3D GIS

Drones for 3D GIS

Welcome to this module on drones, or uncrewed aircraft systems, and their role in 3D geospatial data collection and visualization. We are going to work through this in three parts. First, an overview of the platforms, sensors, and, importantly, the regulatory environment you have to operate within in the United States. Second, flight planning, which is where most of the quality of your final product is actually determined. And third, product generation, where we walk through how overlapping two-dimensional photographs become point clouds, elevation surfaces, orthomosaics, and textured meshes that we can bring into a 3D GIS.

I want to flag one thing at the outset. The regulatory portion of this lecture is the part most likely to be out-of-date by the time you work through it. Treat what I say about rules as orientation, not as legal advice, and always verify against current FAA guidance before you fly.

Module Topics

1. Terminology
2. Drone types
3. Sensor types
4. Positioning considerations
5. Legal considerations
6. Pilot certification
7. Flight planning
8. Structure from motion processing pipeline
9. Structure from motion products: point clouds, DSMs, orthophotographs
10. Integration mesh creation

2

These are the topics covered in this module.

Book Recommendation

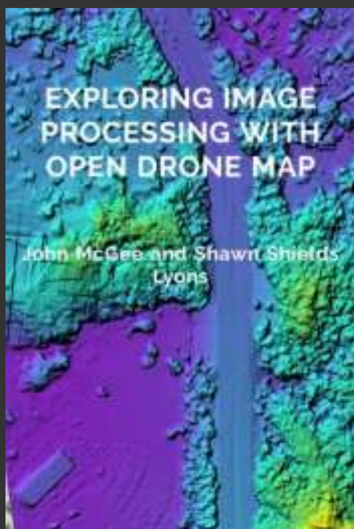

- *Exploring Image Processing with Open Drone Map*
- Virginia View
- John McGee and Shawn Shields Lyons
- Books – WordPress Site for Virginia View

3

Before we get into the material, I want to point you toward a resource that pairs well with this module.

This is "Exploring Image Processing with OpenDroneMap," by John McGee and Shawn Shields Lyons, made available through Virginia View. OpenDroneMap is a free and open-source photogrammetry pipeline, and it is genuinely capable. It will take a folder of overlapping drone images and produce point clouds, digital surface models, and orthomosaics without you ever buying a license.

I mention this for two reasons. The first is practical: if you are a student or working at an institution without a Pix4D or Agisoft Metashape seat, OpenDroneMap removes the cost barrier entirely, and WebODM gives you a browser-based interface that is approachable if you have never run a photogrammetry job before. The second reason is pedagogical. Open-source tools tend to expose their parameters rather than hide them behind a wizard, so working through OpenDroneMap forces you to confront the choices that a commercial package might make quietly on your behalf. That is a good way to actually learn the pipeline rather than just operate it.

Virginia View is part of the AmericaView network, the same network West Virginia View belongs to, and their materials are free to use. I would encourage you to spend some time with this book alongside the lecture.

Drones Overview

Terminology

- Drone
- Unmanned Aerial Vehicle (UAV)
- Unmanned Aircraft System (UAS)
- Unpiloted Aerial Vehicle
- Uncrewed Aerial Vehicle

5

A quick word on terminology, because you will encounter all of these terms and the differences occasionally matter.

"Drone" is the colloquial term. It is what everyone says, including me, and it is perfectly acceptable in conversation and increasingly in writing.

"Unmanned Aerial Vehicle," or UAV, refers specifically to the aircraft itself. "Unmanned Aircraft System," or UAS, is broader. It refers to the whole system: the aircraft, the ground control station, the communications link, and the operator. This is the term the FAA uses in its regulations, so it is the one you will see in any official document. When precision matters, UAS is usually the right word, because the aircraft is genuinely useless without the rest of the system.

"Unpiloted Aerial Vehicle" appears as an alternative, and you will also see "uncrewed" substituted for "unmanned" with increasing frequency in the scientific literature and in agency documents, as a gender-neutral alternative. The acronym stays the same, which is convenient.

My advice for your own writing: define the term you are going to use on first mention, then be consistent. If you are writing anything that touches regulation or that will be read by an agency audience, use UAS.

Types of Drones

- **Multirotor** (quadcopters, hexacopters, octocopters)
- **Fixed-wing**
- **Hybrid** (fixed wing with vertical lift rotors) (tail-sitters, quadplanes, tilt-rotors)

File:2015 Dron DJI Phantom 3 Advanced.JPG - Wikimedia Commons

File:USGS Team Flies Drones for Oregon BLM Forestry Research (14800833373).jpg - Wikimedia Commons

Platforms fall into three broad categories, and the choice among them has real consequences for the kind of survey you can fly.

Multirotor aircraft, which includes quadcopters, hexacopters, and octocopters, are what most people picture. They hover, they take off and land vertically in a small space, and they can fly slowly or stop entirely over a feature of interest. That last capability matters enormously for the oblique and orbital imagery we will need for integrated mesh generation later in this lecture. The tradeoff is endurance. A typical multirotor gives you somewhere in the range of twenty to forty minutes of flight time, and much of that budget gets spent transiting to and from the survey area.

Fixed-wing aircraft trade that flexibility for efficiency. They generate lift from forward motion rather than from constantly spinning rotors, so they cover far more ground per unit of battery. For corridor mapping or a large area survey, a fixed-wing platform may be the only practical option. The costs are that they need space to launch and recover, they cannot hover, and they cannot easily collect the oblique geometry that mesh reconstruction wants.

Hybrids attempt to get both. Tail-sitters, quadplanes, and tilt-rotors take off vertically, transition to wing-borne flight for the survey, then transition back for landing. They are mechanically more complex and generally more expensive, and the transition phases are where things tend to go wrong, but for large-area work from a confined launch site they are compelling.

The images here show a DJI Phantom 3, a very common small multirotor, and a USGS team flying for BLM forestry research in Oregon, which is a good example of exactly the kind of application we are talking about.

Types of Sensors

- RGB
- Multispectral
- Hyperspectral
- NDVI
- Thermal
- Lidar
- SAR
- Sonar

File:Drone with Lidar.jpg - Wikimedia Commons

7

The platform is really just a way to get a sensor to a location, so let's talk about the sensors.

RGB is the default and the workhorse. A standard visible-light camera is what nearly every consumer and prosumer drone ships with, and it is what feeds a structure-from-motion (SfM) pipeline. If your goal is a point cloud, a surface model, an orthomosaic, or a textured mesh, RGB imagery is very likely all you need.

Multispectral sensors add discrete bands beyond the visible, typically red edge and near-infrared, which is what lets you compute vegetation indices. Hyperspectral sensors push that much further, into hundreds of narrow contiguous bands, at correspondingly higher cost, weight, and processing complexity.

I want to note that NDVI is listed here alongside the sensors, and I want to be precise: NDVI is not a sensor. It is a normalized difference index computed from red and near-infrared reflectance. Some multispectral drone packages market themselves as "NDVI cameras," but what they are selling you is a sensor with the necessary bands plus software to compute the index.

Thermal sensors measure emitted longwave infrared, which supports applications in energy auditing, wildlife detection, irrigation assessment, and fire monitoring. Note that thermal sensors generally have much coarser spatial resolution than the RGB sensor sitting next to them.

Lidar deserves particular attention for our purposes. A drone-mounted laser scanner, like the one pictured, gives you a direct three-dimensional measurement rather than one inferred from image matching. The critical advantage over photogrammetry is canopy penetration. Structure from motion can only reconstruct what the cameras can see, which means under forest canopy you get a surface, not a ground. Lidar returns can find gaps in the canopy and reach the ground beneath, which is what makes a true bare-earth terrain model possible. The cost is that drone lidar payloads are considerably more expensive and they require a high-quality inertial navigation system, since every return has to be georeferenced from the aircraft's position and attitude at the instant of the pulse.

SAR and sonar round out the list. Drone-mounted synthetic aperture radar is still relatively specialized. Sonar, obviously, is for bathymetry, deployed from uncrewed surface vessels or in tethered configurations.

Considerations

1. Cost
2. Maintenance
3. Weight
4. Battery life
5. Customization
6. Payload
7. Software requirements

When you are actually specifying a system, whether for a grant budget or a departmental purchase, these are the practical considerations.

Cost is not just the purchase price. Budget for spare batteries, which you will need more of than you think, spare propellers, a case, and quite possibly software licenses. Insurance and, if you are institutionally affiliated, whatever compliance overhead your organization imposes are real costs too.

Maintenance covers propeller replacement, motor wear, firmware updates, and battery health. Lithium polymer batteries degrade with cycle count and with poor storage practice, and they are a meaningful recurring expense over the life of a program.

Weight determines which regulatory regime you fall under. The 0.55 pound threshold triggers registration, and 55 pounds is the upper boundary of Part 107 small UAS operations. Weight also drives battery consumption and wind tolerance.

Battery life determines how much ground you can cover per flight, which cascades into how many batteries you need, how long your field day runs, and whether you can complete a site in one visit. Changing illumination across a long field day is a real problem for photogrammetry, so endurance has quality implications, not just logistical ones.

Customization matters if you intend to swap sensors. Some platforms have open payload bays and published interfaces; others are effectively sealed appliances. If your research plan involves integrating a specific sensor, verify that the platform supports it before you buy.

