Table of Contents
Topic: Products Interpretation

Click to jump to lesson

Back to Top

Interpreting Radial Velocity

Introduction

Notes:

Hello there and welcome to this lesson on interpreting radial velocity!

Recap

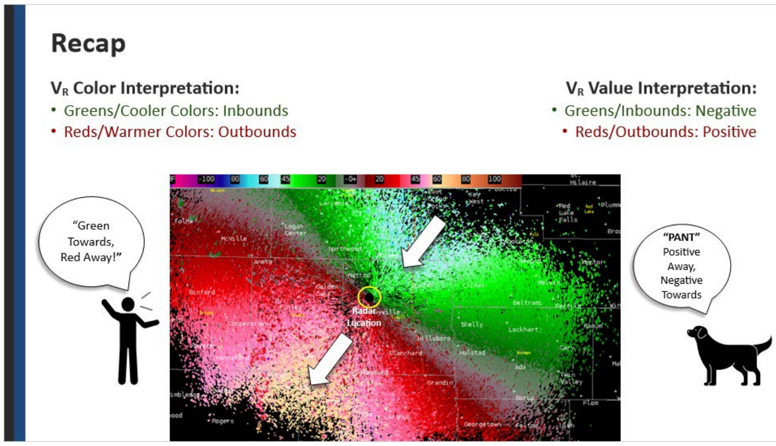

Notes:

Before we get to work learning the process for interpreting a velocity display, let’s quickly recap the groundwork we previously laid in the radial velocity module. If you remember, we just covered interpreting what the colors and values mean. Cooler colors, or greens, indicate winds moving towards the radar and warmer colors, or the reds, indicate winds moving away from the radar. The numeric values listed in the legend are an estimate of the sampled wind speeds along a radial, with negative values representing inbound winds and positive values representing outbound winds. Let’s build on this groundwork to now discuss how you go about determining what the wind direction and wind speed is anywhere on a radar display.

Back to Top

Velocity Interpretation Checks

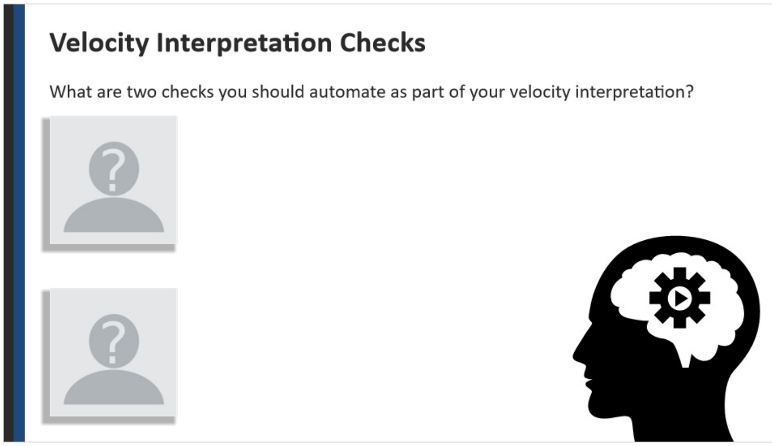

Notes:

Before we cover determining wind speed and direction, there are two super duper important checks that we always need to run through before we begin our interpretation. I want these to be so second nature that they end up in the automated side of your brain - kind of like starting a car and putting it in drive is so natural to us we don't think that these are checkboxes to work through before driving a car. So what should you task your brain with automating when looking at velocity data? Click on either icon to reveal.

Check #1:
Check number 1 - "Where is your radar?" Because all of velocity interpretation relies on the phrases "inbounds" and "towards the radar", we need to know where that thing is that we are comparing everything else to in the first place! Say we're zoomed in looking at this storm. Without knowing where our radar is, we have no idea which direction the inbounds, green colors, or outbounds, red colors, are headed. You can quickly figure this out multiple ways, with the easiest being to zoom out and find where the cone of silence is as shown here in the yellow circled area. In this case, our radar is to the northeast so we know now that for the storm we were looking at, inbounds are moving to the northeast and outbounds to the southwest which means this storm is rotating counter-clockwise, or cyclonically – more on that in a later module!

In addition to finding the cone of silence, you can also turn on cursor sampling and note where your storm is relative to the radar - in this example our storm is 109 nm away and roughly to the southwest, or 241 degrees, of the radar.

A last quick method is to just visually inspect the resolution or width of your radar

Back to Top

range gates. Since the radar beam gets wider the further away you are from radar, range gates will appear coarser or wider. Likewise, they will appear sharper and smaller the closer you get to the radar along a radial. In this example, our radar range gates appear smaller and more defined as we move along a radial to the northeast which tells us the radar is in this direction. Of these three methods, the first involves actively zooming out to establish your frame of reference while the other two can be done without zooming out. Use whatever method is easiest for you as the most important thing is to just always know where your radar is relative to the feature you are looking at.

Check #2:
Check number 2 - “How high up are you sampling?” Remember that because the radar scans azimuthally at a fixed point, the further you get from the cone of silence on a radar display, the higher above ground level you are sampling. In other words, measurements made at close ranges will sample winds at lower heights while measurements made at far ranges will sample winds at higher heights. This is important for your operational interpretation because velocity values can be enhanced aloft but that doesn’t mean they will be that same intensity at the surface.

Note that this and future discussions in this module assume a standard atmosphere. However, in cases of sub-refraction or super-refraction, the beam can get significantly refracted and the actual radar beam heights may vary significantly from what is in the readout.

Back to Top

Check #1 (Slide Layer)

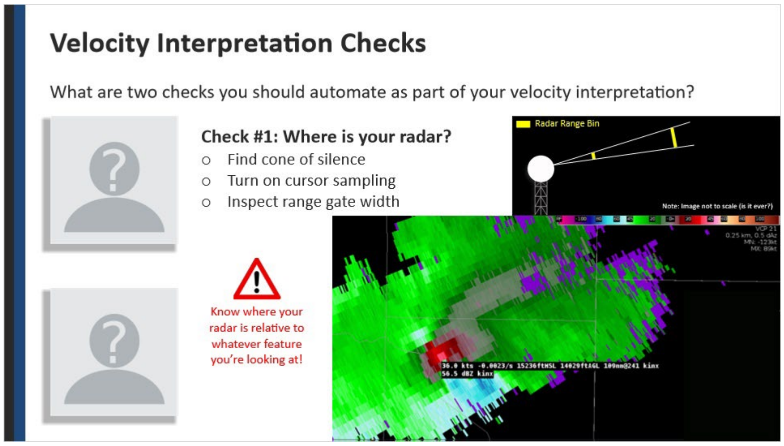

Check #2 (Slide Layer)

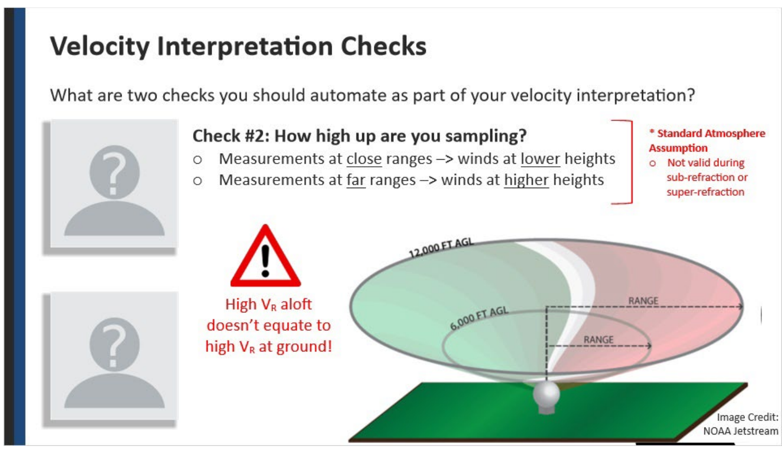

Back to Top

How to interpret wind direction?

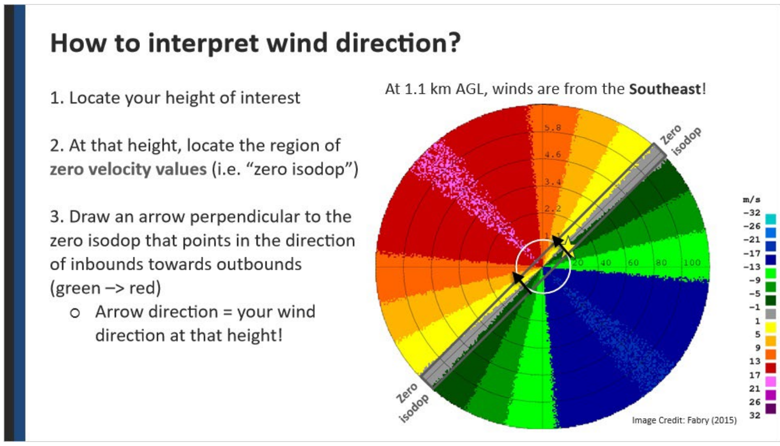

Notes:

Here we are again with our simplified radar display that we've been using in previous lessons. We can now add on range rings that correspond to kilometer distance from the radar as shown on the horizontal line and kilometer height above ground level in the vertical. So at 20 km away from the radar, you will be sampling about 1.1 km above ground level and the further away or more down-radial you go, the higher up you will be sampling. Alright, now let's add velocity - poof, happy colors!

To assess wind direction, the first step is to start at your height of interest. Let's say we want to know what the wind direction is closer to the surface, so let's start at the lowest range ring here, so 1.1 km above ground level.

Next, at that height, locate where the narrow region of zero velocity values are, or the gray radials as shown here. Remember when the wind direction is perpendicular to a radial, no component of that wind is detected so its velocity will be zero. This narrow region of zero velocity values is called the "zero isodop" and will be our best friend in helping us figure out wind direction because it represents where the winds are perpendicular to a radial!

After finding the zero isodop, draw an arrow perpendicular to it that points from inbounds to outbounds, or from the green colors towards the red colors as shown here. The direction that arrow points is your wind direction at that height. So in this example, at 1.1km above ground level, the wind direction is from the southeast.

Back to Top

How to interpret wind speed?

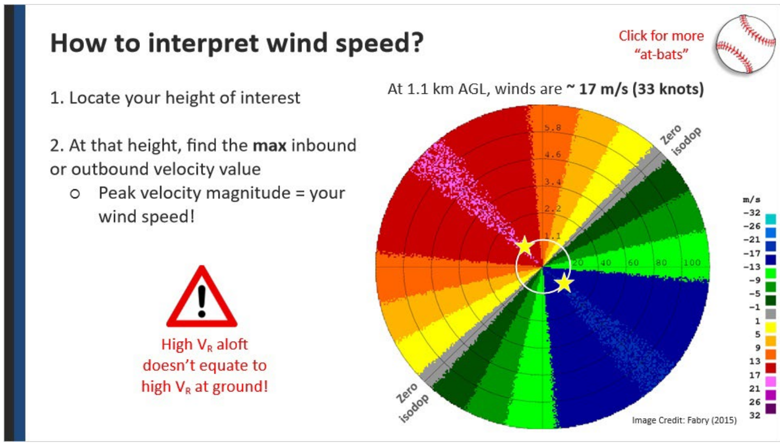

Notes:

The process for determining wind speed is pretty quick. As you did for wind direction, locate your height of interest. So in our example, we were interested in the velocity values at 1.1 km above ground level. At that height, find the max inbound or outbound velocity values as shown. Wind speed here is the magnitude of that peak velocity so in our case, 17 m/s or 33 knots. Note that as part of both determining wind speed and direction, we are assuming the wind is symmetrical at that height.

Just remember our automated check #2 from earlier when it comes to wind speeds where a high wind speed aloft does not equate to that same intensity at the ground. All we can say here is that the wind speed at 1.1 km in height is approximately 17 m/s and that's about it. Now that we've talked through the process for determining wind direction and speed, it's time to get some at bats practicing - click on the baseball icon to begin.

Back to Top

What are some limitations of V?

Notes:

There are a few different situations that can limit our interpretation of radial velocity. Click on each one to learn more.

Absence of Target Returns:
Remember back when we defined radial velocity, it depends on encountering actual moving targets to produce a shift in phase used to calculate radial velocity. Baked within this definition is the limitation that we won't have any radial velocity data available to interpret if there's no moving targets. In this example, note how there are very few moving targets so our display of radial velocity is pretty sparse.

Sampling Height:
One of the biggest challenges when interpreting radial velocity has to do with sampling height. There are many times where you'll sample severe-strength winds aloft as shown in the example here where winds in the circled area are pushing 70 mph but that intensity may not be realized at the surface due to a variety of factors such as the low-level thermal profile, pressure gradient, and surface roughness. To avoid coming off as "this sucks but good luck anyways", I wanted to share a rule of thumb often used by forecasters as a starting point to help with this. If your sampled wind speeds are less than 1 km AGL in height, it's safe to assume these winds will mix down to the surface. In these situations, it's also helpful to really lean into any reports you get or don't get to help boost or lower your confidence in that assumption.

Abnormal Target Movement:
*S*ince most meteorological targets move with the mean wind, radial velocity will in general provide reliable information about wind speed and direction. However,

Back to Top

there are targets out there that do NOT move with the mean wind and these aren't just UAPs, I'm talking about birds, migrating insects, and the sea surface. Birds, for example, add their velocity component to the wind. Conversely, if you are getting returns from ground targets such as buildings, this will bias velocity towards zero. A limitation of this is that whenever you are dealing with non-meteorological targets that move independent of the mean wind or are stationary, this could bias velocity values and there your interpretation.

Practice:
Because interpreting a velocity display requires you to think three-dimensionally and mentally translate colors and values into wind direction and speed, it can be difficult and overwhelming at times to interpret if you don't regularly practice doing so. The glass half full side of this limitation is that it is absolutely in YOUR control to overcome by being intentional to regularly practice velocity interpretation whenever you're looking at AWIPS, GRAnalyst, or RadarScope for example.

Back to Top

Practice (Slide Layer)

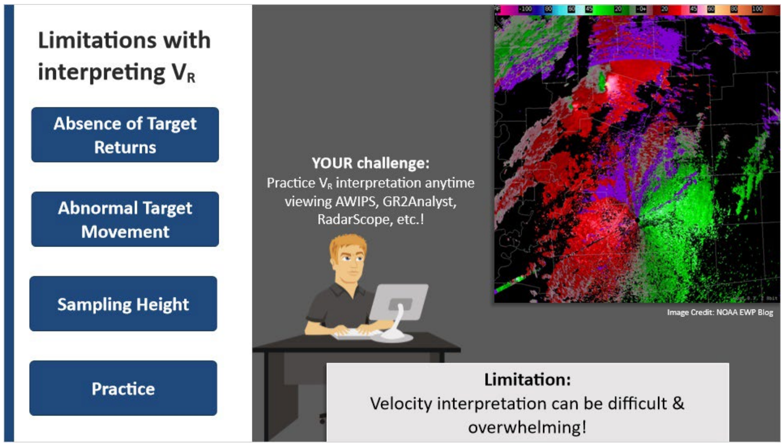

Sampling Heights (Slide Layer)

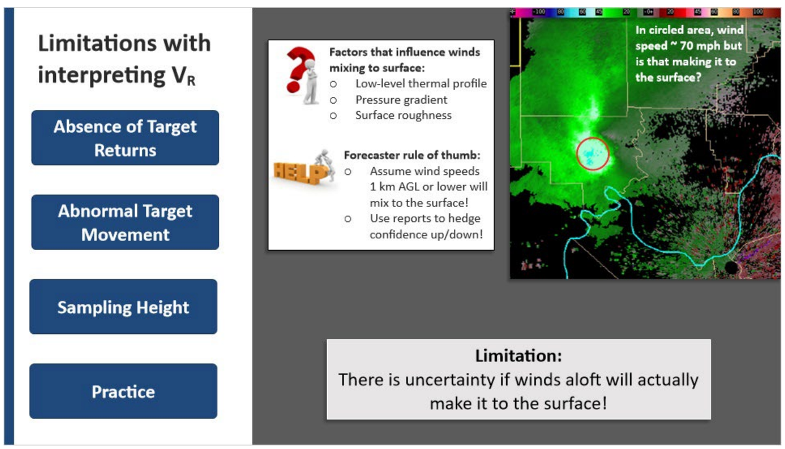

Abnormal Target Movement (Slide Layer)

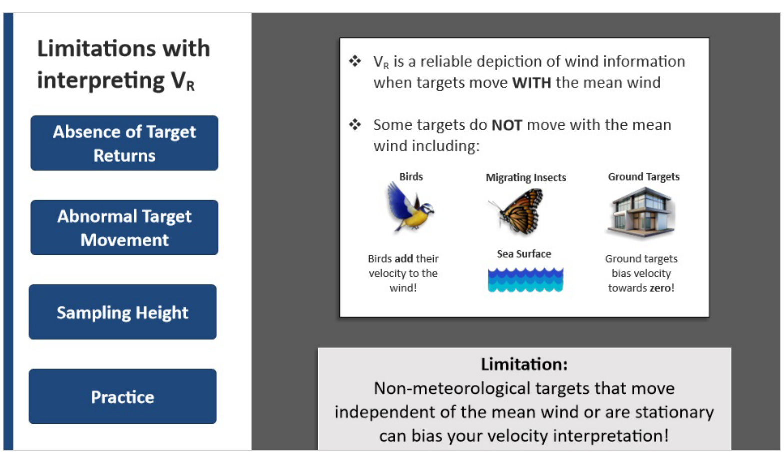

Back to Top

Absence of Target Returns (Slide Layer)

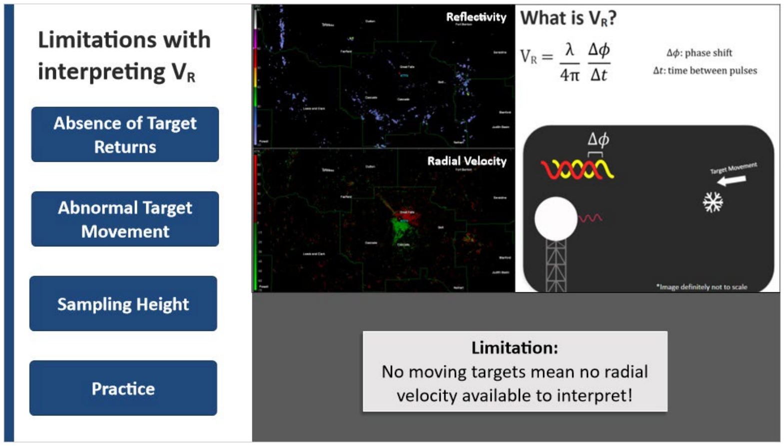

Main Takeaways

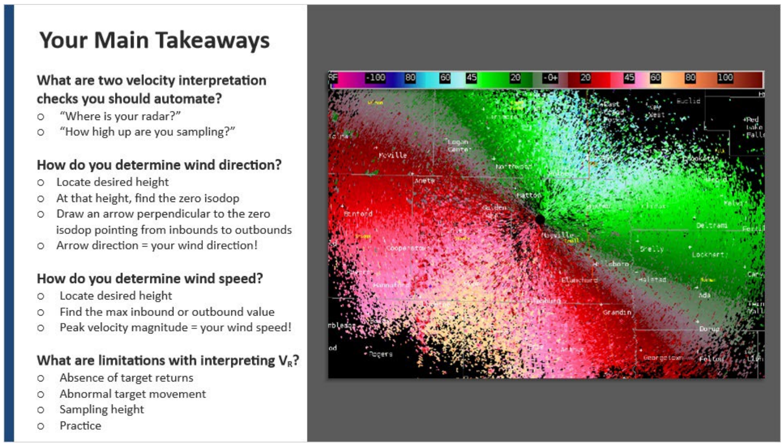

Notes:

Before you take the quiz, here are the main takeaways for your review. Click "next" when ready.

Back to Top

Contact Information

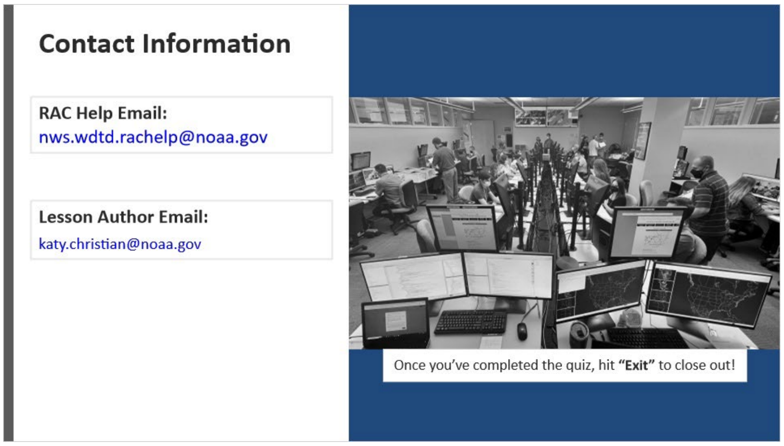

Notes:

If you've got any questions about the material in this training, please feel free to send it to the RAC team or myself at the addresses shown here. If not, click "exit" to close out of this module.

Large-Scale Doppler Velocity Patterns

Introduction

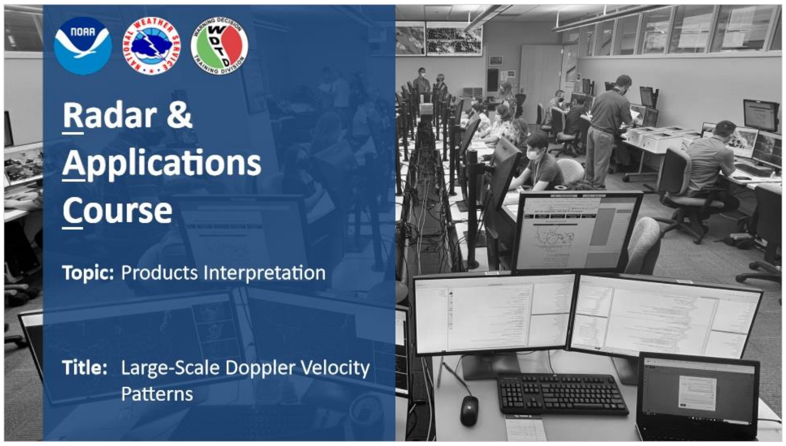

Notes:

Hi, my name is Jessica Bunker and welcome to this Topic on Products Interpretation. This lesson will cover large-scale Doppler velocity patterns, while the next lesson will focus on storm-scale patterns.

Back to Top

Doppler Velocity Patterns

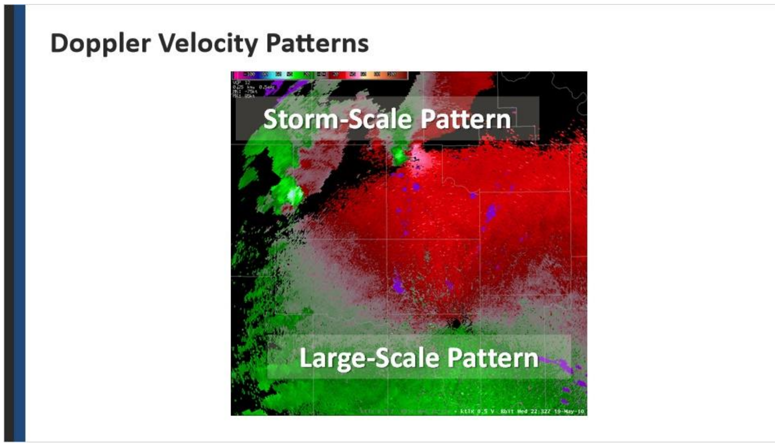

Notes:

In this lesson we will be discussing more complex large scale velocity patterns. In the next lesson will then focus on storm-scale velocity signatures.

Back to Top

“S” Shape Velocity Pattern

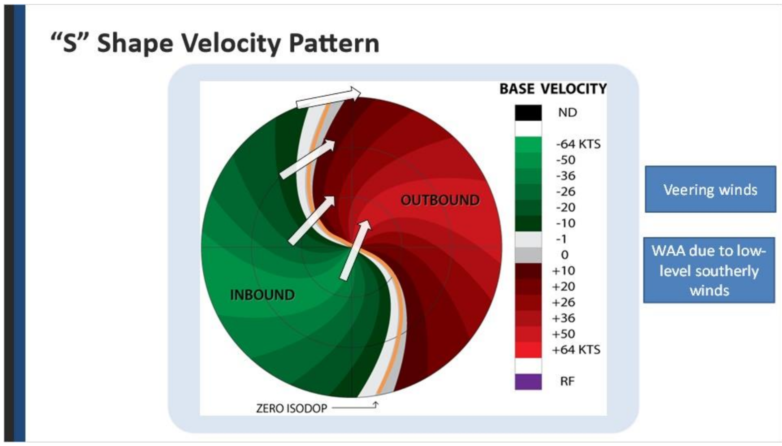

Notes:

Curvature of the zero isodop represents changing wind direction with height. In this example here, the zero isodop is shaped like the letter “S.” The wind direction near the RDA is from the south-southwest while the wind direction is from the west-southwest near the edge of the display. The associated vertical wind profile shown here indicates that the winds are turning clockwise with height. The meteorological term for this is veering. Veering generally indicates that warm air advection is occurring due to the low-level southerly winds.

Back to Top

“S” Shape Velocity Example

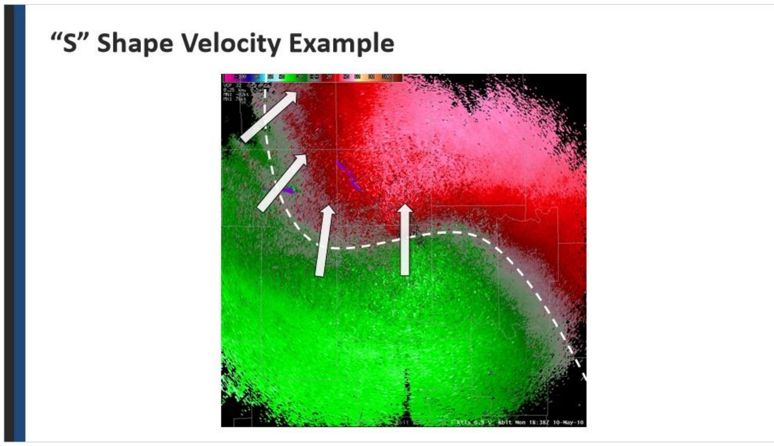

Notes:

Here is a real-world example of veering winds with height. There is a “S” shape to the zero isodop, with winds from the south near the surface and winds from the southwest near the edge of the display. Thus, the winds are veering with height throughout the layer.

Back to Top

Backward "S" Shape Velocity Pattern

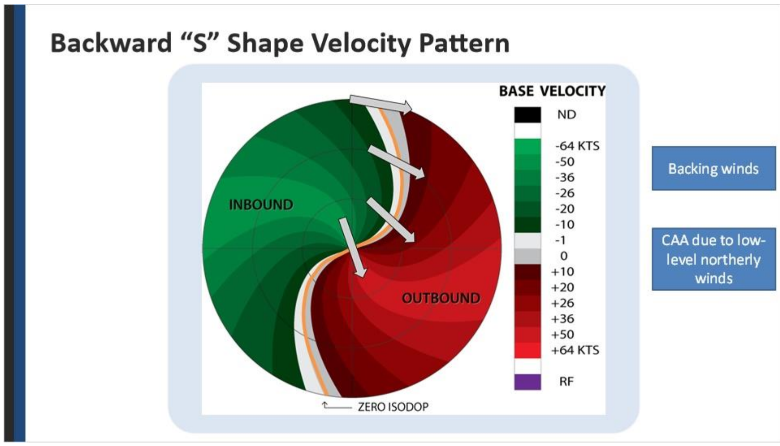

Notes:

Some velocity patterns can exhibit a backward "S" shape pattern. In this case, the winds are turning counterclockwise with height, which is completely opposite of the standard "S" shape profile you saw earlier. In this graphic, the winds are from the north-northwest near the RDA and from the west-northwest near the edge of the display. When the vertical wind profile is turning counterclockwise with height, as seen here, it is referred to as backing. Backing winds are generally associated with cold air advection.

Back to Top

Diffluence

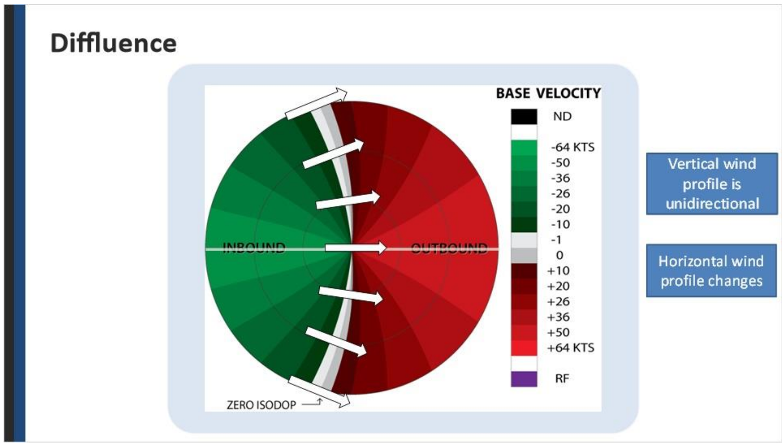

Notes:

We will now move away from pure homogeneous wind fields and talk about diffluent and convergent wind fields. The best way to examine these kinds of wind fields is to split the display into two parts. Starting with the top half, the wind direction changes from the west near the surface to west-southwest at the edge of the display. Looking at the bottom half of the display, the wind direction is westerly near the RDA but changes to west-northwest at the southern edge of the display. This associated pattern shows the air spreading out as it passes over the RDA.

Note that the zero isodop has a bowing shape to it and that the inbound velocities are within the bow. Now, at any one point, the *vertical* wind profile is still going to be unidirectional; however, we now add that complexity of the wind field changing over horizontal distances (not horizontally homogeneous).

Back to Top

Confluence

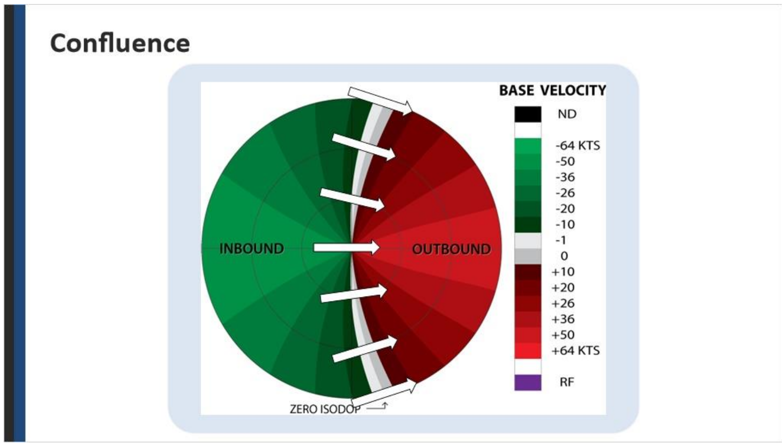

Notes:

Looking at a diagram at what a confluent wind flow would be, you can see that the outbound velocities are now on the inside of the bow-shaped zero isodop. Here, the winds are coming together over the RDA.

Back to Top

Confluence Velocity Example

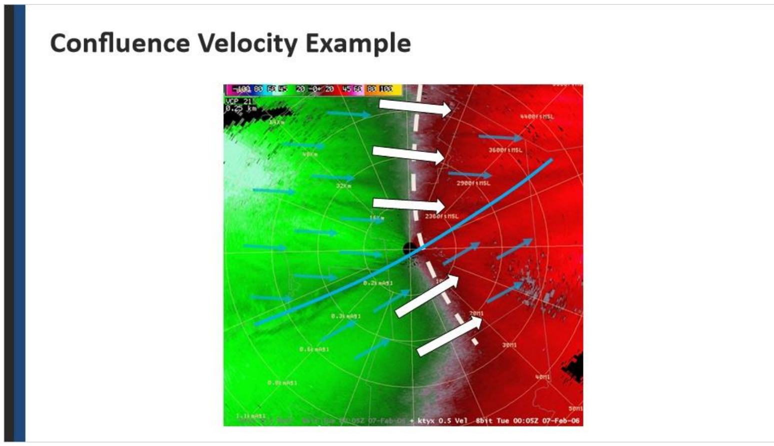

Notes:

Here is a real-life example of a confluence zone. A lake effect snow band is passing over this radar from southwest to northeast. You can see that the zero isodop bends in a way that the outbound velocities are on the inside of the bowing shape.

Back to Top

Sloping Wind Maximum

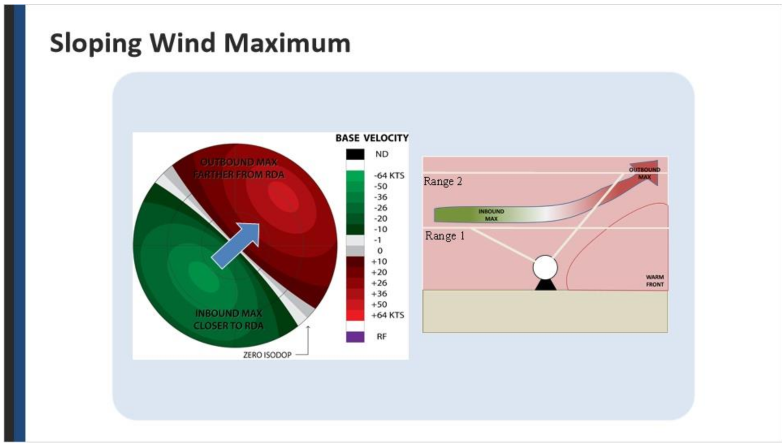

Notes:

Sometimes, you might run into a situation where you have a sloping wind maximum, like a low-level jet moving over a warm front. In this graphic here, the general flow at all levels is from the southwest to the northeast. If you were to look at the location of the velocity maximum for the inbound winds, you can see it is just beyond the first range ring. Now looking at the velocity maximum on the outbound side, you can see it is over the second range ring. This indicates that the wind maximum is increasing with height as it moves across the display.

