus 1. And w dot that dotted line-- I don't know, we'll call it dot dot dot-- plus b equals 0.

So what does this mean? There are a lot of vectors. Well, I mean, we're usually in two-dimensional space, so we can basically just say that there's two components of this w vector, w_1 and w_2 . And they're just two coefficients in a linear equation. So for instance, what we're interested in finding, this dot dot dot line, we'll just call that x , so with nothing on it. Actually, maybe that'll be easier.

This is equivalent to saying w_1x_1 plus w_2x_2 plus b equals 0, where x_1 is this, and x_2 is this. We would possibly call them x and y . So one way to think about it is w_1 , we'll call it a , ax plus-- call w_2 b -- by. Oh, don't call it b . Well, ax plus cy plus b equals-- I'll put this all in parentheses-- this is basically an equation like this. Or y equals negative a over c x minus b over c .

It's basically y equals mx plus b . Does everyone see that? This thing that we're looking for, this w dot x plus b equals 0, is the equation of a line in Cartesian coordinates. It just looks uglier.

Normally, when we're doing all this solving for w and b , we would have to put in tons of equations, plug in all of the support vectors in there. And we'd have to use these little devils called alphas. Alphas essentially-- if it wasn't clear in the lecture, which it usually isn't completely clear to everyone, wasn't clear to me completely-- alphas, the way I like to think about them, the alphas in this problem is they are the weight of how significant any particular point on the graph is towards creating the boundary.

The higher the alpha is, the more that that point narrows in the boundary. The lower the alpha is, the less that that point narrows in the boundary, the wider the road can be. And if that point doesn't do anything, if that point is irrelevant and could be removed and it wouldn't affect the boundary, the alpha is? Everyone?

AUDIENCE: Zero.

PROFESSOR: Zero. Well, that was one person, but you can suffice for everyone. The alpha is 0. And that

means if it's not a support vector, if it's not one of the vectors on the boundary lines, it will always have an alpha of 0 because it doesn't affect. So keeping that in mind, there's a few fun and important equations about alphas that we'll need if we're solving many equations for many unknowns, which hopefully we won't have to do.

The sum over the positive alphas equals the sum over the alphas-- the negative points. And this is true over all the points. But since all of the alphas are 0, except for the support vectors, it also means the alphas of the positive support vectors are equal to the alphas of the negative support vectors.

Additionally, our old buddy, the w vector, is equal to the sum over all i that are plus vectors of $w_i \alpha_i$ minus m over j minus vectors of $w_j \alpha_j$. Now, all of these equations can be used in a bloody mess to figure out the answer to what we're trying to find, which is circles-- well, actually, they can't be used as circle support vectors and draw the dotted line.

But once we do that, all these equations can be used in a bloody mess to give us the next thing that we want, which is w and b . So fortunately, there's another way to get w and b . If you guys really want, at the end of the hour we can also try to derive w and b using the many equations in many unknowns, but it's a bit painful. We'll try to do it the cool way.

So let's start off. This is the one we're looking at. We need to find where the support vectors are. So the first thing we need to do is simply eye it. Fortunately, on the test, there will always be ones that you can eye if you're supposed to circle the support vectors.

There's obviously some number of pluses and some number of minuses. I say obviously, but maybe not. But hopefully obviously, and we'll find out because I'm going to call on random people. So give me a positive support vector.

AUDIENCE: Um, going to the one that looked like [INAUDIBLE]?

PROFESSOR: Which plus sign, [INAUDIBLE]?

AUDIENCE: One on the right.

PROFESSOR: The one all the way on the right. Yeah, that plus sign is a positive support vector. That's good. All right? Excellent. Now, give me a negative support vector. That one? No?

AUDIENCE: Yeah, sorry.

PROFESSOR: Ah, no problem. Give me a negative support vector.

AUDIENCE: I should definitely ask you, what's a support vector?

[LAUGHTER]

PROFESSOR: That is a good question. The question is, what is a support vector? How many other people will admit to having this question? See? You're not alone. OK. Before I go on, I'm going to assume-- you guys make sure I'm correct-- Monday was, just being sure so I can tailor based on this. Monday was the support vector machine lecture. But it was also very difficult to follow. That's what I usually expect.