Payload capacity is the constraint on that. A lidar unit or a hyperspectral sensor may simply exceed what a given airframe can lift.

And software requirements deserve more attention than they usually get. Some manufacturers restrict flight planning to their own app. Some restrict what data formats you can export, or route your data through their cloud. Ask what your data looks like coming off the aircraft, and whether you can get raw imagery with full EXIF metadata, before you commit.

GNSS and Positioning Considerations

Method	Horizontal	Vertical	What it requires
Autonomous / SPP (Single Point Positioning)	1.5–5 m	3–10 m	Nothing. Single- or dual-frequency receiver
SBAS (Satellite-Based Augmentation System) — **WAAS** (Wide Area Augmentation System, US), **EGNOS** (European Geostationary Navigation Overlay Service)	0.5–1 m	1–2 m	Clear view of a geostationary satellite; free
Code DGNSS (Differential Global Navigation Satellite System)	0.3–1 m	0.5–1.5 m	Base or beacon; largely obsolete
PPP (Precise Point Positioning)	5–15 cm	10–25 cm	Global orbit/clock products; 10–30 min convergence
PPP-RTK (Precise Point Positioning – Real-Time Kinematic) / commercial correction service	3–10 cm	5–15 cm	Subscription service; seconds to minutes convergence
RTK (Real-Time Kinematic)	1–2 cm + 1 ppm	2–3 cm + 1 ppm	Base station + live correction link; ambiguity fix
Network RTK / **VRS** (Virtual Reference Station)	2–4 cm	3–5 cm	**CORS** network, cellular coverage, **NTRIP** client
PPK (Post-Processed Kinematic)	1–3 cm	2–4 cm	Raw observable logging at rover and base; post-processing

9

This table is worth spending time on, because positioning quality is the single largest determinant of whether your final products are actually usable for measurement, and it is routinely misunderstood.

The distinction that matters is between where the aircraft "thinks" it is and where it actually is. Every drone has some GNSS receiver, so every image gets a coordinate stamped in its metadata. That does not mean the coordinate is any good.

At the top of the table is autonomous or single point positioning. This is a bare GNSS receiver with no correction of any kind, and it gives you roughly one and a half to five meters horizontally and three to ten meters vertically. That vertical figure is the one that surprises people. GNSS is always substantially worse vertically than horizontally, because the satellite geometry is inherently one-sided; you only ever see satellites above you.

SBAS, which in the United States means WAAS, brings that to sub-meter horizontally by broadcasting wide-area corrections from geostationary satellites. It is free and requires nothing but a clear view of the sky, which is why most consumer aircraft use it.

Code differential GNSS is largely a legacy technology at this point and I mention it mainly so you recognize it in older literature.

Precise Point Positioning gets you to centimeters using global orbit and clock products, but requires a convergence period of roughly ten to thirty minutes. PPP-RTK, offered through commercial correction services, shortens that convergence dramatically in exchange for a subscription.

RTK is what most survey-grade drone work uses. A base station at a known position computes corrections and transmits them to the rover in real time, and once the receiver resolves carrier phase ambiguities you are at one to two centimeters horizontally. Note the "plus one ppm" term. That is a baseline-length-dependent error, so accuracy degrades as you get farther from your base.

Network RTK, or VRS, replaces your physical base with a network of continuously operating reference stations, delivered over cellular using the NTRIP protocol. It is enormously convenient because you do not haul or set up a base, but you are dependent on cell coverage and on network geometry.

PPK, post-processed kinematic, is the one I would encourage you to consider seriously for research work. Instead of a live correction link, you log raw observables at both the rover and the base, and you combine them afterward in software. Accuracy is essentially equivalent to RTK, but you eliminate the radio link as a point of failure, you can process forward and backward in time through brief signal outages, and you can reprocess if something goes wrong. For an operation where you cannot afford to discover a problem after you have left the site, that reprocessing ability is worth a great deal.

One last point. None of this eliminates the need for independent checkpoints. Whatever positioning method you use, you should establish some ground points measured by an independent means and hold them out of the solution, so you have an actual accuracy assessment rather than the software's own optimistic estimate of its internal consistency.

Legal Considerations

- **Part 107**
  - Workhorse for commercial, research, and government-adjacent flying
  - Remote Pilot Certificate (aeronautical knowledge test, renewed every 24 months via free online recurrent training)
  - Baseline limits
    - Under 55 lbs
    - Visual line of sight
    - Daylight (or civil twilight with anti-collision lighting)
    - 400 ft AGL (or within 400 ft of a structure)
    - Max 100 mph
    - One aircraft per pilot
    - No operation over people or moving vehicles (unless the aircraft meets an operational category)
    - Most limits are waivable through the FAA's DroneZone portal
- **Recreational flying**
  - Under 49 U.S.C. § 44809
  - No pilot certificate
  - Must pass TRUST (free, unproctored)
  - Fly strictly for recreation
  - Follow a community-based organization's safety guidelines
  - Stay under 400 ft
- **Public aircraft operations**:
  - Public universities can often operate under a Certificate of Waiver or Authorization (COA) as a public aircraft operator rather than under Part 107, which can be more flexible for research but requires institutional self-certification of airworthiness and pilot competency

10

Now we get into the regulatory environment, and I want to walk through this carefully because it determines what you can legally do.

In the United States, there are essentially three regimes you might operate under. Part 107 is the first, and it is the workhorse. Nearly all commercial, research, and government-adjacent flying happens under Part 107.

To fly under Part 107 you need a Remote Pilot Certificate, which you get by passing an aeronautical knowledge test, and which you maintain by completing free online recurrent training every twenty-four calendar months.

The baseline operating limits are on the slide. The aircraft must be under fifty-five pounds. You must maintain visual line of sight, meaning you can see the aircraft with unaided vision, and note that this is not satisfied by watching the video feed. Operations are limited to daylight or civil twilight with appropriate anti-collision lighting. You are limited to four hundred feet above ground level, with an important exception: you may exceed that when flying within four hundred feet of a structure, which is what makes tall-structure inspection possible. Maximum groundspeed is one hundred miles per hour. One aircraft per pilot, meaning no swarms without a waiver. And no operation over people or moving vehicles unless the aircraft meets one of the operational categories that permit it.

Here is the crucial part: most of these limits are waivable. The FAA's DroneZone portal handles waiver applications for night operations, beyond visual line of sight, operations over people, and multiple aircraft. Waivers require you to demonstrate an equivalent level of safety, and they take time, but they exist. If your research design requires something outside the baseline, the answer is not necessarily no, it is that you need to apply early.

The second regime is recreational flying under Section 44809. No pilot certificate required, but you must pass TRUST, a free and unproctored online test, fly strictly for recreational purposes, follow a community-based organization's safety guidelines, and stay under four hundred feet. The word "strictly" carries the weight here. If you are collecting data for a class, for a thesis, or for anything that furthers a business or institutional purpose, that is not recreational flight, regardless of whether anyone is paying you. Student research is Part 107 work.

The third regime is public aircraft operations, and this is directly relevant to those of us at a public university. Public institutions can often operate under a Certificate of Waiver or Authorization as a public aircraft operator rather than under Part 107. This can be considerably more flexible for research, but it comes with a real obligation: the institution self-certifies airworthiness and pilot competency, which means your university is taking on responsibilities that the FAA would otherwise handle. If your institution has a COA, find out what it covers and who administers it before you assume it applies to your project.

Registration and Identification

- Registration is required for anything over 0.55 lb
- Remote ID
  - Every registered drone must broadcast identity, position, and takeoff location, either via standard built-in Remote ID or a broadcast module
  - The exception is flying within an FAA-Recognized Identification Area (FRIA), which are established at fixed sites through community organizations or educational institutions

11

Two administrative requirements that apply regardless of which regime you are flying under.

First, registration. Anything over 0.55 pounds must be registered with the FAA. That threshold is low enough that essentially every drone you would use for mapping is included. Registration under Part 107 is per-aircraft; recreational registration covers the operator across all their aircraft. The registration number must be displayed on the exterior of the aircraft.

Second, Remote ID, which is the more substantive requirement and the newer one. Every registered drone must broadcast its identity, its position, and its takeoff location. Think of it as a digital license plate that anyone with a receiver, including law enforcement, can read in real time. Most aircraft manufactured recently have standard Remote ID built in. Older aircraft need an add-on broadcast module attached to the airframe.

There is one exception worth knowing about: FAA-Recognized Identification Areas, or FRIAs. These are fixed geographic sites where aircraft without Remote ID may be flown, and they are established through community-based organizations or educational institutions. If your program flies older aircraft, or aircraft you have built yourself for research, a FRIA at your field site or test facility may be worth pursuing. Educational institutions are explicitly eligible to apply, which is a detail a lot of university programs are unaware of.

Airspace

- Class G
  - below 400 ft needs no authorization
- Controlled airspace (B, C, D, and surface-E)
  - Requires authorization
  - Obtained near-instantly through LAANC at pre-approved altitude ceilings from the UAS Facility Maps, or via DroneZone for anything above those grid values.
- National security restrictions over certain federal facilities
- National Park Service lands prohibit launch, land, and operate from within park boundaries

12

Airspace is the thing that determines where you can fly, and this is where a lot of new operators get into trouble simply by not knowing what class of airspace they are standing in.

