REMOTE SENSING

Ander Knudby
University of Ottawa

LibreTexts™

University of Ottawa

Remote Sensing

Ander Knudby

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (https://LibreTexts.org) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.

LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by NICE CXone Expert and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptions contact info@LibreTexts.org or visit our main website at https://LibreTexts.org.

This text was compiled on 05/19/2026

TABLE OF CONTENTS

Licensing

Welcome to Remote Sensing

Chapters

- Licensing
- 1.1: Introduction
- 1.2: Radiometric measurements
- 1.3: Calculations of TOA radiance and TOA reflectance
- 1.4: EMR interactions with the Earth's atmosphere and surface
- 1.5: Visualization and manipulation of satellite images
- 1.6: Classification
- 1.7: Accuracy assessment
- 1.8: Change detection
- 1.9: Vegetation and Fire
- 1.10: Photogrammetry and Structure-from-Motion
- Index
- Glossary
- Detailed Licensing

Index

Glossary

Detailed Licensing

Licensing

*A detailed breakdown of this resource's licensing can be found in **Back Matter/Detailed Licensing**.*

CHAPTER OVERVIEW

Front Matter

TitlePage
InfoPage
Table of Contents
Licensing

University of Ottawa

Remote Sensing

Ander Knudby

This text is disseminated via the Open Education Resource (OER) LibreTexts Project (https://LibreTexts.org) and like the thousands of other texts available within this powerful platform, it is freely available for reading, printing, and "consuming."

The LibreTexts mission is to bring together students, faculty, and scholars in a collaborative effort to provide an accessible, and comprehensive platform that empowers our community to develop, curate, adapt, and adopt openly licensed resources and technologies; through these efforts we can reduce the financial burden born from traditional educational resource costs, ensuring education is more accessible for students and communities worldwide.

Most, but not all, pages in the library have licenses that may allow individuals to make changes, save, and print this book. Carefully consult the applicable license(s) before pursuing such effects. Instructors can adopt existing LibreTexts texts or Remix them to quickly build course-specific resources to meet the needs of their students. Unlike traditional textbooks, LibreTexts' web based origins allow powerful integration of advanced features and new technologies to support learning.

LibreTexts is the adaptable, user-friendly non-profit open education resource platform that educators trust for creating, customizing, and sharing accessible, interactive textbooks, adaptive homework, and ancillary materials. We collaborate with individuals and organizations to champion open education initiatives, support institutional publishing programs, drive curriculum development projects, and more.

The LibreTexts libraries are Powered by NICE CXone Expert and was supported by the Department of Education Open Textbook Pilot Project, the California Education Learning Lab, the UC Davis Office of the Provost, the UC Davis Library, the California State University Affordable Learning Solutions Program, and Merlot. This material is based upon work supported by the National Science Foundation under Grant No. 1246120, 1525057, and 1413739.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation nor the US Department of Education.

Have questions or comments? For information about adoptions or adaptions contact info@LibreTexts.org or visit our main website at https://LibreTexts.org.

This text was compiled on 05/19/2026

TABLE OF CONTENTS

Licensing

Welcome to Remote Sensing

Chapters

- Licensing
- 1.1: Introduction
- 1.2: Radiometric measurements
- 1.3: Calculations of TOA radiance and TOA reflectance
- 1.4: EMR interactions with the Earth's atmosphere and surface
- 1.5: Visualization and manipulation of satellite images
- 1.6: Classification
- 1.7: Accuracy assessment
- 1.8: Change detection
- 1.9: Vegetation and Fire
- 1.10: Photogrammetry and Structure-from-Motion
- Index
- Glossary
- Detailed Licensing

Index

Glossary

Detailed Licensing

Licensing

*A detailed breakdown of this resource's licensing can be found in **Back Matter/Detailed Licensing**.*

1.1: Introduction

Hopefully you have taken this course because you already have an idea of what remote sensing is, and you want to learn more about it. But what exactly is remote sensing?!?

What is remote sensing?

Some very broad definitions go something like ‘remote sensing is studying something without touching it’. While true in the most basic sense of the words, such definitions aren’t really very useful. Without going into philosophical discussions about what ‘studying’ and ‘touching’ mean anyways, I think we can agree that you are not *doing* remote sensing right now as you are reading this document on a computer screen… despite learning what this document says without touching it. More appropriately for this course, we can define remote sensing as using an instrument (the sensor) to collect information about the Earth’s surface (or other parts of the Earth, its oceans or atmosphere, or other planets for that matter) at large extents and at some distance. Typical examples include use of satellite imagery or aerial photography, but acoustic sensing (e.g. of the seafloor) and other technologies that do not create data in the form of an ‘image’ are also part of the broad subject of remote sensing. This kind of definition means that another term, ‘Earth Observation’, is often used interchangeably with ‘remote sensing’.

Review of remote sensing technologies

Passive sensing

There are many ways to split existing remote sensing technologies up into categories. Some technologies rely on ambient energy (e.g. energy that is naturally present in the environment, such as sunlight). These are called ***passive sensing*** technologies.

Active sensing

Others provide their own source of energy, emitting an energy pulse toward a target area and recording the part of it that is reflected back to the sensor (think about the sonar used on submarines). These are called ***active sensing*** technologies.

