The wedge product of differential forms is kind of like the cross product of vectors; however, instead of trying to interpret it as another vector (or a scalar), we view it as another differential form of higher 'rank' than the original forms. The ordinary differential forms that we're used to are rank 1, and they can be evaluted at a point and a vector; to evaluate a differential form of rank 2, you need a point and 2 vectors. If you keep going with more wedge products, then you get differential forms of even higher rank: to evaluate a differential form of rank n, you need a point and n vectors.

The wedge product also involves subtracting one thing from another (again like the cross product); if α and β are 1-forms (differential forms of rank 1, as we've been using so far), P is a point, and $\mathbf{v}$ and $\mathbf{w}$ are vectors, then

$$(\alpha \wedge \beta)\big|_{\substack{R=P \\ dR=\mathbf{v},\mathbf{w}}} = \alpha\big|_{\substack{R=P \\ dR=\mathbf{v}}} \beta\big|_{\substack{R=P \\ dR=\mathbf{w}}} - \alpha\big|_{\substack{R=P \\ dR=\mathbf{w}}} \beta\big|_{\substack{R=P \\ dR=\mathbf{v}}}.$$

For example, if $\alpha = x^2\,dx + xy\,dy$, $\beta = y^2\,dx - xy\,dy$, $P = (2,3)$, $\mathbf{v} = \langle 0.01, 0.04 \rangle$, and $\mathbf{w} = \langle -0.01, 0 \rangle$, then

$$\left((x^2\,dx + xy\,dy) \wedge (y^2\,dx - xy\,dx) \right)\Big|_{\substack{(x,y)=(2,3) \\ d(x,y)=\langle 0.01,0.04\rangle,\langle -0.01,0\rangle}}$$

$$= (x^2\,dx + xy\,dy)\big|_{\substack{(x,y)=(2,3) \\ \langle dx,dy\rangle=\langle 0.01,0.04\rangle}} (y^2\,dx - xy\,dy)\big|_{\substack{(x,y)=(2,3) \\ \langle dx,dy\rangle=\langle -0.01,0\rangle}}$$

$$- (x^2\,dx + xy\,dy)\big|_{\substack{(x,y)=(2,3) \\ \langle dx,dy\rangle=\langle -0.01,0\rangle}} (y^2\,dx - xy\,dy)\big|_{\substack{(x,y)=(2,3) \\ \langle dx,dy\rangle=\langle 0.01,0.04\rangle}}$$

$$= \left((2)^2(0.01) + (2)(3)(0.04) \right)\left((3)^2(-0.01) - (2)(3)(0) \right)$$

$$- \left((2)^2(-0.01) + (2)(3)(0) \right)\left((3)^2(0.01) - (2)(3)(0.04) \right)$$

$$= (0.28)(-0.09) - (-0.04)(-0.15) = -0.0312.$$

A few basic properties of the wedge product follow immediately (where α, β, γ are 1-forms and u is a 0-form, that is an ordinary non-differential quantity):

$$\alpha \wedge (u\beta) = (u\alpha) \wedge \beta = u(\alpha \wedge \beta);$$
$$(\alpha + \beta) \wedge \gamma = \alpha \wedge \gamma + \beta \wedge \gamma;$$
$$\alpha \wedge (\beta + \gamma) = \alpha \wedge \beta + \alpha \wedge \gamma;$$
$$\alpha \wedge \beta = -\beta \wedge \alpha;$$
$$\alpha \wedge \alpha = 0.$$

So if you treat the wedge product as a kind of multiplication, then you can use the ordinary rules of algebra, so long as you keep track of the order of multiplication in the wedge product.

To define a wedge product between forms of higher rank, you have to add and subtract all possible permutations of the possible orders in which to write the vectors at which the result is evaluated. Keeping track of all of this in a general formula is complicated, but the important point for our calculations is that the rules above continue to apply, and additionally we have an associative law for wedge products:

$$(\alpha \wedge \beta) \wedge \gamma = \alpha \wedge (\beta \wedge \gamma).$$

We will not actually need to evaluate these higher-rank forms in this course; what's necessary is to work with them algebraically.