Class G is uncontrolled airspace. It is most of the country by area, particularly in rural regions, and below four hundred feet in Class G you need no authorization at all. You still have to follow all the other Part 107 rules, but you do not need permission from anyone.

Controlled airspace is different. Class B around the largest airports, Class C around mid-size ones, Class D around towered airports, and surface-level Class E all require authorization before you fly. The good news is that the process is largely automated now. LAANC delivers authorization near-instantly up to pre-approved altitude ceilings that are published in the UAS Facility Maps. Those ceilings vary by grid cell: near a runway threshold the ceiling might be zero, meaning no automatic approval at all, while a few miles out it might be two or three hundred feet. If you need to exceed the published grid value, you apply through DroneZone, and that is a manual review that takes weeks, not seconds.

Beyond airspace class, there are national security restrictions over certain federal facilities, which are permanent and not waivable through the normal process, and there is a specific rule worth remembering for anyone doing environmental work: National Park Service lands prohibit launching, landing, and operating from within park boundaries. Note the framing. This is a land-use rule, not an airspace rule, which is precisely why the Park Service can impose it. It also means the rule attaches to your point of launch and recovery and to your operation on the ground, so launching from outside the boundary and flying in is not the clean workaround people sometimes assume it is. If you need to work over park land, contact the park's research permitting office well in advance.

What is LAANC?

- Low Altitude Authorization and Notification Capability
- A collaboration between FAA and Industry
- Automates the application and approval process for airspace authorizations
- Requests are checked against multiple airspace data sources in the FAA UAS Data Exchange such as UAS Facility Maps, Special Use Airspace data, Airports and Airspace Classes, as well as Temporary Flight Restrictions (TFRs) and Notices to Airmen (NOTAMs).
- If approved, pilots can receive their authorization in near-real time
- Unless specifically requested in an authorization, drone pilots do not need to notify the tower before they fly
- Provides
  - Drone pilots with access to controlled airspace at or below 400 feet
  - Awareness of where pilots can and cannot fly
  - Temporary Flight Restrictions (TFR) information and updates to the Special Use Airspace (SUA) schedule
  - Air Traffic Professionals with visibility into where and when drones will operate
- LAANC provides airspace authorizations only. Pilots must still check NOTAMs, weather conditions, and abide by all airspace restrictions

LAANC CONNECTED

Federal Aviation Administration

UAS Data Exchange (LAANC) | Federal Aviation Administration

Let's spend a moment on LAANC specifically, because it is the tool you will use most often.

LAANC stands for Low Altitude Authorization and Notification Capability. It is a collaboration between the FAA and industry that automates what used to be a slow manual application process for airspace authorization.

Mechanically, here is what happens. You submit a request through an approved app, specifying your location, your altitude, and your time window. That request gets checked against multiple authoritative data sources in the FAA UAS Data Exchange: the UAS Facility Maps that define approved altitude ceilings by grid cell, Special Use Airspace data, airport and airspace class boundaries, Temporary Flight Restrictions, and Notices to Airmen. If everything checks out, your authorization comes back in near-real time, often in under a minute.

Note the point on the slide that unless an authorization specifically requests it, you do not need to call the tower before you fly. LAANC has already handled that coordination. Calling the tower unprompted is not helpful and generally not welcome.

What LAANC provides is threefold. For drone pilots, it gives access to controlled airspace at or below four hundred feet, awareness of where you can and cannot fly, and current TFR and Special Use Airspace scheduling information. For air traffic professionals, it gives visibility into where and when drones are operating in their airspace, which is the whole point of the system from a safety standpoint.

And now the critical caveat, which is on the slide in the last line and which I want to emphasize. LAANC provides airspace authorization only. It does not check the weather for you. It does not release you from checking Notice to Airmen (NOTAMs). It does not clear you of any other airspace restriction. An approved LAANC authorization means the FAA has no airspace objection to your operation; every other element of preflight responsibility remains yours.

Foreign-Made Drones and Components

- The FY2025 National Defense Authorization Act (NDAA) required the FCC to add DJI and Autel equipment to the Covered List by December 22, 2025 absent a national security determination that they posed no undue risk. On December 21, an interagency body issued a determination covering not just DJI and Autel but all foreign-made UAS and UAS critical components, and the FCC updated the Covered List accordingly on December 22.
- DJI and Autel challenged the designations in February 2026, both at the FCC and in the Ninth Circuit. Importantly, the Covered List does not let the FCC ground or disable drones already in service (operators can keep flying aircraft they legally own under existing FAA rules). The restriction targets future availability, and reversing it would require new legislation, regulatory change, or a completed security review paired with policy relief. Reporting from June 2026 indicates DJI remains on the Covered List with new models blocked from import, though the FCC extended firmware and security update permissions to January 1, 2029.
- Separately, the NDAA "Blue UAS" restrictions have long barred DoD and many federally funded programs from procuring covered foreign drones. If any of your work runs on federal grants, check the specific award terms (several agencies now flow procurement restrictions down to subrecipients).

This slide covers a rapidly moving situation, and I want to be clear that this is a snapshot. Verify the current state before you make a purchasing decision.

The background: the FY2025 National Defense Authorization Act required the FCC to add DJI and Autel equipment to the Covered List by December 22, 2025, unless a national security determination concluded they posed no undue risk. On December 21, an interagency body issued a determination that covered not just those two manufacturers but all foreign-made UAS and foreign-made UAS critical components, and the FCC updated the Covered List accordingly the following day.

Both companies challenged the designations in February 2026, at the FCC and in the Ninth Circuit.

Now, the practical implications, because there was a lot of alarm about this that was not entirely warranted. The Covered List does not give the FCC authority to ground or disable aircraft already in service. If you legally own a DJI aircraft, you can keep flying it under existing FAA rules. What the Covered List restricts is future availability, primarily by blocking equipment authorization for new imports. Reversing it would require new legislation, a regulatory change, or a completed security review paired with policy relief. Reporting from June 2026 indicates DJI remains on the list with new models blocked from import, though the FCC extended permission for firmware and security updates through January 1, 2029.

There is a separate and older restriction that has arguably more day-to-day effect on academic work. The NDAA "Blue UAS" provisions have long barred the Department of Defense and many federally funded programs from procuring covered foreign drones. That matters directly for us: if any part of your work runs on a federal grant, read your specific award terms. Several agencies now flow these procurement restrictions down to subrecipients, which means the restriction can reach you even if you are two steps removed from the funding agency. Discovering that after you have bought an aircraft is an unpleasant conversation with your sponsored programs office.

Blue UAS

- Blue UAS Cleared List
  - Complete aircraft platforms approved for government procurement
- Blue UAS Framework
  - Cleared critical components (flight controllers, radios, cameras, GNSS modules), so integrators can build compliant aircraft from vetted parts.
- NDAA-compliant ≠ Blue UAS
  - A drone can be NDAA compliant without being on the Cleared List; Blue UAS is a higher level of vetting intended specifically for federal government use
- Green UAS
  - Association for Uncrewed Vehicle Systems International (AUVSI) parallel certification for the commercial market.
  - It now functions as a fast-track
  - Platforms completing Green UAS certification can move into Blue UAS Cleared without a separate full assessment

15

Since Blue UAS came up on the previous slide, let's define the terms, because they get used loosely and they do not mean the same thing.

The Blue UAS Cleared List is a list of complete aircraft platforms approved for government procurement. If a platform is on the Cleared List, a federal agency can buy it without additional vetting.

The Blue UAS Framework is a parallel list of cleared critical components: flight controllers, radios, cameras, GNSS modules. The purpose is to let integrators assemble compliant aircraft from vetted parts rather than requiring every finished platform to go through a full independent assessment.

Now the distinction that trips people up. NDAA-compliant is not the same as Blue UAS. A drone can be NDAA compliant, meaning it does not use covered components from covered countries, without appearing on the Cleared List. Blue UAS represents a higher level of vetting, specifically intended for federal government use. Vendors sometimes blur this, so if your requirement is genuinely Blue UAS, check the actual list rather than taking marketing language at face value.

Green UAS is AUVSI's parallel certification aimed at the commercial market rather than the federal one. Its role has evolved: it now functions as a fast track, in that platforms completing Green UAS certification can move into Blue UAS Cleared status without going through a separate full assessment. That is a meaningful development, because it means smaller manufacturers have a more tractable path into federal eligibility than they once did.

State and Local

- Federal law preempts airspace regulation
- States and localities retain authority over land use, privacy, trespass, harassment, and voyeurism
- Many states have drone-specific statutes on surveillance of private property, interference with first responders, and use over correctional facilities
- A growing number of states have also enacted their own restrictions on Chinese-made drones for government agencies

16

We have been talking about federal rules, but that is not the whole picture.

The foundational principle is that federal law preempts regulation of the airspace itself. A city cannot set its own altitude limits or create its own drone licensing scheme; that authority belongs to the FAA.

But states and localities retain substantial authority in adjacent areas: land use, privacy, trespass, harassment, and voyeurism. This is the same mechanism we saw with the Park Service rule. A jurisdiction cannot tell you how high you may fly, but it can absolutely tell you that you may not launch from its public park, and it can prosecute you under a privacy statute for what you recorded while airborne.

Many states have enacted drone-specific statutes. The common categories are surveillance of private property, interference with first responders at an incident scene, and use over correctional facilities. And, connecting back to the previous slides, a growing number of states have enacted their own restrictions on Chinese-made drones for use by state government agencies, which for those of us at public universities can mean a state-level restriction layered on top of the federal one.