Optical sensing

Another way to categorize remote sensing technologies is by the kind of energy they use. Some technologies rely on visible light (e.g. traditional aerial photography), while other technologies expand the range of wavelengths sensed to the ultraviolet and infrared (which is also the kind of ‘light’ your TV’s remote control uses). When measuring reflected sunlight in these wavelengths, these are called ***passive opticalsensing*** technologies – optical because the instruments employ classic optics in their design. The satellite images seen on Google Earth, or used as a background on televised weather forecasts, are images created with passive optical remote sensing. A form of ***activeoptical sensing*** is called lidar, light detection and ranging. It involves emitting a laser pulse (the active part) toward a target and measuring the time it takes before that pulse has hit something and part of it returns to be measured at the sensor. The intensity of the return pulse is typically also measured.

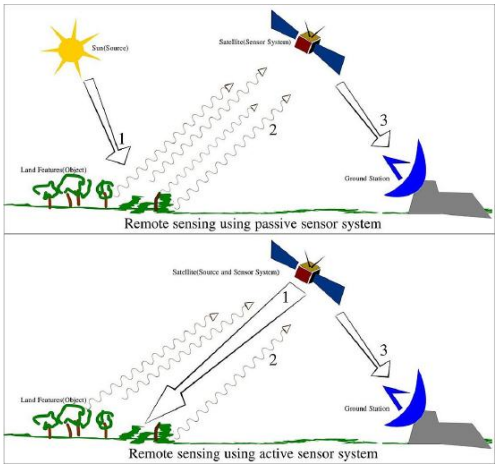

1: The difference between passive and active remote sensing systems. Source: Remote Sensing Illustration by Arkarjun, Wikimedia Commons, CC BY-SA 3.0.

2: NASA Satellite Captures Hurricane Earl on September 1, 2010 by NASA Goddard Photo and Video, Flickr, CC BY 2.0.

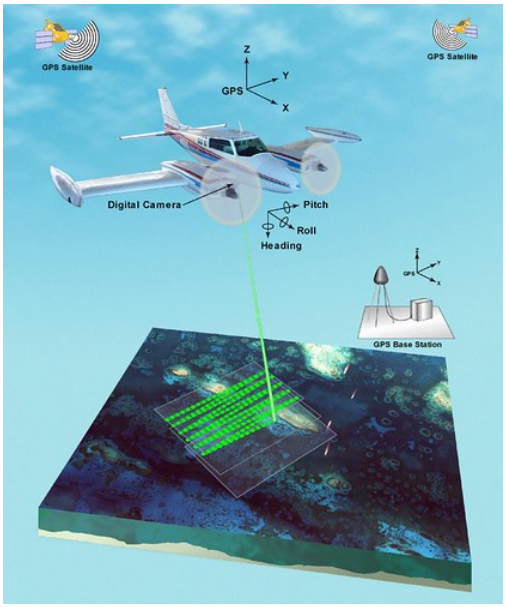

3: Concept diagram for airborne lidar. The location and orientation of the instrument is provided by the GPS and IMU instruments on the aircraft. With this knowledge, distances measured between the instrument and objects on the ground can be converted to find the locations of those objects, together providing a 3D image of the Earth's surface. fis01335 by NOAA Photo Library, Flickr, CC BY 2.0.

Thermal sensing

Thermal sensing is similar to passive optical sensing except the energy measured by the sensor is in the 'thermal region' of the wavelength spectrum. This is the kind of passive thermal sensing you do yourself when putting your hands close to the campfire to feel its heat, so we can feel thermal radiation but we can't see it with our eyes like we can see sunlight. The Sun produces much more thermal radiation than the Earth does (it is much hotter, after all), but the Sun is also very far away, so if you build a thermal sensor and put it in an airplane or on a satellite and point it towards Earth, the thermal radiation emitted by the Earth overwhelms

the small amounts of thermal radiation that has reached Earth from the Sun and been reflected by the Earth's surface into the sensor. As a result, thermal remote sensing is useful for measuring the temperature of the Earth's surface.

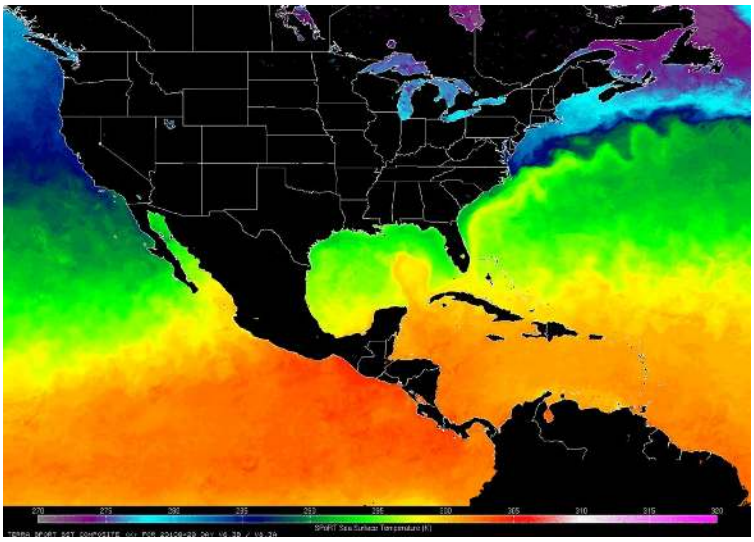

4: This map shows the distribution of Sea Surface Temperature for one day in 2010. Sea Surface Temperature is routinely measured by several satellites, including MODIS Terra that produced the data shown here. Sea Surface Temperature Map (NASA, SPoRT, 04/23/10) by NASA's Marshall Space Flight Center, Flickr, CC BY-NC 2.0.