So what is a support vector? Well, all these pluses and minuses, if we were me, and if, I guess-- yeah, if we were me and if I was describing this problem, the one that we work out in class, I would call them points because they're on the graph. They're points. They're data points. But however, in more difficult versions of this problem that have n dimensions, where n is some ridiculous number of dimensions that you're never going to graph.

Like say some of the research I'm doing now, I could use support vector machines on some of these articles that I'm reading about cyber events to try to figure out if there's a real event or if it's just someone complaining about how we're really vulnerable or something like that and no event actually happened.

So the reason why they call these guys vectors is when you're not able to graph them on a Cartesian plane, there's still this long vector of many different dimensions. Right now, though, these points represent the vectors. This is very simple. It's easier to view them this way. But for instance, that plus at negative 1, 2 represents the fact that there is a vector going in the direction of negative 1, 2 with a magnitude such that it reaches negative 1, 2.

So all these points are just a point representation of a vector. You probably, in any class that worked with vectors, saw this, saw vectors being represented as points. Question?

AUDIENCE: Always from the respect to the origin?

PROFESSOR: Yes. The question is always with respect to the origin. The answer is canonically, when vectors are represented as points, yes, it's always with respect to the origin.

So that's the basic idea is that all these points are vectors. So what are support vectors? Well,

you could call them support points for this case. But the reason we call them support vectors is again, in the generalized case that you might be doing in the real world with real AI, you're going to have a giant vector. And it's not just going to be points on a graph. Well, usually.

So the support vectors, the support points, we found one of them correctly. It's this guy. They're going to be the ones that again, they don't have an alpha of zero. They're the ones that bind in the, as Petra calls it, the road, the boundary lines. They're going to be on the edge of plus. Whichever direction we draw it, this plus is the edge of the plus region. If we made this the edge of the plus region and everything on this side is plus and everything on this side is minus, we'd be screwed because there's two pluses on the other side of that.

Generally, when trying to find a support vector, you do something a little bit similar to my crazy method of doing nearest neighbors, and try to find a plus-minus pair that's close to each other. Sometimes though, it's not just two points because sometimes if you try to draw the simple-minded thing, which is the perpendicular bisector of the two points, you get screwed because there's another point in your way.

So now that I've given away a clue, let's go-- and hopefully that made sense to you guys. The support vectors are the ones on the edges that are just barely a plus for sure, or just barely a minus for sure. Let's go back. Can you give me a negative support vector?

AUDIENCE: The top one?

PROFESSOR: Hmm?

AUDIENCE: The top negative point?

PROFESSOR: The one on the top left? Yes. And does anyone think that there's a third support vector? Well, let's simple-mindedly try the thing that-- remember, support vectors always attempt to have the widest possible space between the pluses and minuses that they can. So let's simple-mindedly try to do the perpendicular bisector and see if screws us over.

So when we simple-mindedly do the perpendicular bisector, it goes through here like this. And it's just fine. So these are the only two support vectors. And there's our divider line.

So we're on the home stretch. But we have to find w and b . In olden days, we would find w and b by plugging in w dot the plus support vector, plus b equals 1. Oh, that's very crucial. These w dot x plus x minus are only true equaling 1 or negative 1? Or only true for support vectors?

It's always true that $w \cdot \text{any positive point} + b$ will be some positive number. But it won't always be 1. In fact, it will always be greater than 1 up over here. It will always be less than -1 down over there.

In olden days, we would plug in -1, 2 into this equation. We would plug in 3, -2 into this equation. We'd plug in α plus equals α minus in its sums. And since there's only 1 plus 1 minus, we'd know they were equal. And then we'd fidget around with this w equation.

However, there is a better way to do it. And so let's use this cheap strategy to solve this version of the SVM. Here's how. First, and I know I didn't draw these completely straight. Sorry. But can anyone, by looking at-- this is three, -2. 2 And this is -1, 2. Can anyone tell me what the equation-- you can do y equals mx plus b . Can anyone tell me what the equation of the dotted line is supposed to be if I was good at drawing?