Back to Top

Discontinuities and Fronts: Front Approaching RDA

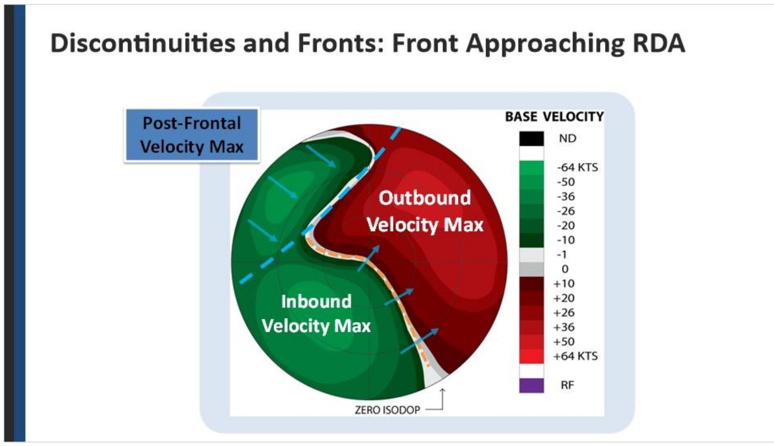

Notes:

Now let's move on to an even more complicated wind pattern. The following examples will show a frontal boundary moving through the display area, and the expected wind patterns at various stages of the passage. First, we will start with the front approaching the RDA, in this case, from the northwest. Here, you can see that the southeastern 2/3 of the display has an "S" shaped pattern, with velocity maxima located to the northeast and southwest of the RDA. Behind the front, located here, you can see a secondary inbound wind maximum to the northwest, which is not "connected" with any of the other two maxima on this display.

Back to Top

Front Approaching RDA Example

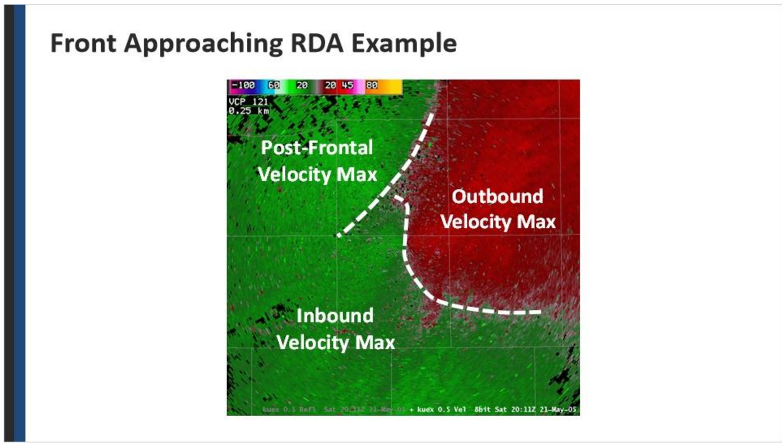

Notes:

This real world example also has a front approaching the RDA. In the real world, it is not always easy to see frontal boundaries, especially in just one volume scan. Using the static image here, you can see the where the front is located via the cutoff between the inbounds and outbounds to the north, and the velocity minimum to the west. A backing wind pattern is seen south of the boundary. The two connected wind maxima are located to the west-southwest and east-northeast of the RDA, while the disconnected post-frontal velocity maximum is located to the northwest of the RDA. Note that this boundary is aloft in this display. That is, the radar beam is sampling the elevated portion of the boundary.

Discontinuities and Fronts: Front Over RDA

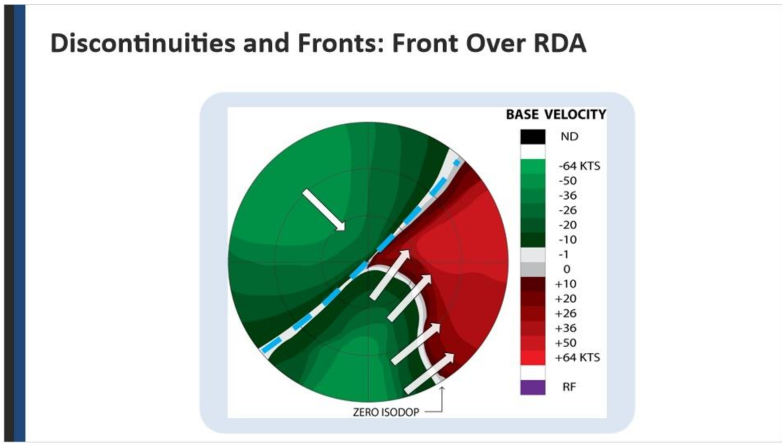

Notes:

Moving forward in time, here is a graphical representation of a frontal boundary now located over the RDA. The boundary is still oriented from southwest to northeast, as seen by the zero isodop. The winds are from the northwest behind the boundary, while they are generally from the southwest in a veering pattern ahead of the boundary.

Back to Top

Front Over RDA Example

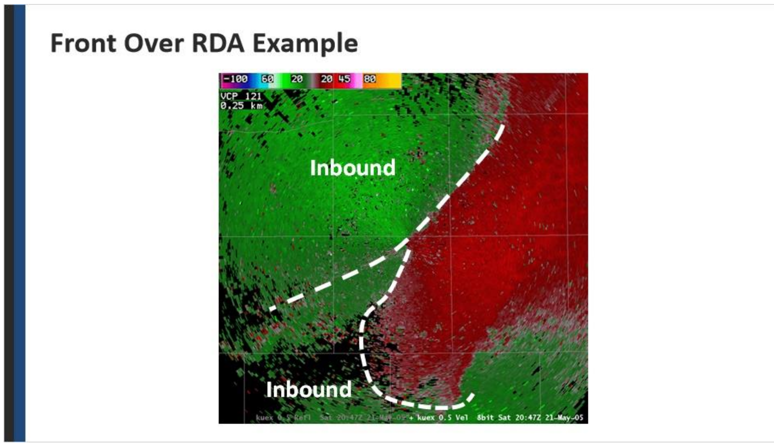

Notes:

Now moving forward in time with the real-world example, the frontal boundary is now at the RDA. The inbound velocities are located to the southwest of the RDA ahead of the boundary and to the northwest of the RDA behind the boundary. A sharp change in speeds indicates the location of the boundary. Also, in this example, note the backing wind profile ahead of the front.

Discontinuities and Fronts: Front After Passing RDA

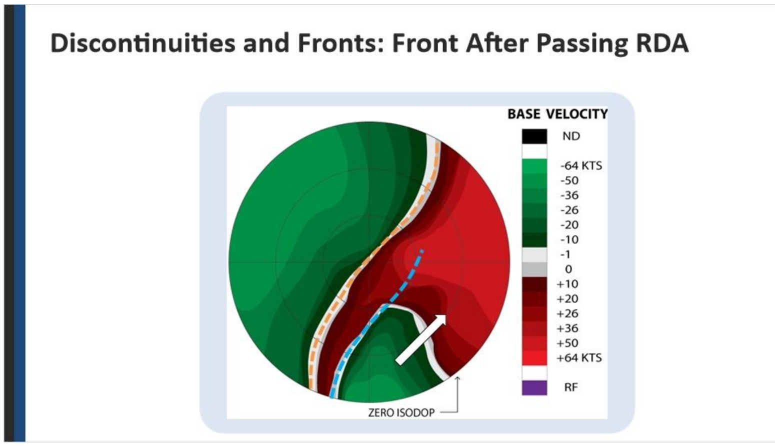

Notes:

The front has now passed over the RDA and is now located to its southeast. The winds ahead of the front are still from the southwest. Behind the front, the winds are now backing with height.

Back to Top

Front After Passing RDA Example

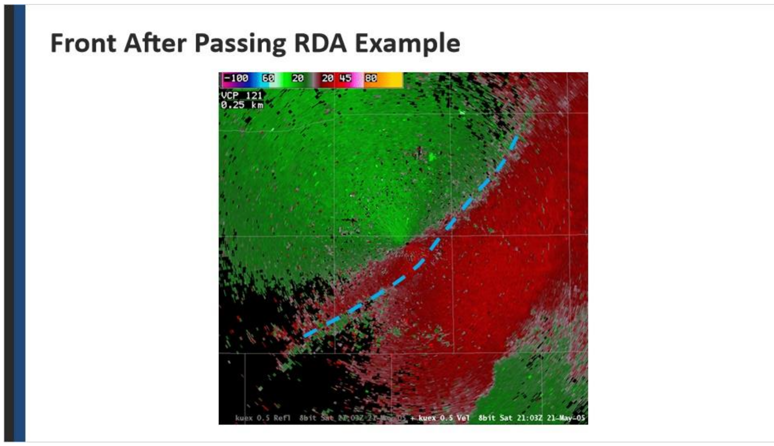

Notes:

Back to the real world example, the front is now located to the southeast of the RDA, which is denoted by the blue line here. Notice that this front is losing definition with time. Only to the southwest of the RDA is there a leading edge of inbound velocities. Otherwise, use the subtle leading edge of the enhanced outbounds as the frontal interface.

Back to Top

Summary

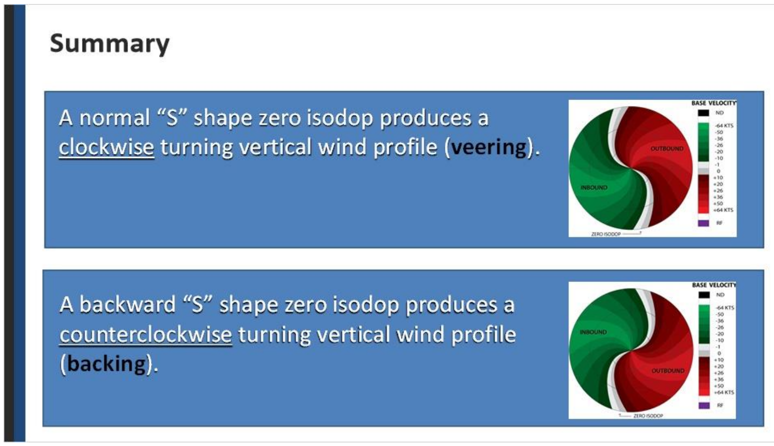

Notes:

Lets summarize what we learned. A normal "S" shape zero isodop produces a clockwise turning vertical wind profile (veering), while a backward "S" shape zero isodop produces a counterclockwise turning vertical wind profile (backing).

Back to Top

Summary (cont.)

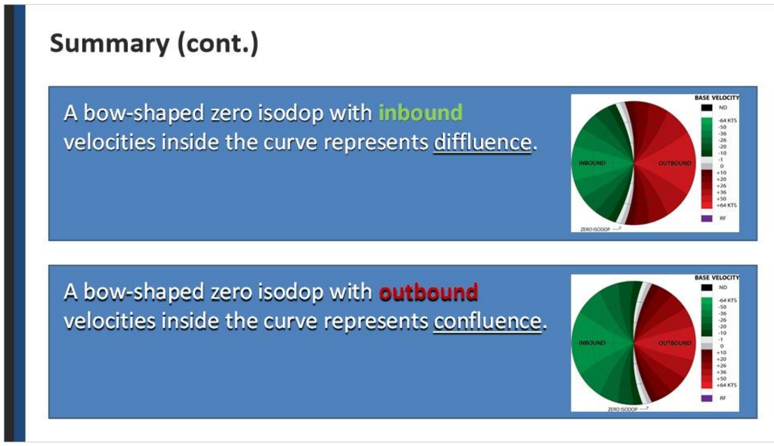

Notes:

Finally, a “bowed” shape zero isodop with inbound velocities inside the curve represents diffluence, while a “bowed” shape zero isodop with outbound velocities inside the curve represents confluence.

Back to Top

Contact Info

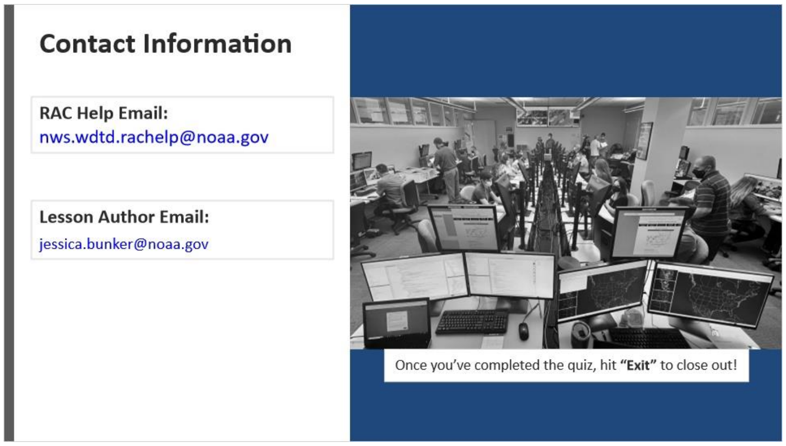

Notes:

If you have any questions or if you would like more information on topic, please contact the emails listed here.

Storm-Scale Doppler Velocity Patterns

Introduction

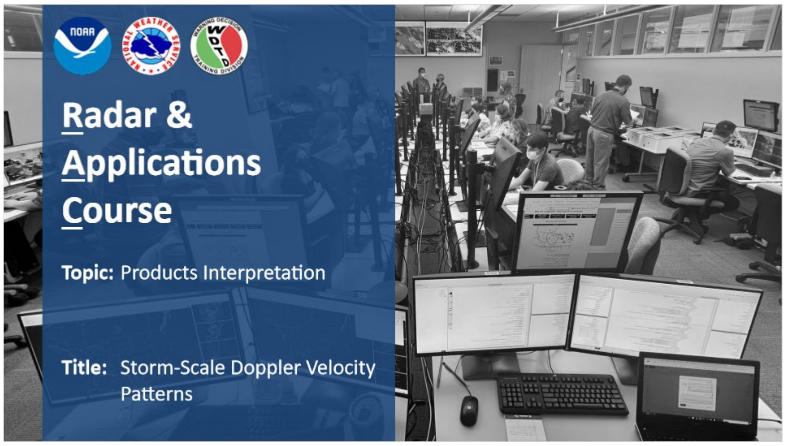

Notes:

Hi again! Its Jessica Bunker and welcome to this lesson on storm scale doppler velocity patterns.

Back to Top

Doppler Velocity Patterns

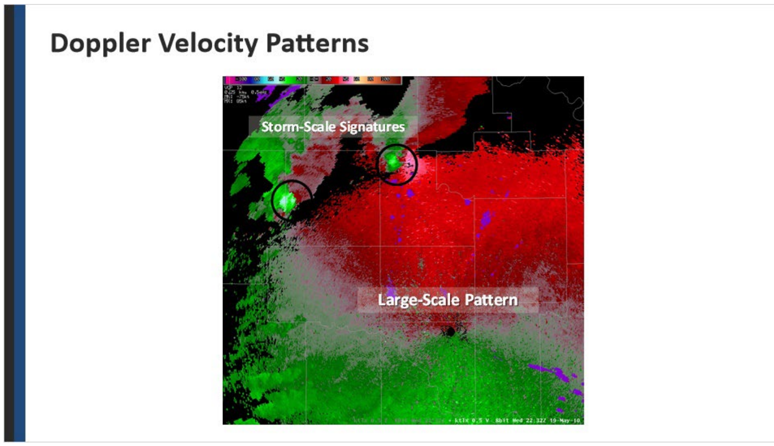

Notes:

In the first lesson, you saw a variety of factors and examples that influence the large-scale velocity field and its display in the AWIPS environment. Now, we will go ahead and focus on the small-scale phenomena, which cover only a few range gates, and therefore, have a relatively small change in elevation.

Back to Top

Locating the RDA

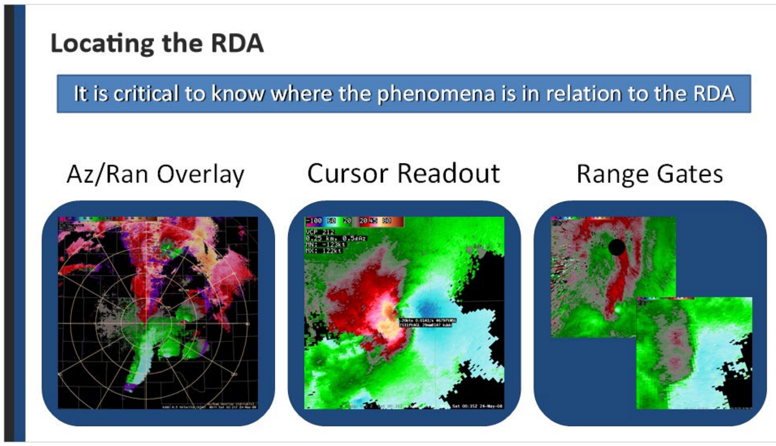

Notes:

When examining data on this scale, you will be zooming in AWIPS in order to see small scale rotation and/or convergence and divergence. It is critical to know where the phenomena is in relation to the RDA. Here, you can no longer assume that the RDA is in the center of the display, or on the display at all. The following three actions, either used separately or in combination, will help in locating the RDA.

--You can select the Azimuth and Range (or Az/Ran) Overlay from the Tools menu to help determine the location of the RDA by overlaying a polar grid centered on the RDA.

--Or you can place your cursor at the point of interest and hold down the left mouse button. The cursor readout will give the azimuth and range (in statute miles) from the RDA. To see this readout all the time, you can right-click and turn on Sampling, as well.

--Or you can visually analyze the range gates in your velocity display. Range gates increase in width along each radial as they increase in distance from the RDA. This is one advantage to an unsmoothed radar display.

Back to Top

Small Scale Pattern: Convergence and Divergence

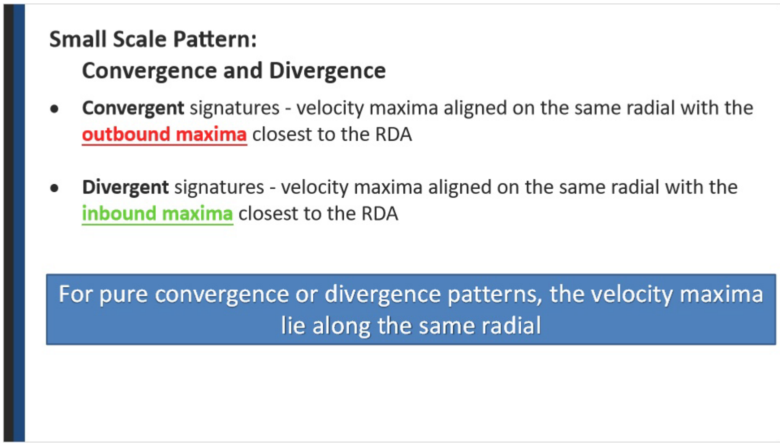

Notes:

When interpreting pure convergence or divergence patterns, the velocity maxima lie along the same radial. Whether the pattern is convergent or divergent is dependent upon which maximum is closest to the RDA.

With a convergent signature, the outbound maxima is closest to the RDA.

With a divergent signature, the inbound maxima is closest to the RDA.

Back to Top

Convergence

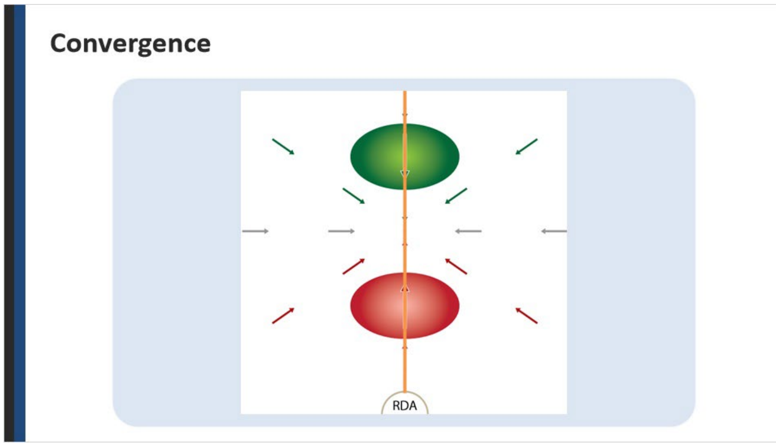

Notes:

Here is a basic diagram of a convergent velocity signature. As you can see here, both the maxima lie along the same radial with the outbound velocity maximum closest to the RDA.

Back to Top

Divergence

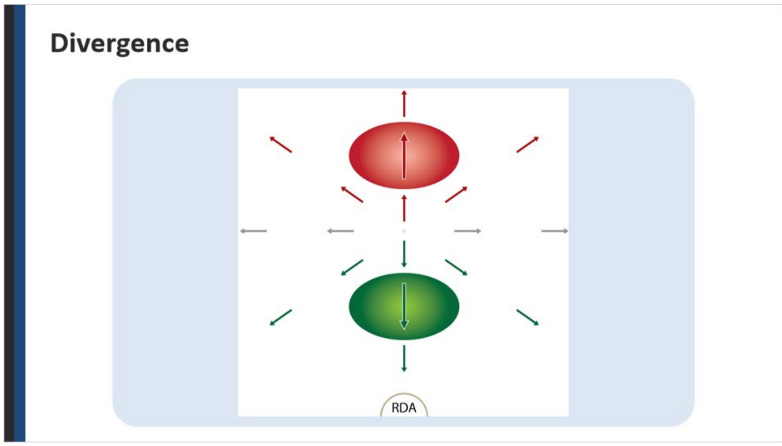

Notes:

Same thing here for the divergent velocity signature, except that the inbound velocity maximum is now closest to the RDA.

Back to Top

Divergent Velocity Example

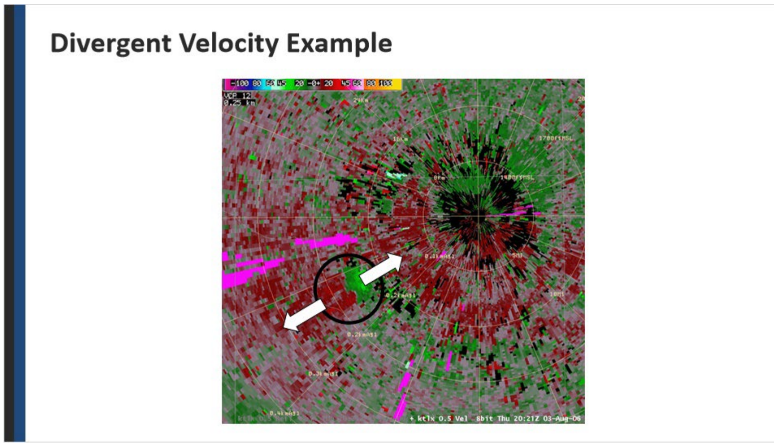

Notes:

In this real-world example, there is a divergence signature located to the west-southwest of the RDA. This occurred just after a downburst from a thunderstorm. Note that the maximum inbound velocity is closer to the RDA than the maximum outbound velocity.

Back to Top

Linear Convergence and Divergence

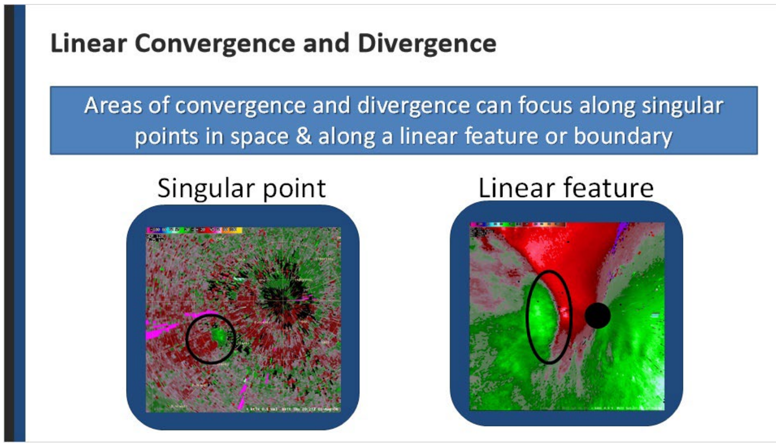

Notes:

The examples shown in the last few slides were of pure convergence and divergence on a single point in space. However, areas of convergence and divergence can also focus along a linear feature.

Back to Top

Linear Convergence

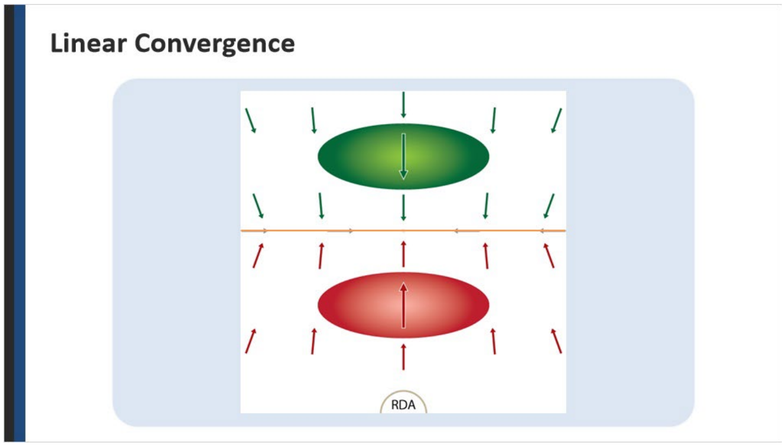

Notes:

Here is a basic diagram of a convergent velocity signature focused along a linear feature. As you can see here, both the maxima still lie along the same radial with the outbound velocity maximum closest to the RDA. The exception is that these maxima are elongated across a number of radials at about the same range from the RDA.

Back to Top

Linear Divergence

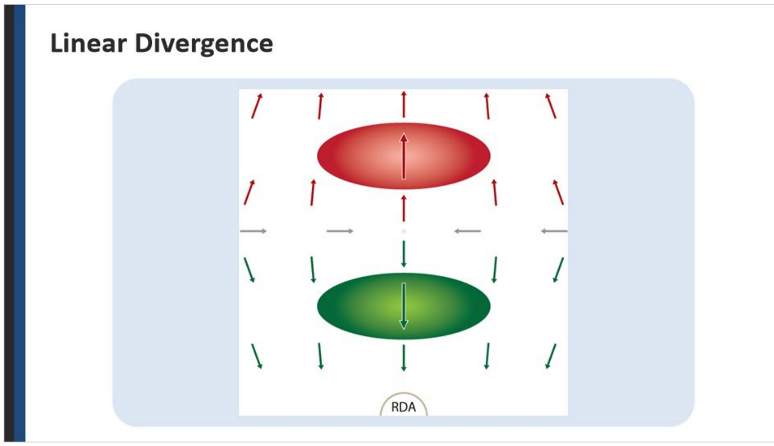

Notes:

Same thing here for the divergent velocity signature. Again, both maxima are elongated across a number of radials, but now the inbound maximum is closest to the RDA.

Back to Top

Linear Convergence Example

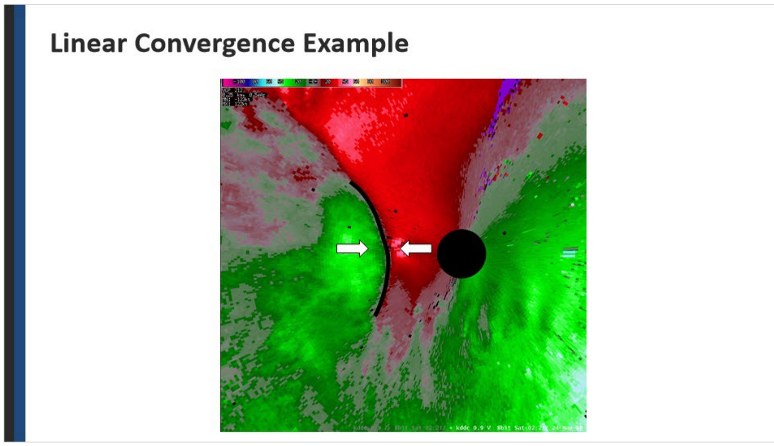

Notes:

Here is a real-world example of a linear storm-scale convergence signature. In this case, a QLCS is approaching the RDA from the west, and a segment of the line is bowing out at this point, creating an enhanced convergence signature. Areas of enhanced convergence along a line segment could lead to mesovortex formation, which can enhance the wind threat and increase the probabilities of a tornado in areas of vorticity that are generated.

Back to Top

Small Scale Pattern: Rotation

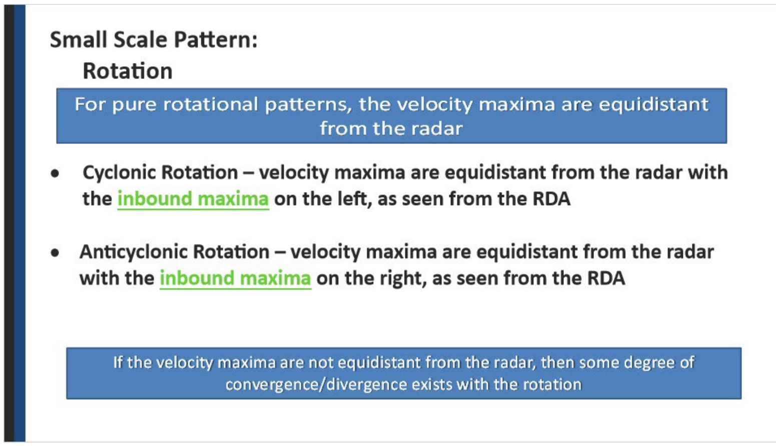

Notes:

So now let's consider times when there's a rotational component.

When examining pure rotational patterns, the velocity maxima are equidistant from the radar. Whether the pattern is cyclonic or anticyclonic is dependent upon whether the inbound maximum is on the left side or the right side of the signature, as seen by the RDA.

--With cyclonic rotation, the inbound maximum is on the left hand side, while the inbound maximum is on the right hand side with anticyclonic rotation.

--Velocity maxima oriented any other way means some combination of rotation and convergence or divergence is occurring. We'll talk more about that later in the lesson.

Back to Top

Cyclonic Rotation

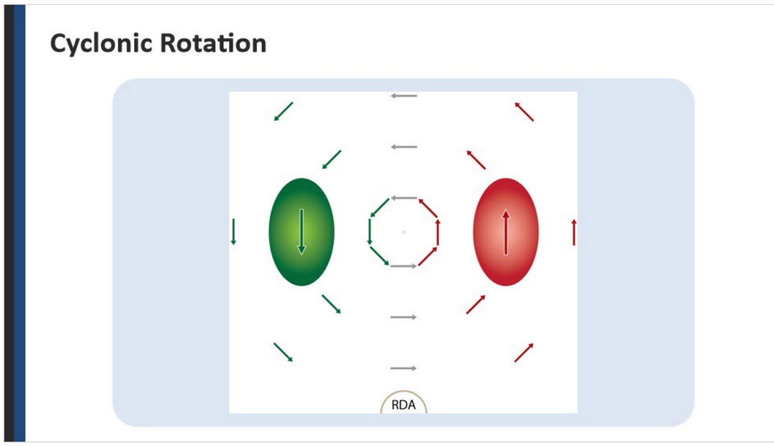

Notes:

Again, for the following examples, the RDA is located to the south of the velocity signature.

Here is a basic example of pure cyclonic rotation. As you can see here, both of the velocity maxima are equidistant from the RDA with the maximum inbound velocities on the left side of the signature.

Anticyclonic Rotation

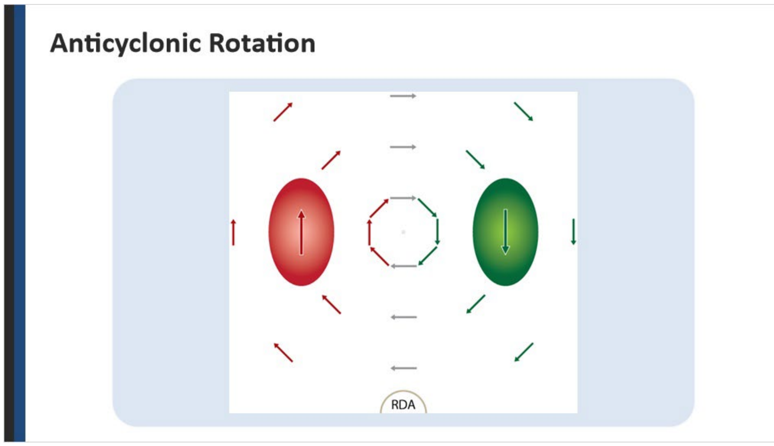

Notes:

Now with pure anticyclonic rotation, both velocity maxima are again equidistant from the radar and the inbound velocity maximum on the right side of the signature.