Microwave sensing

Moving into the kinds of radiation we can neither see nor feel ourselves, ***microwave sensing*** is an important part of remote sensing. The Earth constantly emits small amounts of microwave radiation, and ***passive microwave sensing*** can be useful in a way somewhat analogous to thermal remote sensing, to measure characteristics about the Earth's surface (e.g. where there is snow and ice). Active microwave remote sensing, also called ***radar remote sensing***, uses emitted radar pulses just like radars on ships and at airports. Radars are mounted on aircraft or satellites and pointed toward Earth, radar pulses are emitted, and the reflected part of each pulse measured to provide information about aspects of the Earth's surface, most of which cannot be measured with optical or thermal technologies. These include information on the Earth's terrain (for making digital elevation models), soil moisture, vegetation density and structure, and the presence of sea ice and ships. Radar remote sensing has the advantage that microwave radiation moves through the atmosphere more or less unimpeded, and because it is an active technology it can thus be used anywhere at any time, including during the night and in bad weather.

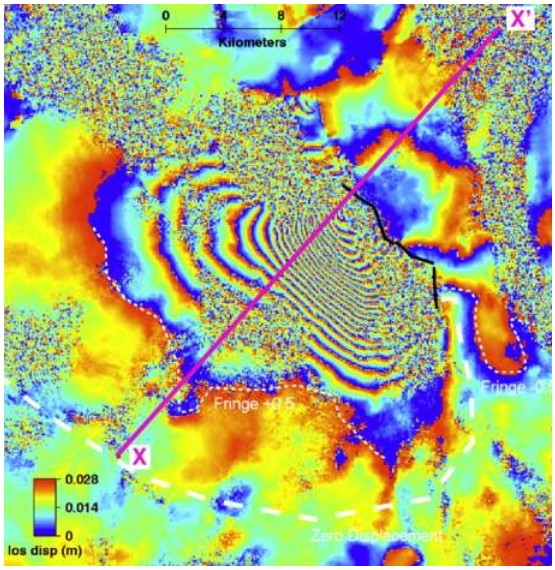

5: An amazing application of radar remote sensing is a technique called interferometric synthetic aperture radar, or InSAR. When a radar instrument passes the same area during two different orbits, changes in the shape of the Earth's surface can be mapped with millimeter precision. For example, when earthquakes happen the change in the Earth's surface can be mapped, as shown here for the 2010 earthquake in Dinar, Turkey. Smaller changes that result from e.g. buildings sinking into the substrate or bridges expanding and contracting in hot and cold weather, can also be mapped with InSAR. Interferogram of the Dinar, Turkey 1995 earthquake by Gareth Funning, GEodesy Tools for Societal Issues (GETSI), CC BY-NC-SA 3.0.

Gravimetric sensing

Another kind of instrument that has been put on satellites with great success measured the Earth's gravity field and how it changes spatially. The Earth's gravitational pull as sensed by an orbiting satellite is stronger when the satellite passes over an area with relatively large mass (e.g. the Himalayas), as opposed to an area with less mass (e.g. deep oceans). These differences can be measured very accurately and used to not only map the general shape of the Earth's gravity field, but also to measure changes in it through time. Such changes are caused by tides (more water equals more mass equals more gravity), melting glaciers, and even lowered groundwater tables from unsustainable use of groundwater.

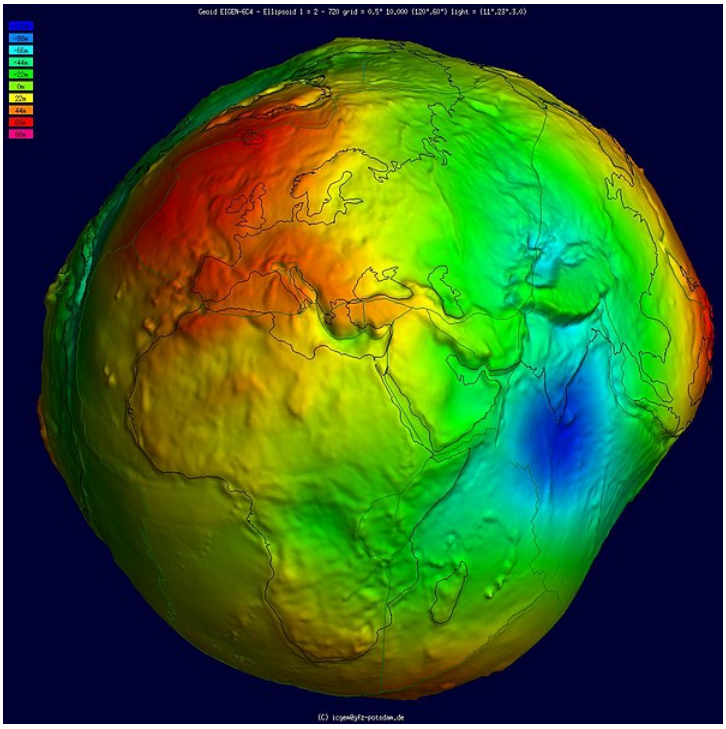

6: Earth shaped and coloured according to gravitational pull as measured by the European Space Agency's GOCE satellite. Geoid undulation 10k scale by the International Centre for Global Earth Models, Wikimedia Commons, CC BY 4.0.

Acoustic sensing

Acoustic sensing is fundamentally different from the previously listed technologies because it uses acoustic waves (i.e. waves that form as a result of compression and expansion of water masses) while all the others use electromagnetic radiation (i.e. waves that have a magnetic and electric component and can propagate through empty space). While there are many uses of acoustic sensing, primarily mapping water depth, seafloor and even sub-seafloor structure, the instruments and its physical basis are sufficiently different from the previously listed technologies that we will not look at them further in this course.