AUDIENCE: [SEVERAL ANSWERS]

PROFESSOR: People say y equals x minus 1. And I say yes, y equals x minus 1. So therefore, the pluses would be y is greater than or equal to x minus 1 indeed. So we've already seen that $w \cdot x$ plus b somehow can be converted into this form. Right? So therefore, if we have y equals x minus 1, then we know that we have $w \cdot x$ plus b equals 0. Let's do that here. So we know that $w_1 x_1$. We can even call it x and y . I think it'll be fine. No one will come after us. $w_2 y$ plus b equals 0. But we also know that y equals x minus 1, which means that if y equals x minus 1, then according to this thing we have over here, then negative w_1 over w_2 equals--

So we know that negative w_1 over w_2 -- and we have $-b$ over w_2 . So y equals x minus 1. And if we solve this equation to make it look like this, we would have y equals negative w_1 over w_2 Minus b over w_2 . So we know that in some way, shape, or form-- we know that then therefore, w_1 over w_2 is some scalar multiple of minus 1. And we know that b over w_2 is, in fact, some scalar multiple of positive 1.

Scalar multiple, what's a scalar multiple? Well, why is it a scalar multiple? Why isn't it just going to be negative 1 or positive 1? Just because in this equation, we can multiply the entire equation by any number and it will still have the same boundary line. You guys see that? Oh, there's an x here. If we multiplied everything, since it's all divided by w_2 . If we double w_2 , but also doubled b and w_1 , it would be the exact same equation. Do you guys agree?

So there's, in fact, infinitely many possible equations. You say, well, great, Mark. You've

figured out what form it is. So you figured out that w_1 over w_2 equals some scalar multiple of negative 1. So it's negative 1 times-- what's everyone's favorite letter?

AUDIENCE: k.

PROFESSOR: k. Negative 1 times k. And we figured out that b over w_2 is-- I guess we can just do negative -- is positive k. But what's k? How are we going to figure it out? Well, it's a good question. And I will tell you how.

I will assert the following fact as true without proof. Then I will not prove it. 1 over the magnitude of w , which is this vector here with w_1 and w_2 , equals this where this is that line that I just drew, the line from here to this point. 1 over the magnitude of w equals this. Therefore, since 1 over the magnitude of w equals this, and this equals, I believe, $2\sqrt{2}$, because we're going over to, down to, Pythagorean Theorem, $2\sqrt{2}$. So therefore, flip everything over. Magnitude of w equals 1 over $2\sqrt{2}$. So therefore, magnitude of w equals $\sqrt{2}$ over 4 .

But why are we OK? Well, how do we calculate the magnitude of w ? Do people know, in general, magnitudes of vectors? Generally, for these vectors, we do it by the square root of the sum of the components squared. So the square root of w_1 squared plus w_2 squared equals $\sqrt{2}$ over 4 . But that's not all. That's not all, we say, because we know from this over here that the ratio of w_1 and w_2 is--

AUDIENCE: [SEVERAL ANSWERS]

PROFESSOR: Yeah, the ratio of w_1 and w_2 is going to be-- actually, sorry. I shouldn't put a k here. I realize I probably have been confusing you guys a lot. w_1 over w_2 is just -1 . b over w_2 is just 1 . That's just a fact. There's no k . The k is to determine what w_1 and w_2 are. So w_1 equals $-k$. And w_2 equals positive k . And b equals also positive k .

By the way, here's a question for you. Could I have put the negative sign on w_2 and b instead of on w_1 ? So many people said yes. That's a very smart answer. Actually, no, because of the fact that the pluses are on the negative x -axis. It's just a little trick I picked up. When one of them is negative and the other one isn't, follow the pluses.

So we know that w_1 is $-k$, w_2 is positive k , and b is positive k . w_1 over w_2 is -1 . b over w_2 is positive 1 . So what do we know about the ratio of w_1 and w_2 ? It's equal to -1 . And that means

that when we square it, w_1 squared equals w_2 squared. So therefore, this is the square root of $2 w_1$ squared, which equals $\sqrt{2} w_1$. Well, actually no, it doesn't equal $\sqrt{2} w_1$ because w_1 is actually negative. So it's negative $\sqrt{2} w_1$. It doesn't matter.