Back to Top

Right-Hand Rule for Cyclonic/Anticyclonic

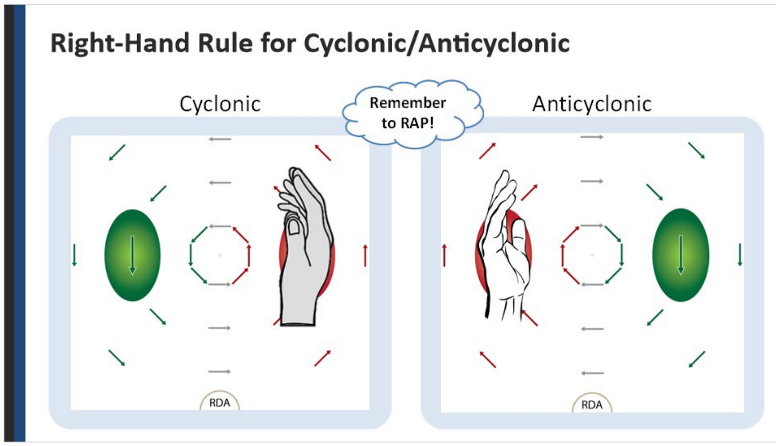

Notes:

Another way to remember cyclonic and anticyclonic (that's much easier for me) is to use the right-hand rule. First, I RAP to find the outbounds (red, away, positive). Then I align my right hand in that direction, with the base of my hand closest to the RDA and my fingertips pointing outbound. Then I curl my fingers in the direction of the inbounds (or green). When my thumb faces towards me, it's cyclonic. When it faces away from me (or into the screen), it's anticyclonic.

Back to Top

Cyclonic Convergence

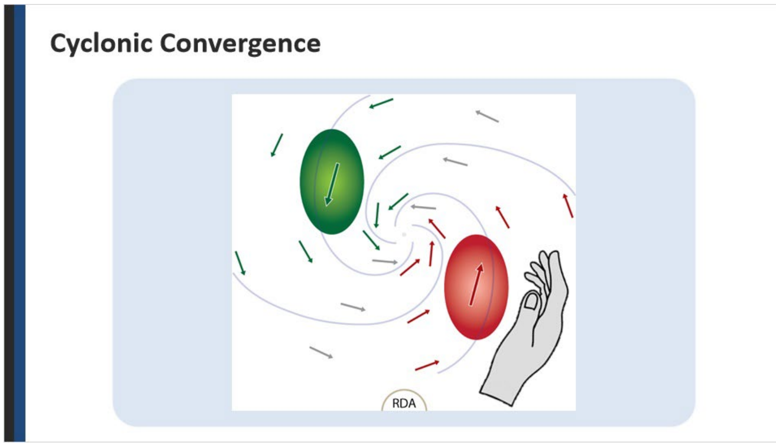

Notes:

So let’s put it all together and show combinations of both rotation and convergence or divergence. This is an example of cyclonic convergence. First note that both the maxima are not on the same radial and not equidistant from the radar. So it can’t be pure rotation.

Here, the outbound maximum is closest to the RDA, signifying convergence, and the inbound maximum is to the left, signifying cyclonic rotation. Similarly, the right hand rule would have your hand oriented like this, curling towards the inbounds, and giving you a thumbs up, also signifying cyclonic rotation.

Back to Top

Cyclonic Divergence

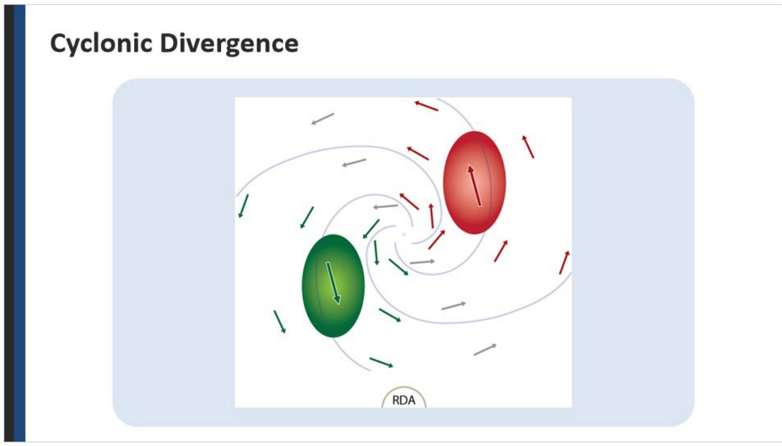

Notes:

Here is an example of cyclonic divergence. The inbound maximum is closest to the RDA, signifying divergence, and the inbound maximum is still to the left, signifying cyclonic rotation.

Back to Top

Anticyclonic Convergence

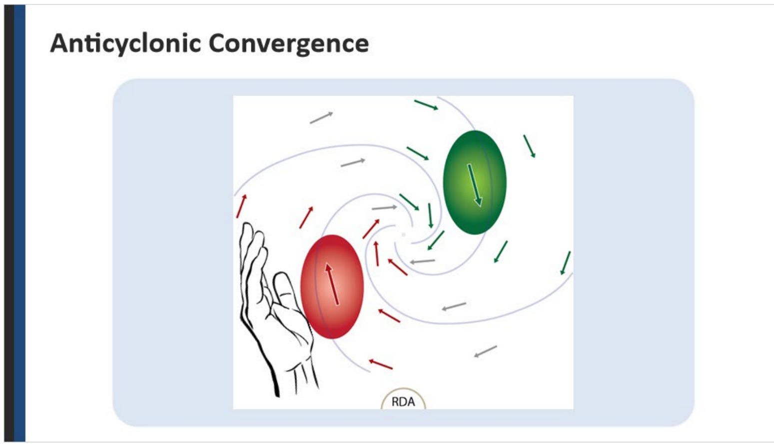

Notes:

Here is an example of anticyclonic convergence. The outbound maximum is closest to the RDA, signifying convergence, and the inbound maximum is to the right, signifying anticyclonic rotation.

Once again, the right hand rule would have your hand oriented like this, curling towards the inbounds, and giving you a thumbs down into the screen, signifying anticyclonic rotation.

Back to Top

Anticyclonic Divergence

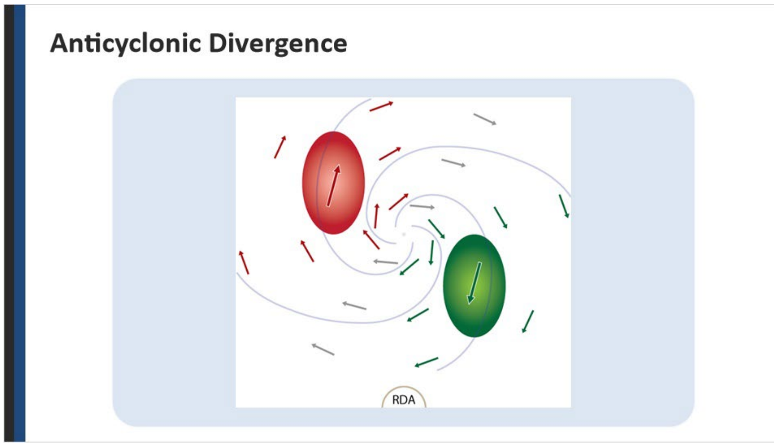

Notes:

Finally, here is an example of anticyclonic divergence. The inbound maximum is closest to the RDA, signifying divergence, and the inbound maximum is to the right, signifying anticyclonic rotation.

Back to Top

Storm Scale Rotation - Example

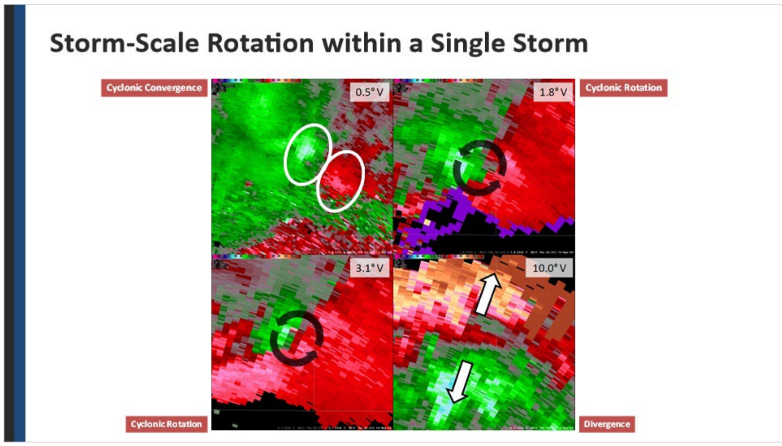

Notes:

Here is a real-world example of storm-scale signatures through various tilts of what would become a tornadic supercell. Using the range gate method, we can note that the RDA is located to the south-southwest of the storm.

--In the upper-left panel (0.5° tilt), you see a cyclonic convergence signature with the storm. The red outbound maximum is slightly closer to the RDA, signifying convergence, and the inbound maximum is to the left, signifying cyclonic rotation.

--The next two elevation scans up (1.8° tilt in the upper-right panel and 3.1° tilt in the lower-left panel) are close to “pure” cyclonic rotation, since the maxima are equidistant from the radar.

--Finally, the highest tilt (in the lower-right panel) is an example of storm-top divergence, with the maxima oriented along the same radial.

Move onto the next slide when you are ready.

Back to Top

Remember RDA location!

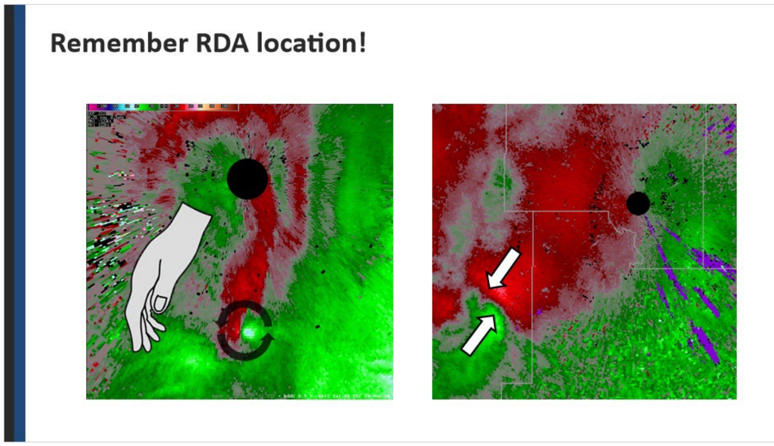

Notes:

One final note…Don’t forget to keep in mind where the RDA is located when identifying storm-scale signatures. Til now, we’ve only shown examples where the RDA is to the south. But in this example, the RDA is north of the storm-scale feature.

This is why I like the right-hand rule, because you don’t have to remember left and right. Simply find the outbound maximum, orient your right hand such that your wrist is closest to the radar, and curl your fingers towards the inbound maximum. In this example, the thumb is pointing towards you, so it’s cyclonic rotation. Since the maxima are equidistant from the radar, there’s no convergent or divergent signatures.

But this example on the right shows cyclonic convergence because the outbound maximum is closer to the RDA.

Back to Top

Summary Table

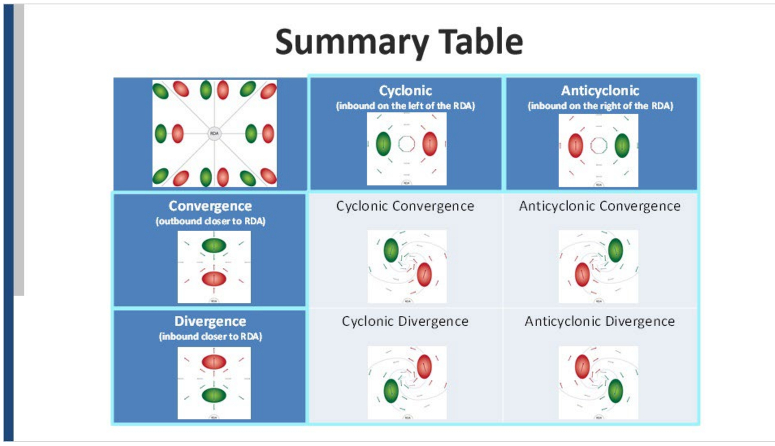

Notes:

Let’s go ahead and summarize storm-scale velocity signatures. Convergence signatures have the velocity maxima lie along the same radial with the outbound maximum closest to the radar. Divergence signatures have the velocity maxima lie along the same radial with the inbound maximum closest to the radar.

Cyclonic rotation signatures have the velocity maxima equidistant from the radar with the inbound maximum to the left, as seen from the radar. Anticyclonic rotation signatures have the velocity maxima equidistant from the radar with the inbound maximum to the right, as seen from the radar.

Velocity maxima oriented any other way means some combination of rotation and convergence or divergence is occurring.

Move onto the next slide when you are ready.

Back to Top

You are done with this topic!

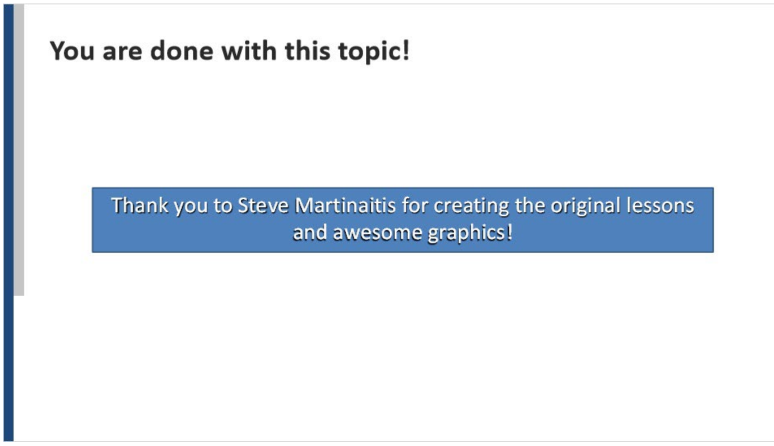

Contact Info

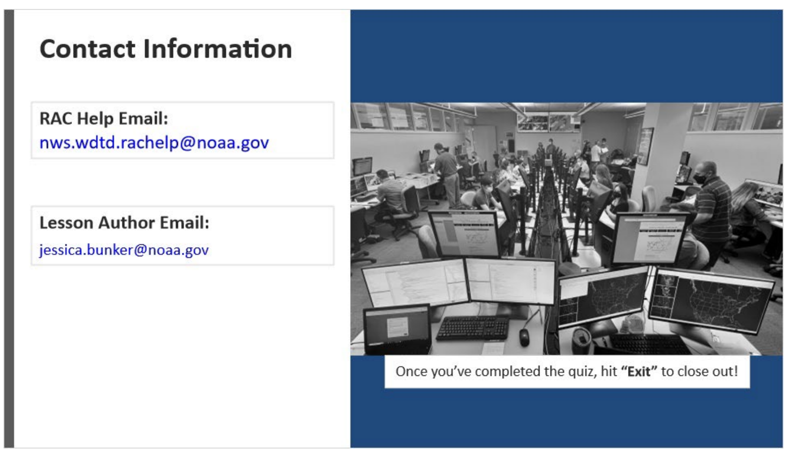

Notes:

If you have any questions or if you would like more information on this topic, please contact the emails listed here.

Back to Top

Identifying Elevation Sidelobe Contamination

1.1 Introduction

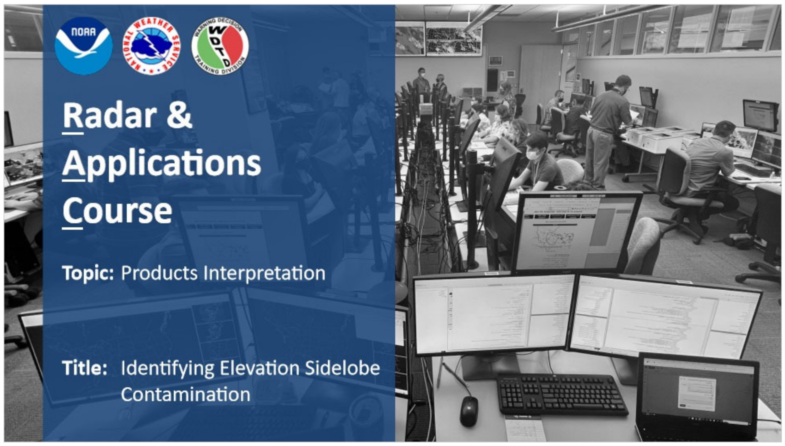

Notes:

Welcome to this lesson on Identifying Elevation Sidelobe Contamination. Let's get started.

1.3 Motivation for This Training

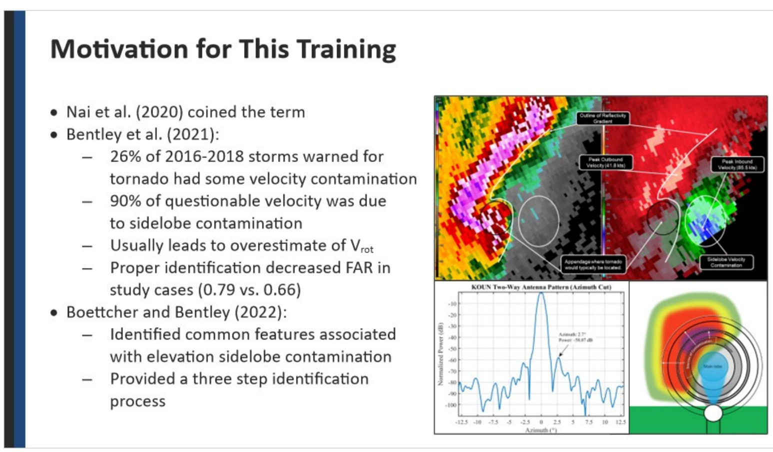

Notes:

If you have been taking the full (or Hydro Track) of the Radar & Applications Course, you know there's a lesson dedicate to data quality issues. So, why does elevation sidelobe contamination have it's own lesson? Well, this topic has grown in visibility over the last several years. Nai et al. (2020) coined the term elevation sidelobe contamination to better describe what was previously called vertical sidelobe contamination. Bentley et al. (2021) found 26% of low-level couplets from 2016 to 2018 had some velocity contamination, and approximately 90% of that was due to sidelobe contamination. When elevation sidelobe contamination occurs, it generally leads to an overestimate of low-level rotation velocity (or Vrot). Their study found that properly identifying elevation sidelobe contamination in those cases could have led to a reduction of False Alarm Ratio during those years from 0.79 to 0.66. Likewise, Boettcher and Bentley (2022) built off of the initial elevation sidelobe contamination work and identified common features associated with these artifacts. They also provided a three step identification process to help operational forecasters. Armed with this information, hopefully you see how this issue deserves a stand alone lesson.

Back to Top

1.4 Reminder: What is a Sidelobe?

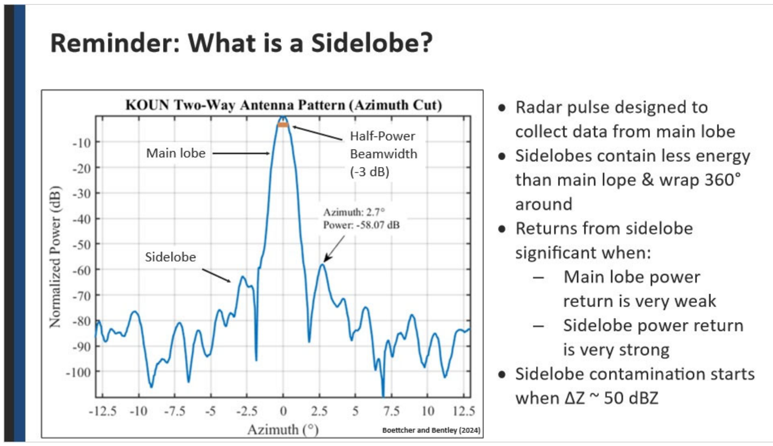

Notes:

Before we go any further, let's do a quick review on what is a sidelobe. The image on the slide shows a sample wave form from a WSR-88D. Weather radars are designed to transmit a radar pulse that collects data from the main lobe. The main lobe is centered on an azimuth of 0 degrees and has a normalized power of 0 dB in the graph. The other local maxima in the graph are the sidelobes. The first sidelobe peak in power around 2.5 degrees of azimuth away from the center of the main lobe, wraps 360 degrees around the main lobe, and contain substantially less energy. However, this sidelobe can interact with targets and return a significant amount of energy when the radar pulse in the main lobe has very weak power returns and the sidelobe power returns are very strong. When the absolute, cumulative power return for all of the sidelobes is around 50 dBZ higher than what is in the main lobe, then sidelobe contamination starts to be visible.

1.5 How Do Sidelobes Impact Radar Data Collection?

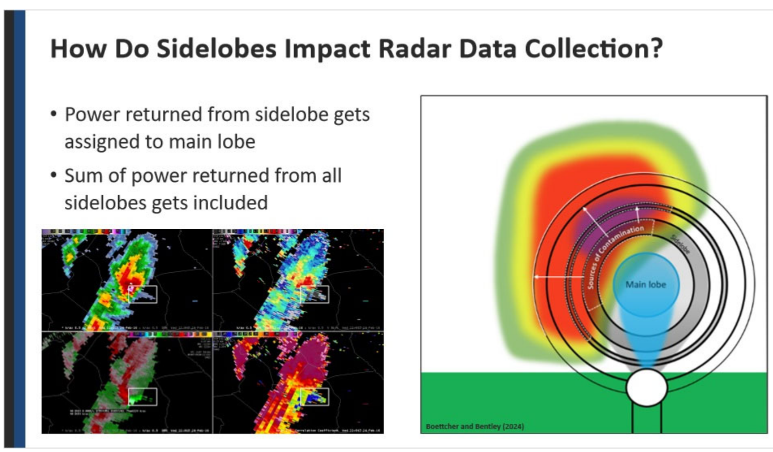

Notes:

When sidelobe contamination becomes visible, it happens because power return from the sidelobe gets assigned to the mainlobe. In other words, the radar thinks these strong targets (which could be weather, terrain, or anthropogenic in nature) are actually located where the main lobe is pointing. The image on the right-hand slide from Boettcher and Bentley (2024) shows a conceptual model of what this might look like in reality while the four-panel in the lower left shows how it might appear in your radar products. The sum of all the power returned in the sidelobe (labeled as the sources of contamination in the graphic), which is often quite weak, will be assigned to the range gate where the main lobe is sampling. The result in your radar data is often a region (or even a spike) of weak reflectivity values, negative differential reflectivity, very low correlation coefficient, and velocity values that may or may not make meteorological sense on first glance. Because the velocity data may not be obviously bad at first glance, it results in what looks like a velocity couplet, but is actually an imposter.

Back to Top

1.6 Sidelobe Contamination Example

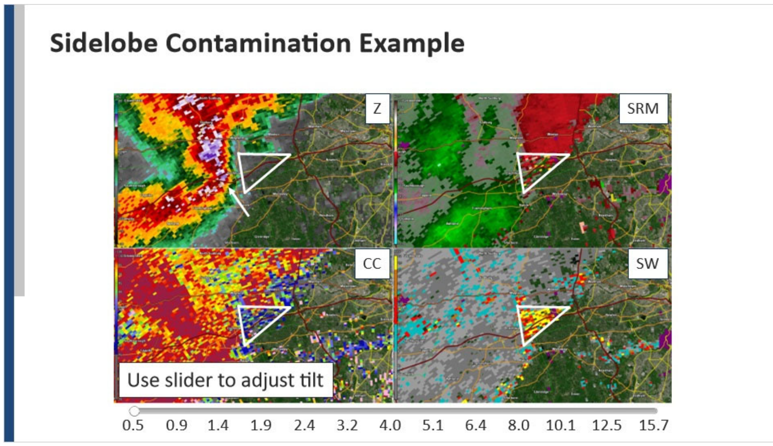

Notes:

Here's an example mentioned by Boettcher and Bentley (2022) of elevation sidelobe contamination on the 0.5 tilt indicated by the white triangle in each product. The 0.5 degree tilt I include was a SAILS tilt collected between the 3.2 and 4.0 scans, and you can use the slider bar on the bottom of the slide to view each tilt. Forecasters commonly make a mistake when diagnosing elevation sidelobes by only focusing on the tilts a specific tilt away from 0.5 (usually about 3 degrees) and directly above the affected area. That's one reason for using the newer term elevation (and not vertical) sidelobe contamination in this lesson. This particular example likely shows sidelobe contamination that is cumulative. So, it's not only from directly above, but also from the strong reflectivity returns just to the west (and at the same range) of the highlighted area on all of the tilts at and above 0.5 degrees (as indicated by the white arrows). It should be noted that some of the contamination may also be coming from sidelobe power returned between 0.5 degrees and the surface. We will explain why that is important later in the lesson.

15.7 (Slide Layer)

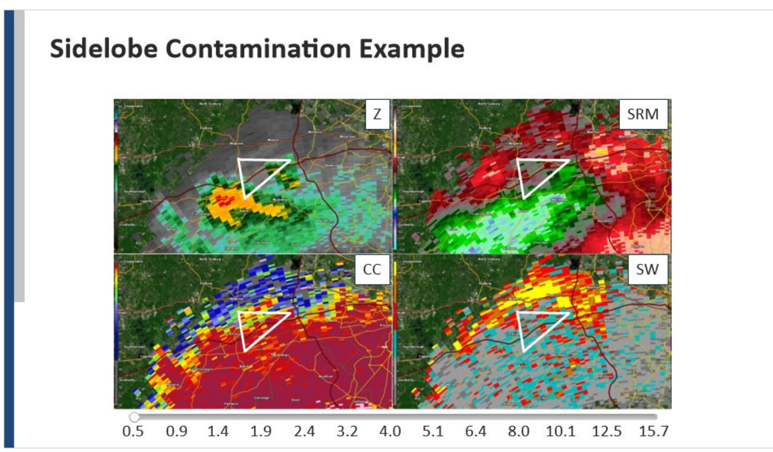

Back to Top

12.5 (Slide Layer)

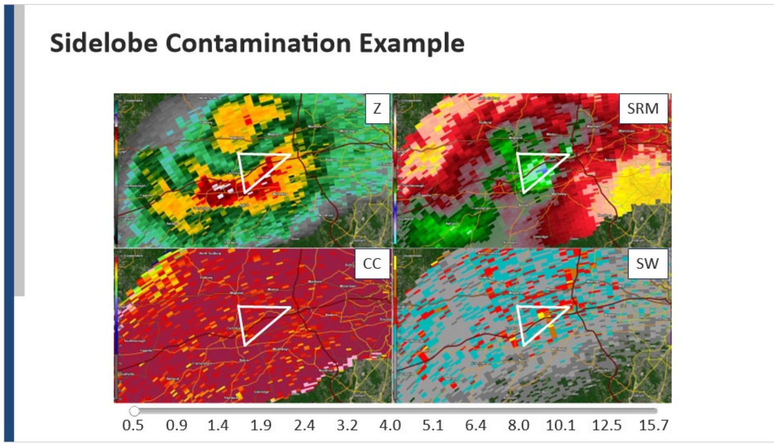

10.1 (Slide Layer)

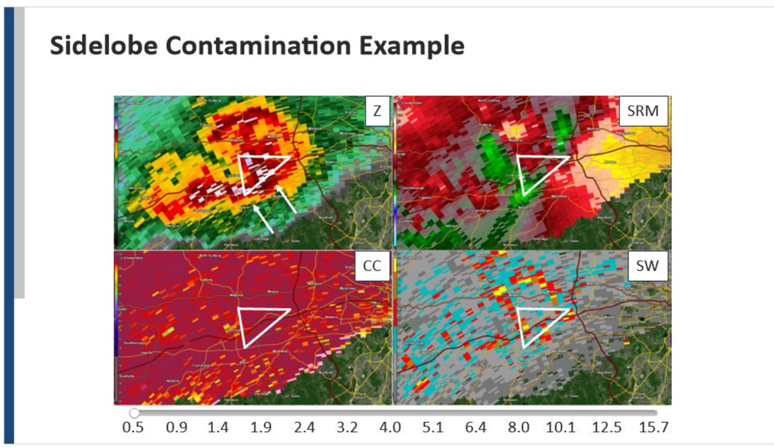

8.0 (Slide Layer)

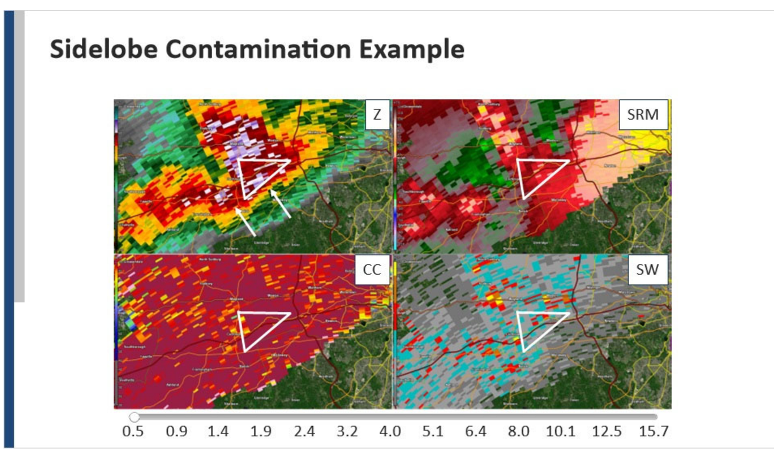

Back to Top

6.4 (Slide Layer)

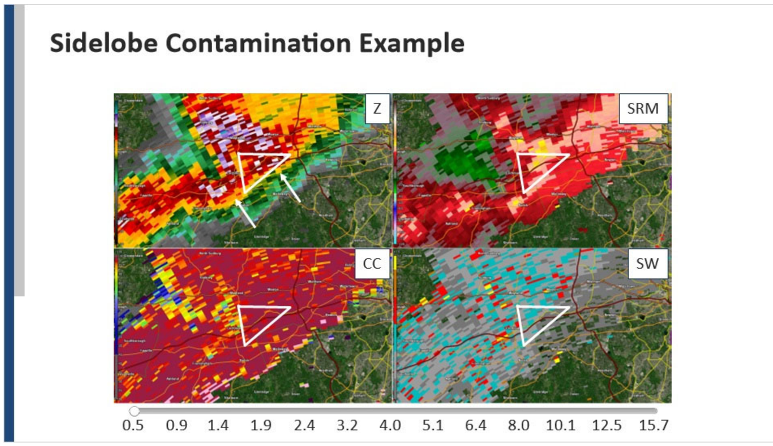

5.1 (Slide Layer)

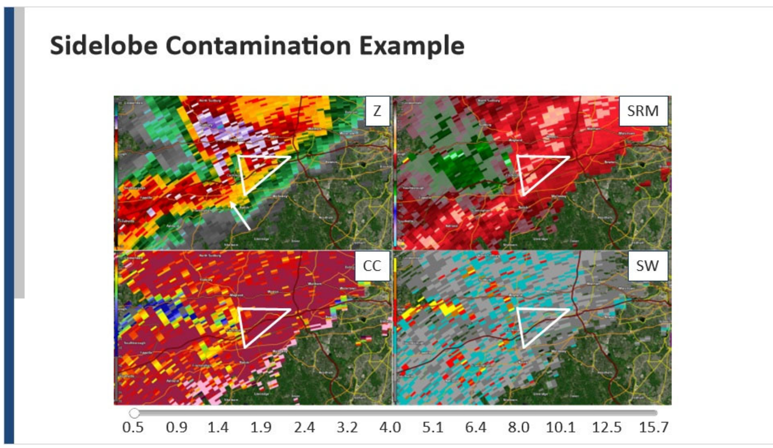

4.0 (Slide Layer)

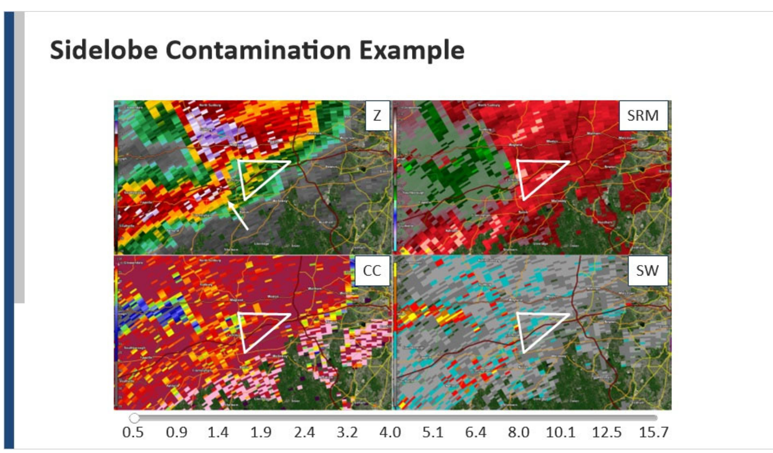

Back to Top

3.2 (Slide Layer)

Sidelobe Contamination Example

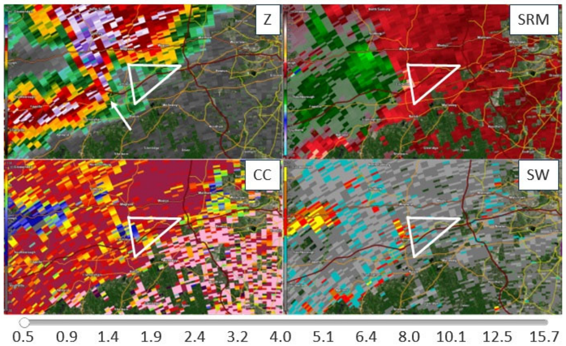

2.4 (Slide Layer)