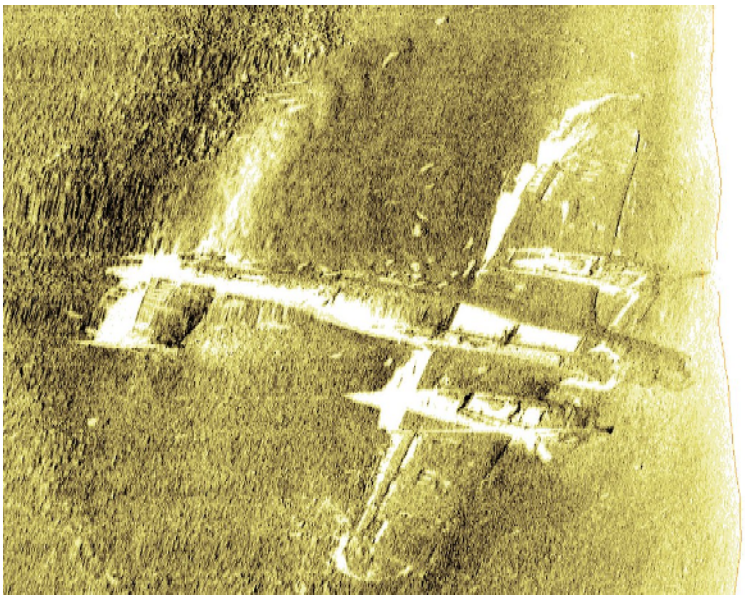

7: Acoustic sensors can both provide information on water depth and substrate composition over large areas, and provide detailed information on seafloor structure in smaller targeted areas. Seen here is a 3D reconstruction, based on acoustic data, of a World War II aircraft located off the UK coast. Aircraft on the seabed by Wessex Archaeology, Flickr, CC BY-NC-SA 2.0.

Land-based sensors

All remote sensing technologies are initially developed and tested in labs on land, where the fundamental science is done, prototypes are produced and tested, initial data are analyzed, and the technology is matured. While some are eventually reconstructed for use in an aircraft or on a satellite, these technologies typically retain their use on land. Well-known examples include the cameras that are in smartphones, the thermal cameras electrical inspectors use to look for faulty circuits, the radars police use to check how fast you are driving, and the lidars some self-driving cars use for situational awareness. While land-based sensors are useful for collection of very precise information from a small area or object, they are not useful for obtaining the kinds of data needed for mapping. Nevertheless, the data produced by land-based instrument are often useful in the remote sensing process (e.g. by comparing what you see in a satellite image to what you measured on the ground).

Airborne sensors

Eventually, someone will take a land-based technology and wonder ‘what would happen if I took that instrument, put it in an aircraft, pointed it down, and turned it on?’. The advent of aerial photography followed quickly after the development of the camera when someone decided to bring a camera up into a hot-air balloon, and airborne lidar and radar systems also followed the development of their respective technologies without much delay. There are at least two challenges that must be overcome when taking a remote sensing instrument airborne. 1) The instrument must be able to collect lots of data quickly, and store this data for later use (or transmit it directly to a storage device elsewhere), and 2) each collected data point must typically be georeferenced – in other words a geographic coordinate pair (e.g. latitude and longitude) must be associated with each data point. These are by no means insurmountable challenges today, when a typical smartphone can act as both data storage and GPS, and putting sensors on airborne platforms has recently become an area in explosive growth because cheap and easy-to-fly drones can now replace manned aircraft as the platform that carries the instrument.

Space-borne sensors

After an instrument has been proven useful when mounted on an aircraft, eventually someone will suggest that ‘we should totally take this and put it on a satellite, so we can collect data continuously for years without having to land, refuel, submit flight plans and so on’. In fact, (to my knowledge) all Earth-observing satellites carry instruments whose prototypes were tested on aircraft, because not only does this help fine-tune the instrument hardware before launching it into space, it also gives remote sensing scientists a preview of what the data from the instrument will be like, so they can start writing programs that process the instrument data and convert it into useful information about the Earth. There are important advantages and disadvantages to the use of

satellites compared to drones and manned aircraft. The great advantage is that satellites typically last for many years and can collect data continuously over that period. They are thus incredibly cost-effective. Just imagine what it would cost to provide daily updated information on sea ice cover in the Arctic (and Antarctic for that matter) without satellites! However, satellites have limited fuel, and space is a rough environment so the instruments on board degrade with time and may ultimately fail entirely. For example, in 2012 the multi-billion-dollar European satellite ENVISAT stopped communicating with its command centre, and despite the best efforts by the European Space Agency to resume communication the satellite with all its instruments was declared 'dead' two weeks later. Landsat 6, a general-purpose American land observation satellite failed to reach its orbit after launch, and never produced any data at all. One important problem with satellites is that when these things happen, they can't land and be fixed, as could be done if a problem developed with an instrument sitting in an aircraft.

While the large majority of space-based sensors are located on satellites, other spacecraft have also carried a few important remote sensing instruments. A radar system was placed on a Space Shuttle in 2000 and used to produce a near-global digital elevation model at 30-meter spatial resolution during its 11 days of operation. Recently, a Canadian company called UrtheCast operated two passive optical sensors from the International Space Station – one was an ultra-high-definition video camera that produced video at three frames per second with a spatial resolution of 1.1 meters, the other is a more typical imaging instrument that produced colour imagery at 5-meter spatial resolution.