The point is that if that equals $\sqrt{2}$ over 4, then w_1 is-- everyone?

AUDIENCE: Negative $1/4$.

PROFESSOR: Negative $1/4$. Bingo. And if w_1 is $-1/4$, everything else falls into place. What are w_2 and b ? Everyone?

AUDIENCE: Positive $1/4$.

PROFESSOR: Positive $1/4$. We got it. We're done with this part of problem. However, bonus. Let's come to the alphas, which they didn't ask you to calculate. Actually, you know what? We'll do the alphas if we have enough time, since they didn't actually ask you to calculate them. However, my recommendation is since there's only one alpha-plus and one alpha-minus, they must be equal from this equation, since the sum of the alpha-plus equals the sum of alpha-minus. And so therefore, w equals the sum of w -- sorry, this should be an x . Of course, there's not a million w 's in this equation. The sum of the positive data points times their alphas is minus the negative data points times their alphas.

So we're looking at here $-1/4$, $1/4$ equals-- what do we got here? Positive point negative 1, 2? So we've got alpha, alpha of that point negative 1, 2 minus alpha of that minus point. And what is that? It's 3, -2. 3, -2. So if both of the alphas which are equal were 1, we'd have -4, 4. But we want $-1/4$, $1/4$. So actually both of the alphas are $1/16$. And that's the answer.

We'll do that more in depth if we have time. But we won't. So let's do number two. So let's go into faster mode. Number two, very similar to number one in many ways. But as you can see, one of the main things that they added an extra minus sign at 2, -1. So I think we can all agree that this will still be our plus-- Actually, they added another plus sign there, too. So maybe this plus sign is a support vector. But it's not. This plus sign is a support vector. What do you guys think about the new negative sign? Will it become a support vector since it is strictly closer to the pluses? Yep, you're right.

OK, so this is a very beautiful division because if I do this correctly, which I didn't, but if we pretend that I did. Then the dotted line is--

AUDIENCE: [INAUDIBLE].

PROFESSOR: y equals x . OK, so with the dotted line at y equals x , then just like we did up here, we know that if y equals x plus 0, we know that first of all, b equals 0. Second of all, we know that if y equals x , then we know that $-w_1$ over w_2 equals 1. The pluses are still on the left and up, so we know that w_1 is some negative number, $-k$, and w_2 is some positive number k . Great.

How are we going to figure it out? Well, let's call this d for distance, or whatever you want to call it. So 1 over w equals d . d In this case is not 2 over 2 . Can everyone tell what d is here?

AUDIENCE: [INAUDIBLE]

PROFESSOR: It's actually-- so it goes over 2 and 1 . So it should be $1 \frac{1}{2}$ root 2 since this width distance, which is twice as much, goes over 3 and 3 , which is 3 root 2 . So it's $1 \frac{1}{2}$ root 2 . I don't like putting in decimals and stuff there. So we'll say that 2 over magnitude of w equals $2d$ equals 3 root 2 . So therefore, right, Pythagorean, one, two, three, one, two, three, 3 root 2 . So therefore, magnitude of w equals-- let's see. Switch them over. We should get root 2 over 3 .

And if magnitude of w is root 2 over 3 , we can do our same trick from before, square root of 2 w^2 squared equals root 2 over 3 . And this is just root 2 times w^2 equals root 2 over 3 . So therefore, w^2 is? $1/3$. And w_1 is? $-1/3$. Bingo. We've got w_1 . We've got w_2 . We know that b was zero because obviously it's 0 . It's y equals x . And we're done. That was fast. The alphas might taken longer.

Actually, the alphas on this one are more of a pain in the ass than anywhere else because let's take a look at this one if you can see it. We've added in yet some new points. We've got this point up here and this point down there. So I think pretty clearly this plus and minus are the closest to each other. But what happens if we take the perpendicular bisector between these two and do like this? This plus is in the middle. So therefore, this plus is going to have to also be a support vector.

So we can't just draw this line. We have to include this. What's our best division? Vertical lines, that's right. Vertical lines just so. And that means that the equation of our boundary, the dotted line here, y -axis. So the equation of the boundary with the y -axis, then b equals 0 . And hell, w_2 equals 0 . So the only thing that is not--

AUDIENCE: Wait, w_1 [INAUDIBLE].