The basic example that shows how to do this is the transformation between rectangular and polar coordinates. Given

$$x = r\cos\theta,$$
$$y = r\sin\theta,$$

we differentiate to get

$$\mathrm{d}x = \cos\theta\,\mathrm{d}r - r\sin\theta\,\mathrm{d}\theta,$$
$$\mathrm{d}y = \sin\theta\,\mathrm{d}r + r\cos\theta\,\mathrm{d}\theta.$$

Given this, the algebra of the wedge product determines this calculation:

$$
\begin{aligned}
\mathrm{d}x \wedge \mathrm{d}y &= (\cos\theta\,\mathrm{d}r - r\sin\theta\,\mathrm{d}\theta) \wedge (\sin\theta\,\mathrm{d}r + r\cos\theta\,\mathrm{d}\theta) \\
&= (\cos\theta\,\mathrm{d}r) \wedge (\sin\theta\,\mathrm{d}r) + (\cos\theta\,\mathrm{d}r) \wedge (r\cos\theta\,\mathrm{d}\theta) + (-r\sin\theta\,\mathrm{d}\theta) \wedge (\sin\theta\,\mathrm{d}r) + (-r\sin\theta\,\mathrm{d}\theta) \wedge (r\cos\theta\,\mathrm{d}\theta) \\
&= \cos\theta\sin\theta(\mathrm{d}r \wedge \mathrm{d}r) + r\cos^2\theta(\mathrm{d}r \wedge \mathrm{d}\theta) - r\sin^2\theta(\mathrm{d}\theta \wedge \mathrm{d}r) - r^2\sin\theta\cos\theta(\mathrm{d}\theta \wedge \mathrm{d}\theta) \\
&= \cos\theta\sin\theta(0) + r\cos^2\theta(\mathrm{d}r \wedge \mathrm{d}\theta) - r\sin^2\theta(-\mathrm{d}r \wedge \mathrm{d}\theta) - r^2\sin\theta\cos\theta(0) \\
&= 0 + r\cos^2\theta(\mathrm{d}r \wedge \mathrm{d}\theta) + r\sin^2\theta(\mathrm{d}r \wedge \mathrm{d}\theta) + 0 \\
&= r(\cos^2\theta + \sin^2\theta)\,\mathrm{d}r \wedge \mathrm{d}\theta \\
&= r\,\mathrm{d}r \wedge \mathrm{d}\theta.
\end{aligned}
$$

With experience, you can do this sort of thing much faster; for example, you can immediately recognize the terms that will become zero and skip them.

This is an example of changing coordinates in two variables; we can also use two variables to parametrize a surface in three-dimensional space. For example, on the surface of the unit sphere (the sphere of radius 1 centred at $(x, y, z) = (0, 0, 0)$), if we write x and y using r and θ above, then we can further write

$$r = 1\sin\phi,$$
$$z = 1\cos\phi,$$

where the 1 indicates the radius of the sphere and the angle ϕ varies from 0 to π. Differentiating,

$$\mathrm{d}r = \cos\phi\,\mathrm{d}\phi,$$
$$\mathrm{d}z = -\sin\phi\,\mathrm{d}\phi.$$

Thus,

$$
\begin{aligned}
\mathrm{d}x \wedge \mathrm{d}y &= r\,\mathrm{d}r \wedge \mathrm{d}\theta = \sin\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta, \\
\mathrm{d}x \wedge \mathrm{d}z &= (\cos\theta\cos\phi\,\mathrm{d}\phi - \sin\phi\sin\theta\,\mathrm{d}\theta) \wedge (-\sin\phi\,\mathrm{d}\phi) = 0 + \sin^2\phi\sin\theta\,\mathrm{d}\theta \wedge \mathrm{d}\phi, \\
\mathrm{d}y \wedge \mathrm{d}z &= (\sin\theta\cos\phi\,\mathrm{d}\phi + \sin\phi\cos\theta\,\mathrm{d}\theta) \wedge (-\sin\phi\,\mathrm{d}\phi) = 0 - \sin^2\phi\cos\theta\,\mathrm{d}\theta \wedge \mathrm{d}\phi.
\end{aligned}
$$

However,

$$\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z = (\sin\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta) \wedge (-\sin\phi\,\mathrm{d}\phi) = \sin^2\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\phi \wedge \mathrm{d}\theta = 0.$$

This makes sense if $\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$ represents something like a volume, since the volume of the *surface* of a sphere is zero.

To see how $\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$ indeed represents something like a volume, I should explain how to integrate higher-rank differential forms. You typically integrate a differntial form over a shape (or 'manifold') whose dimension (as given by the number of parameters used to parametrize the manifold) matches the rank of the form. We have already seen this with rank-1 forms integrated over 1-dimensional curves, which can be parametrize by 1 parameter t.