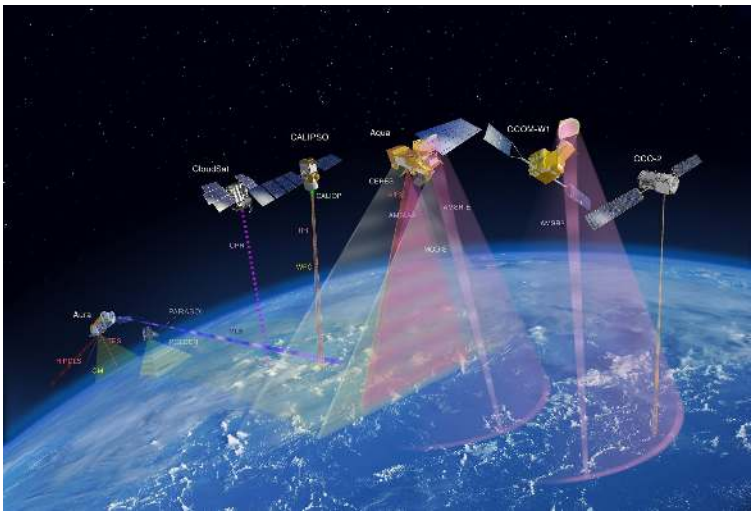

8: There is today a remarkable number of Earth-observing satellites in orbit. This image shows NASA's 'A-train' satellite constellation, a series of satellites whose orbits closely follow each other, all passing overhead during the early afternoon, local solar time. Constellations provide the opportunity to use data from one satellite to aid interpretation of data from another satellite, thus improving the quality of many satellite products. Atrain-879×485 by Shakibul Hasan Win, Wikimedia Commons, CC BY-SA 4.0.

Electromagnetic radiation and its properties

Now it is time to look at bit more closely at how remote sensing works, and what is its physical basis. That will provide you with a better understanding of how remote sensing data are created, and how they may be used in creative ways to extract exactly the kind of information needed for a specific purpose.

The energy that is measured by a remote sensing instrument (other than acoustic instruments), and that is used to produce an image, is called electromagnetic radiation (often abbreviated to EMR). It can be difficult to understand exactly what electromagnetic radiation is, but a useful way to think about it is to consider that light is a specific kind of electromagnetic radiation, one that our eyes and brain are good at detecting. Other kinds of electromagnetic radiation include the harmful ultraviolet (UV) rays that sunscreen provides some protection from, the thermal radiation that we feel when moving close to a campfire, and the radar waves used to detect aircraft and ships (and to map some surface properties on Earth). Physically, electromagnetic radiation can be visualized as waves that propagate through space. The waves have two components, one electrical and one magnetic, which are at 90-degree angles both to each other and to the direction of propagation (Figure 9).

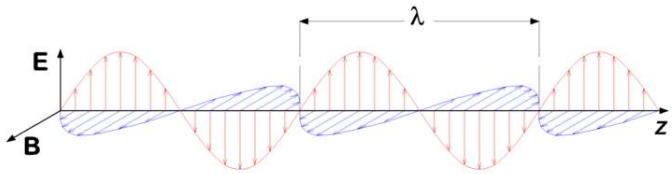

9: Electromagnetic radiation consists of a transverse wave with an electric field (red, vertical plane) and a magnetic field (blue, horizontal plane), propagating (in this figure) in the horizontal direction. Electromagnetic wave by P.wormer, Wikimedia Commons, CC BY-SA 3.0.

Wavelength

EMR waves have certain properties that we can measure and use to describe it. EMR waves can be characterized by their *wavelength*, which is measured as the physical distance between one wave peak and the next, along the direction of propagation. Human eyes are able to detect electromagnetic radiation with wavelengths roughly in the range between 400 and 700 nanometers (a nanometer is 10-9 meter, or a billionth of a meter), which is therefore what we call visible light. EMR waves propagate at the speed of light, and there is a simple and direct relationship between the wavelength and the frequency of EMR waves, typically expressed as $c = \nu\lambda$, where c is the speed of light, v (the Greek letter nu) is the frequency, and λ (the Greek letter lambda) is the wavelength. The frequency is defined as the number of wave peaks passing a fixed point per unit of time (typically per second). Note that some people, including engineers and physicists working with radiation, sometimes use the wavenumber instead of wavelength or frequency. The wavenumber ($\tilde{\nu}$) is defined as $\tilde{\nu}=1/\lambda$. We will not use the wavenumber in the rest of these notes, but it is good to be aware of its use for later studies.

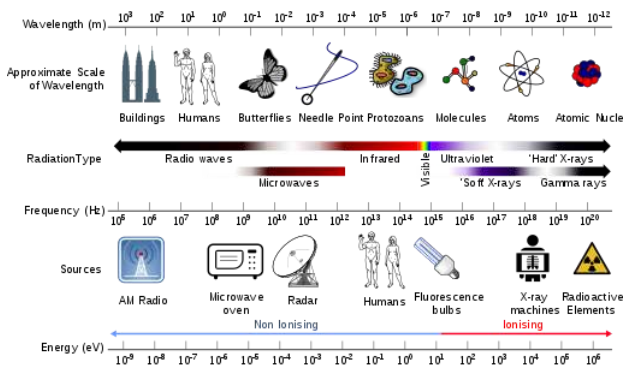

10: An overview of different kinds of EMR, their wavelengths, frequencies, and some common sources. Electromagnetic spectrum with sources by Dinksbumf, Inductiveload and NASA, Wikimedia Commons, CC BY-SA 3.0.