The practical takeaway is that compliance is layered. Federal authorization is necessary but not sufficient. Before flying at an unfamiliar site, you need to know the state statute, any local ordinance, and the land manager's own rules for the parcel you are launching from.

Pilot Certification

- Requirements
  - 16 or older
  - Able to read, speak, write, and understand English
  - In a physical and mental condition to safely operate
  - Pass the initial knowledge test
- UAG (Unmanned Aircraft General – Small) exam
  - 60 multiple-choice questions
  - 70% to pass,
  - $175
  - Two hours allotted
- Must get a FTN (FAA Tracking Number) before taking the exam
- Content
  - Regulations (~25%)
  - Operations (~20%)
  - Other = airspace classification and chart reading, weather and weather sources, loading and performance, radio communication procedures, emergency procedures, crew resource management, aeronautical decision-making, and airport operations
- You must complete the free online recurrent course ALC-677 every 24 calendar months to remain qualified to operate under Part 107

Let's talk about getting certified, since most of you who want to do this work will need to.

The eligibility requirements are modest. You must be at least sixteen years old; able to read, speak, write, and understand English; in a physical and mental condition to operate safely; and you must pass the initial knowledge test. Note that there is no medical certificate requirement and no flight hour requirement. There is no practical flight test. The entire barrier is the written exam.

The exam is the Unmanned Aircraft General – Small, abbreviated UAG. It is sixty multiple-choice questions, you need seventy percent to pass, it costs one hundred seventy-five dollars, and you have two hours. You must obtain an FAA Tracking Number, or FTN, through the IACRA system before you can schedule it. Do that first, because it catches people off guard on test day.

The content breakdown is roughly a quarter regulations and a fifth operations, with the remainder distributed across airspace classification and chart reading, weather and weather sources, loading and performance, radio communication procedures, emergency procedures, crew resource management, aeronautical decision-making, and airport operations.

I will offer some candid advice on where people struggle. The regulations are memorization and most people handle them fine. The section that consistently causes failures is sectional chart reading. You have to interpret aeronautical charts: identify airspace classes from the symbology, read the altitude figures, recognize special use airspace, decode the airport data blocks. If you have no aviation background, this is genuinely unfamiliar and it will take real study time. Do not leave it until the end.

Weather is the second common weak point, particularly density altitude, stability, and reading METARs and TAFs.

Finally, certification is not permanent. You must complete the free online recurrent course, ALC-677, every twenty-four calendar months to remain qualified to operate under Part 107. It is free, it is online, and it is not proctored, but if you let it lapse, you are not legally current.

Studying Resources

- FAA Remote Pilot – sUAS Study Guide (FAA-G-8082-22)
  - The FAA's own study guide, mapped directly to the test topics
  - Remote Pilot - Small Unmanned Aircraft Systems Study Guide
- Airman Certification Standards for Remote Pilot (FAA-S-ACS-10)
  - The actual specification of what's testable.
  - Remote Pilot – Small Unmanned Aircraft Systems Airman Certification Standards
- Airman Knowledge Testing Supplement (FAA-CT-8080-2H)
  - Computer Testing Supplements | Federal Aviation Administration
- Pilot's Handbook of Aeronautical Knowledge
  - Pilot’s Handbook of Aeronautical Knowledge | Federal Aviation Administration

If you are going to sit for the exam, here are the resources I would use, and I would emphasize that all of them are free and published by the FAA.

The *FAA Remote Pilot Small UAS Study Guide*, document number FAA-G-8082-22, is the FAA's own study guide, mapped directly to the test topics. Start here.

The *Airman Certification Standards for Remote Pilot*, FAA-S-ACS-10, is arguably the more valuable document, and it is the one most people skip. This is the actual specification of what is testable. Every knowledge element the exam can draw from is enumerated. If you work through the *ACS* and can address each element, you have covered the exam by construction rather than by hoping your study guide was comprehensive.

The *Airman Knowledge Testing Supplement*, FAA-CT-8080-2H, contains the figures, charts, and legends you will be given during the exam. This is important for a practical reason: you should practice with the exact materials you will have in the testing center. The chart excerpts you will be asked to interpret come from this supplement, so working through it in advance removes a great deal of test-day friction.

And the *Pilot's Handbook of Aeronautical Knowledge* is the broad reference for crewed aviation. It is far more than you need for the UAG exam, but it is where I would go for a genuine explanation of weather, aerodynamics, or airspace, as opposed to the compressed treatment in the study guide.

There are also commercial test prep courses, and some of them are good. But I would not buy one before working through the free materials first. The *ACS* in particular is the document that tells you what you actually need to know.

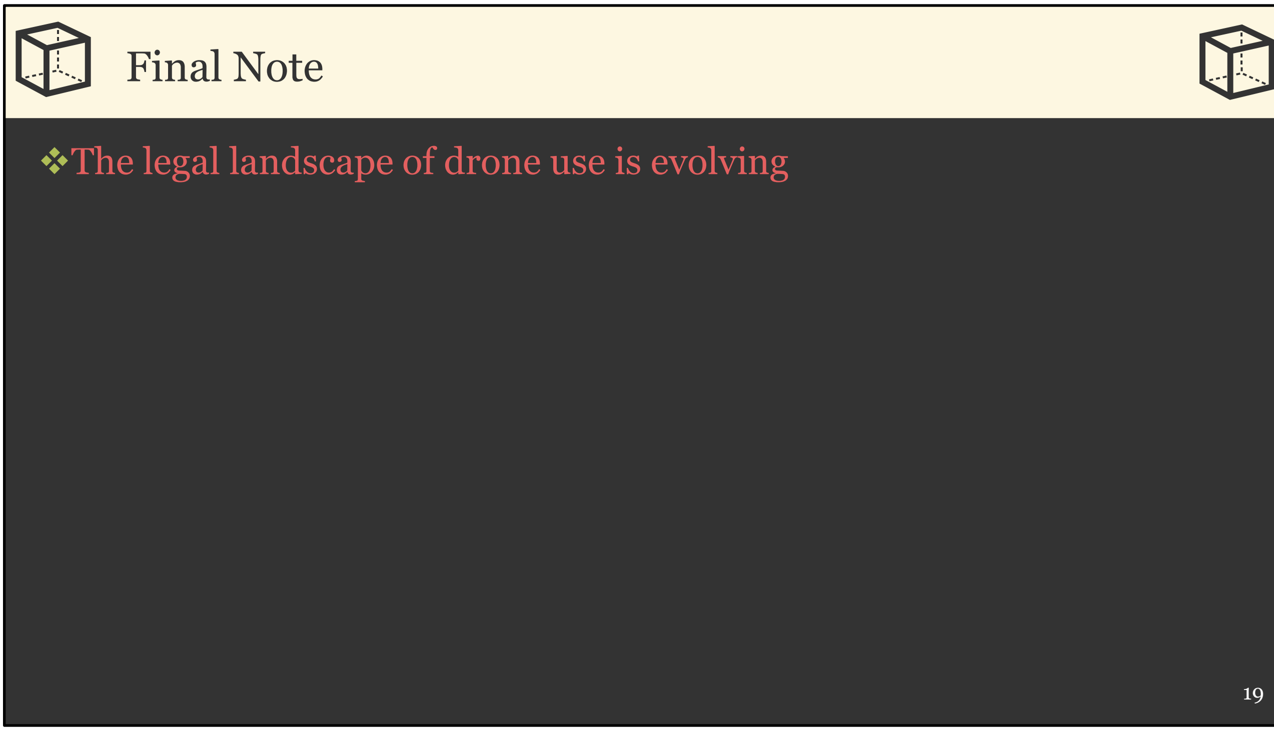

I want to close this regulatory section with a caveat, and I want it on its own slide because it is that important.

The legal landscape of drone use is evolving, and it is evolving quickly. The material I just presented reflects the situation as of the preparation of this lecture. Some of it may already have changed by the time you work through this material.

The areas most likely to shift are the ones I flagged as being in motion: the Covered List situation and the litigation surrounding it, the ongoing rulemaking around beyond-visual-line-of-sight operations, which would be a significant expansion of what is routinely possible, and state-level legislation, which is being introduced continuously.

The stable core is unlikely to change soon: the certification requirement, the basic operating limits, the airspace authorization framework, and the layered federal-state-local structure of authority. Learn that structure, because it is what lets you situate a new rule when you encounter one.

The practical habit I would encourage is this: before any flight at an unfamiliar site or under unfamiliar conditions, check current guidance rather than relying on what you learned once. The FAA's UAS pages and DroneZone are authoritative. Secondhand summaries, including this lecture, are not.

Flight Planning

Obtaining Desired Spatial Resolution

- Ground sampling distance (GSD)
- Depends on:
  - Flight altitude (AGL)
  - Lens focal length
  - Pixel pitch (physical center-to-center spacing between adjacent photosites on the imaging sensor, expressed in micrometers; the sensor's physical resolution unit; sensor width divided by pixel count along that dimension)
  - (pixel pitch × flying height) / focal length
- Since elevation changes, need to adjust flight height during flight to maintain a constant GSD

Ground sampling distance, or GSD, is the ground dimension represented by one pixel in your imagery. It is the resolution figure that actually matters, as opposed to the megapixel count on the box.