Sidelobe Contamination Example

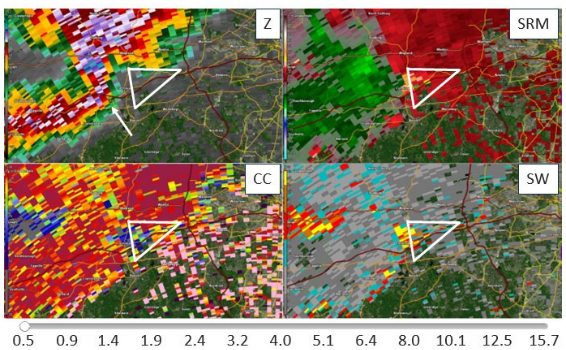

1.9 (Slide Layer)

Sidelobe Contamination Example

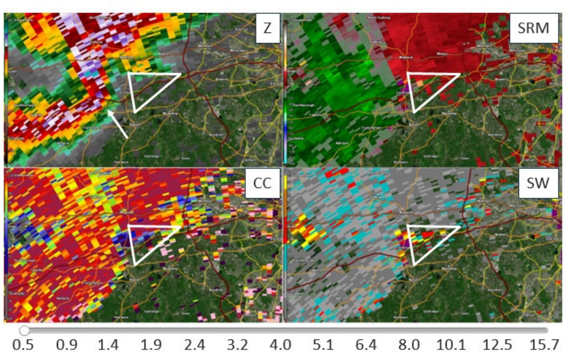

Back to Top

1.4 (Slide Layer)

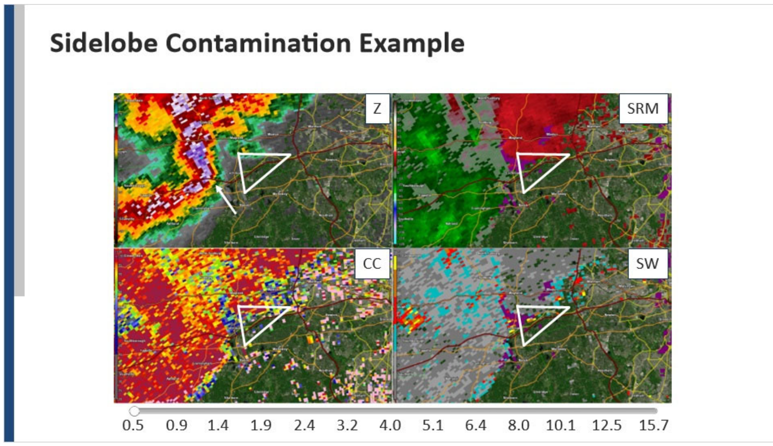

0.9 (Slide Layer)

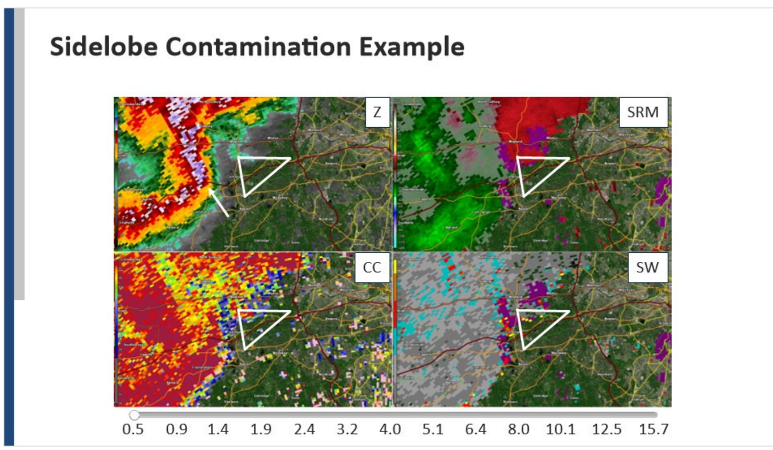

1.7 Key Goal: Identify Imposters

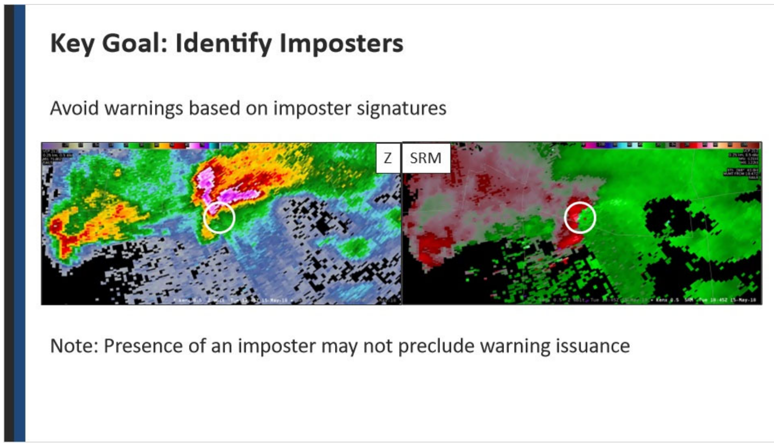

Notes:

Now that you have seen some examples, you should have a better understanding about what we mean about elevation sidelobe contamination. We highlight this issue in order to achieve an important goal: Identifying velocity contamination that may result in imposter velocity couplets. Identifying imposter circulations has critical importance to reducing Tornado Warning issuance on storms that would otherwise not warrant it. Before we discuss specifics on how to identify these imposters in the remainder of this lesson, I need to point out an important point. Just because a storm contains an imposter circulation, that shouldn't prevent you from issuing a warning if the remaining evidence indicates a tornado is likely present or imminent. You still need to follow the preponderance of the evidence and make the best scientifically defensible decision you can.

1.8 Common Feature #1: Blocky Velocity

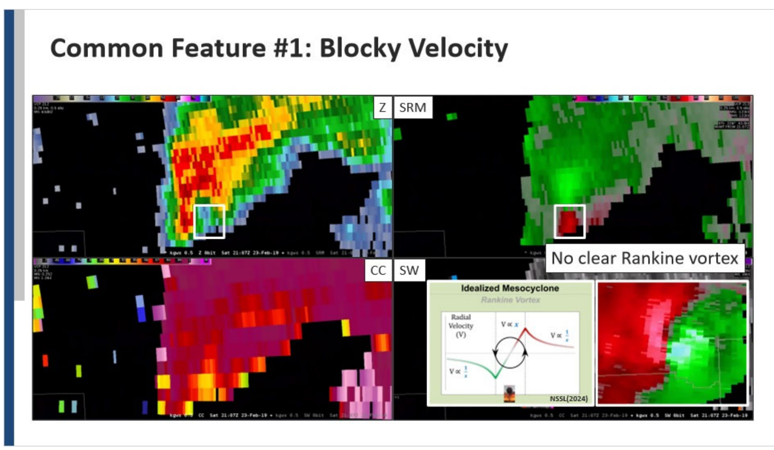

Notes:

For the first few features, we will focus on the velocity data. In this example, we see an apparent velocity couplet in SRM. However, a quick look suggests a problem in the outbound velocities on the south side of couplet. See how the velocities look blocky with no distinguishable velocity gradient or maximum near the center of the circulation. Typically, velocity data in a mesocyclone appears as a Rankine vortex like in the graphics I show on the slide. What we typically see in an authentic mesocyclone is a gradual increase in velocities as we approach the center of the circulation with a rapid switch in direction as you move across the center of circulation.

Back to Top

1.9 Common Feature #2: Unrealistic Shear

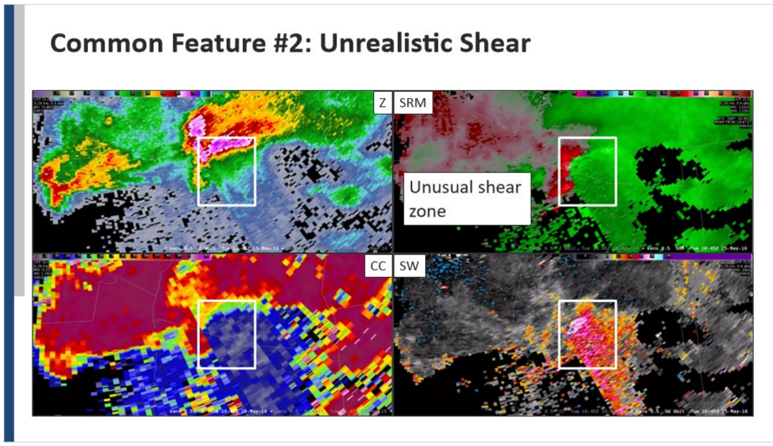

Notes:

Another questionable aspect of this potential circulation is the unrealistic region of azimuthal shear. The shear region in this box extends to approximately 3 nautical miles in length. Legitimate mesocyclonic shear rarely gets that large. This particular example has unrealistic cyclonic shear and some questionably strong convergence along the forward flank due to a combination of elevation sidelobe contamination and a three-body scatter spike over a portion of the contamination region.

1.10 Common Feature #3: Wrong Location

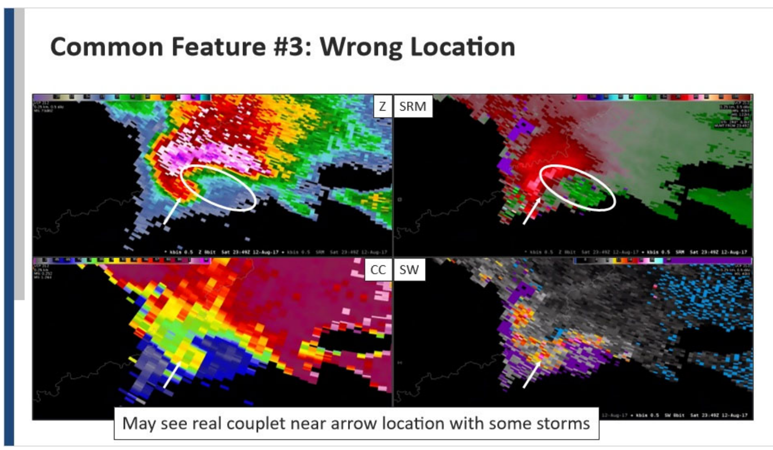

Notes:

Another quirk with imposter circulations is their location. The oval in the graphic highlights an apparent circulation in the velocity data. This location puts the circulation more along the storms forward flank than near the hook like appendage in Reflectivity indicated by the arrow further to the south. While this may seem a subtle difference, it’s an important one. In some storms with a potential imposter like this one, you will notice a second circulation in the area where you would expect it. Being able to spot the difference helps you gauge the true rotational velocity for the storm and not get sucked in by the imposter.

Back to Top

1.11 Common Feature #4: Unbalanced Vrot

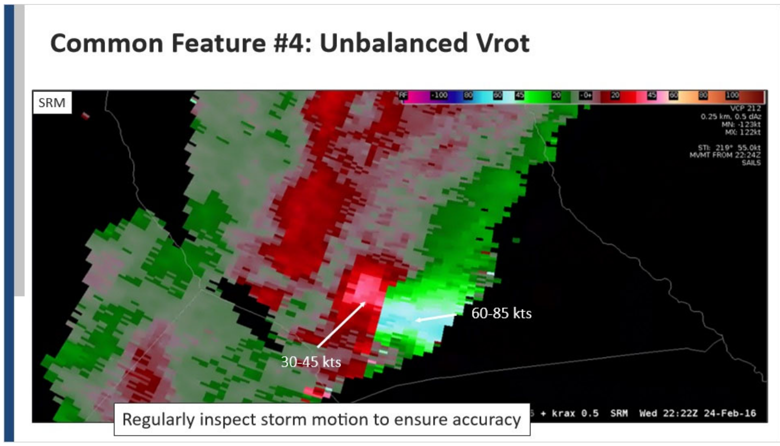

Notes:

Let's go back to the SRM data for the next feature common with imposter circulations. Assuming you have a good storm motion, most low-level mesocyclones in SRM will have relatively balanced velocity couplet maxima. I'm not saying they will be the same, but the difference will be relatively small. In this circulation, the difference in the maxima are generally 30 kts different, but as much as 50 kts different in the broader mesocyclonic circulation. Identifying this feature depends on the radar operator using a good storm motion. Consider this a reminder to regularly inspect your SRM storm motion when investigating couplets.

1.12 Common Feature #5: At Lower Levels

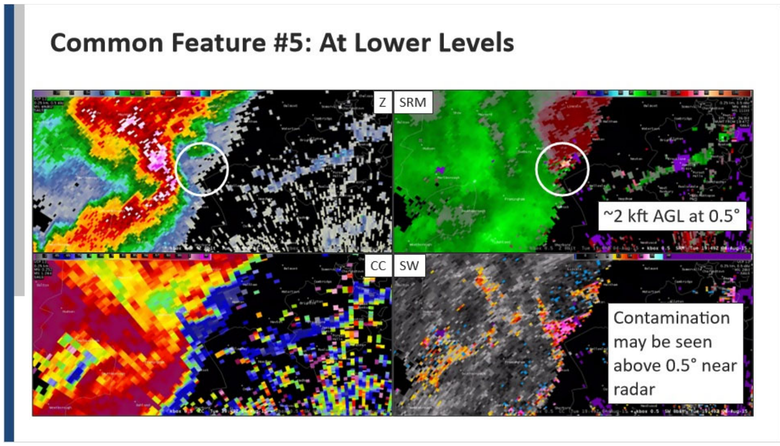

Notes:

Velocity couplet imposters tend to be seen in the lower parts of the storm in the weak echo region underneath the midlevel reflectivity overhang. For classic supercells, you will usually see them in the lowest 6 km AGL of a storm. For mini-supercells, that height would be even lower. In other, more intense, deeper storms, you might see velocity contamination even higher in the storm. In other words, typical heights don't mean a hard and fast rule.

Back to Top

1.13 Common Feature #6: Weaker Returns

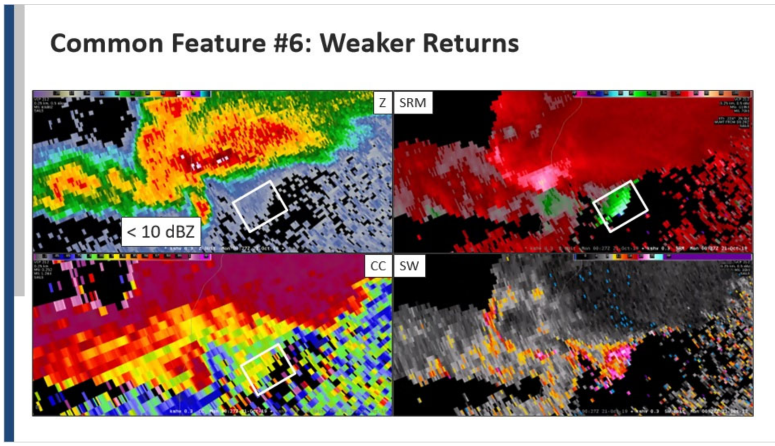

Notes:

The last common feature is that the contaminated velocity values tend to be found in areas with weaker reflectivity returns. Expect most Reflectivity values to be 20-25 dBZ or less with corresponding CC values below 0.9 (suggesting non-meteorological echoes). In this particular example, the CC values are a little higher than what you may typically see in comparable situations. In this case, the elevation sidelobe returns likely dominate the signal in parts of the highlighted area, resulting in CC values more comparable to Mie scattering from larger hail than from typical non-meteorological returns.

1.14 Other Potential Features #1: Noisy V/SRM & Spectrum Width

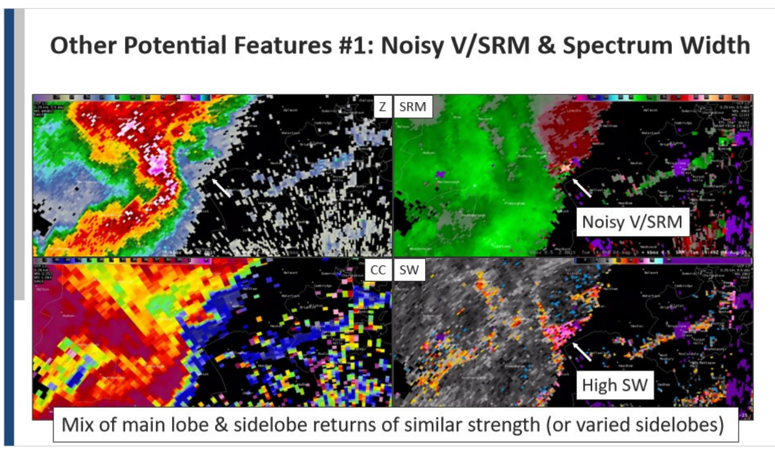

Notes:

You may see some other potential features with elevation sidelobe contamination depending on how the main lobe and sidelobes contribute to the returned signal. In one scenario, you might see noisy velocity data associated with high Spectrum Width values. When this happens, the returned power likely contains a fairly even mix of power return from the main and side lobes. The same scenario might occur if the sidelobes sample different parts of the storm with significantly different velocity values (or if you get a mix of TBSS and sidelobe velocity contamination).

Back to Top

1.15 Other Potential Features #2: Smooth V/SRM & Spectrum Width

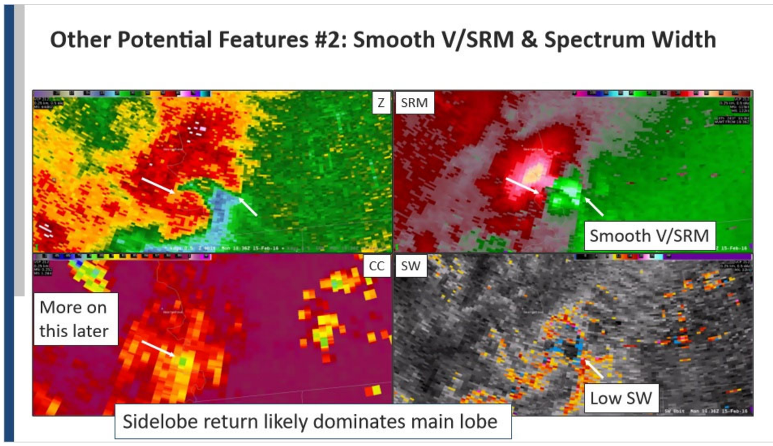

Notes:

You could also observe fairly smooth velocity data associated with low Spectrum Width values in an imposter circulation like the inbound velocities highlighted here. These situations occur when the side lobe return likely dominates the signal from the main lobe and the sampled velocity is fairly uniform. Before I move on, I need to point out one more thing in this particular example. While the area I highlighted shows a clear imposter, another area to the southwest appears to be part of a legitimate couplet. Even though a clear imposter circulation exists, this storm still warrants a Tornado Warning due to the remaining features visible in this storm.

1.16 Imposter Identification Process

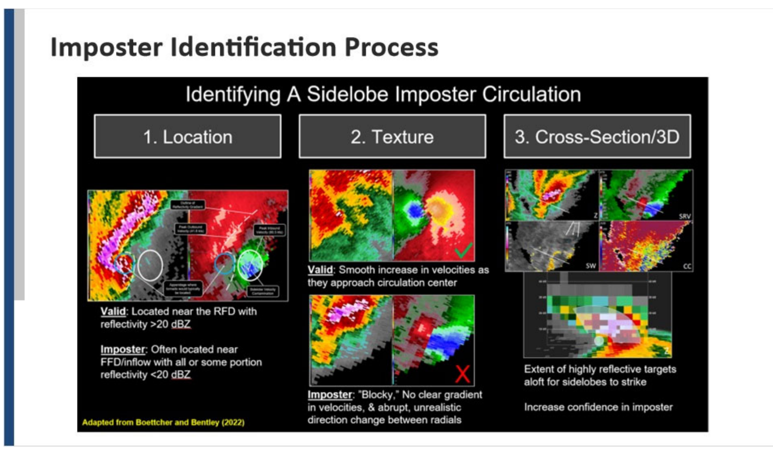

Notes:

At this point, we have explained many details on elevation sidelobe contamination. Before we wrap things up, I want to walk through the methodology that Boettcher and Bentley (2022) proposed in their study and shown on this slide. When forecasters interrogate a storm with a circulation that might be an imposter, they should focus on three steps to confirm that the circulation likely is an imposter.

Back to Top

1.17 Step #1: Circulation Location

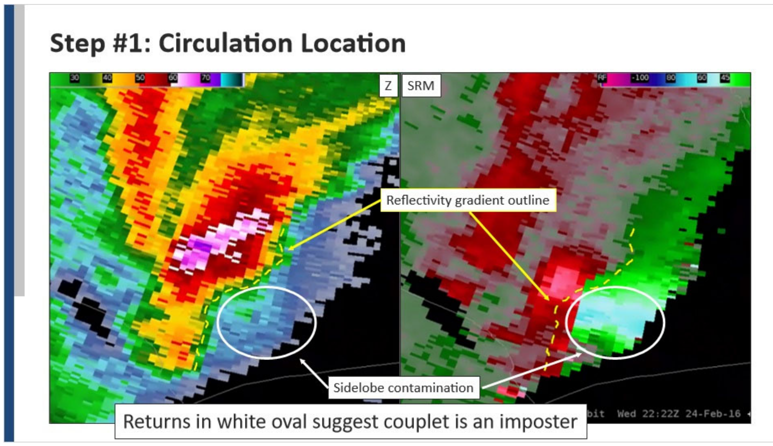

Notes:

First, look at the potential circulation's location to identify if it is an imposter. Specifically, is the circulation located near the RFD and/or hook echo? Or is it located more in the forward flank and with a significant amount in weak echo returns like the area in the white oval in the image? The latter location would suggest the circulation is an imposter.

1.18 Step #2: Velocity Texture

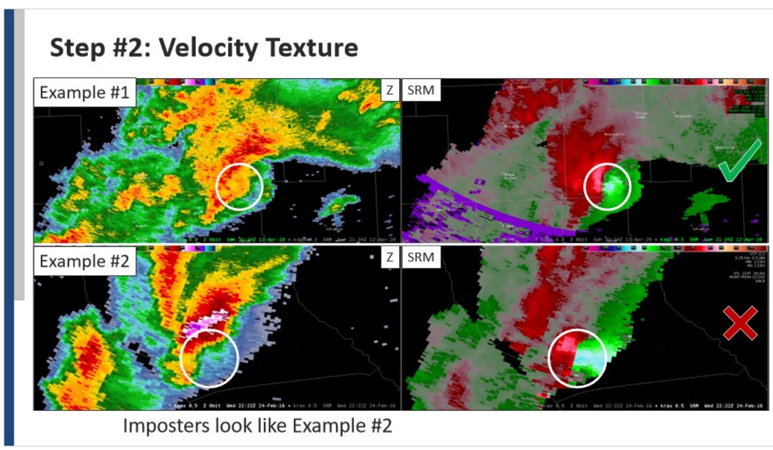

Notes:

The second step involves looking at the velocity texture in the potential imposter. Real circulations (like the example marked by the green check mark) generally have a Rankine vortex appearance with a smooth increase in velocities to well-defined velocity maxima. Imposter circulations (like the one marked by the red X) have a blockier appearance, with no clear gradient in velocities visible.

Back to Top

1.19 Step #3: Vertical Examination

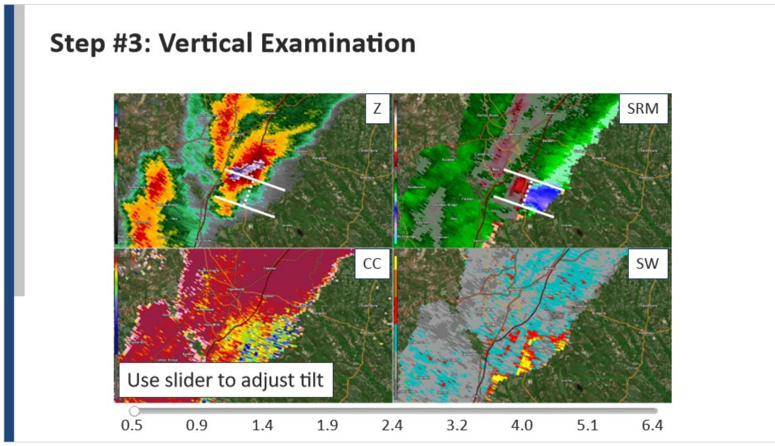

Notes:

In the last step, you want to look at the vertical structure of the storm to identify highly reflective targets that could be struck by sidelobes. Remember that those sidelobe targets need to be located at the same range as the main lobe, so it's more of an arc than a true line normal to the beam. Ideally, you would use a vertical cross-section oriented perpendicular to the radar beam to perform this task. However, radar cross-section tools in NWS operations have some significant limitations at the moment. As a result, the better option is to analyze the radar data in an all-tilts display to look for strongly reflective cores. I have added some annotations on the slide to help on the pertinent tilts. It's the blue, inbound radial velocities on the 0.5 degree tilt that appear to be corrupted. Looking aloft, you can see several tilts have strong reflectivity values perpendicular to the beam just to the west of the impacted area. I need to make one more important point. The velocity values in the sidelobe contaminated area do not need to match your possible source region. It's likely that multiple areas in the storm (including regions below the lowest tilt) are contributing to the contamination each with their own velocity values. The multiple sources could cause the radar to estimate a velocity different than what is seen in any one source region.

6.4 (Slide Layer)

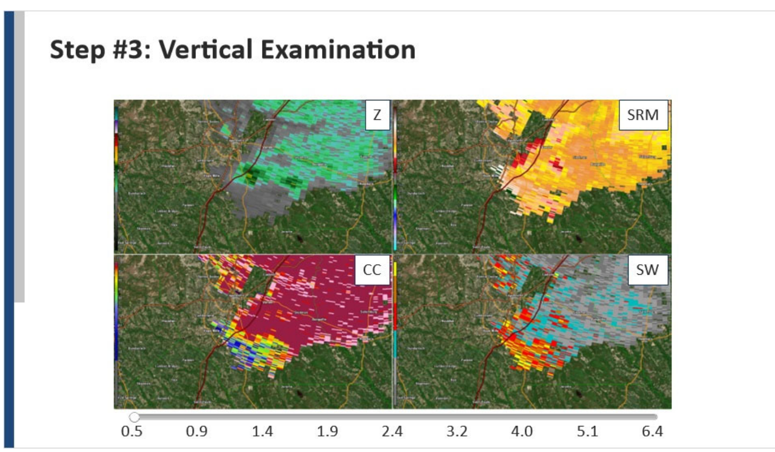

Back to Top

5.1 (Slide Layer)

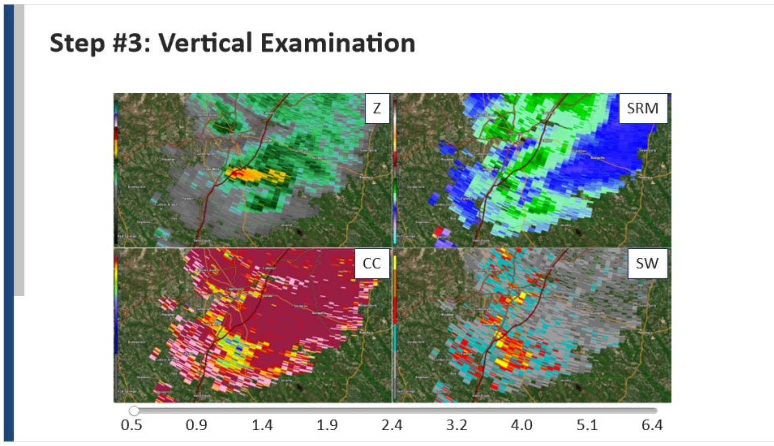

4.0 (Slide Layer)

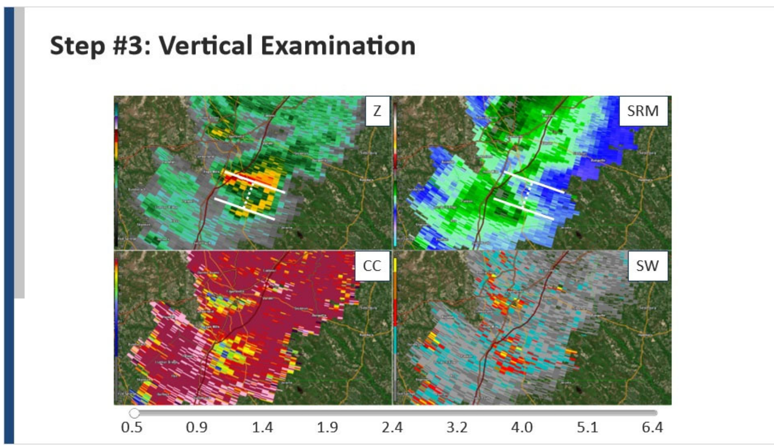

3.2 (Slide Layer)

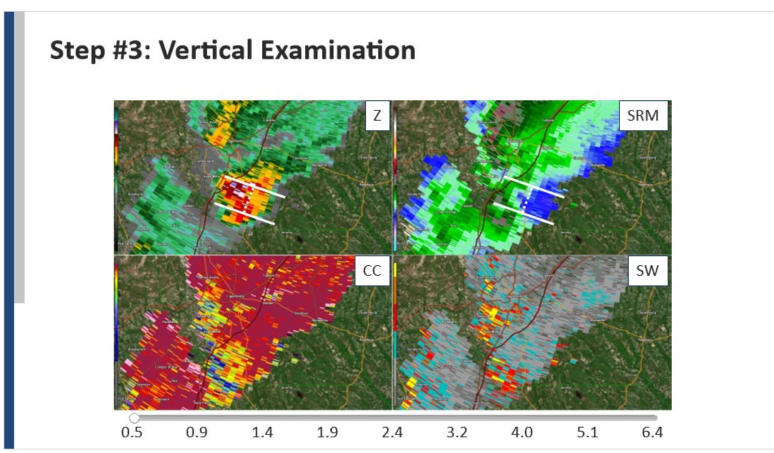

Back to Top

2.4 (Slide Layer)

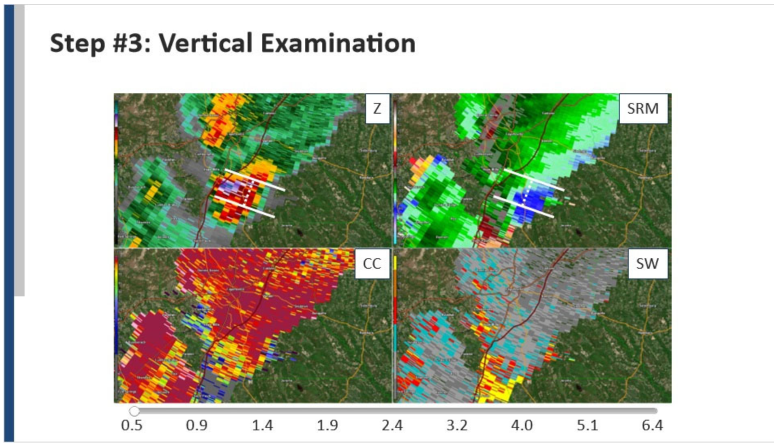

1.9 (Slide Layer)

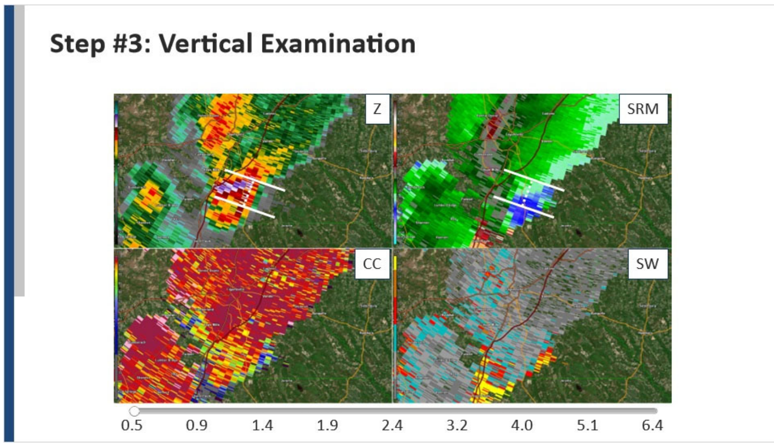

1.4 (Slide Layer)

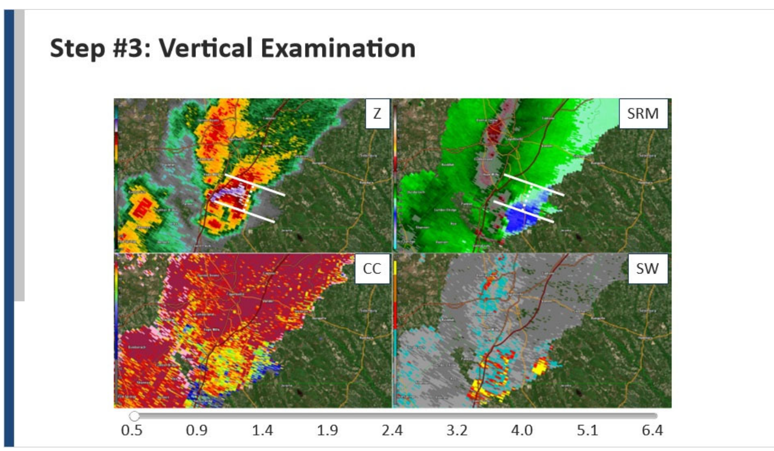

Back to Top

0.9 (Slide Layer)

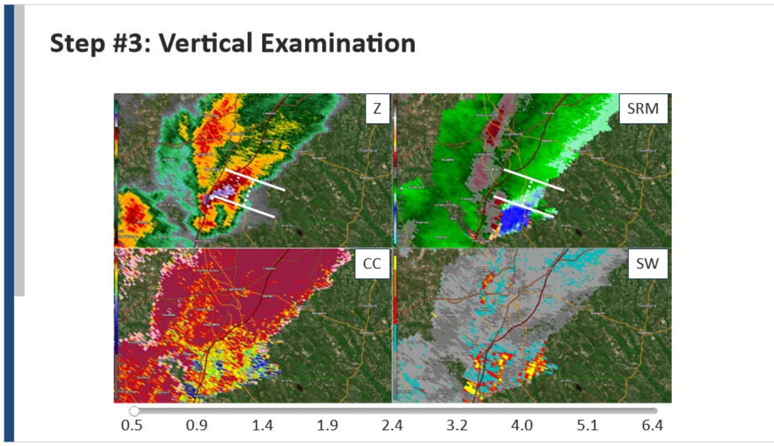

1.20 End Result: Ignore the Bad, Lean into the Good, Make Best Decision You Can

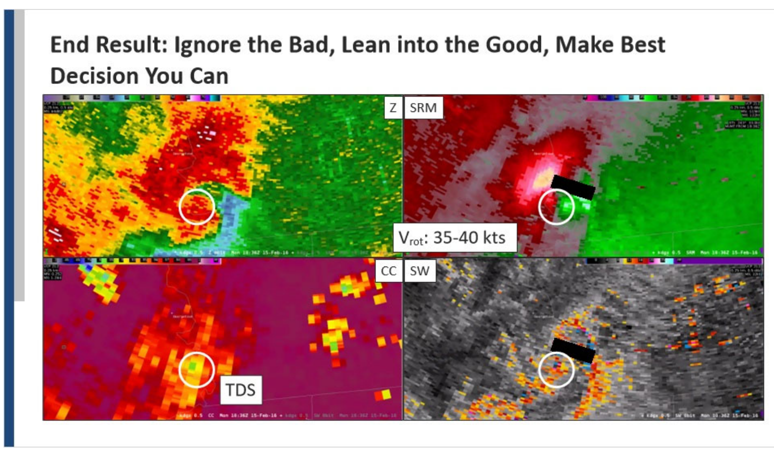

Notes:

So, what is the end result of this process? You should use the three step process to identify a region (or regions) of questionable data. Once you know where that is, you can mentally block it out (as I have visually done on the slide). That will allow you ignore the data you know is bad and the lean into the good data that remains. In this case, blocking out that bad data helps us see that we actually have a circulation visible in the hook echo with 35-40 kts of rotation velocity and a co-located tornado debris signature. In other words, we make the best decision we can with the data that remains.