According to quantum theory, EMR can also be considered as consisting of a stream of individual energy packages called photons. Each photon contains an amount of energy that is proportional to its frequency, a relationship expressed as $E=h\nu$, where h is Planck's constant (roughly 6.6 * 10-34 Js). The relationship between wavelength, frequency and energy per photon is shown in Figure 10. Also found in Figure 10 are the common names for EMR of a certain wavelength. You can see, for example, that X-rays have very short wavelengths, very high frequencies, and therefore very high energy content per photon, which is an important reason to limit your exposure to them.

Polarization

In addition to wavelength, EMR can be characterized by its *polarization*, which can be thought of as the orientation of oscillations of the electric field. For example, in Figure 9, the electric field oscillates exclusively in the vertical plane, so the EMR shown in the figure is said to be vertically polarized. Many sources of EMR, like the Sun, produce electromagnetic waves whose polarization changes so quickly in time that, when measured, the combination of many waves appears to have no distinct polarization. In other words, the orientation of the oscillations of the electric field (and by extension of the magnetic field) is vertical, horizontal, and anything in between. Such EMR is called unpolarized. Other sources of light, such as that created by most instruments, is polarized, meaning that the orientation of the oscillations of the electric field does not change with time. A familiar example of

polarization is sunlight reflecting off a water surface. When the Sun's unpolarized light hits the water surface, the waves whose electric field is oriented vertically are absorbed or refracted into the water more easily than those whose electric field is oriented horizontally, which are more likely to be reflected off the water surface. When you look at a water surface, the horizontally polarized parts of the light you see have thus predominantly been reflected off the water surface, while the vertically polarized parts of the light you see have predominantly been reflected from inside the water itself, or from the bottom if you are looking at shallow water. Therefore, if you use polarized sunglasses that have been made to effectively remove horizontally polarized light, the light that reaches your eyes is predominantly that part of the light field that comes from within the water itself. An example of this is shown in Figure 11, where you can see that the use of a polarization filter enables more of the detail from the seafloor to be visible in the image.

11: An example of looking at shallow water, without (left) or with (right) a polarization filter. Reflection Polarizer2, by Amithshs, Wikimedia Commons, public domain.

How electromagnetic radiation is created

Now that we know how to quantify different properties of electromagnetic radiation, we will take a brief look at how it is created in the first place. Physicists say that EMR is produced when charged particles are accelerated, this happens all the time, so it is not helpful in the context of remote sensing. Of greater use in remote sensing, we can look at the *spectral radiance* created by a surface, i.e. the amount of radiation created and its distribution across different wavelengths. This depends primarily on the temperature of the surface, and is neatly described for blackbodies by Planck's Law:

$$L_{\lambda,T} = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\left(\frac{hc}{\lambda k_B T}\right)} - 1}$$

In Planck's Law, $L_{\lambda,T}$ is radiance, with the subscripts indicating that it is dependent on wavelength and temperature (of the surface). h is Planck's constant, c is the speed of light in a vacuum, λ is wavelength, k_B is the Boltzmann constant (not to be confused with the Stefan-Boltzmann constant!), T is the surface temperature in degrees Kelvin. While Planck's Law can initially look intimidating, keep in mind that h, c and k_B are all constants, so all it says is that the radiance emitted by a surface ($L_{\lambda,T}$) can be calculated from the temperature (T) of that surface for a given wavelength (λ).

To take a look at the result of Planck's Law, you can imagine an object with a certain surface temperature, like the Sun whose temperature is roughly 5800K, and calculate the radiance it emits at, say, 400nm, by simply plugging the temperature and wavelength into the equation. You can then repeat for 401nm, 402nm, and so on, to produce what is called a Planck curve. An example of several Planck curves is provided in Figure 12. Note that very little radiation is produced at very short wavelengths (e.g. 0-200nm), and each curve has a distinct peak that shows the wavelength at which the maximum amount of energy is emitted. The location of this peak can be calculated using Wien's Law, which is based on Planck's Law, and which says that the wavelength of peak emission (λ_{max}) can be calculated as λ_{max}=bT, where b is Wien's displacement constant: 2,897,767 nm K. For the Sun, this comes to approximately 500 nm, which is blueish-greenish light.

One important note on Planck's Law is that it applies to blackbodies, which are imaginary entities that are defined as perfect absorbers and perfect emitters. While real-world objects, like the Sun and the Earth, are not blackbodies, they come close enough to being blackbodies most of the time for Planck's Law to be a useful approximation that can be used to quantify their emission of electromagnetic radiation. How close or far a surface is to being a blackbody is defined by its emissivity, which is the amount of radiance produced at a given wavelength compared to that which would be produced by a blackbody. Natural materials tend to have high emissivity values, in the range of 0.98-0.99 (water), 0.97 (ice), and 0.95-0.98 (vegetation). However, some natural

materials can have low emissivity, like dry snow (appr. 0.8) or dry sand (0.7-0.8). What this means, for example, is that dry snow at a given temperature only produces 80% of the radiance that it should according to Planck's Law. A further complication is due to the fact that emissivity changes with wavelength, but that's a bridge to cross another time.

Planck's Law can be used to determine the temperature of a surface whose spectral radiance is measured. This is done simply by inverting the equation to isolate temperature:

$$\mathrm{T} = \frac{\mathrm{hc}}{\ln\left(\frac{2hc^2}{\lambda^5 L} + 1\right)\lambda k_B}$$

While actual determination of the Earth's surface temperature from space-based measurement of (spectral) radiance is slightly more complicated because both emissivity and the influence of the atmosphere need to be taken into account, the inversion of Planck's Law as outlined above is the basic principle upon which such mapping rests.