GSD depends on three things. Flight altitude above ground level, which is the one you control. Lens focal length, which is fixed for a given camera. And pixel pitch.

Let me define pixel pitch carefully, because it is the term people are least comfortable with. Pixel pitch is the physical center-to-center spacing between adjacent photosites on the imaging sensor, expressed in micrometers. It is the sensor's physical resolution unit. You compute it by dividing the sensor's physical width by the pixel count along that dimension. So, a sensor thirteen millimeters wide with five thousand pixels across has a pitch of about 2.6 micrometers.

The relationship is simple: GSD equals pixel pitch times flying height, divided by focal length. Keep your units consistent and it works out directly. Double your altitude and your GSD doubles, meaning your resolution gets coarser by a factor of two. Halve your altitude and you get twice the detail, at the cost of covering a quarter of the area per image and needing roughly four times as many images.

That tradeoff is the fundamental tension in flight planning. Detail costs coverage, flight time, battery, and processing time, quadratically.

Now the point at the bottom of the slide, which is the one people miss. Since elevation changes across a site, flying at a constant altitude above your takeoff point means your height above ground varies, which means your GSD varies. If your site has a hundred meters of relief, and you flew at a hundred meters above the launch point, your imagery over the ridge tops is at a very different resolution than your imagery in the valley bottom. In an extreme case you fly into the terrain.

The solution is terrain following, sometimes called terrain awareness, where the flight plan uses an existing elevation model to modulate the aircraft's altitude and hold a constant height above ground. Most modern planning software supports this, usually by ingesting a Shuttle Radar Topography Mission (SRTM) or user-supplied DEM.

Overlap

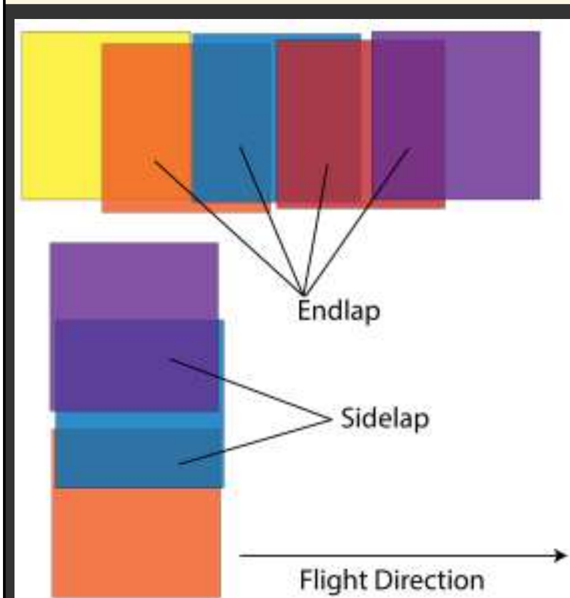

- 75/65 forward/side to support structure from motion
- Higher for forest canopy (85/80 forward/side)

22

Overlap is the second parameter you set, and it is the one that structure from motion is most sensitive to.

The reason we need overlap at all is fundamental to how photogrammetric reconstruction works. To determine the three-dimensional position of a point, that point must be visible in multiple images taken from different positions. The algorithm identifies the same physical feature across images, and the geometric relationship between those observations is what yields depth. A point seen in only one image has no depth information at all.

The conventional starting point is seventy-five percent forward overlap and sixty-five percent side overlap. Forward overlap, sometimes called endlap, is the overlap between successive images along a flight line, and you control it through the interval between shutter triggers relative to ground speed. Side overlap, or sidelap, is the overlap between adjacent flight lines, controlled by line spacing.

Those numbers feel high to people used to traditional aerial photography, where sixty and thirty were common. The difference is that classical photogrammetry used a small number of precisely calibrated images with known exterior orientation. Structure from motion instead solves everything simultaneously from a large number of redundant, uncalibrated observations.

Redundancy is the mechanism, so it needs much more of it.

Now the important part on this slide: increase overlap for forest canopy, to something like eighty-five forward and eighty side. The reason is that canopy is a difficult matching surface. It is geometrically complex, it is self-similar, leaves move between exposures, and shadowing changes the appearance of the same feature from different angles. Feature matching fails much more often, so you need substantially more redundancy to get an acceptable reconstruction. The same logic applies to any low-texture or repetitive surface: water, sand, snow, fresh asphalt, a uniform agricultural field.

One practical note on cost. Because overlap compounds in two dimensions, going from seventy-five/sixty-five to eighty-five/eighty is not a modest increase. It roughly doubles your image count, and processing time scales worse than linearly with image count. Plan your batteries and your compute accordingly.

Other considerations

- Sensor-specific considerations
- Sun angle (generally 10 to 2 PM is best)
- Wind speeds at flight height
- Air density at flight altitude
- Battery life
- Obstacles

Beyond resolution and overlap, here are the other factors that go into a flight plan.

Sensor-specific considerations come first, because they modify everything else. A thermal sensor has a much coarser native resolution than the RGB camera beside it, so achieving a given thermal GSD requires a much lower flight. A multispectral array may need a calibrated reflectance panel imaged before and after the flight. A lidar payload has its own scan geometry and pulse rate that determine point density, and it has different, generally more forgiving, requirements about illumination.

Sun angle is one of the most important and most frequently ignored. Roughly ten in the morning to two in the afternoon is the standard window. The reasoning is that low sun angles cast long, hard shadows, and shadows are a serious problem for photogrammetry. A shadow edge is a strong image feature, but it is a feature that moves as the sun moves, so the algorithm can match it across images and produce geometry that does not exist. Shadowed areas also lose detail entirely, so the matcher has nothing to work with. Related to this: try to complete a survey within a reasonably tight time window, since illumination changing substantially across a long collection introduces radiometric inconsistency between the early and late images.

Wind speed at flight altitude, not at ground level, is what matters. Wind is typically stronger aloft, and a wind that is unremarkable where you are standing may exceed your aircraft's capability at altitude. Headwinds and tailwinds along flight lines produce inconsistent ground speed, which means inconsistent forward overlap, and crosswinds force the aircraft to crab, which changes the image footprint from what you planned.

Air density at flight altitude affects lift. This is the density altitude concept from the exam, and it matters practically: high elevation, high temperature, and high humidity all reduce air density, which reduces the thrust your rotors produce. A flight profile that works in the valley on a cool morning may not work on a ridge on a hot afternoon.

Battery life determines how you break a site into flights. Plan conservatively. Manufacturer flight time figures assume hover in still air at moderate temperature, and none of those conditions apply to an actual survey in wind and cold. I would plan to land with a substantial reserve.

And obstacles. Trees, towers, power lines, and terrain. Power lines are particularly hazardous because they are difficult to see and difficult for obstacle avoidance sensors to detect. Survey the site visually before you fly it, and build a mental picture of where the vertical hazards are.

A brief survey of the software you use to build and fly a mission.

On the desktop side, UgCS is a commercial planning package notable for strong terrain-following support and for handling complex mission geometries, including facade and corridor missions, and it works across a range of airframes. Mission Planner and QGroundControl are both open-source ground control stations built around the MAVLink protocol. Mission Planner is associated with the ArduPilot ecosystem, QGroundControl with PX4, though there is considerable overlap. If you are working with a research or custom-built platform rather than a sealed consumer product, these are the tools you will use, and they expose essentially every parameter of the flight.

On mobile, DJI Pilot 3 is the first-party app for DJI's enterprise line. Map Pilot Pro, DroneDeploy, Dronelink, and Litchi are third-party apps that add mission planning capability to consumer aircraft. Each has its own strengths. DroneDeploy pairs planning with a cloud processing pipeline. Dronelink is particularly good for complex, scripted missions including orbital and facade work, which is directly relevant to the mesh collection we will discuss shortly. Litchi has long been popular for waypoint missions and has a strong following for cinematic work.

Two practical points. First, be aware of the dependency you are accepting.

Third-party apps depend on manufacturer software development kits (SDKs), and when a manufacturer changes its SDK or its policy, an app can lose functionality on short notice. Second, look at what the software does with your data. Some of these platforms are designed to push your imagery into a proprietary cloud pipeline. That may be exactly what you want, but for research work you should know whether you can get raw imagery with complete metadata out of the system, and where your data is physically stored.