Back to Top

1.21 Summary

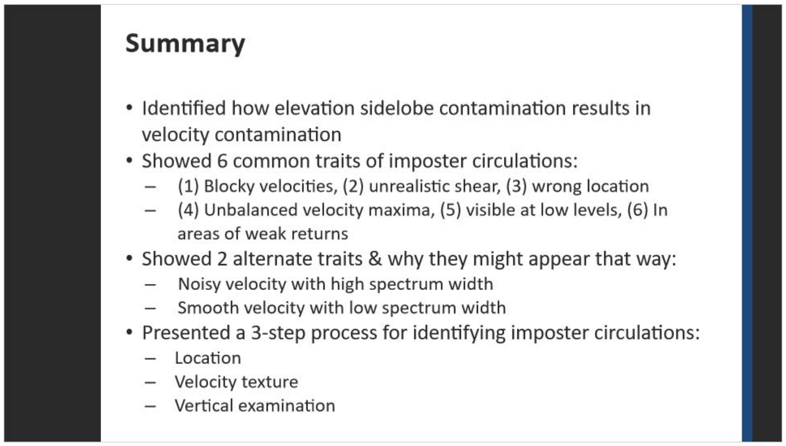

Notes:

In summary, we discussed various aspects of the causes of and ways to identify elevation sidelobe contamination artifacts. When these artifacts result in a potential imposter circulation, we discussed 6 common traits the imposter might have. These traits include blocky velocities, unrealistic shear, wrong location, unbalanced velocity maxima, seen at lower levels, and in weak returns. In some cases, velocity data may be noisy with high spectrum width values if the main lobe and sidelobe returns are similar. In other cases, the velocity data may be smooth with lower spectrum width values if the sidelobe contamination dominates the signal. Forecasters can follow a simple three step process to identify potential imposter circulations. First, make sure the circulation is located in the correct location. Second, examine the velocity texture to make sure it appears realistic. Third, examine the vertical storm structure to see if there's highly reflective sidelobe returns near and above the low-level storm inflow.

When you are ready, please proceed to the next slide to start the quiz.

1.27 Contact Information

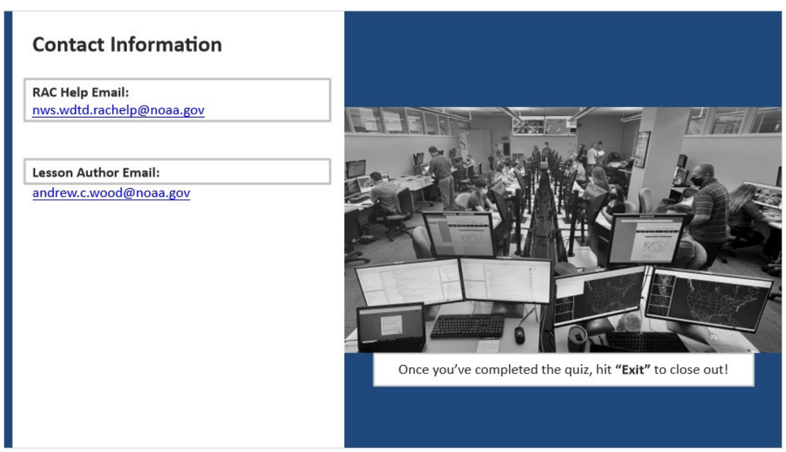

Notes:

Thank you for completing this lesson. If you have any questions, please use the e-mail addresses on the slide to contact us. Use the Exit button in the upper right-hand side of the lesson player to close the window.

Back to Top

Impacts of Wind Farm Contamination

1.1 Introduction

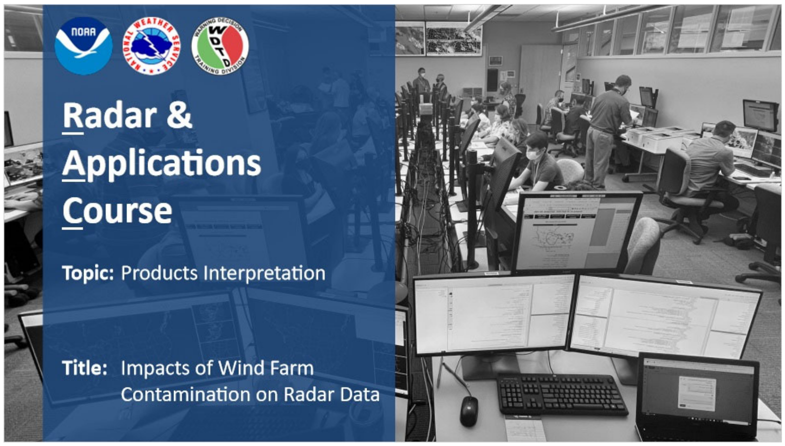

Notes:

Welcome to this lesson on the Impacts of Wind Farm Contamination on Radar Data. Let's get started!

1.3 Motivation for This Lesson

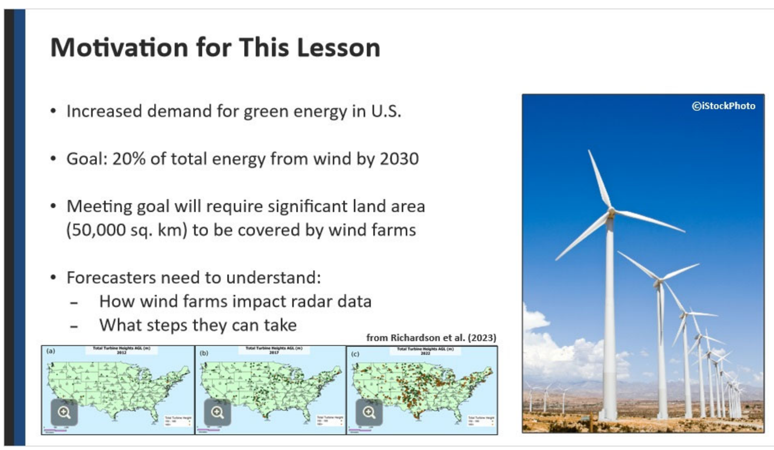

Notes:

As various utilities, municipalities, and other seek out alternatives to fossil fuels, there has been an increased demand for green energy in the United States. A goal was set over a decade ago to increase wind energy production to 20% of total energy by 2030 (Dept. of Energy, 2008), and the U.S. wind production has increased dramatically since then. To meet that goal, some analysts believe as much as 50,000 square km of land mass will need to be covered by wind farms in order to meet that goal (Toth et al, 2011). Regardless of whether that happens or not, wind farms are proliferating in desirable locations at a rapid rate. These images on the bottom of the screen show how the number of turbines exceeding 150 m has proliferated dramatically from 2012 to 2022. NWS forecasters need to be aware of how these wind farms impact radar data quality and what steps, if any, they can take to address these data quality issues.

Back to Top

1.4 Relevant Attributes of Wind Turbines/Farms

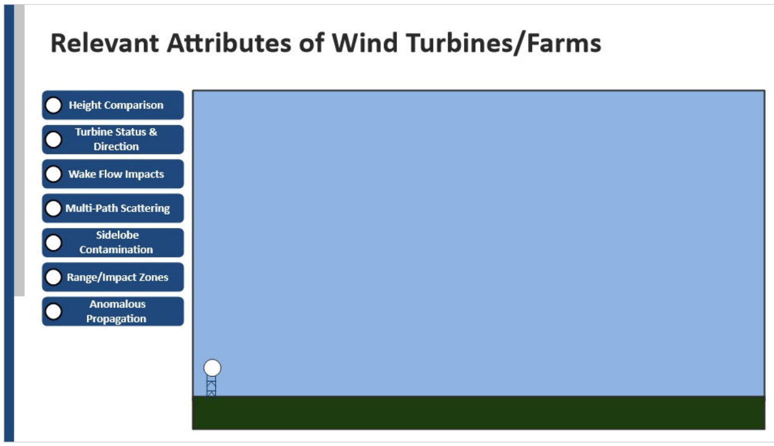

Notes:

Let's start this conversation by discussing the relevant attributes of wind turbines and farms that can impact radar data quality. Choose the buttons on the left to learn more about each topic.

Anomalous Propagation (Slide Layer)

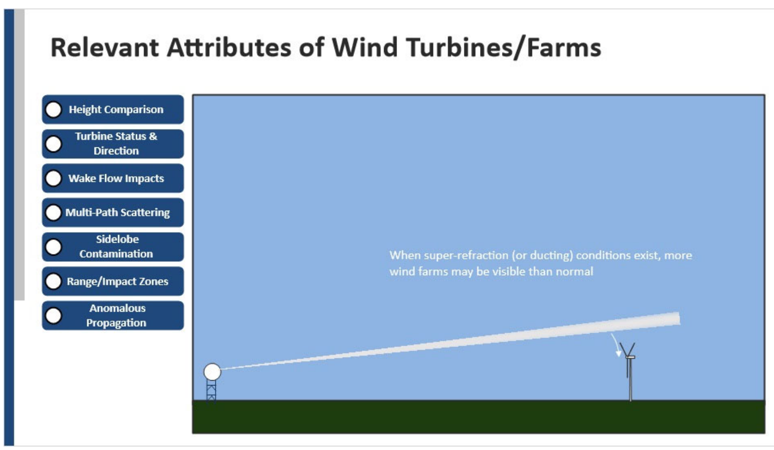

Notes:

When wind turbines and farms are located further away from a radar site, they may not impact the lowest tilt of radar data during typical atmospheric conditions. However, these sites may be visible in radar data during super-refractive (or ducting) conditions. In these situations, wind farms may appear along side other areas of anomalous propagation in your radar products.

Back to Top

Range/Impact Zones (Slide Layer)

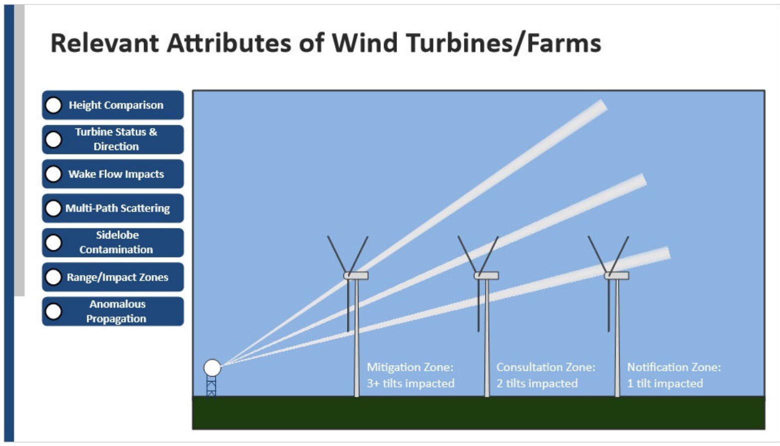

Notes:

A lot of previous research on wind farms focus on the range a wind turbine (or farm) is from a site to describe their impact on radar data (Zhang, 2024; Tang, 2024; Radar Operations Center, 2024). A newer way to describe impacts involves impact zones. Impact zones focus on the number of tilts impacted by the wind farm (Ward, 2024). Using impact zones eliminates the extra variable of turbines being located on higher or lower terrain relative to the radar. We will talk about range and impact zones in more detail on a later slide.

Sidelobe Contamination (Slide Layer)

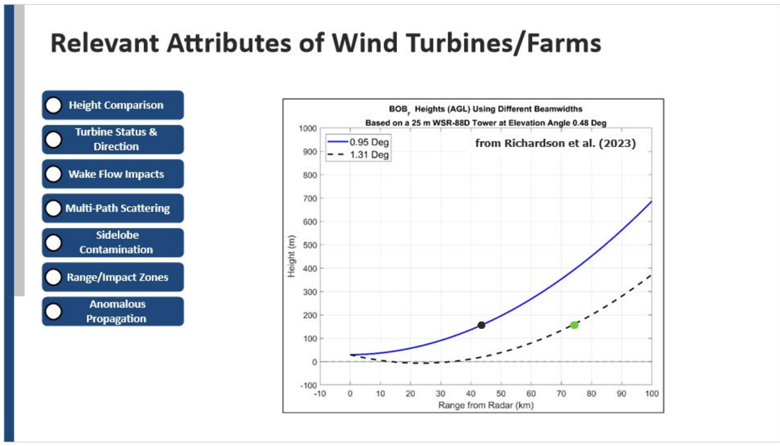

Notes:

Wind turbines can be highly reflective and result in energy from sidelobes contributing to returns on higher and adjacent tilts (Norin, 2014). Hall et al. (2015) found sidelobe returns visible for an off-shore turbine (which tend to be the tallest in use) was visible up to 5 km above the wind farm in normal propagation, clear-air conditions. These returns are typically weak and only noticeable when no significant meteorological returns are present in those areas (Toth et al., 2011). This graphic from Richardson et al. (2023) show how wind turbine clutter can impact radar data at significantly further ranges when you factor in sidelobe contamination. For example, a 150 m tall wind turbine may no longer impact the main radar lobe after 45 km, but may contribute returns in the first sidelobe out to 75 km.

Back to Top

Multi-Path Scattering (Slide Layer)

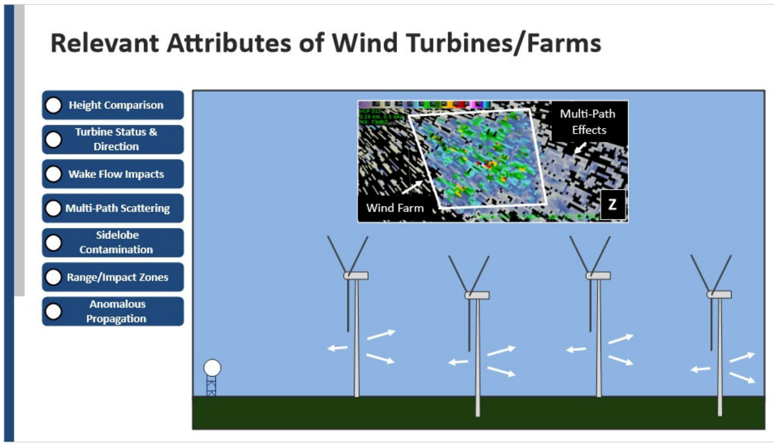

Notes:

When turbines are relatively close to a radar, multi-path scattering effects may be observed. Multi-path scattering occurs due to the radar pulse being reflected off of the various wind turbines (often multiple times) and then being reflected back to the radar. These reflections will make it appear as if some weaker returns are occurring down radial of the wind farm. In some sense, the physical process of multi-path scattering is similar to three-body scattering. When precipitation is occurring, multi-path effects less likely to be visible even if the turbines remain visible in base data (Toth et al., 2011).

Wake Flow Impacts (Slide Layer)

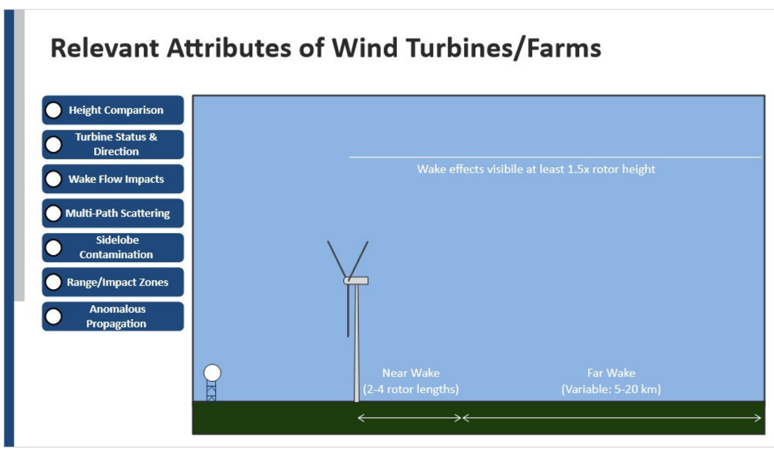

Notes:

When wind turbines are operating, the winds in the turbines wake are significantly altered by the turbine (Isom et al., 2009; Hood et al., 2010; Tang, 2024). The impacts can be broken into near and far wake effects (Porte`-Agel et al., 2020). Near wake effects occur in a fairly narrow zone only a few rotor-lengths long down radial from the turbine. The wake effects then begin to dissipate in the far wake zone. The far wake zone is larger, but also more variable in size. Far wake effects can be seen in velocity and spectrum width data anywhere from 5-20 km downwind from a turbine (or wind farm). Additionally, wake effects occur at levels above the turbine to a height of at least 1.5 times the peak rotor height.

Back to Top

Turbine Status and Direction (Slide Layer)

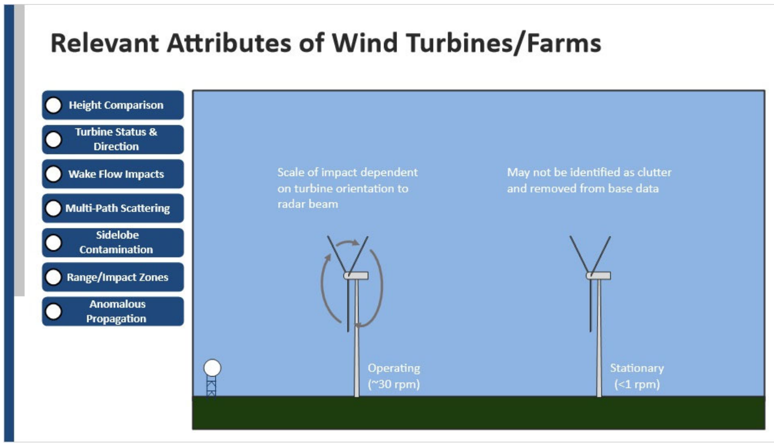

Notes:

Wind turbine appearance in radar data will vary depending on their operational status. When turbines actively generate electricity, the movement of the blades impacts velocity and spectrum width data, specifically. Other products may also be impacted as well. The scale of those impacts will vary based on how the turbine is oriented relative to the radar beam. If the turbine blades have been feathered, or are basically stationary, the turbines will appear to the radar more like traditional ground clutter. However, they still may not be properly identified by CMD as clutter where those returns could be removed from the base data.

Height Comparison (Slide Layer)

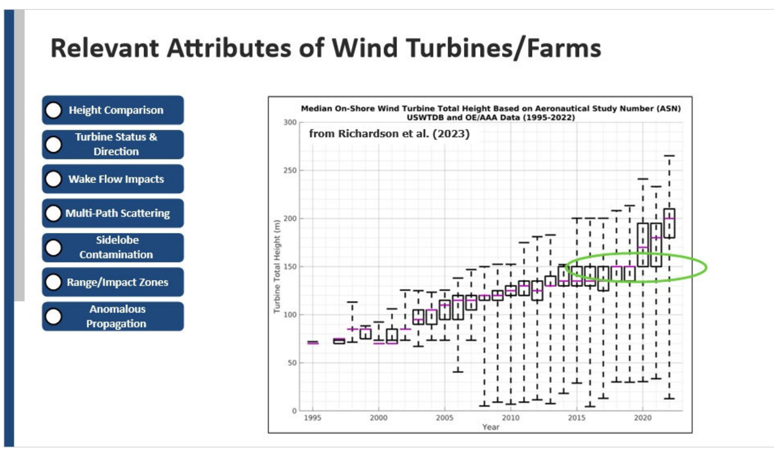

Notes:

Wind turbines have grown in size over the years. Back in the 1980s, turbines were approximately 20 meters tall on average. Nowadays, they are over 100 meters tall on average. And that's just on land. Turbines located offshore average over 150 meters tall. Remember, that's just the average. Recent projects have turbines over 700 feet tall onshore and over 1100 feet tall offshore (Ward, 2024; Albano and Auten, 2023). This box and whiskers chart from Richardson et al. (2023) shows it another way. The purple line indicates the median total turbine height for newly built on-shore turbines over time. Since 1995, the median has increased three fold. Turbines that were 75th percentile in height in the 2010s are only 25th percentile as of the 2020s. As turbines become larger in height, it increases the chance that a wind farm will be in the line of sight of one or more radar elevation angles.

1.5 Base Data Impacts

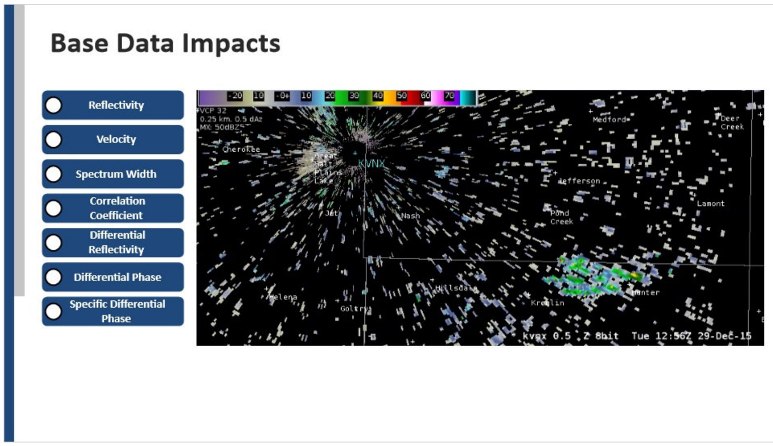

Notes:

Let's look at some typical radar returns from an area of wind farm contamination with no precipitation returns to see how products are impacted. As we discuss the features visible in this example, not all wind farms case will produce the same signatures as what's seen in this example.

Specific Differential Phase (Slide Layer)

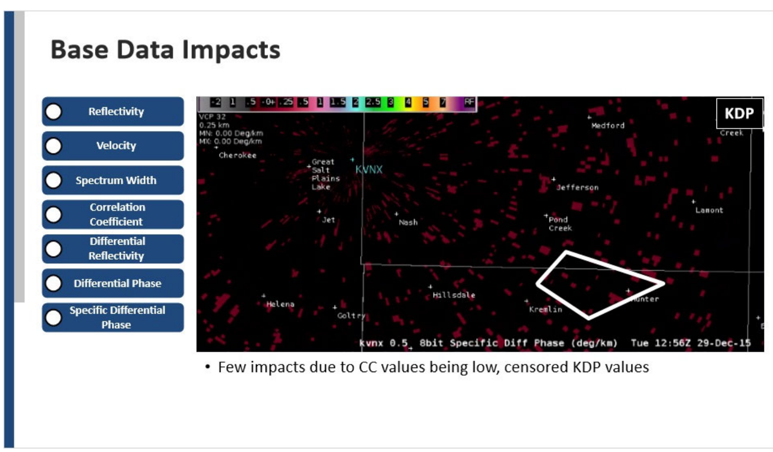

Notes:

Since Specific Differential Phase reports no data when Correlation Coefficient is below 0.9, it is the one product wind farms have little impact on as you can see from the example shown.

Back to Top

Differential Phase (Slide Layer)

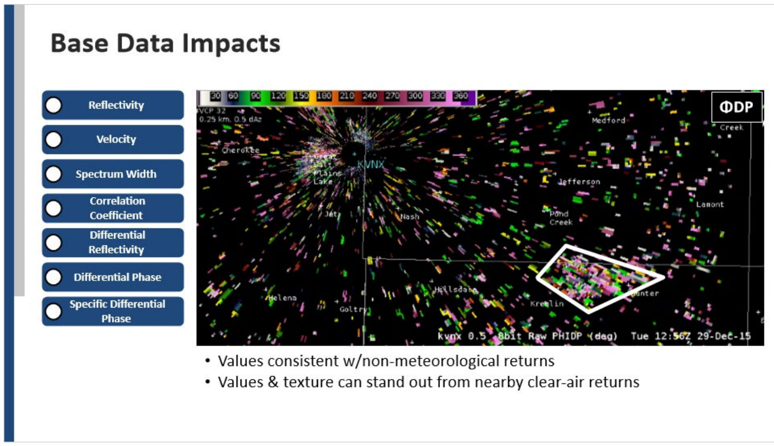

Notes:

Most users don't look at Differential Phase on a regular basis. I wanted to mention this product because wind farm data can stand out in this product. Much like Correlation Coefficient, the PhiDP data from the wind farm are consistent with non-precipitation. However, the values (and texture of those values) often stand out as appearing different than nearby, clear-air returns (Frech and Seltmann, 2017).

Differential Reflectivity (Slide Layer)

Notes:

Differential Reflectivity in the vicinity of wind farms will often have a general area of negative values with some positive areas sprinkled in. This matches observations from Frech and Seltmann (2017) of land-based wind farms in Germany. Radar observations of a marine wind farm in the UK found positive ZDR values from the wind farm (Hall et al., 2015). So, it's possible that marine wind farms will have a different appearance in ZDR data than those over land. The ZDR texture was also less varied than for Reflectivity and Correlation Coefficient in the marine case.

Back to Top

Correlation Coefficient (Slide Layer)

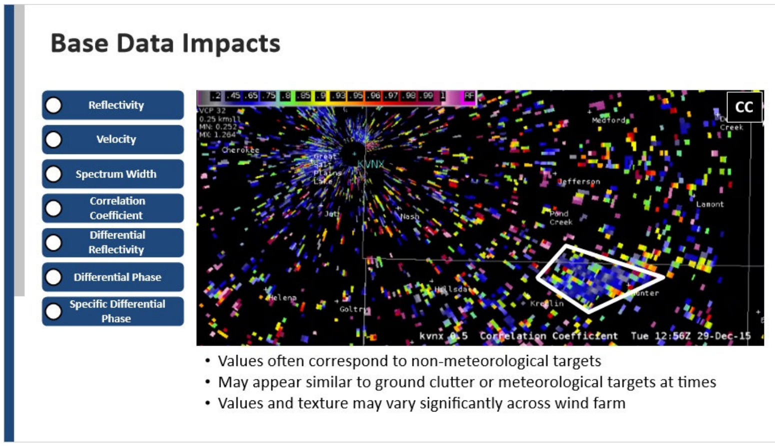

Notes:

Correlation Coefficient values clearly correspond to non-meteorological returns in this case. In many cases, the CC values in the wind farm area will have noticeably lower values than nearby precipitation echoes and possibly even the clear air returns that surround it. However, CC values can be closer to that of precipitation than other ground clutter targets. A mix of high and low CC values may also occur across the wind farm area (Zhang, 2024). As a result, the texture of the CC field can vary significantly across the wind farm region (Hall et al., 2015).

Spectrum Width (Slide Layer)

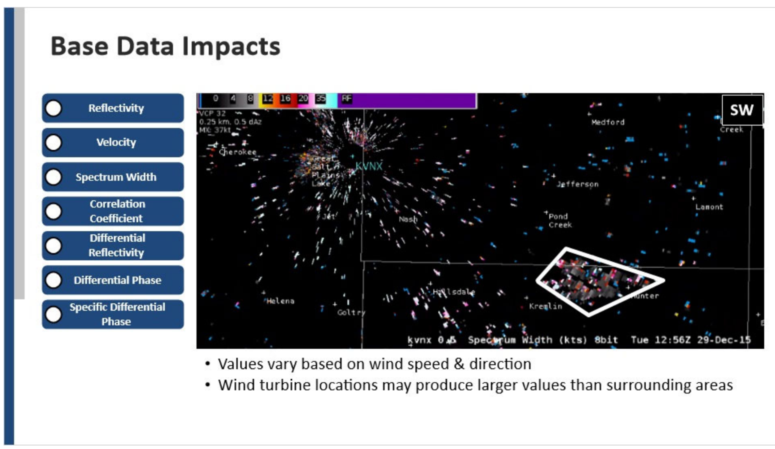

Notes:

Like velocity, the magnitude of spectrum width will vary (Norin, 2014). In this example, most spectrum width values are low likely due to most cross-radial winds at this level of the atmosphere. There are some bins with higher values of spectrum width that could correspond to individual turbines.

Velocity (Slide Layer)

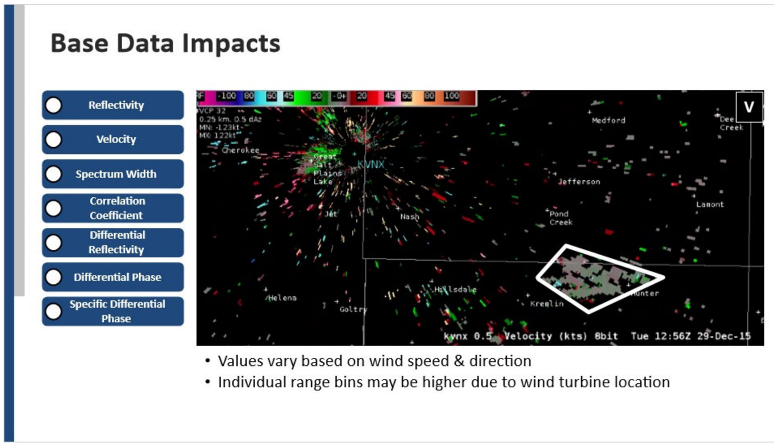

Notes:

The magnitude of velocity returns will vary depending on atmospheric wind speeds and direction (Toth et al., 2011; Norin, 2014). In this case, the surface winds are mostly perpendicular to the beam. So, most range bins are +5 to -5 knots. A few range bins have higher values that could correspond to individual turbines (or their near wake regions).

Reflectivity (Slide Layer)

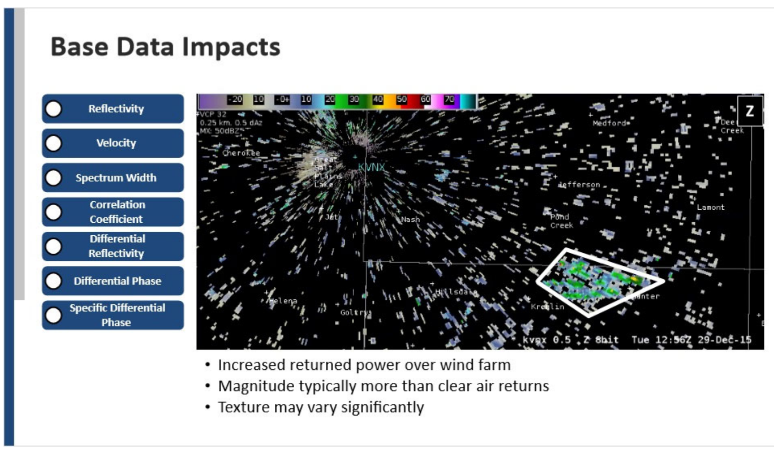

Notes:

When wind turbines intersect the radar beam, radar reflectivity shows an increased power return in that area. The magnitude of the returned power return will vary but should be noticeably stronger than nearby clear air returns. Texture in data may vary significantly from volume scan to volume scan and site to site.