PROFESSOR: w_2 equals 0. So w_2 equals 0. b equals 0. But w_1 is not equal to zero because it's just the equation of the y -axis. So we therefore know that the equation is just w_1 times x equals 0. So it's w_1 times x equals 0. And we know that that just means essentially that x is going to be some k . It's also going to be negative because of the fact that the pluses are still on the left.

Then we're going to have to figure out what that k is. We'll use our old trick-- by this point old, hopefully. One over magnitude of w equals d . This time d is just 1. So therefore, magnitude of w equals 1. There's only one component in w . So therefore w_1 is

AUDIENCE: -1.

PROFESSOR: -1 because the plus is on the left. Do people see that? Not too bad. This one's easy to calculate the w . But it's not as easy to get all the alphas. But let's move on to a new and even more fun-- maybe not-- question, which is this guy. As you can see-- well maybe not. This is a one dimensional vector. These vectors only have a single dimension. So it just looks like a number line here. That dimension varies from it looks like -9 to positive 9. It just has one component. You don't have to worry about any of these crazy magnitudes with two components, just everything as a single component.

However, it's obvious that a linear basis line is going to completely screw us up here, since lines at this point are just like, grunk, all these are pluses. All these are minuses. Well, great, that doesn't get them all. So how are we going to do it? Well, we're going to use what is usually perhaps the hardest thing in SVMs, but in this case is not going to be too bad for us. We're going to use a kernel.

Now, based on how little I understood kernels the first time I took this class, I'm guessing that you guys would like to have some explanation on these kernels. You probably saw them. You remember the kernels from Patrick's lecture vaguely? There's this ϕ . And then there's this k . And then they get really complicated.

OK, so here's how the kernel works. The basic idea is this. And I'll write it over here. Oh, wow, there's more stuff. OK. I'll write it right here. The basic idea is this. We're taking the normal space, which is this number line or it could be any kind of normal space, and we're going to take a vector, we're going to put into it a function called ϕ . And ϕ of vector x brings x into some new dimension.

ϕ , or "phee" if you like it better, is usually a nasty piece of work and something you never,

ever want to look at. Sometimes it's not too bad. Phi is the function that brings it into the new dimension. OK? And when you brought the data into a new dimension, sometimes you can just cut a straight line in that dimension and you'll just be happy.

However, something that was noted by the very, very smart inventor of support vector machines is that you don't actually need to work with function phi, even if phi is an absolutely horrible monstrosity, because of the fact that you never need to know what all these vectors x actually are in the new space, at least not directly. In none of these equations up here do we ever use x by itself. However, we do use x being dot product with something else. So he figured out a very sneaky and excellent shortcut.

OK, so-- oh, I shouldn't use x_1 and x_2 . I'll use x and z . So if you have two vectors, x and z , which are in a regular space, you put them into this function called the kernel. Then it will tell you $\phi(x)$ dotted with $\phi(z)$. And if you have that, you don't need phi. Does everyone see that? Does everyone see why we don't need phi? Look at all these equations up here. We have never looked at x by itself in these vector equations at least.

Now, calculating alphas, yeah, that gets a little bit fuzzy. Also, you may ask, why would you do this? You can't calculate the alphas. It turns out that actually, other than for these very simple linear problems, human minds cannot calculate the alphas. In fact, you run a very complicated quadratic optimization. In fact, finding out the best alphas is the thing that you will climb on when you're doing SVMs in the real world. You say, all right, I'll run my algorithm when I know there's only one peak, which is very, very good because it's quadratic optimization. Let me figure out the alphas.

So in fact, it doesn't matter that you can't use these alpha equations to figure out the alphas if you only know the kernel function and not the phi function because normally, the computer figures out the alphas for you with quadratic optimization. Just in these simple problems, we know you can calculate the alphas.