Product Generation

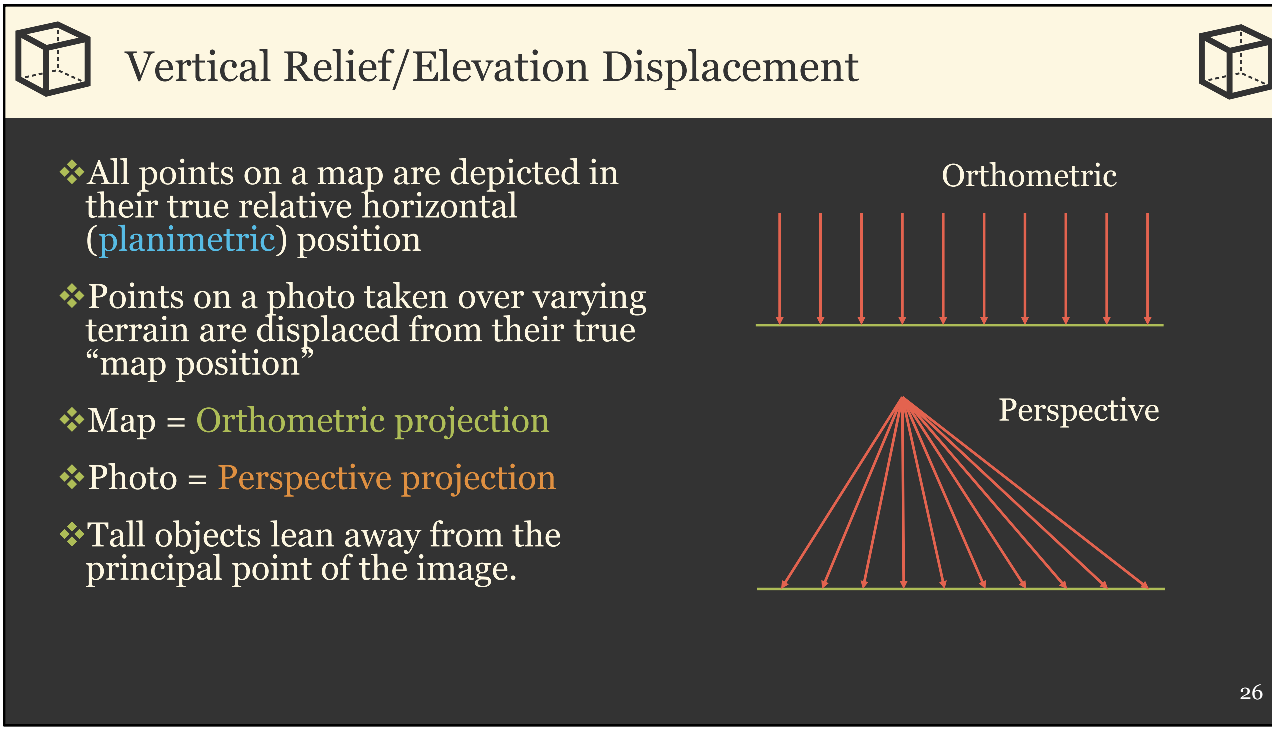

This slide, and the following two slides, were taken from our *Remote Sensing* course as a review of vertical relief displacement and orthophotography.

Images represent a perspective view of the landscape as opposed to an orthometric view; the data are collected from a vantage point as opposed to directly above each individual location in the extent.

This results in a perspective projection as opposed to an orthographic projection, as demonstrated by the graphics on this slide. Practically, this viewing geometry results in some impacts on representation of the landscape in the resulting images.

Most notable, is the impact of terrain. Points on a photo taken over varying terrain are displaced from their true map position. Also, tall objects tend to lean away form the center of the image, or image principal point.

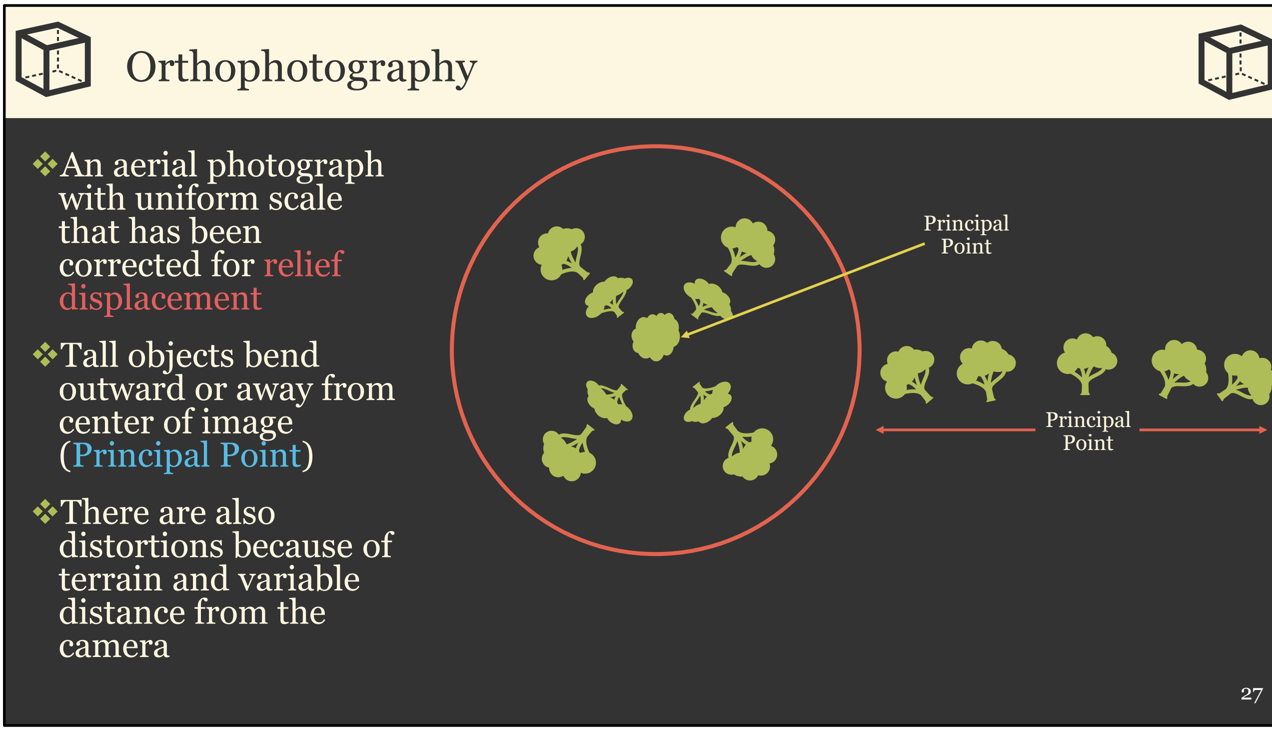

This slide further discusses the concept of vertical relief displacement resulting from the perspective view.

Again, tall objects, such as buildings and trees, will bend or lean away from the center or principal point of the image.

Also, variability in the terrain surface means that the actual landscape surface will occur at different distances from the sensor. So, the location and scale of the landscape surface is also impacted or distorted.

Orthophotography

Orthorectification = A process used to remove vertical relief displacement and give image uniform scale

To do this, you need to know information about:

- The flight
- The camera
- The terrain

Only orthophotography should be used as reference information to digitize and make maps

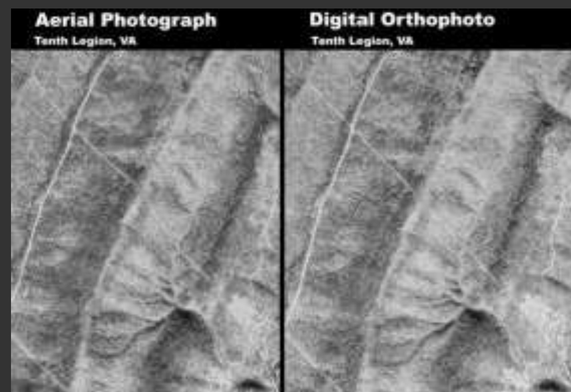

https://www.usgs.gov/media/images/aerial-photograph-vs-orthoimage

28

Orthorectification is the process used to remove vertical relief displacement and give an image uniform scale. The goal here is to transform the image from a perspective to orthometric projection, which will alleviate the impact of topography, height, and viewing geometry. The resulting orthophotograph will have consistent scale throughout.

In order to perform orthorectification, you need to have information about the flight, camera, and terrain.

Once orthorectification has been performed, the image has been converted to an orthophotograph.

The difference is highlighted in the provided images. The first image is the raw image with scale distortions. Note that the long, linear clearing seems to “zig-zag” as it goes up and down the topography. This is because the ground surface-to-sensor distance varies due to the terrain. Once orthorectification has been performed, the linear feature appears to be straighter. This is because the scale distortion was removed.

The most important point here is that only orthophotography should be used as reference information to digitize and make maps in a GIS. If orthorectified imagery is used, the scale distortions will be incorporated into the products.

If you have spent any time viewing aerial imagery, such as in Google Earth, you may have noticed that many buildings still appear to lean in the orthophotography. That is because a terrain surface or digital elevation model is commonly used in the orthorectification process. Since no above ground features are represented in the elevation surface, only relief displacement associated with the topography is fixed. Above ground features, such as buildings, are still distorted.

Structure from Motion (SfM) Requirements

Parameter	Typical target	Notes
Forward overlap	75–80%	85%+ over canopy or low-texture surfaces
Side overlap	65–70%	80% over canopy; 60% is the practical floor
GSD	2–4× finer than desired DSM cell	A 5 cm DSM wants ~1.5–2 cm imagery
Images per ground point	≥5, ideally 9+	Below 5, depth estimates are unreliable
Motion blur	<0.5 pixel	Ground speed × exposure < GSD/2
Base-to-height ratio	~0.3–0.5	Too small degrades depth precision; too large breaks matching
Oblique coverage	20–35° tilt, cross pattern	Not optional for accurate DSMs

This table gives you the actual numbers for a structure-from-motion collection. I would treat this as a reference to come back to when you are planning a flight.

Forward overlap of seventy-five to eighty percent, going to eighty-five or higher over canopy or low-texture surfaces. Side overlap of sixty-five to seventy percent, going to eighty over canopy. Note that sixty percent is listed as the practical floor. Below that, you will get gaps and failures, and no amount of processing effort will fix it.