Back to Top

1.6 How Important is Wind Farm Location from Radar?

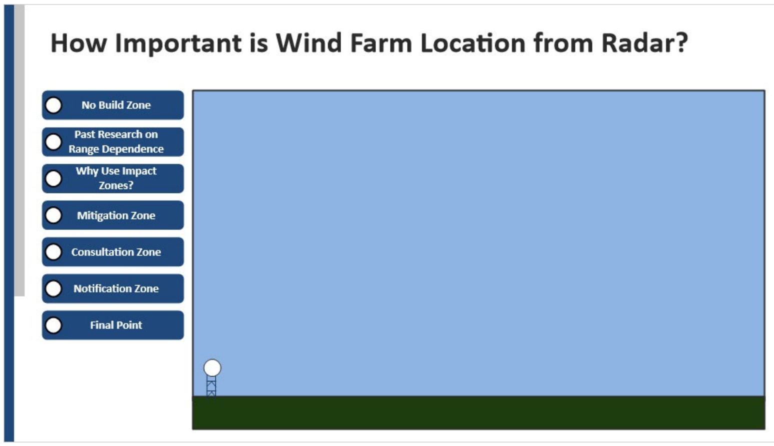

Notes:

Since we have talked about how the base data gets impacted, now let's discuss how the wind farm's location relative to the radar plays a crucial role in data contamination impacts (Zhang, 2024; Tang, 2024; Radar Operations Center, 2024). Choose the buttons on the left to learn more.

Final point (Slide Layer)

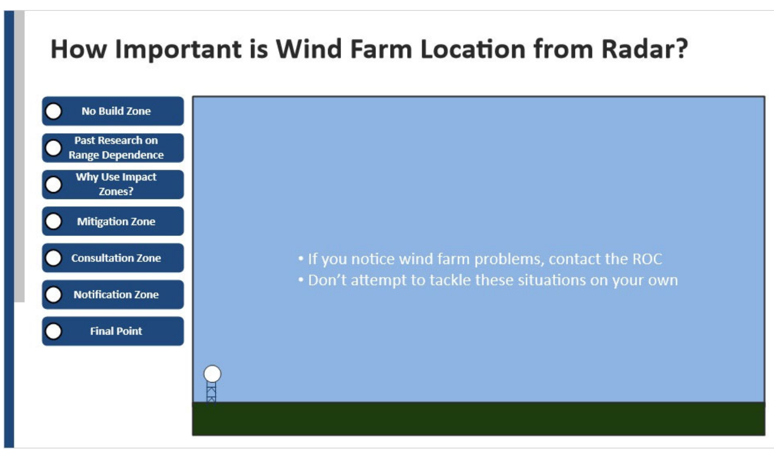

Notes:

If wind farms cause operational problems for your office, contact the ROC if you haven't already. Likewise, contact the ROC if you hear about a prospective wind farm being built in your area. They work on wind farm issues quite a bit, so don't try and tackle these issues on your own.

Back to Top

Notification zone (Slide Layer)

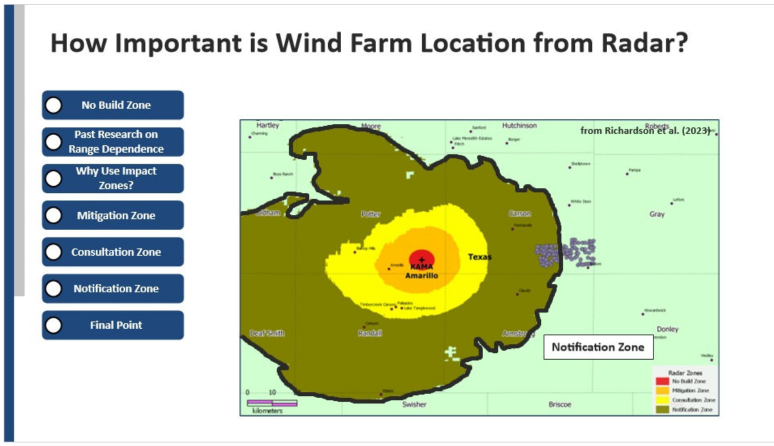

Notes:

The notification zone indicates the areas where only one tilt will typically be impacted. Wind farms in this region (shown in the darker green) still produce data quality artifacts, but they tend to be more easily worked around by NWS forecasters.

Consultation zone (Slide Layer)

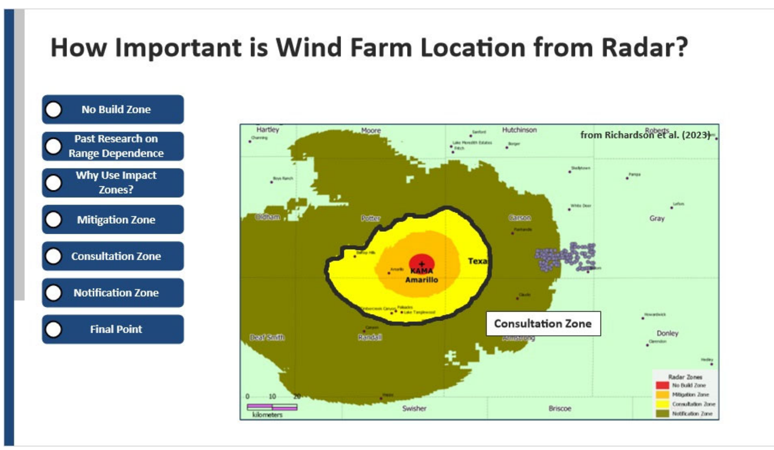

Notes:

The consultation zone indicates the areas where two tilts will typically be impacted. In this region (shown in yellow in the graphic), data quality impacts are likely to occur. The scope of the impacts will depend on the size of the wind farm, the height of the turbines, and related factors.

Back to Top

Mitigation zone (Slide Layer)

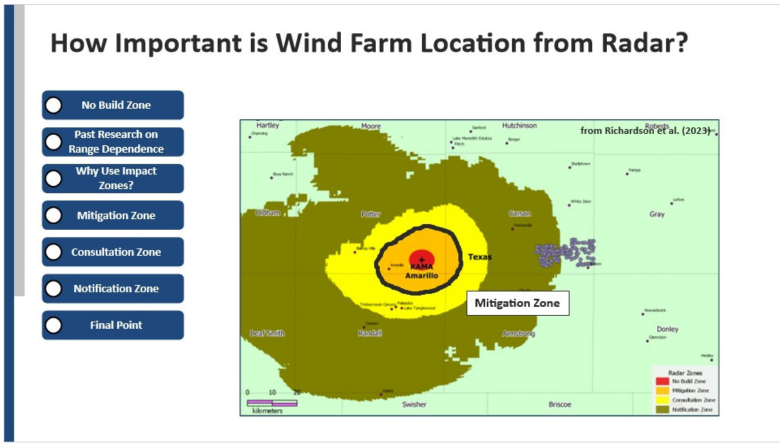

Notes:

The mitigation zone indicates the area where three or more tilts are typically impacted. This region (shown in orange in the map on the screen) effectively extends the no build zone for future wind farm projects. Some radar sites have wind turbines located in the mitigation zone already and the impacts of those wind farms on radar data quality can be significant.

Why use impact zones? (Slide Layer)

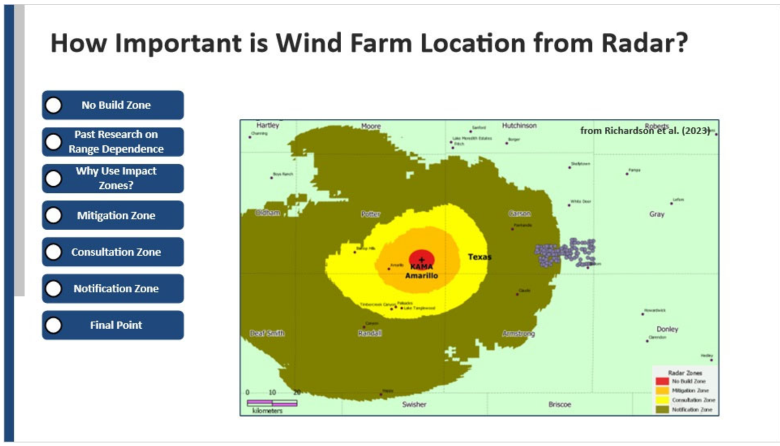

Notes:

Impact zones improve on this concept by incorporating other factors beyond simple line of sight calculations (Ward, 2024). Examples of additional factors included in impact zones also factor are the surface elevation around the radar and proposed turbine heights. Maps can then be generated to show the number of tilts likely to be impacted under standard atmospheric conditions that help the Radar Operations Center (ROC) more effectively characterize the issues than using range alone.

Back to Top

Past Research on Range Dependence (Slide Layer)

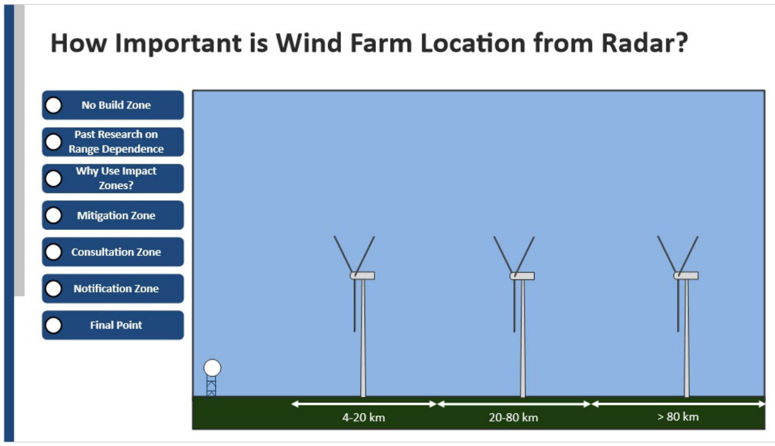

Notes:

Past research tended to characterize wind farm impacts on radar data quality primarily based on the range of turbine from the radar. In this past work, wind farms located 4-20 km from the radar had the biggest impacts. These impacts included multi-path effects, elevation sidelobe contamination, & wake flow impacts (Radar Operations Center, 2024). These same impacts could occasionally be seen further away (say 20-80 km from the radar), but it usually required atmospheric conditions that supported ducting. Even in standard atmospheric conditions, some data quality artifacts are routinely visible due to returns from the wind farm at these distances. Outside of 80 km, artifacts are not common during normal atmospheric conditions. However, ducting conditions may make some wind farms visible at much further ranges.

No Build Zone (Slide Layer)

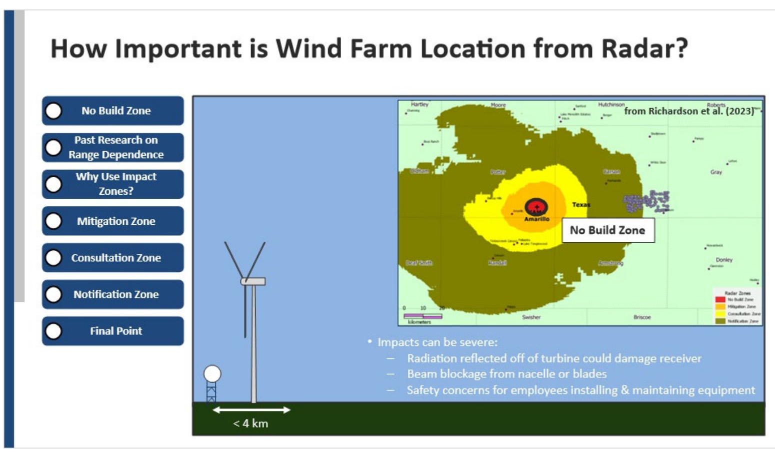

Notes:

Wind turbines are not permitted to be built within 4 km of a WSR-88D site as the impacts to radar data quality can be severe. These impacts can include receiver damage from excess power return, beam blockage by the turbine nacelle or blades, and even personnel safety of installation and maintenance crews for the wind farm. (Radar Operations Center, 2024)

Back to Top

1.7 Keys Ways That Wind Farms Impact NWS Operations

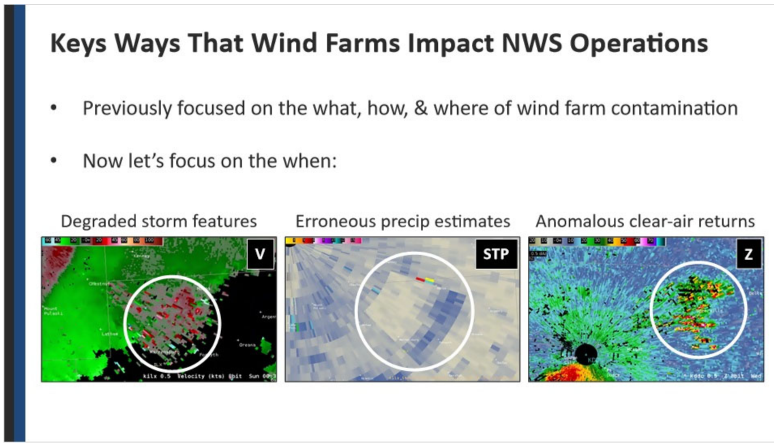

Notes:

Up to this point, we have focused on what kinds of artifacts might be present, how those artifacts might look like in your base data, and where the location of the wind farm relative to the radar might impact the significance of the data contamination. Now that you have that background, let's discuss the three primary ways when wind farms will likely impact your operations. Those three ways are: Degraded storm feature interrogation, erroneous precipitation estimates, and anomalous clear-air returns. We will discuss each in more detail on subsequent slides.

Back to Top

1.8 Key Issue #1: Degraded Storm Feature Interrogation

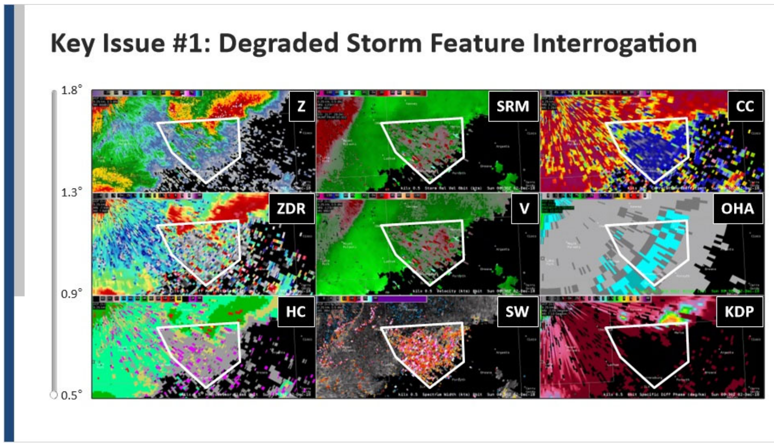

Notes:

The first situation when wind farms will impact your operations is when a storm is located over the wind farm and its features are degraded in the radar data (Porte`-Agel et al., 2020; Albano and Auten, 2023; Ward, 2024). On the base layer of this slide, I have overlaid the rough outline of the Radford's Run wind farm that is clearly visible on the lowest radar tilt fro KILX. This wind farm has turbines located over a large area approximately 15-40 km from the Lincoln radar. If you use the slider bar on the left, you can move between the lowest four tilts of radar data to see how the wind farm's impact on storm features identification lessens as you scan higher up. The arrows on the various tilts indicate the location of a mesocyclone near Maroa that produced a tornado starting around this time. Even though features are more visible on the 1.8 degree tilt, there are still signs of potential data contamination due to elevation sidelobe contamination and wakeflow effects at this level in Correlation Coefficient, Spectrum Width, and the velocity products.

I should note that this example shows where a radar signature is obscured by the wind farm. Albano and Auten (2023) shows an example where signature obscuration and false rotation signatures are present as a supercell moves over the same wind farm during a different event in January 2023.

1.8 (Slide Layer)

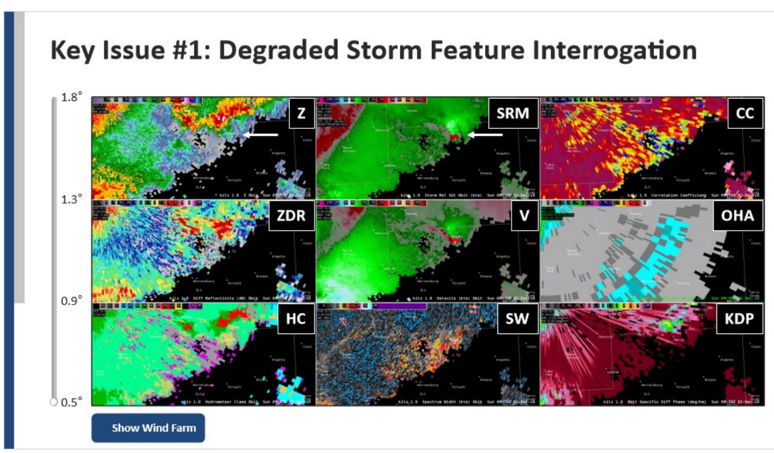

Back to Top

1.3 (Slide Layer)

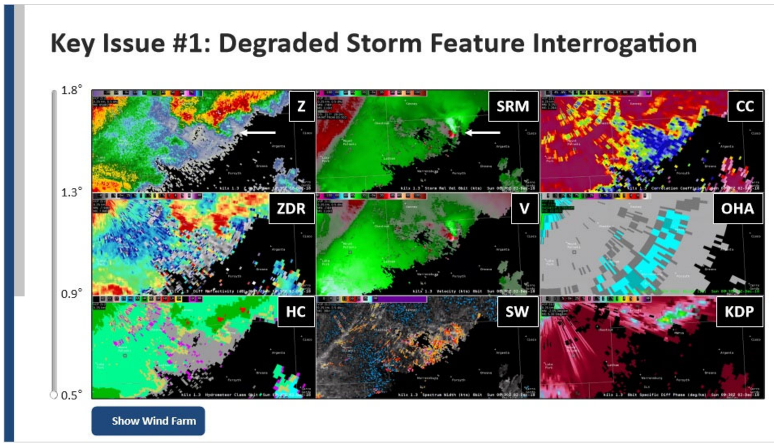

0.9 (Slide Layer)

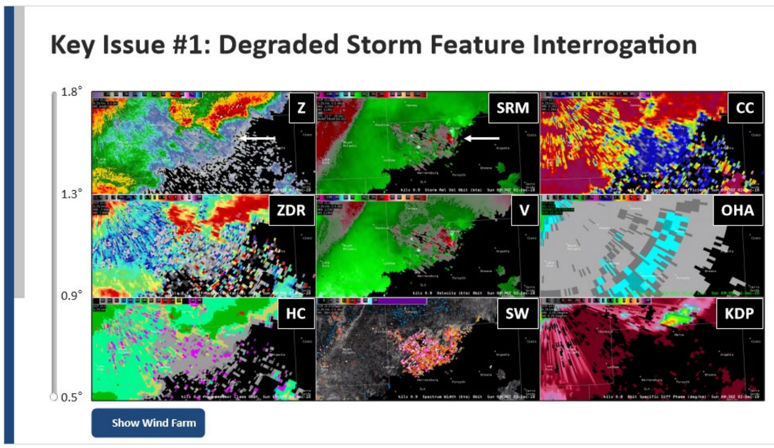

0.5 (Slide Layer)

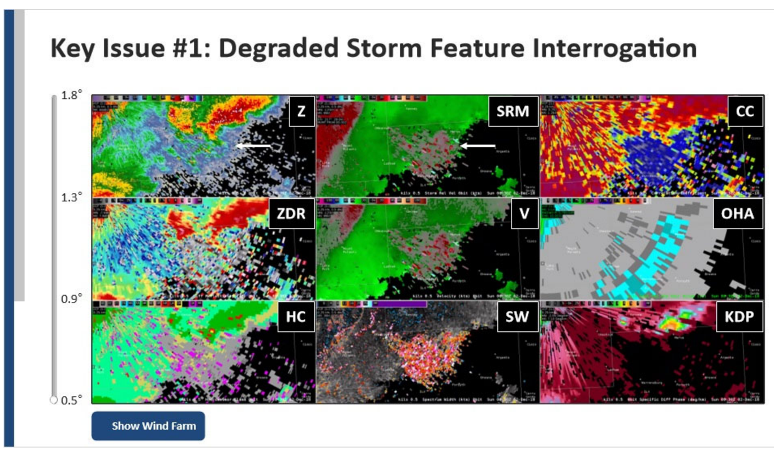

Back to Top

1.9 Feature Degradation Example Continues

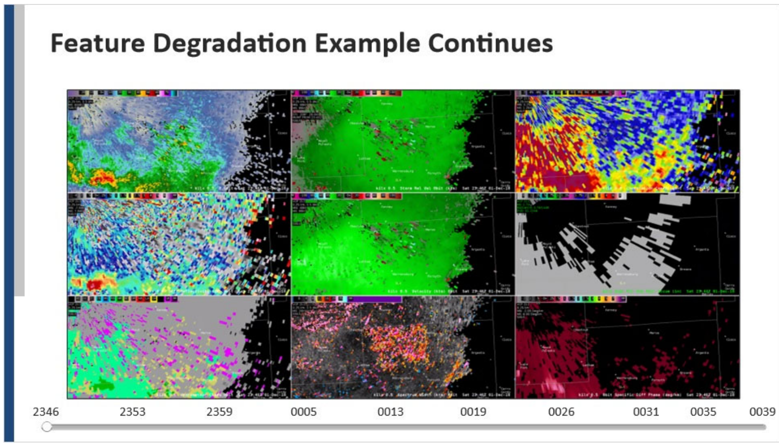

Notes:

We continue the feature degradation from the previous slide, but in this case we are showing how the lowest tilt gets impacted as the storm moves over the wind farm. Use the slider bar to move through the nearly hour long period of data. Take your time and look at how the base products change over time in the highlighted area. What do you think is the primary factor causing that change? When you think you have an answer to that question, advance to the next slide.