In general, to integrate a rank-p form over a manifold parametrized by p parameters, you divide the manifold up into pieces along the level curves of the parameters, for each piece evaluate the differential form at a point in that piece and at the vectors along the sides of the piece (evaluating $\mathrm{d}u \wedge \mathrm{d}v$ first at the vector along the side where u increases and then at the vector along the side where v increases), multiply by ± 1 according to the orientation of the manifold, and then add these pieces up, taking the limit as the size of the largest piece goes to zero.

Here the **orientation** of the manifold indicates directions along it. In the case of a curve, there are two ways to go along the curve, giving two orientations. In the case of a surface, if we start going in some direction, then we can turn from that direction in one way or the other. In general, every small piece of a manifold has two orientations, no matter what the dimension. A form such as $\mathrm{d}u \wedge \mathrm{d}v$ matches the orientation if moving in the direction in which u increases and then turning in the direction in which v increases matches the turning given by the orientation; if not, then we must use -1 for that piece.

In practice, we don't evaluate an integral as a limit of such sums; instead, we evaluate it as an integrated integral in the parameters. To do this, we simply set up limits of integration over the values that the parameters can take and write down an iterated integral that makes sense, inserting a factor of -1 if the orientation of the differential form is opposite that of the manifold. For example, to integrate $\mathrm{d}x \wedge \mathrm{d}y = \sin\phi\cos\phi\,\mathrm{d}\phi \wedge \mathrm{d}\theta$ on the top half of the unit sphere, oriented to turn clockwise when viewed from outside the sphere, we start with

$$\int_{\theta=0}^{2\pi} \left(\int_{\phi=0}^{\pi/2} \sin\phi\cos\phi\,\mathrm{d}\phi \right) \mathrm{d}\theta = \int_{\theta=0}^{2\pi} \frac{1}{2}\,\mathrm{d}\theta = \pi;$$

but then, because we turn *counterclockwise* to move from a direction in which ϕ increases to a direction in which θ increases, the actual value is $-\pi$.

In the textbook, you'll never be given directly differential forms to integrate. In Section 15.6, you integrate vector fields through surfaces; to integrate the vector field $\mathbf{F}$, you integrate the differential form $\mathbf{F}(x,y,z) \cdot \mathrm{d}\mathbf{S}$, where

$$\mathrm{d}\mathbf{S} = \frac{1}{2}\,\mathrm{d}\mathbf{r}\,\hat{\times}\,\mathrm{d}\mathbf{r} = \langle \mathrm{d}y \wedge \mathrm{d}z, \mathrm{d}z \wedge \mathrm{d}x, \mathrm{d}x \wedge \mathrm{d}y \rangle = \frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v}\,\mathrm{d}u \wedge \mathrm{d}v$$

(where $\mathbf{r} = \langle x,y \rangle z$ as usual). This requires the use of the right-hand rule for the cross product; we accordingly use the right-hand rule to convert between a direction through the surface (which is technically called a *pseudo-orientation*) and an orientation on the surface. So for example, to integrate the vector field $\mathbf{F}(x,y,z) = \langle 0 \rangle 01 = \mathbf{k}$ downwards through the top half of the unit sphere is the same as integrating $\mathrm{d}x \wedge \mathrm{d}y$ on that hemisphere oriented clockwise (which I calculated above to be $-\pi$).

In Section 15.5, you integrate scalar fields through surfaces; to integrate the scalar field f, you integrate the differential form $f(x,y,z)\,\mathrm{d}\sigma$, where

$$\mathrm{d}\sigma = |\mathrm{d}\mathbf{S}| = \sqrt{(\mathrm{d}y \wedge \mathrm{d}z)^2 + (\mathrm{d}z \wedge \mathrm{d}x)^2 + (\mathrm{d}x \wedge \mathrm{d}y)^2} = \left| \frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v} \right| |\mathrm{d}u \wedge \mathrm{d}v|.$$

Now orientation is irrelevant, instead, simply make sure that all parameters are increasing in the iterated integral.

As for Chapter 14, here we are simply integrating scalar fields on the flat surface of the plane, using

$$\mathrm{d}A = |\mathrm{d}x \wedge \mathrm{d}y|$$

and on all of three-dimensional space, using

$$\mathrm{d}V = |\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z|.$$

So for example, in polar coordinates,

$$\mathrm{d}A = |\mathrm{d}x \wedge \mathrm{d}y| = |r\,\mathrm{d}r \wedge \mathrm{d}\theta| = |r|\,|\mathrm{d}r \wedge \mathrm{d}\theta| = r\,|\mathrm{d}r \wedge \mathrm{d}\theta|,$$

where the last step is valid if we only use non-negative values of r.