GSD should be two to four times finer than the cell size of the surface model you intend to produce. So, if you want a five centimeter DSM, you want imagery at roughly one and a half to two centimeters. This is a point people get wrong constantly. The intuition is that if you want five centimeter output you should collect five centimeter imagery, but photogrammetric reconstruction is a statistical process. Depth estimates are noisy at the pixel level, and the way you get a clean surface is by having many redundant, finer observations to aggregate within each output cell.

Images per ground point should be at least five, ideally nine or more. Below five, depth estimates are unreliable. This is the parameter that actually matters, and overlap percentage is just the means by which you achieve it. Notice that this is also why edges of your survey area are always weaker than the middle: points near the boundary are seen in fewer images by construction. Always fly a buffer beyond your area of interest.

Motion blur under half a pixel. The rule is that ground speed times exposure time must be less than half your GSD. This is a real constraint that limits how fast you can fly, and it is why low light is a problem for mapping: a longer exposure at a given ground speed means more smear. Blurred images do not just produce blurry output, they produce failed feature matches.

Base-to-height ratio around 0.3 to 0.5. This is the ratio of the distance between camera positions to the flying height, and it captures a genuine tension. Too small a baseline means the two viewing rays to a point are nearly parallel, so the intersection angle is shallow and depth precision is poor. Too large a baseline means the images look too different from each other and feature matching starts to fail. The listed range is the sweet spot.

And finally, oblique coverage: twenty to thirty-five degrees of camera tilt, flown in a cross pattern. Note the annotation, which I want to underline: this is not optional for accurate DSMs. A pure nadir block produces a systematic vertical error, a doming or bowl effect, that arises from a correlation between the estimated lens distortion parameters and the surface shape. Oblique imagery breaks that correlation. If you have ever produced a surface model with an inexplicable dome in the middle, this is almost certainly why.

SfM Camera Considerations

- Fixed focal length
- Focus locked at infinity
- Mechanical or global shutter (rolling shutter introduces a systematic distortion proportional to ground speed)
- Manual exposure or at minimum fixed aperture and ISO (auto-exposure varying across a block introduces radiometric inconsistency that hurts matching at block edges)
- Aperture around f/4–f/5.6 on typical drone sensors (enough depth-of-field, before diffraction starts eating the pixel pitch)
- Diffuse, overcast light is ideal (hard shadows create features that move between frames and produce spurious geometry; specular highlights on wet surfaces do the same)
- Shoot RAW if the pipeline supports it, though JPEG is usually acceptable for geometry

Camera configuration for structure from motion. Photogrammetric reconstruction wants boring, consistent imagery, which is nearly the opposite of what a camera's automatic modes are designed to give you.

Fixed focal length. If you have a zoom lens, lock it. The bundle adjustment solves for a single interior orientation across the block, so a focal length that changes between images breaks that assumption.

Focus locked at infinity, for the same reason. Autofocus changes the effective focal length subtly between exposures, and it introduces variation the calibration cannot account for.

Mechanical or global shutter is a meaningful advantage. A rolling shutter reads the sensor progressively rather than all at once, so if the aircraft is moving, different rows of the image correspond to different camera positions. That introduces a systematic distortion proportional to ground speed. Good software can model it, but the correction is imperfect, and if you are choosing a payload, a mechanical shutter is worth paying for.

Manual exposure, or at minimum fixed aperture and ISO. Auto-exposure varying across a block introduces radiometric inconsistency, and where that hurts most is at block edges and between flight lines flown at different times, which is exactly where reconstruction is already weakest. Set exposure for the scene before you launch and leave it alone.

Aperture around f/4 to f/5.6 on typical drone sensors. This is a genuine optimum, not a rule of thumb. Wider apertures give you insufficient depth-of-field and worse optical performance at the edges of the frame. Narrower apertures run into diffraction, where the physical size of the diffraction pattern starts to exceed your pixel pitch, and you lose actual resolution regardless of how many pixels the sensor has.

Diffuse, overcast light is ideal, which is the opposite of what people expect. Hard shadows create high-contrast features that move between frames as the sun moves, and the matcher can lock onto them and produce spurious geometry. Specular highlights on wet surfaces do the same thing, since a highlight is a function of viewing angle rather than of the surface itself. A flat overcast day is a photogrammetrist's friend.

And shoot RAW if the pipeline supports it. For geometry alone, JPEG is usually acceptable, since compression artifacts are small relative to the features being matched. RAW matters more if you care about radiometry, if you intend to do any spectral analysis, or if you are producing an orthomosaic where color consistency across the block is important.

Processing Pipeline

1. Trajectory processing (PPK only) = GNSS post-processing for camera positions at exposure time
2. Image preparation = cull unusable frames, preserve original geometry
3. Feature detection and matching = identify corresponding keypoints across overlapping images
4. Sparse reconstruction/bundle adjustment = solve camera poses, tie point coordinates, and lens calibration together
5. Georeferencing and re-optimization = apply control, re-solve, refine
6. Dense reconstruction = depth maps per image, fused into a dense point cloud
7. Point cloud conditioning = noise removal, blunder cleanup
8. Surface generation = rasterize to DSM (DTM would require point classification to separate out just ground
9. Orthomosaic generation
10. Validation = checkpoint residuals and spatial pattern (back to 5 or 4 if it fails)

This is the processing pipeline. I want to walk through each stage so that when you run software and watch a progress bar, you know what it is actually doing.

Trajectory processing applies only if you flew PPK. This is the GNSS post-processing step that combines your rover and base observables to compute accurate camera positions at the exact moment of each exposure. It happens before anything else, because everything downstream uses those positions.

Image preparation means culling unusable frames, meaning the blurred ones, the ones shot during turns, the ones of the inside of your car. The critical instruction here is to preserve the original geometry. Do not crop, do not resize, do not apply lens correction in a photo editor. The pipeline needs the original image geometry and the original EXIF metadata to solve for camera calibration.

Feature detection and matching identifies distinctive keypoints in each image and finds correspondences across overlapping images. This is where texture matters, and it is where a smooth water surface or a uniform snowfield fails.

Sparse reconstruction and bundle adjustment is the mathematical core. It simultaneously solves for camera poses, the three-dimensional coordinates of the tie points, and the lens calibration parameters, by minimizing reprojection error across the entire block at once. The word "together" on the slide is doing a lot of work. These parameters are correlated with each other, which is exactly why the oblique imagery we discussed matters so much: it decorrelates the lens parameters from the surface geometry.

Georeferencing and re-optimization applies your control, whether that is RTK or PPK camera positions or surveyed ground control points, and then re-solves. The re-solve matters. Applying control as a rigid transformation after the fact is not the same thing as incorporating it into the adjustment.

Dense reconstruction computes a depth map for each image and fuses them into a dense point cloud. This is typically the most computationally expensive stage and the one where GPU capability matters most.

Point cloud conditioning is noise removal and blunder cleanup. There will be blunders, particularly around water, around vegetation edges, and in shadow.

Surface generation rasterizes the cloud into a digital surface model. And note the parenthetical: a DTM, a bare-earth model, requires point classification to separate ground returns from everything else, and that is a distinctly harder problem with photogrammetric data than with lidar, because photogrammetry has far fewer genuine ground points under any kind of vegetation.

Orthomosaic generation produces the corrected, uniform-scale image product we discussed earlier.

And validation, which is the stage most often skipped. Compute residuals at independent checkpoints, points that were held out of the solution entirely. And look at the spatial pattern of those residuals, not just the RMSE. A systematic pattern, a dome, a tilt, a gradient across the site, tells you something quite different from randomly distributed error, and it points at a specific cause. Note the instruction on the slide: if validation fails, you go back to stage five or stage four and re-solve. It is a loop, not a straight line.