0039 Z (Slide Layer)

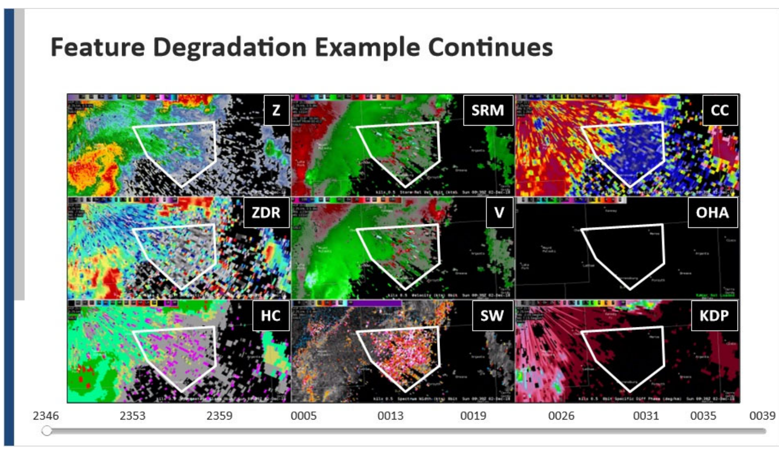

Back to Top

0037 Z (Slide Layer)

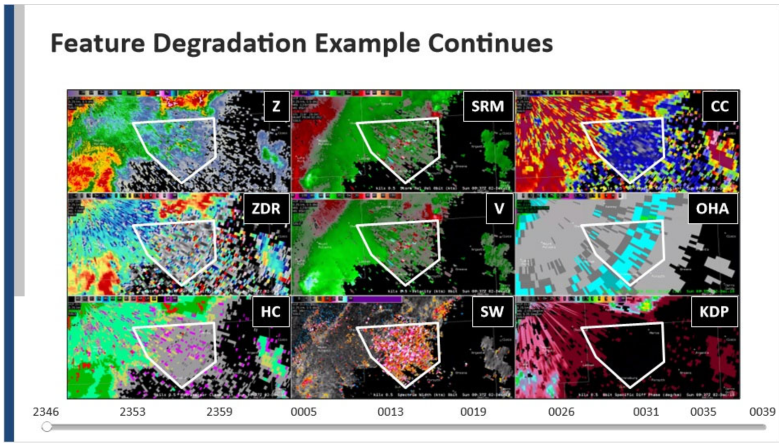

0035 Z (Slide Layer)

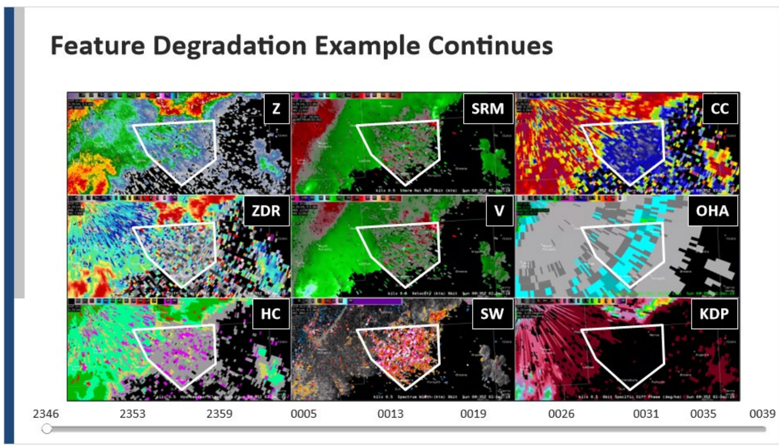

0033 Z (Slide Layer)

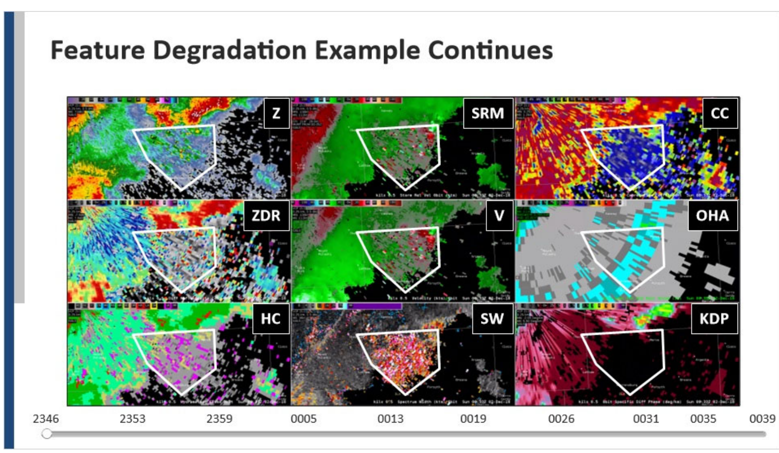

Back to Top

0031 Z (Slide Layer)

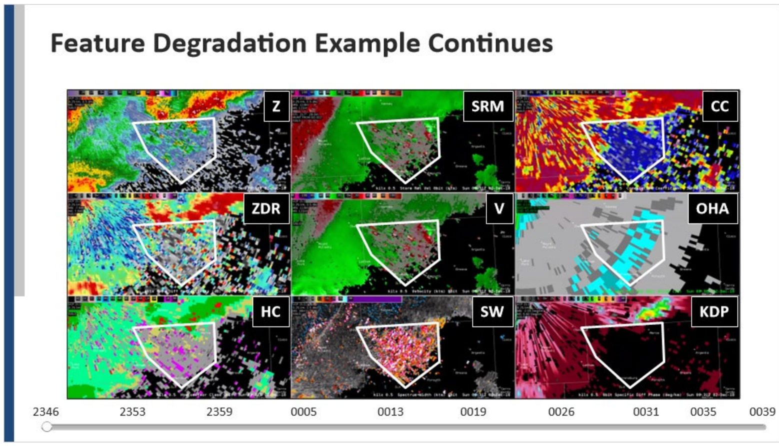

0030 Z (Slide Layer)

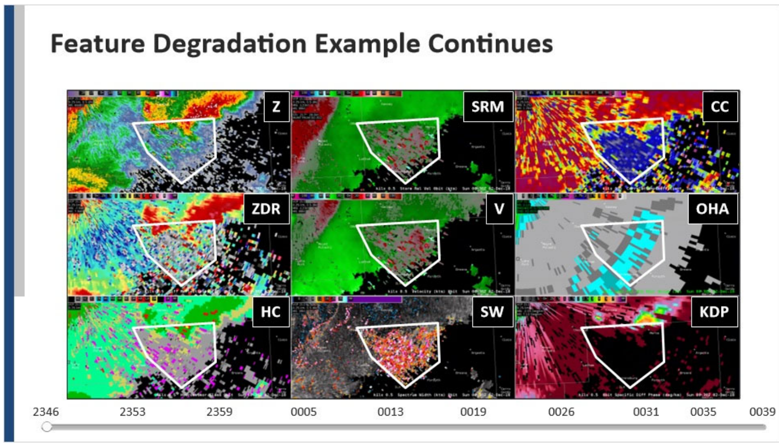

0027 Z (Slide Layer)

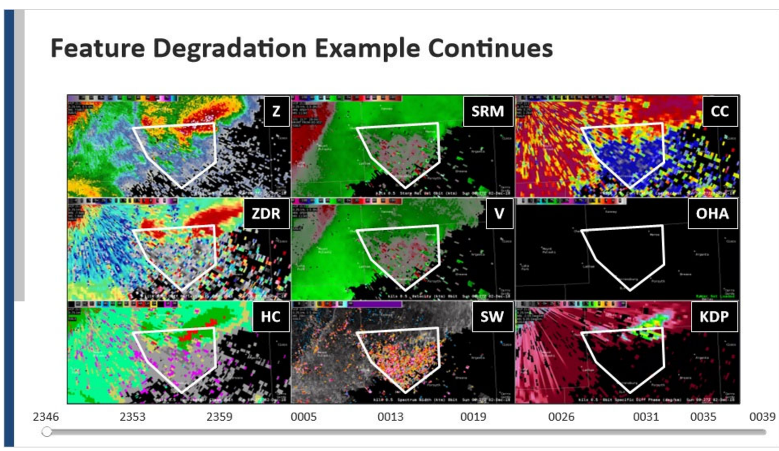

Back to Top

0026 Z (Slide Layer)

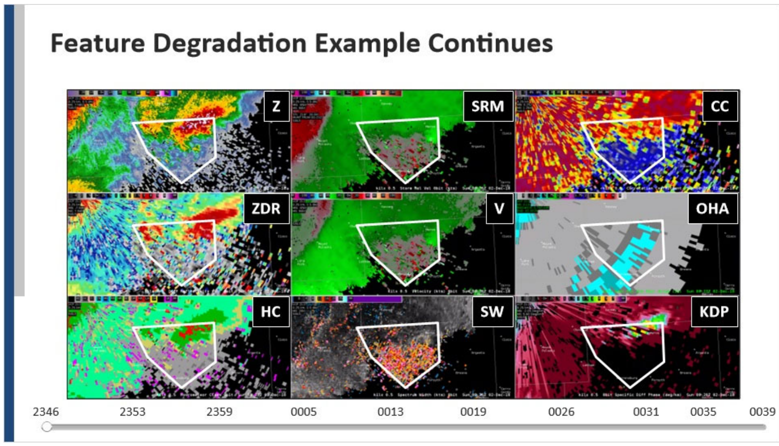

0023 Z (Slide Layer)

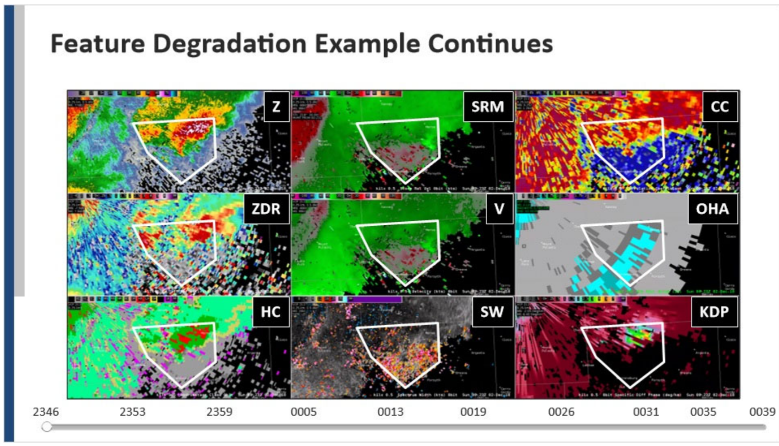

0021 Z (Slide Layer)

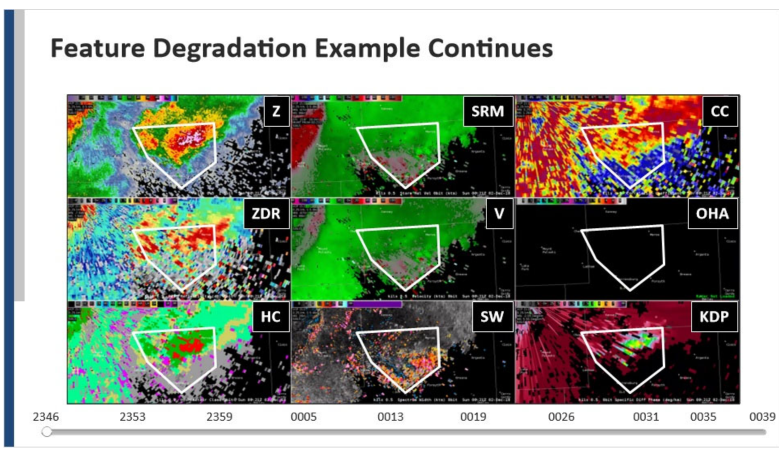

Back to Top

0019 Z (Slide Layer)

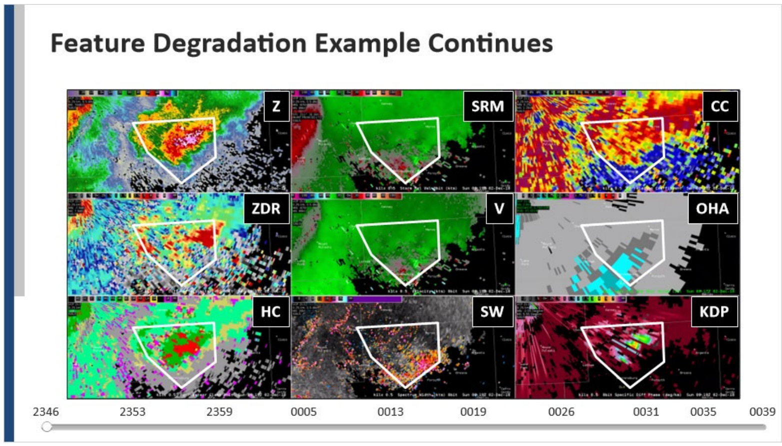

0017 Z (Slide Layer)

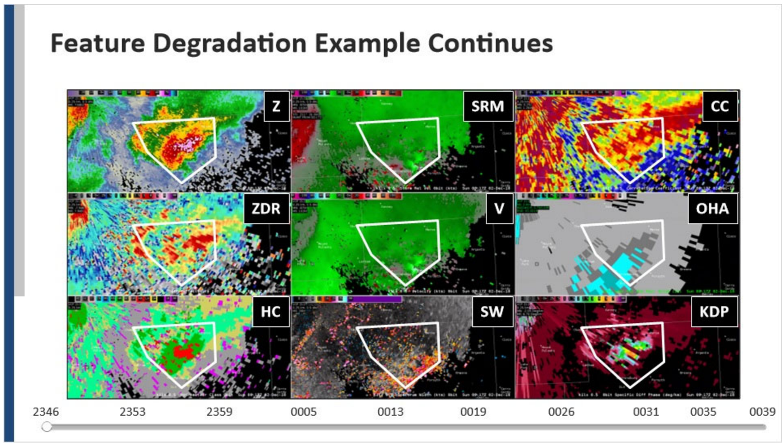

0015 Z (Slide Layer)

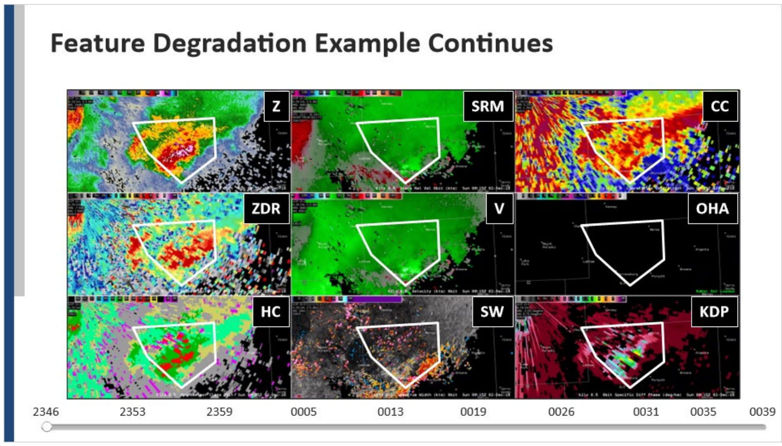

0013 Z (Slide Layer)

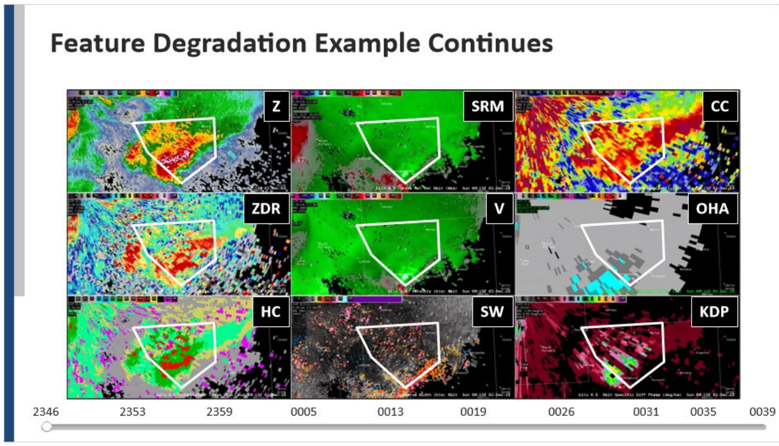

0011 Z (Slide Layer)

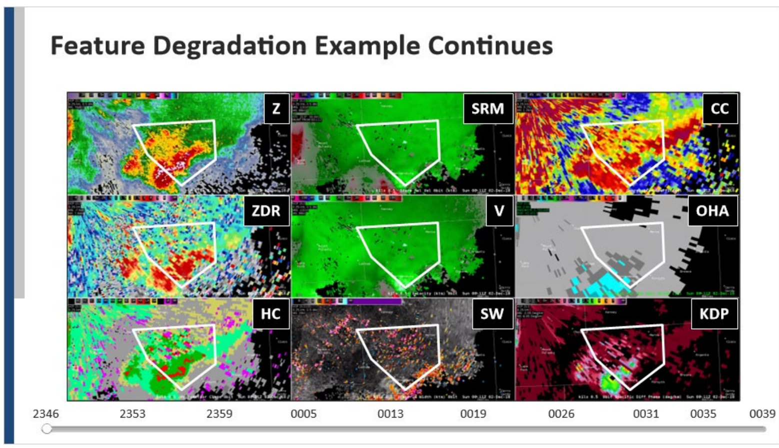

0007 Z (Slide Layer)

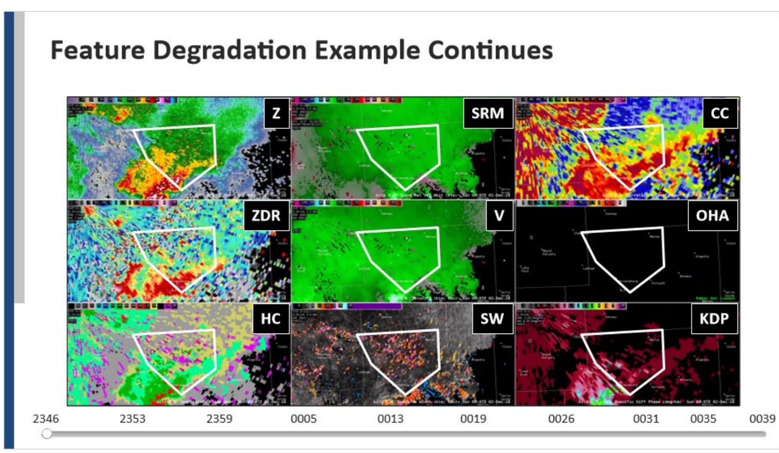

Back to Top

0005 Z (Slide Layer)

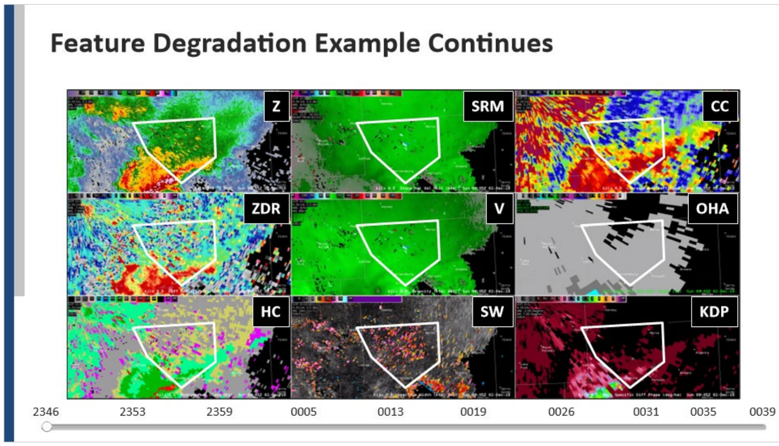

0003 Z (Slide Layer)

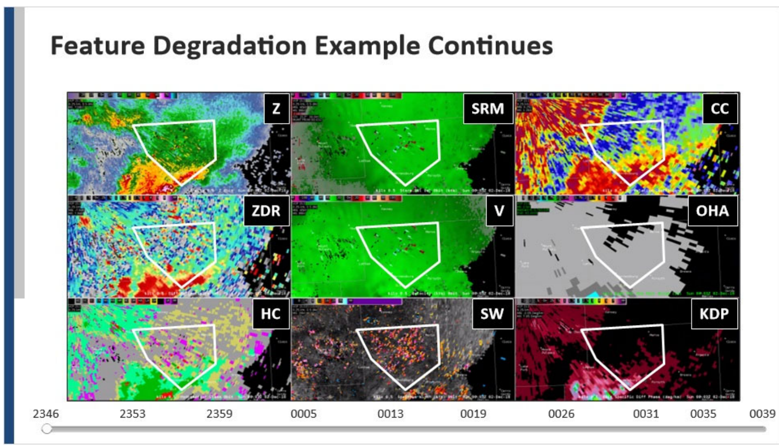

0001 Z (Slide Layer)

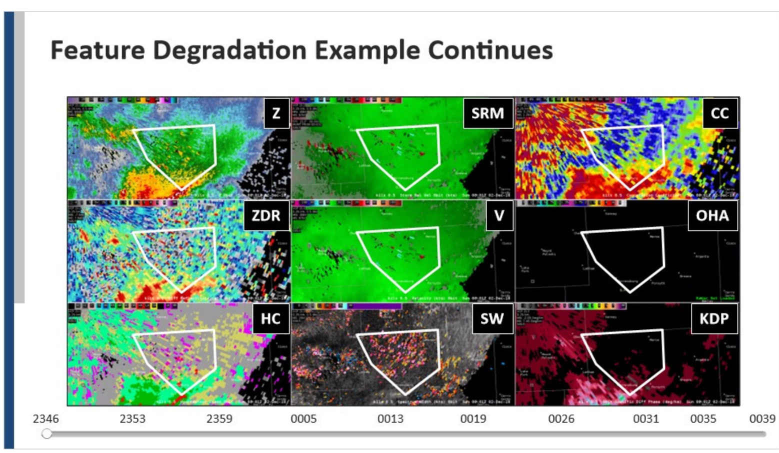

Back to Top

2359 Z (Slide Layer)

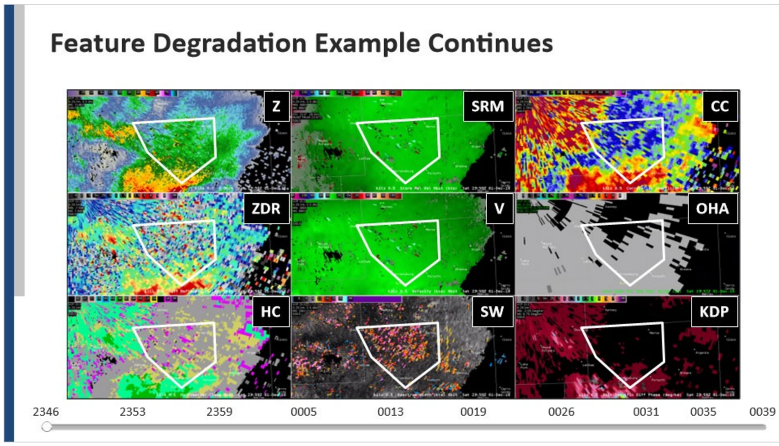

2357 Z (Slide Layer)

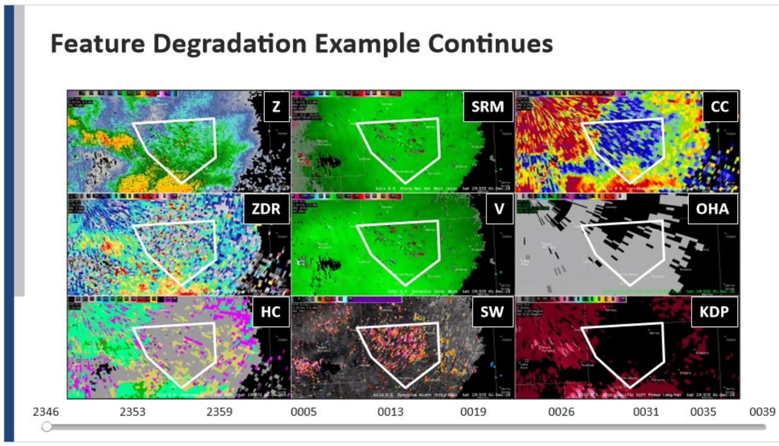

2354 Z (Slide Layer)

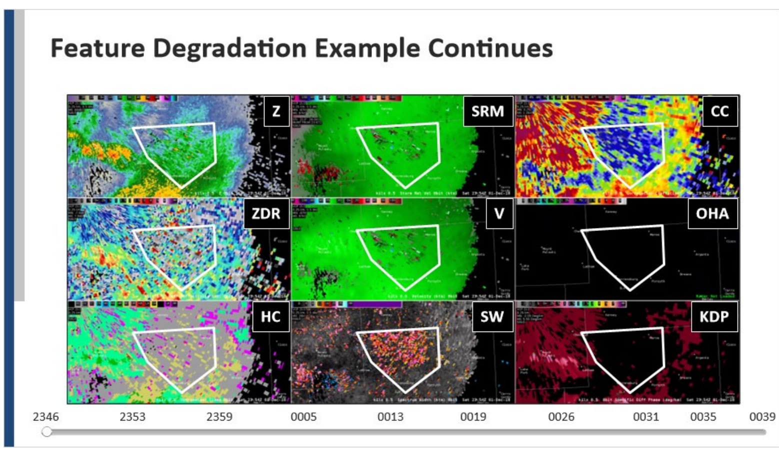

Back to Top

2353 Z (Slide Layer)

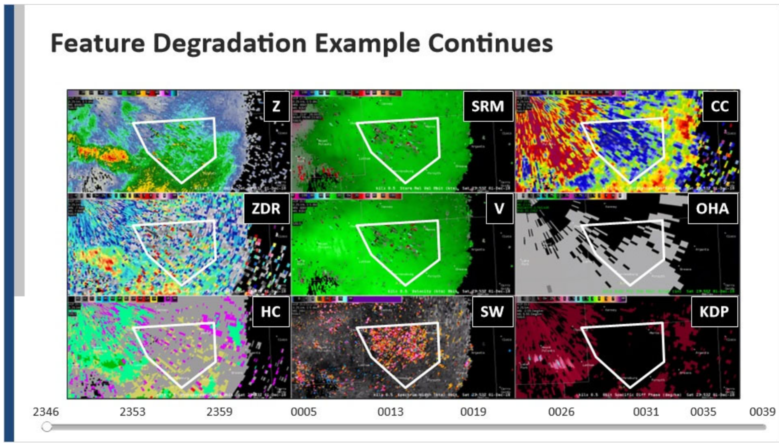

2350 Z (Slide Layer)

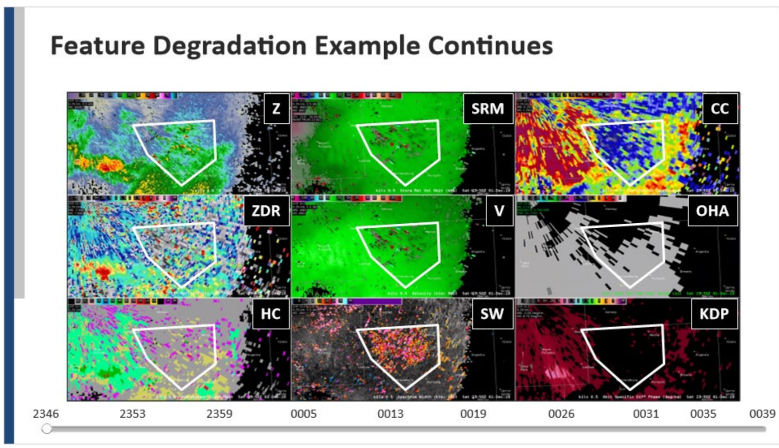

2348 Z (Slide Layer)

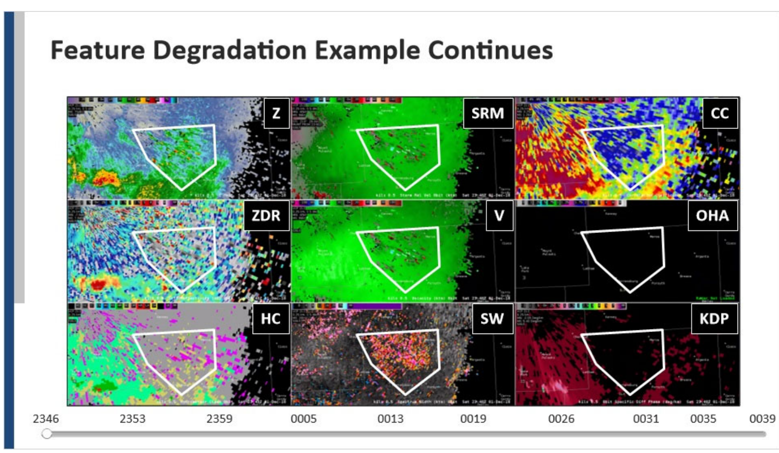

Back to Top

2346 Z (Slide Layer)

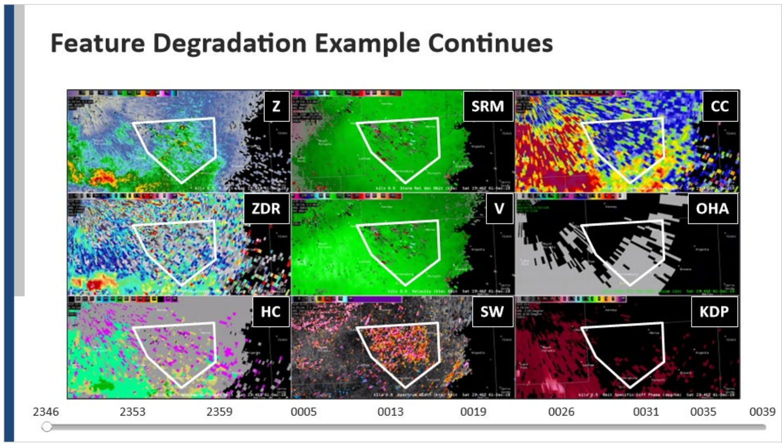

1.10 Weather Signal Vs. Wind Farm Signal

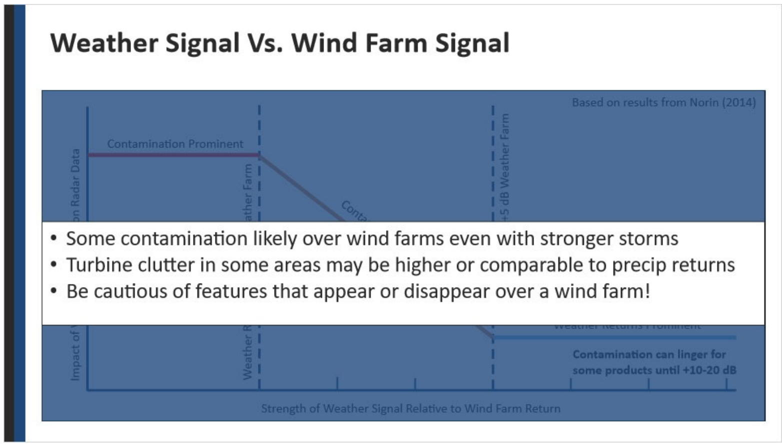

Notes:

I will use this chart based on results from work in Norin (2014) to begin the explanation of what is going on in the previous series of images. Early on, the atmosphere over the wind farm has no significant weather returns, so the turbine clutter dominates the return power. Eventually, a supercell approaches the wind farm resulting in some weather returns passing over that region. The Reflectivity returns gradually increase with time until the storm core is over the wind farm. As this happens, the returns from the wind farm gradually become less visible until the weather returns are most prominent in that window of 0007-0015 Z (Zhang, 2024; Norin, 2014). After 0015 Z, the leading edge of the supercell starts leaving the area over the wind farm. The wind turbine returns become more visible again (especially on the eastern edge of the polygon in the supercell inflow where power returns from the storm or lower). So, the process shown in the graph goes in reverse. It's important to note that for some products, like CC and ZDR, the wind farm contamination may linger until the weather power return is 10 (or even 20) dBZ stronger than the turbine clutter (Klein, 2024).

NOTE: Even during that time window when weather returns are most prominent, some wind farm contamination remains visible in most base data products. In a large wind farm, there will almost always be some range gates where the power return from the turbines will be more than 5 dB above that from the weather signal. Additionally, if the turbine clutter and weather returns are similar in strength, the base data may be still be contaminated but look OK. So, you still need to be cautious about significant storm features that appear or disappear as storm moves over a wind farm.

Back to Top

1.11 Key Issue #2: Erroneous Precipitation Estimates

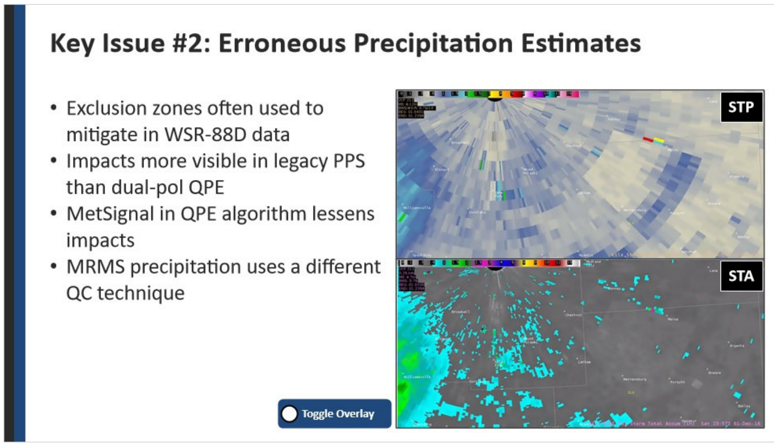

Notes:

Another key issue that can occur with wind farms is erroneous precipitation estimates. Since wind turbines move and can often have non-zero velocity values, the CMD algorithm doesn't detect them as ground clutter. Exclusion zones can be used to mitigate impacts of wind farms for precipitation products, and are especially important for the legacy PPS algorithm. The MetSignal algorithm (Krause, 2016) helps identify turbine clutter as non-meteorological and, hence, cans lesson wind farm impacts in these products without the use of exclusion zones. However, exclusion zones can be applied to both PPS and QPE algorithm outputs. If you use the toggle button visible on the slide, you can display and hide an overlay of the exclusion zone implemented for the example shown. The outline is easily visible in the legacy PPS product on the top of the screen along with a few gates along the edges where turbine clutter has bled through and led to likely overestimation errors. The outline of the exclusion zone is less visible, but still there, in the dual-pol QPE product on the bottom of the screen.

One last thing. The Multi-Radar/Multi-Sensor (or MRMS) precipitation products use a different QC technique to identify wind farms (Tang et al., 2020), so their output will not show the exclusion zone outlines.

Exlusion Zone On (Slide Layer)

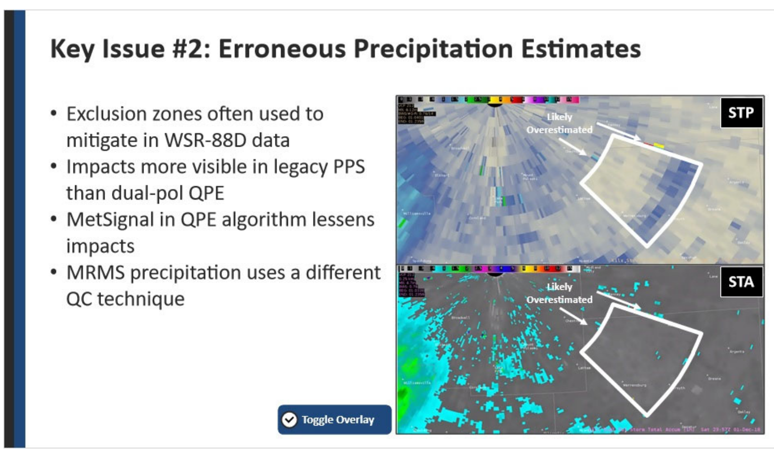

Back to Top

1.12 Example of Precipitation Product Impacts

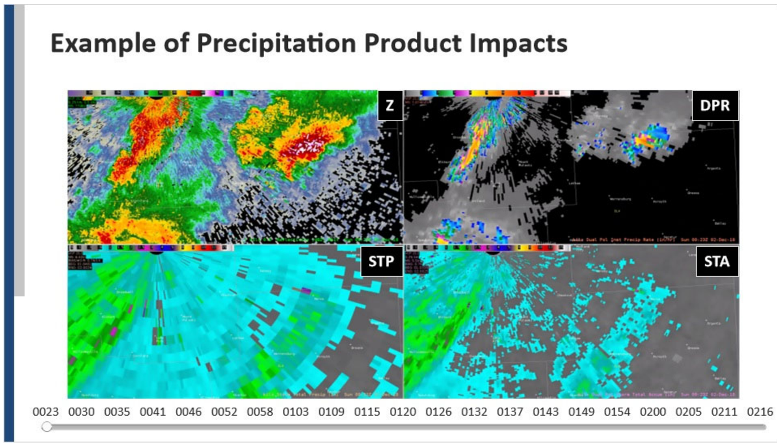

Notes:

Let's take a longer look of the example I showed on the previous slide. I have Reflectivity in the upper left with legacy PPS storm total accumulation in the lower left. The dual-pol instantaneous rate and storm total are in the upper and lower right, respectively. Unlike the last slide, I am using the same color table for the two accumulation algorithms for an "apples to apples" comparison over time. I outlined the exclusion zone on the 0023 slides so you can see where it is located, but can then see all of the actual data on the rest of the time steps. The impacts of the exclusion zone are still more visible in the STP product, but it's not nearly as obvious when using the same color table as the STA. So, the legacy PPS color table definitely has an impact on visibility.

The application of the exclusion zone can be seen in both products, but is less visible in the STA product due to the use of MetSignal. Two small precipitation maxima occur on the fringes of the exclusion zone on its northwest and north side in PPS that don't appear to be meteorological in nature. I labeled both of them on the 0216 Z frame of both the STP and STA products. The north maxima can be observed in the STA as well, but its magnitude is significantly less than in STP. This example highlights the utility of the available tools for preventing wind turbine clutter contamination as well as the challenges of removing all of it from precipitation product estimates.