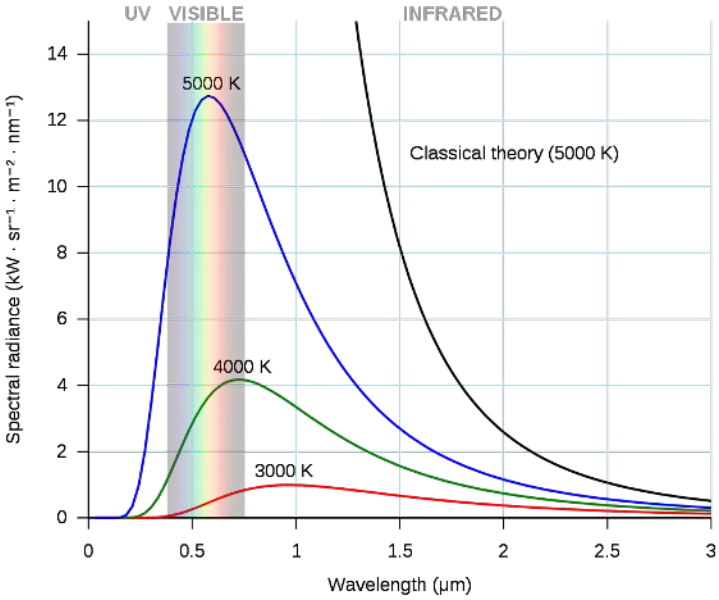

12: Example of three Planck curves, for surfaces at 3000-5000 degrees Kelvin. Black body by Darth Kule, Wikimedia Commons, public domain.

This page titled 1.1: Introduction is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.2: Radiometric measurements

2

While much remote sensing data can be visualized as imagery, and interpreted directly as such to locate e.g. cities and ships and fires, it is important to understand that most satellite sensors can also be thought of as precisely calibrated instruments used to measure the characteristics of electromagnetic radiation that arrives from a given direction at the time of measurement. In this chapter we will focus on such use of satellite imagery. First we need to define some concepts related to electromagnetic radiation and how it is measured. After that we will look at how to employ those concepts to extract useful information from satellite imagery.

Concepts related to EMR intensity

Several other properties of EMR fields relate to their intensity – the amount of energy contained in them. The fundamental SI unit of energy is the Joule (J). Because radiation moves as the EMR waves propagate, an important measure of the intensity of an EMR field is the amount of energy moving onto, through or from a surface or volume ***per unit time***. This is called *radiant flux*, and is measured in Joules per second, or Watts (W). In remote sensing, we are almost always interested in measuring the intensity of radiation over a finite area (e.g. one represented by a pixel in a satellite image). This is called the *radiant flux density*, and is measured as radiant flux per unit area, for example W m^{-2}. For practical use, the term *irradiance* is used to describe radiant flux density incident upon a surface, and *exitance* is used to describe radiant flux density leaving a surface (Figure 13).

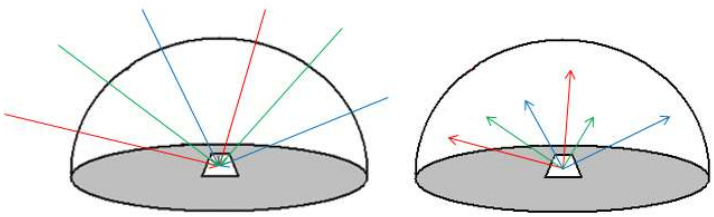

13: Irradiance and Exitance. On the left, irradiance is the radiant flux incident on a surface per unit area. On the right, exitance is the radiant flux leaving a surface per unit area. Both are typically measured in W m-2. The arrows of different colour indicate that these two radiant flux densities are measured for all wavelengths present in the radiation field. By Anders Knudby, CC BY 4.0.

While irradiance and exitance are useful concepts, we cannot measure them with instruments on flying platforms because they encompass radiation incident upon or leaving a surface in any direction. In other words, to measure exitance for a 1 m^2 part of the Earth's surface, we would need to place a hemispherical sensor over the area in question, thus measuring the outgoing radiation in all upward directions (if measuring sunlight reflected from the surface, shading by the instrument would obviously make the measurement useless). Another unit of EMR intensity commonly used in remote sensing is therefore something more closely related to what is actually measured by instruments, it is called *radiance* and is defined as the radiant flux density per unit of projected source area, in a particular direction defined by a solid angle (Figure 14). A solid angle can be thought of as a cone – in most remote sensing cases this cone is incredibly narrow as it stretches from the observed area to the relevant detecting element in the sensor.

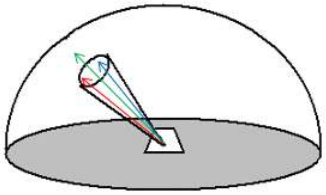

14: Radiance. Compared to irradiance and exitance, radiance is calculated per unit of projected surface area, and within a specified solid angle. Radiance is typically measured with units of W m-2 sr-1, where sr stands for steradian, which is a measure of solid angle. By Anders Knudby, CC BY 4.0.

When measuring EMR with instruments on aircraft or satellites, measurements of radiance are usually made for discrete intervals of wavelengths. For example, 'Band 1' on many satellites has been designed to measure only the radiance of electromagnetic radiation with certain wavelengths that appear blue to human eyes, such as between 420 and 470 nm. Assuming that this range of wavelengths are all measured perfectly by 'Band 1', we can find the *spectralradiance* by dividing the radiance measured by the range of wavelengths measured (in this case 50 nm). The spectral radiance is a very commonly used measure in remote sensing, and is often directly related to the raw values (i.e. Digital Numbers) found in each band in remote sensing images.