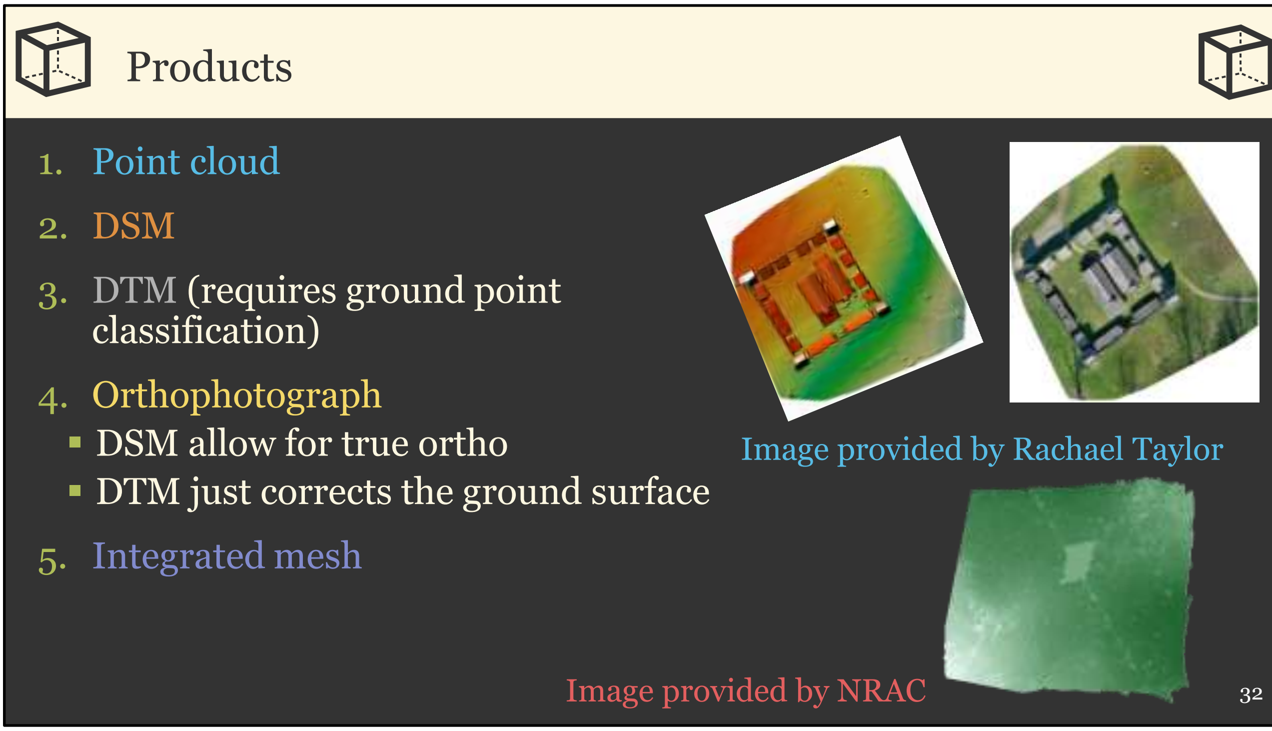

These are the products that come out the other end, and this is where drone data connects directly to everything else in this course.

The point cloud is the fundamental product, and it is the one I would encourage you to keep. It is a dense collection of three-dimensional points, usually with color attached from the source imagery. Everything else on this list is derived from it, so if you retain the cloud you retain the ability to regenerate other products later with different parameters. Store it in LAS or LAZ so it interoperates with the lidar tooling we have used elsewhere in this course.

The DSM, digital surface model, is a raster of the top of everything: canopy, rooftops, and bare ground where nothing else is present. This is what photogrammetry gives you naturally, because the cameras only see the top surface.

The DTM, digital terrain model, is the bare-earth surface, and note the qualifier: it requires ground point classification. In open terrain this works well. Under forest canopy, photogrammetric DTMs are unreliable, sometimes badly so, because there simply are not enough points that reached the ground for the classifier to work with. This is the single strongest argument for a lidar payload if forested terrain is your application.

The orthophotograph is the corrected image product. Now look at the two sub-points, because they explain a distinction that confuses a lot of people. If you orthorectify using the DSM, you get what is called a true orthophoto, in which buildings and trees are corrected along with the terrain, so structures appear in plan view with no lean. If you orthorectify using only the DTM, you correct the ground surface but above-ground features retain their relief displacement, which is why buildings still lean in most imagery you have looked at, including what you see in web mapping services. True orthos have their own difficulty: correcting a building to plan view leaves a hole where the building was leaning, and that occluded area has to be filled from other images, which is why true orthos sometimes have visible artifacts along building edges.

And the integrated mesh: a continuous triangulated surface with photographic texture draped over it. This is the product that most resembles the real world visually, and it is what we care about most for 3D geovisualization. It is what feeds an integrated mesh scene layer, and it is what gives you a walkable, viewable model rather than a set of raster surfaces. We will talk about how to collect for it on the next slide.

The image here is provided by NRAC, the National Resources Analysis Center.

Integrated Mesh Considerations

- Oblique coverage is important
- Nadir grid, then four oblique blocks at roughly 40–45° camera pitch flown in orthogonal directions
- Each surface should be seen from at least three substantially different directions, and vertical surfaces need near-horizontal rays
- Orbital passes around any structure of interest, with the camera pointed inward, at two or three altitudes
- Multiple altitudes generally (a high pass for coverage and context, a low pass for detail on the subject).
- Overlap of 80/80 or higher (Mesh quality degrades visibly before geometric accuracy does)
- Generally, always some occlusion (especially when vegetation is present)
- Will likely require some manual clean up

33

Integrated mesh generation deserves separate treatment because the collection requirements are meaningfully different from what you would fly for a DSM or an orthomosaic.

The core issue is this. A DSM is two and a half dimensional. It has one elevation value per horizontal location, so a nadir flight is sufficient by construction. A mesh is genuinely three-dimensional. It has to represent vertical surfaces, overhangs, and undersides, and none of those are visible from directly overhead. If your only imagery is nadir, your building walls will be smeared vertical stretches of roof texture, which looks exactly as bad as it sounds.

So oblique coverage is not a refinement here, it is the requirement. The standard approach is a nadir grid for overall coverage, followed by four oblique blocks flown at roughly forty to forty-five degrees of camera pitch, in orthogonal directions. Four directions because a wall facing north is only well-imaged by the block flying so as to look north.

The general principle underneath that: every surface should be seen from at least three substantially different directions, and vertical surfaces specifically need near-horizontal rays. A ray arriving at a wall at a steep angle carries almost no information about the wall's geometry.

For a specific structure of interest, add orbital passes with the camera pointed inward, at two or three altitudes. This is where mission planning software with orbital support, something like Dronelink, earns its keep, because flying a consistent orbit manually is difficult.

Fly multiple altitudes generally: a higher pass for coverage and context, a lower pass for detail on the subject. This gives the reconstruction both a stable global geometry and the fine detail where you want it.

Overlap of eighty by eighty or higher. Note the parenthetical, which is a genuinely useful diagnostic: mesh quality degrades visibly before geometric accuracy does. You can have a survey whose checkpoint residuals look perfectly acceptable while the mesh looks obviously wrong, with melted edges and blobby vegetation. The visual quality of a mesh is a more demanding standard than positional accuracy.

Expect occlusion, always, especially with vegetation present. There will be surfaces no camera saw, and the mesh generator will interpolate across them, sometimes plausibly and sometimes not.

And plan for manual cleanup. This is the honest note on the slide, and I want to end on it. Vegetation produces noisy geometry. Water produces nonsense, since a specular, moving surface violates every assumption the matcher makes. Thin structures like railings, antennas, and power lines are frequently lost entirely. Reflective surfaces, particularly glass, generate spurious geometry behind them. A publication-quality or client-quality mesh essentially always involves time in a mesh editing tool afterward. Budget for that when you plan a project.

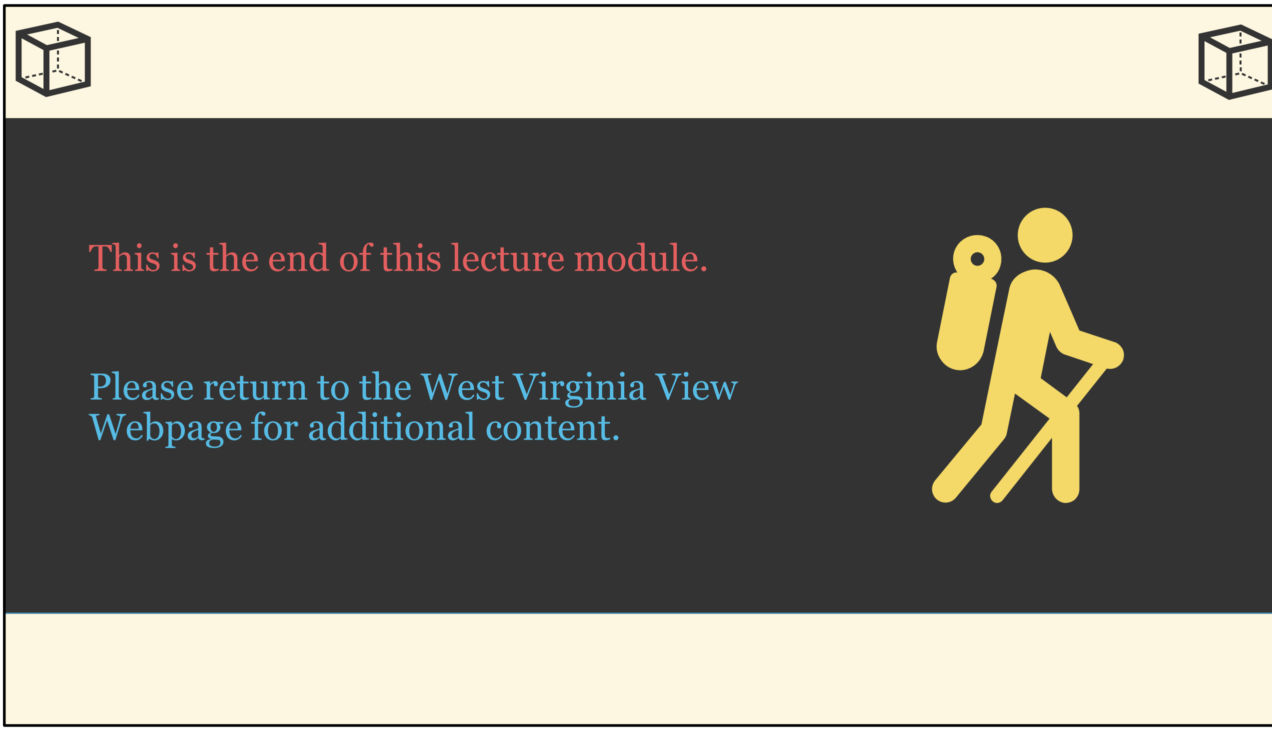

Thanks! Hope you found this useful.