0216 Z (Slide Layer)

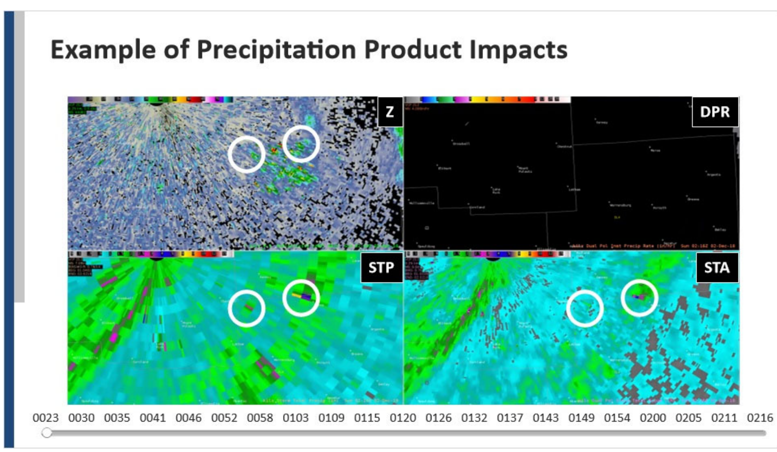

Back to Top

0211 Z (Slide Layer)

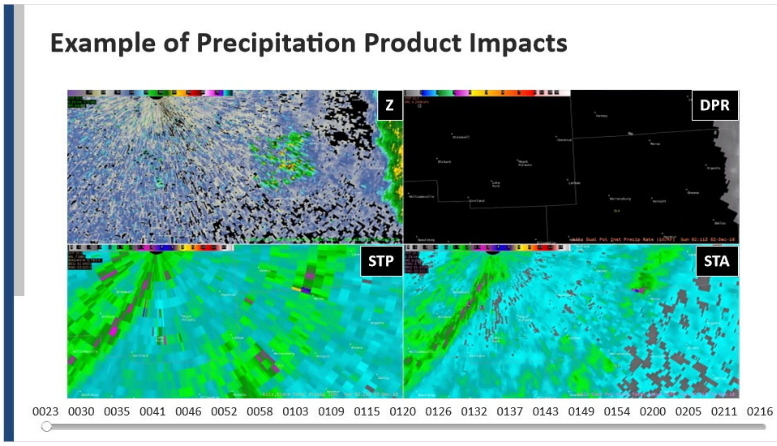

0205 Z (Slide Layer)

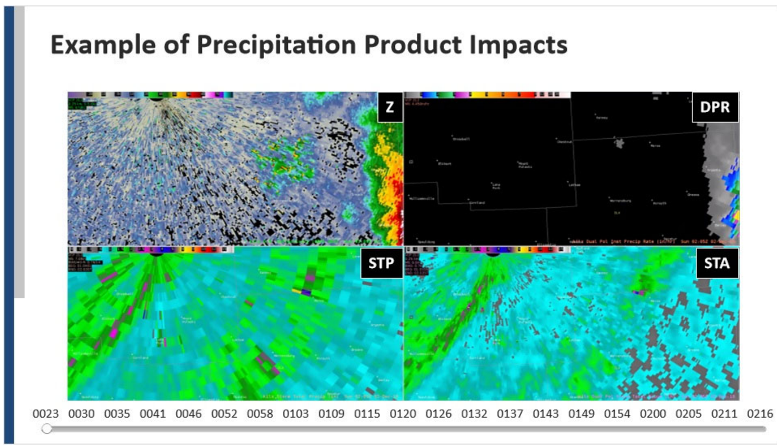

0200 Z (Slide Layer)

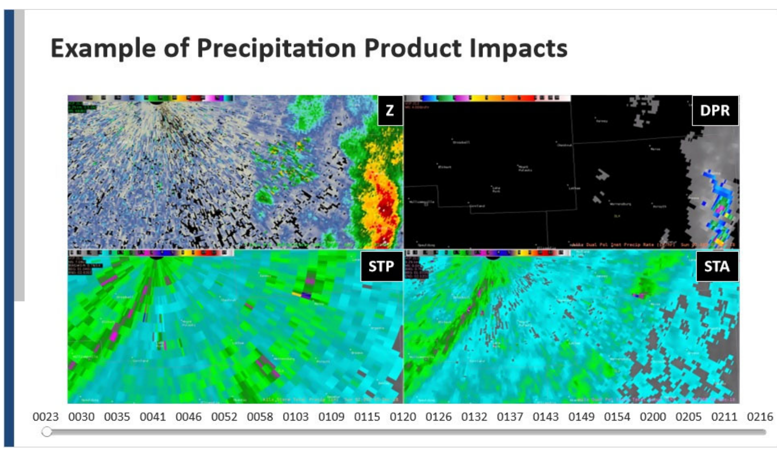

Back to Top

0154 Z (Slide Layer)

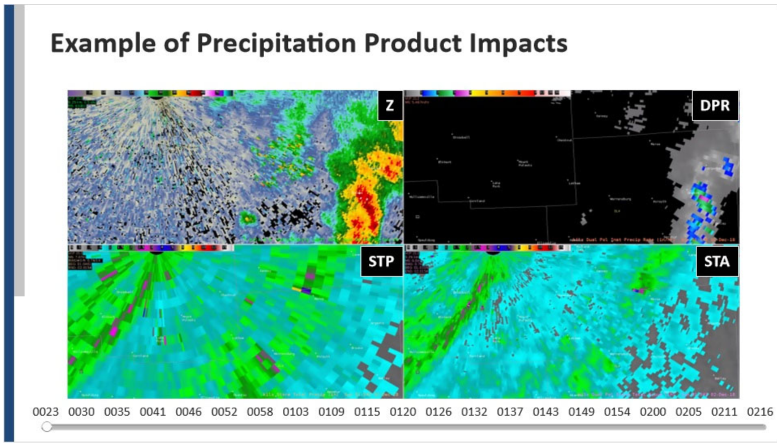

0149 Z (Slide Layer)

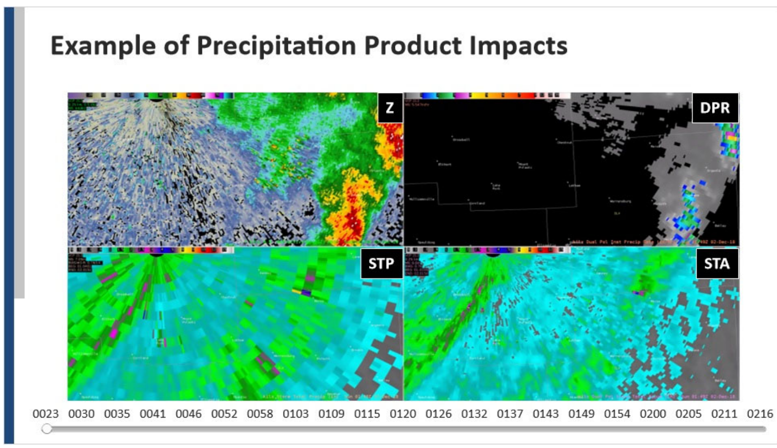

0143 Z (Slide Layer)

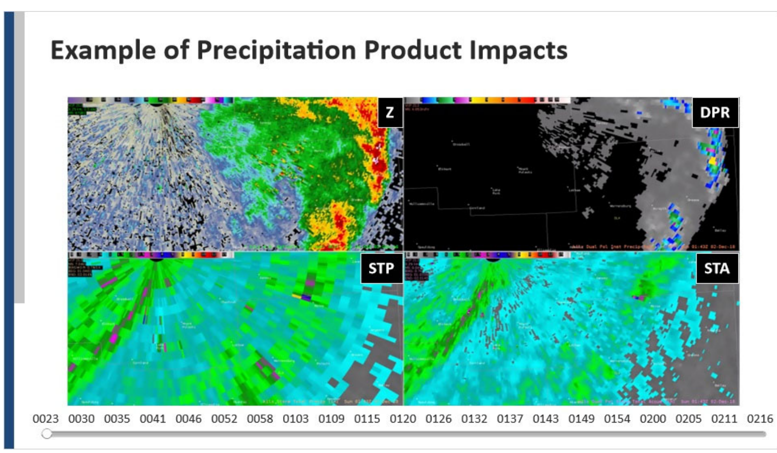

Back to Top

0137 Z (Slide Layer)

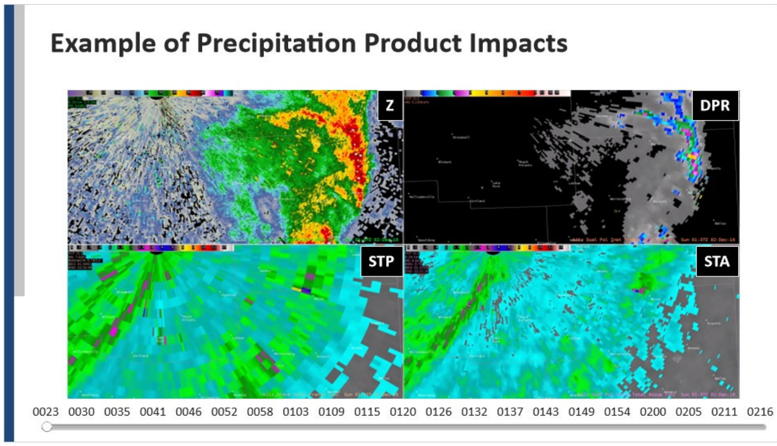

0132 Z (Slide Layer)

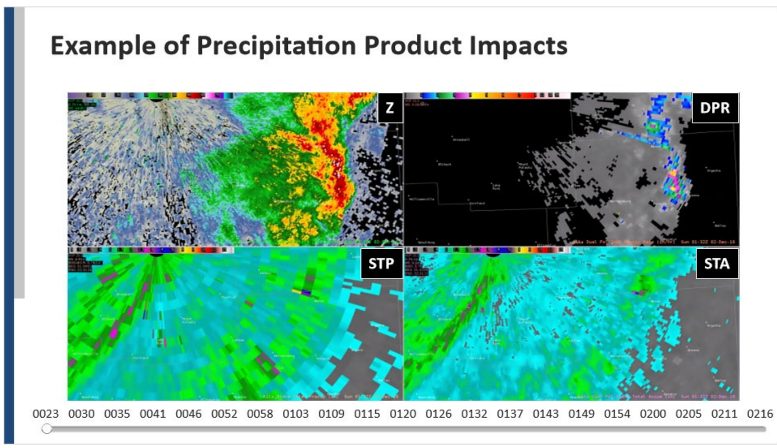

0126 Z (Slide Layer)

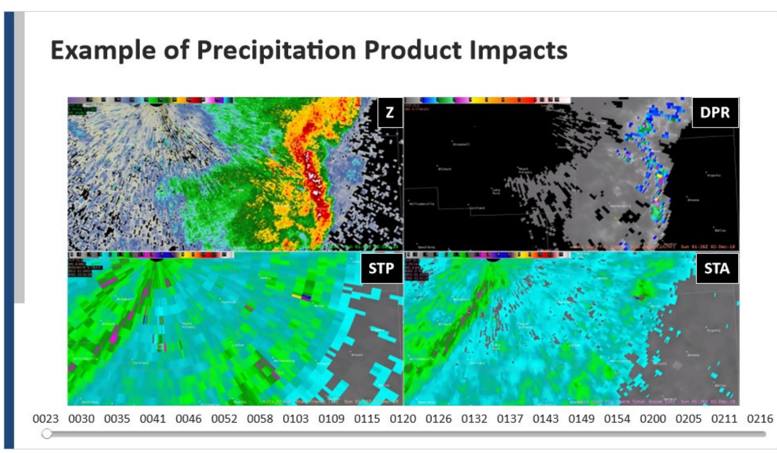

Back to Top

0120 Z (Slide Layer)

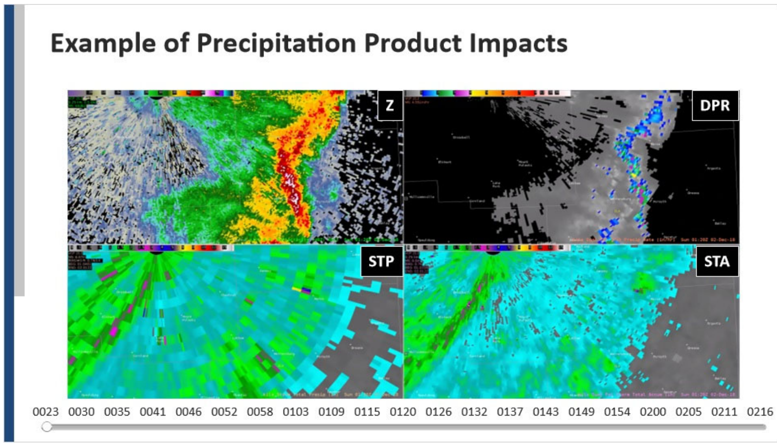

0115 Z (Slide Layer)

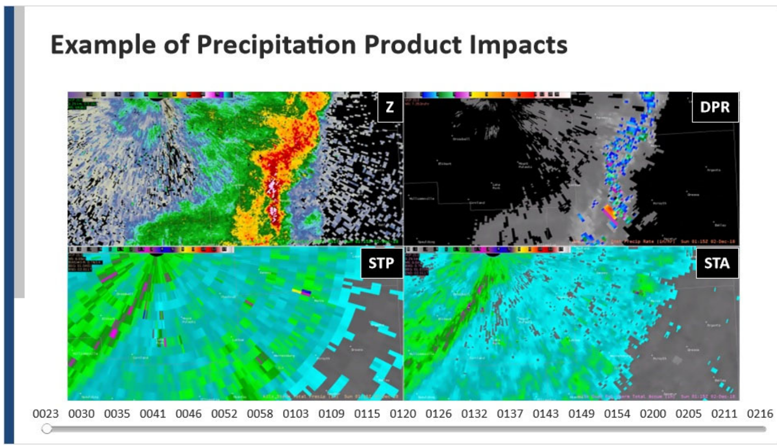

0109 Z (Slide Layer)

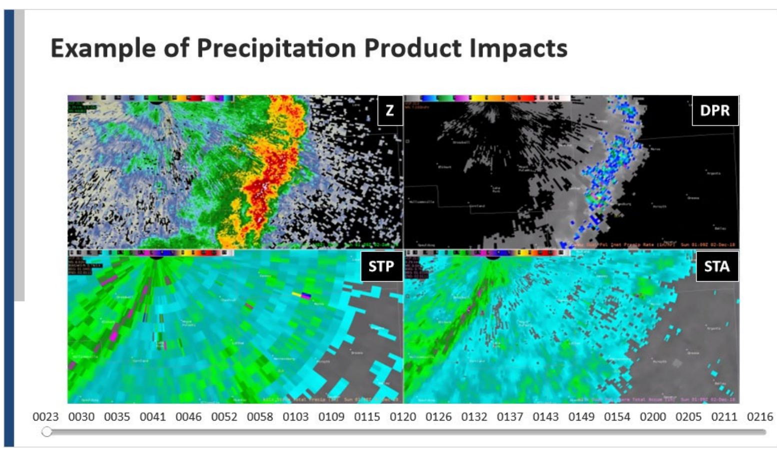

Back to Top

0103 Z (Slide Layer)

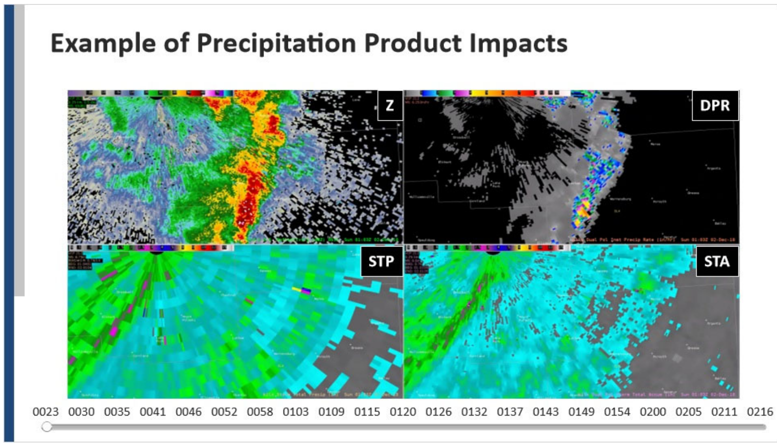

0058 Z (Slide Layer)

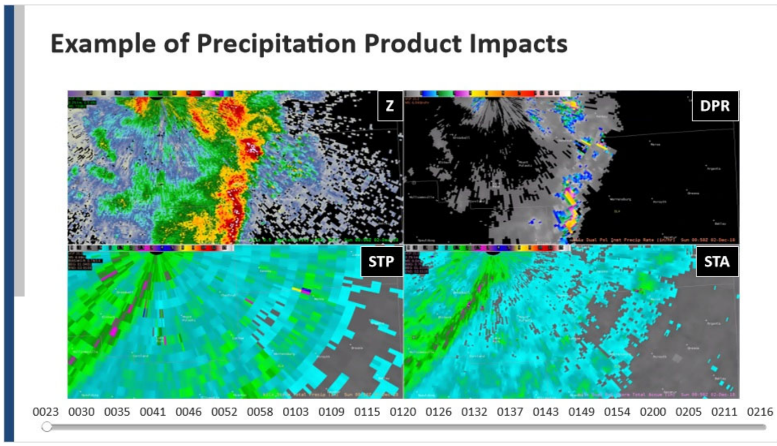

0052 Z (Slide Layer)

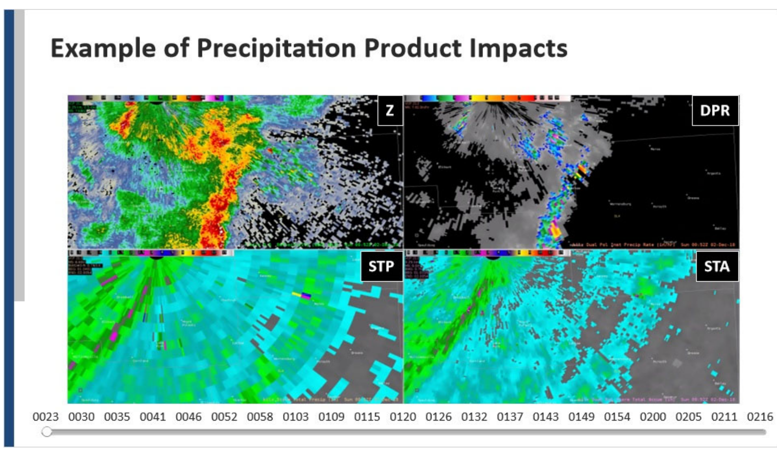

Back to Top

0046 Z (Slide Layer)

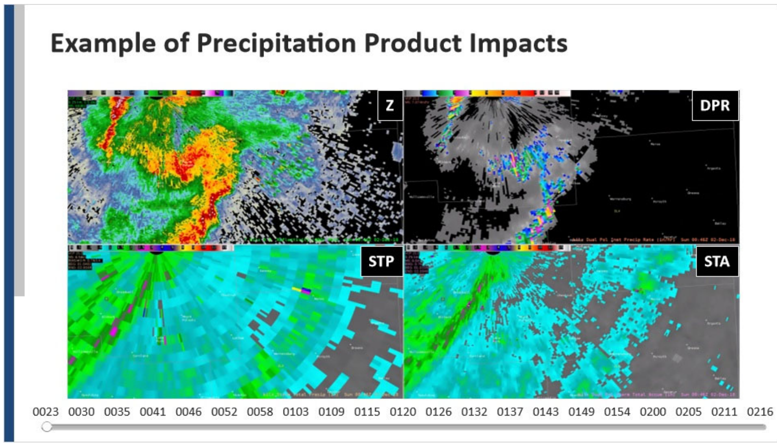

0041 Z (Slide Layer)

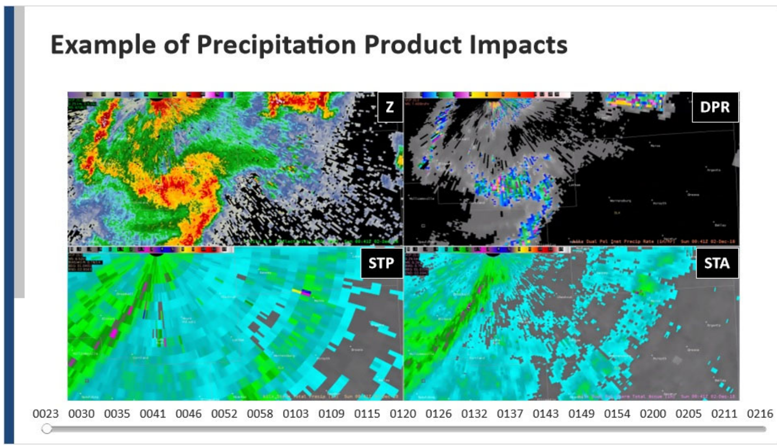

0035 Z (Slide Layer)

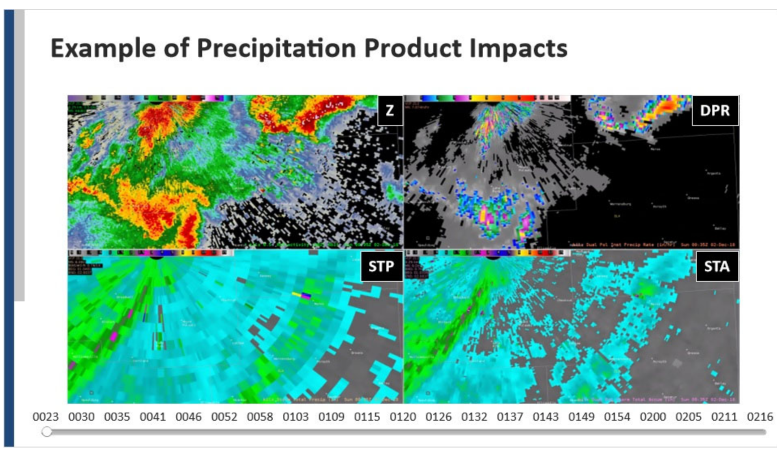

Back to Top

0030 Z (Slide Layer)

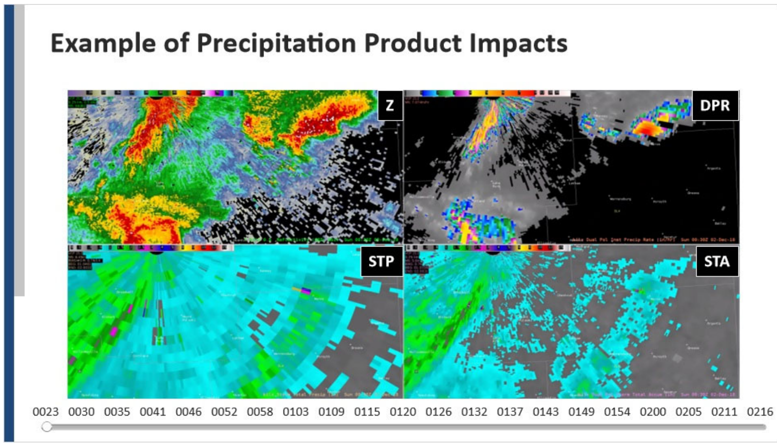

0023 Z (Slide Layer)

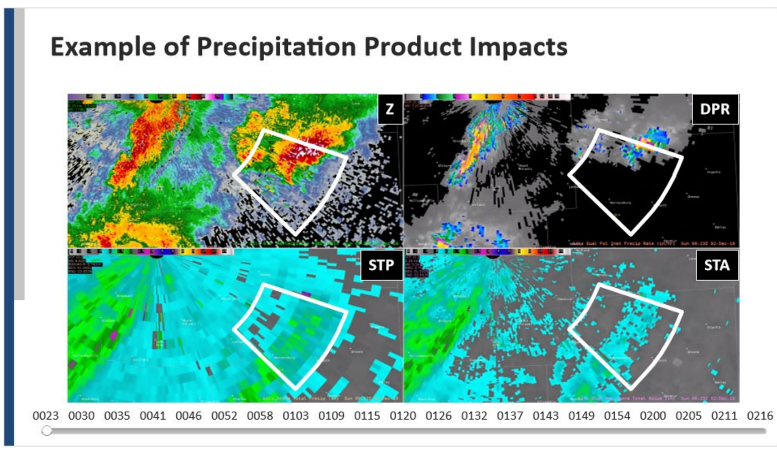

Back to Top

1.13 Key Issue #3: Anomalous Clear-Air Returns

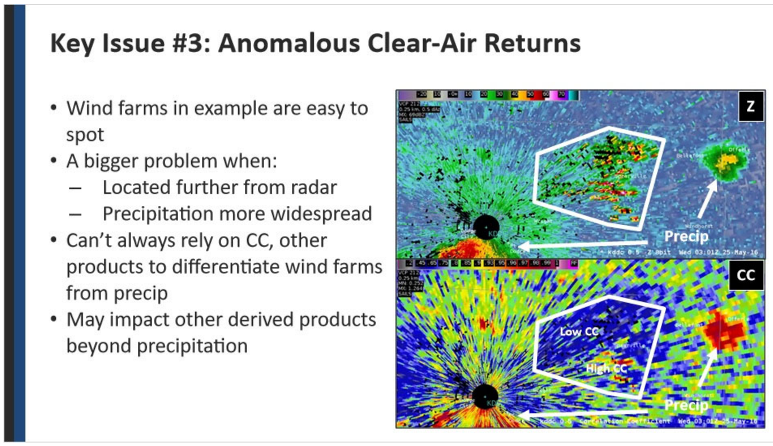

Notes:

Now let's talk about the third potential issue. I have an example on the slide where a wind farm is clearly visible in the Reflectivity data. There is also some precipitation echoes near by. Most forecasters with some experience spotting wind farms in radar data can pick that out. These situations can be a bigger problem when the wind farm is located further away from the RDA than this example or when the surrounding precipitation is more widespread. Part of the problem is that you can always rely on Correlation Coefficient and other base data to easily differentiate the wind farm returns from the meteorological echoes. See how some of the returns in the polygon highlighting the wind farm have low CC values, but other returns have high CC values. While the reason for the difference is unclear, it's possible that there is some difference in the operating status of the wind farms in the two parts of the polygon. The turbines in the higher CC region do look more like other nearby ground clutter rather than the other turbines in the wind farm. Situations like this one could lead to impacts not only to precipitation products, but other derived products as well.

Back to Top

1.14 Example of Anomalous Clear-Air Returns

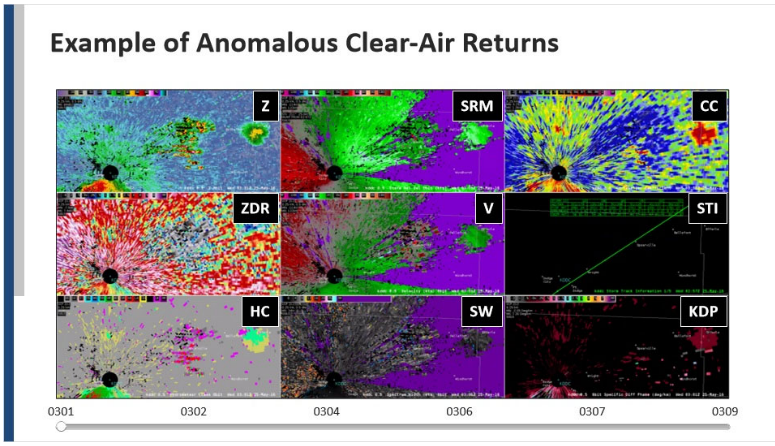

Notes:

Here is a series of images from the example we showed on the previous slide. You can see the consistency in Reflectivity over the wind farm throughout all of the images. Parameters like Differential Reflectivity and Spectrum Width are both consistent, too. Correlation Coefficient is generally low over the wind farm, but there's an area with values from 0.9 through 1.0 in the southern part of the wind farm. That region remains throughout all of the time steps. If you compare that region with the returns just north of the radar, the CC values are fairly similar in the two regions. So, you can't just rely on CC to determine what is meteorological and what isn't.

0309 Z (Slide Layer)

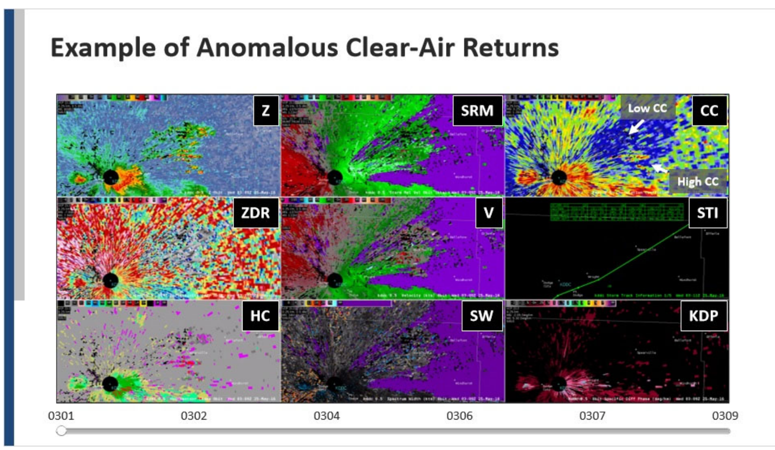

Back to Top

0307 Z (Slide Layer)

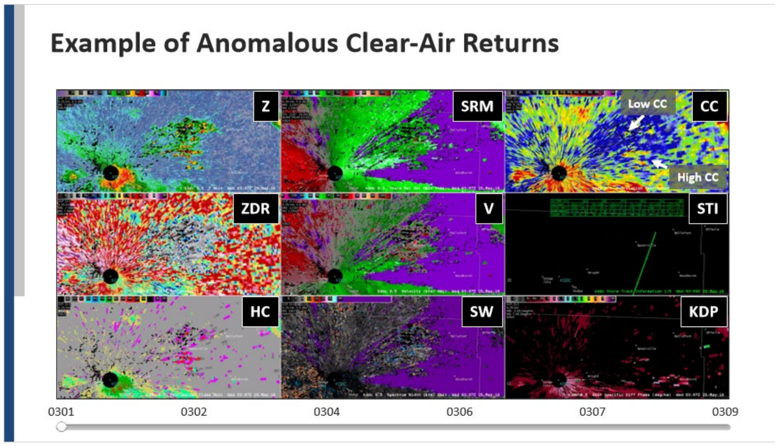

0306 Z (Slide Layer)

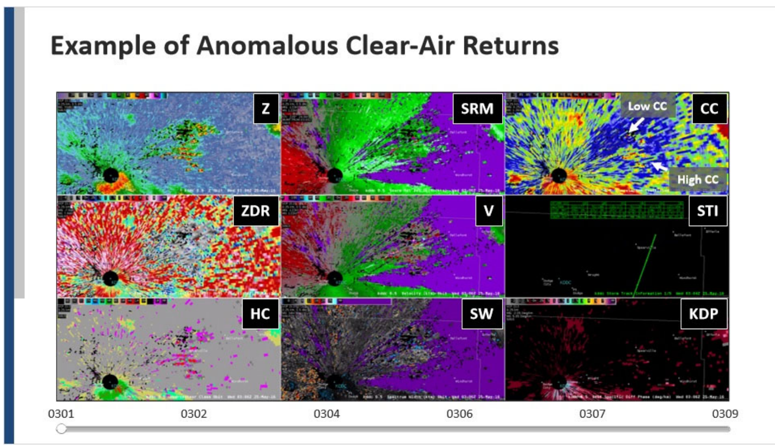

0304 Z (Slide Layer)

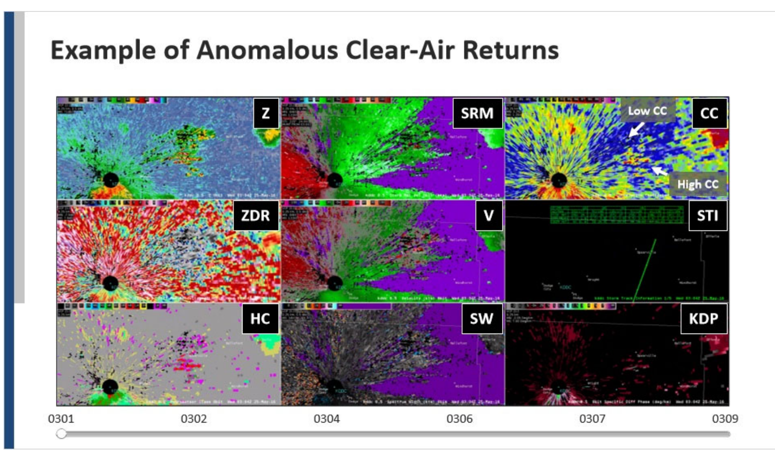

Back to Top

0302 Z (Slide Layer)

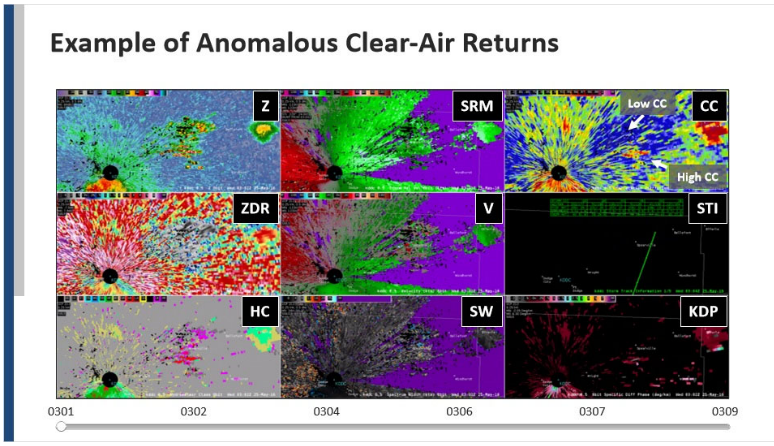

0301 Z (Slide Layer)

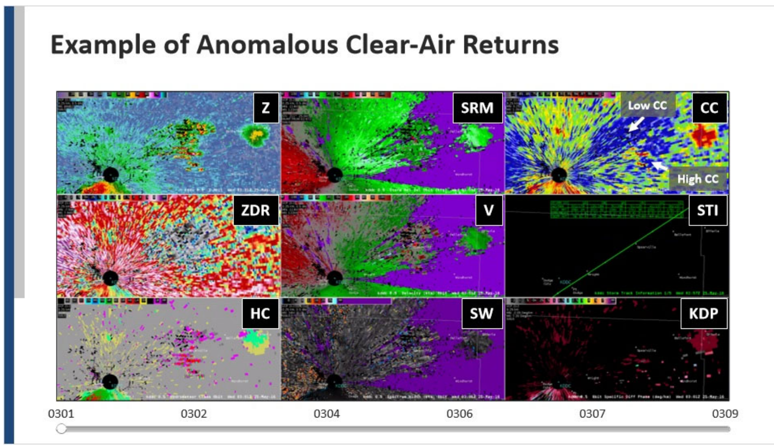

1.15 What Else Can Forecasters Do?

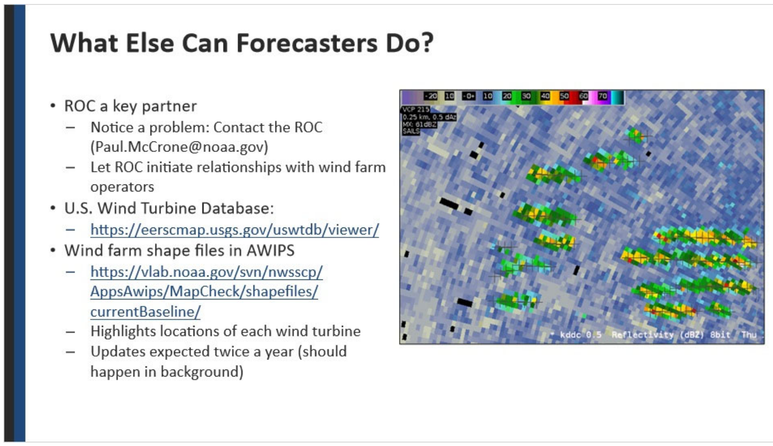

Notes:

Now that we have talked about all the ways wind farms can contaminate radar data, is there anything else forecasters can do? Most importantly, talk to the Radar Operations Center when you notice a new data quality issue related to wind farms. Even if the issue is not so new but it causes significant problems, reach out to the ROC. Let them initiate conversations and build the relationship with wind farm operators when it comes to WSR=88D impacts. A second resource that might help is the U.S. Wind Turbine Database (US Geological Survey, 2024). This web site provides a lot of information on wind farms around the U.S., including sites currently under development. This web site is not perfect as some projects under development may not appear on the site right away. But this database can help. Likewise, there is a shape file of wind turbine locations available that currently requires a manual install on your AWIPS display (Ward, 2024). Your office may have it installed already. An example of the shape file is visible in the image on the slide, and it highlights the location of each wind turbine. The plan is for this map to be included automatically in future AWIPS builds, and should get updated approximately twice a year.

Back to Top

1.16 Summary: Wind Farm Contamination

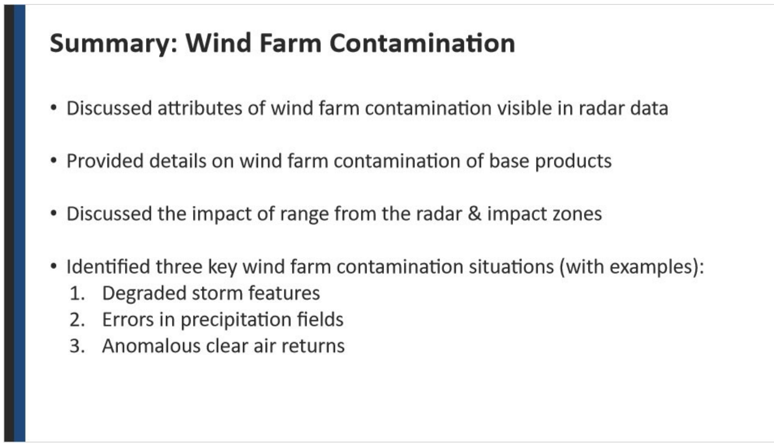

Notes:

In summary, This lesson discussed multiple aspects of wind farm contamination in radar data. We discussed several attributes of wind farm contamination that may impact base and derived radar products. We discussed how wind farm contamination may impact each of the base radar products by showing an example. When other impacts to products are possible that differed from the example, we highlighted those differences. We also discussed how the range between the wind farm and radar can define the scope of the impacts and how the concept of impact zones might better qualify them. Lastly, we discussed three general situations where wind farm contamination is most likely to impact forecasters performing their job duties. Those situations are the degradation of storm features as it passes over the wind farm, errors in precipitation fields, and anomalous clear air returns that might be misidentified as precipitation echoes.

1.26 Contact Information

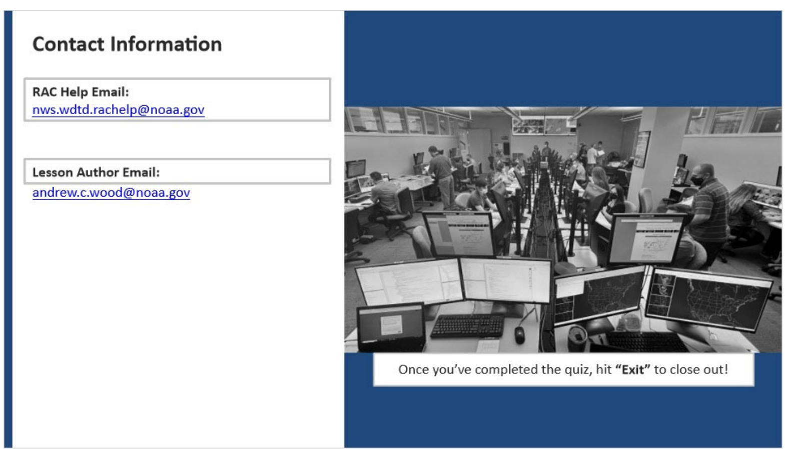

Notes:

Thank you for completing this lesson. If you have any questions, please use the e-mail addresses listed on the slide. Otherwise, use the Exit button in the upper right-hand corner of the lesson player to close the window.