So we have the kernel, which basically gives us the dot product of the things in the new space. So being that as it may, I'll give you the kernel here. I'd like you give me phi. Someone got an idea, whose name was Susan Q. Random Student, apparently. She got an idea that if we had a kernel for x and z -- actually, they're not vectors, I guess. There just single components. And the kernel equals $\cos(\frac{\pi}{4x}) \cos(\frac{\pi}{4z}) + \sin(\frac{\pi}{4x}) \sin(\frac{\pi}{4z})$.

So that is the new dot product. Oh, wait, sorry. I put one of the z's not inside the parentheses. That was silly of me. So cosine of the quantity π over $4x$ times cosine of the quantity π over $4z$ plus sine of quantity π over $4x$ plus sine of quantity π over $4z$ is the new product. So that begs the question. This is an easy one so we can calculate the phi. What is phi of x ?

We're actually taking it from one dimension and we may be playing around with it a lot to get this. And this thing has become a new dot product. It replaces dot product. And remember, the dot product for scalars would have just been multiplying two numbers together. So it actually makes it a little bit more complicated. Does anyone think they know the phi? Oh, we got one. What do you think?

AUDIENCE: [INAUDIBLE]

PROFESSOR: You mean two common vectors, two dimensional?

AUDIENCE: The two points.

PROFESSOR: Absolutely. That's exactly correct. How would you have solved this on the actual quiz if you're not our brave volunteer?

Well, that k , if you squint at it-- not very much actually-- is pretty much a dot product between cosine of π over 4 and sine π over 4 . I mean, look at it. Remember, if the dot product of x and z vectors is x_1, z_1 plus x_2, z_2 -- so that basically is x_1, z_1 plus $x_2 z_2$. Oh, this should have been a times. Yeah, this should have been a times. Sorry. There's a plus there. Anyone who missed it because of that, my bad. That's should have been a times. That should have been a times up there. It's cosine π over $4x$ cosine π over $4z$ plus sine π over $4x$ times sine π over $4z$. So yeah, it's basically the dot product between cosine π over $4x$ and sine π over $4x$. Bingo.

All right, last thing. Well, we're not done yet because we're going to maybe ask some questions. And then we're going to see if we can calculate those alphas. But last thing, let's graph in this new dimension all the points. So obviously, cosines and sines, so we're going to get results between 1 and -1 . Let's see. Maybe I can graph it-- did I write on all these? Wait, maybe this one. No, people drew weird stick figures there. OK. Oh, yeah, this one's kind of messy. But we'll do it on this.

OK, so this is $1, -1, 1/2, -1/2, 1, -1, -1/2, 1/2$. OK? So given that, let's try to graph all these points on this number line into this brave new dimension by using their cosine times π over 4 .

So all right, great. So let's do the pluses first. The plus at 0 is cosine 0, sine 0. So what is that?

AUDIENCE: 1, 0.

PROFESSOR: That's 1, 0. That's right. In fact, the 8 and the -8 are also that times 3. The 8 and the -8 are also that because then it's just 2π minus 2π , which both cosine and sine are periodic. OK, great.

What about the 1? Well, that's cosine π over 4, sine π over 4. And what's that?

AUDIENCE: [SEVERAL ANSWERS]

PROFESSOR: Yeah, it's root 2 over 2, root 2 over 2. So that's something like here, we'll say. And in fact, and the 9 and the -7 are also that. So there's three of these two. What about the -1? That's cosine negative π over 4, sine π over 4.

AUDIENCE: [SEVERAL ANSWERS]

PROFESSOR: That's right. The x value is positive root 2 over 2. And the y value is negative. And again, there's three of them. All right, great. Now let's do the minuses. There's the minus at 3, which is also the same as the minus at -7. The minus at 3 is cosine 3π over 4, sine 3π over 4. Which one is that?

AUDIENCE: [INAUDIBLE].

PROFESSOR: Yeah, that's going to be in the second quadrant. The cosine is going to be negative. But the sine is going to be positive. And so we get 3 points here. And as you may have predicted, the other one, the 5π over 4, is in the third quadrant. We get 3 points here, Where are the support vectors?

AUDIENCE: Question.

PROFESSOR: Question?

AUDIENCE: I understand where you're getting at the three quantities of pluses in the first of the quadrants. But according to the [INAUDIBLE] line there, you want that the total of four-- maybe the values.

PROFESSOR: Oh, you're right. There's only two. Good call. There's two negatives here. And there's two negatives here. Good call. It doesn't change the problem. In fact, if we just graph more points, there might have been more. But that's a very subtle and important distinction. There are two

negatives. But otherwise, yeah, these are graphed correctly. Does anyone see where the support vectors are?

AUDIENCE: The first two.

AUDIENCE: Maybe the top two.

PROFESSOR: So the top two, the minus and plus. We'll try to do the perpendicular bisector. Let's see it. That works. But guess what? These guys are on the same line. So we'd better circle them. So actually, the question is what isn't a support vector? Only this. Only those three. Question?

AUDIENCE: Couldn't you have just done this in one dimension? I mean, you just showed that those ran on the same lines. So you really didn't need the cosine term and a sine term. You could have proved all this with just the cosine.

PROFESSOR: All right, the question is couldn't we have done this in one dimension. All we do is only the cosine. So if we did only the cosine, then they would've still been easily divisible. The answer is, absolutely, we could have. However, the question did not because Susan Q. Random Student decided to do cosine and sine.

But yes, if we had said you, [INAUDIBLE] student, find a ϕ that will work for this, you could have found a ϕ that was just cosine. That would have been easier. However, it's important to be a little work with what somebody else gives you. In this case, they gave you that transformation, which yeah, was wasteful with an extra dimension. You didn't need the sine because you didn't need the y-axis really here. You just needed the x.

Does everyone see this, how this works? You can maybe transform dimensions? The main hardest part is they'll usually give you a k and ask for a ϕ or give you a ϕ and ask for a k . But it's not too bad. Just remember, if they give you a ϕ , do a dot product with it. And if they give you a ϕ that's just one component, dot product of one component, just multiply them together. Easy enough.

If they give you a k , treat it as a dot product and try to reverse engineer. It's usually something like this that's easy to reverse engineer. I really haven't seen it where it's not. So it often looks like the scariest problem. But it's usually not too bad to go between ϕ 's and k 's. Does anyone have any questions on anything that we did on support vector machines? Question.

AUDIENCE: So what's the intuition behind the w ? We solved it and figured out numbers and integrations

with it. But what is it in relation to--

PROFESSOR:

Question is, what is intuition? What is w ? W is the dividing line. It is the drop dead dividing line. When I say the drop dead dividing line, you like those big, bold solid lines over there. Those are your pretty certain lines. Everything past that was a minus. In your training, it's that everything past the big bold line there was a plus in your training stuff. But the dotted line is the one you're really going to use in the test data.

In the test data, when push comes to shove, you might get something if it's inside of that gutter. And if it's on the, say, of that one up there, if it's on the upper left side of that dotted line, you're going to call it a plus. So that dotted line is your decision boundary. And that is basically the idea.

And in fact, the way that the algorithm would do it on the computer is it would quadratically optimize the alphas, which messes around with the dotted line. And by quadratically maximizing the alphas-- you see how the alphas add up to a w . It just checks it around. And eventually, it finds, oh, making the alpha of this one 0 makes it a better optimization. You're trying to get the widest possible road. It would eventually come out to this.

This is trivial for a human to eyeball. But some real problems with 200 data points that have to get one or two of them wrong, classified, and you may be using a quadratic kernel or something that, you can't do that. You just can't. Well, maybe can, in which case you should be getting a MacArthur Fellowship or something like that. But the computer can. And the basic idea is when it comes down to it, it figures out the alphas, that the best w for the widest road. And the w is your decision boundary. Good question. Any other questions about our old friend, SVM?

I have a question for you. After seeing this-- and let's pretend that they only asked to solve for w 's, b 's, these kind of kernels and ϕ , which are the typical things-- how many people now think that you can go through and work an SVM problem? All right, we've got a few. We've got a happy few. Band of brothers. Maybe eight people raised their hand there. That's good.

How many people know what a support vector is now? That's really good. Because if that's all you learned from today's recitation, it's still good. It really is. I'm telling you. I had to take two classes on this and then TA it before I really, really understood it. So you guys are ahead of me.

All right, take care. Have a great weekend. And we'll see you for boosting and vampires next week.