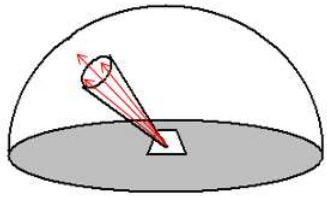

Figure 14, the arrows indicating electromagnetic radiation have here been coloured with a single colour, to emphasize that spectral radiance is measured for a specific range of wavelengths. By Anders Knudby, CC BY 4.0.

While spectral radiance is the unit of EMR intensity that is most closely related to what is actually measured by remote sensing instruments, it has the unfortunate quality of being dependent on the illumination of the observed area. This is obvious if you imagine measuring the amount of light coming off your local neighbourhood parking lot. During a sunny day, lots of light comes off the asphalt surface, sometimes to the extent that you need to squint to look at it. In other words, the surface has a high spectral radiance in the visible wavelengths. On an overcast day, less light comes off the surface – lower spectral radiance. And during the night, obviously, very little light comes off the surface – very low spectral radiance. Measuring the spectral radiance coming off a surface thus cannot directly tell us much about what that surface is – whether it is a parking lot or a lake or a forest. The thing we would prefer to measure instead is the reflectance of a surface – the amount of radiation it reflects per unit of radiation incident upon it. This is a real physical property of a material that is in most cases completely independent of illumination, and therefore something that can be used to identify the material. A high reflectance is what makes white surfaces white, and a low reflectance is what makes black surfaces black. A commonly used kind of reflectance is the albedo, which is what remote sensing people call diffuse reflectance. It is calculated as the exitance divided by the irradiance, and is a unitless measure ranging between 0 (for completely black surfaces that have no exitance) and 1 (for completely white surface that reflect all the radiation incident upon them).

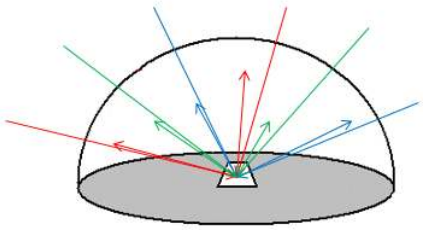

16: Diffuse reflectance. Defined as exitance (outgoing arrows) divided by irradiance (incoming arrows). By Anders Knudby, CC BY 4.0.

Diffuse reflectance cannot be measured directly in remote sensing – in fact neither of the two necessary terms can be measured. However, it can be estimated quite accurately at the top of the atmosphere (more on that later) by employing two tricks. First, we know quite well how much radiation is produced by the Sun at various wavelengths, so we can estimate the irradiance with good accuracy. Second, if we assume that the radiance is a known function of the direction of propagation, then we can convert radiance measured in one direction to exitance. Together, this gives us the information necessary to convert a measurement of radiance to an estimate of diffuse reflectance.

However, we need one last step (for now) to arrive at a measurement that can be used to identify what we are looking at, because diffuse reflectance is a single value that tells us something about how bright a surface is, but it does not tell us anything about its chromacity (which is the aspect of colour that does not include brightness, so for example light red and dark red can have the same chromacity, while light red and light green have different chromacities). However, as outlined above in the paragraph on spectral radiance, measurements made with remote sensing instruments are actually measurements of spectral radiance. We can thus convert

these measurements to estimates of spectral exitance, and divide by spectral irradiance, to arrive at a measure of spectral diffuse reflectance. Without accounting for interactions between electromagnetic radiation and the atmosphere, this is as close as we will get to producing a measure that tells us something about what is covering the Earth's surface in the pixel we measured.

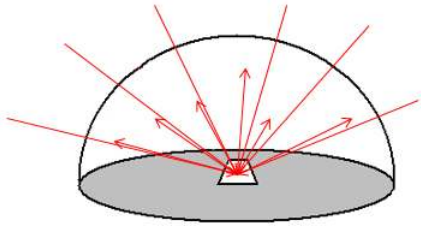

17: Spectral diffuse reflectance. Defined as spectral exitance divided by spectral irradiance. By Anders Knudby, CC BY 4.0.

How electromagnetic radiation is detected and measured

To measure the intensity of incoming electromagnetic radiation, remote sensing instruments rely on special materials that are efficient absorbers of radiation with the desired range of wavelengths, and that generate an electrical current when exposed to radiation – a current which is then measured. While different materials are employed to absorb and thus detect radiation at different wavelengths in different kinds of sensors, the basic principle is well illustrated with passive optical remote sensing instruments, most of which rely on CCD or CMOS sensors, just like the camera in your phone.

Imagine a camera attached to a satellite in space, pointing down on Earth. For a very short period of time, the aperture of the camera is opened, letting in light that has been produced by the Sun (according to Planck's Law) and reflected by the Earth's atmosphere and surface back into space, exactly in the direction of the satellite. Inside the camera is a 2-dimensional CCD, a kind of checkerboard with 3 x 3 individual little detecting elements, each of which is capable of absorbing sunlight and thereby producing an electrical current proportional to the intensity of incoming light. Detecting elements exposed to more intense sunlight produce a current with greater voltage, so the voltage is a direct indication of the amount of sunlight each element was exposed to. An example is provided in Table 1.

Table 1: Example of voltages generated by individual detecting elements in a CCD housed in a spaceborne camera. By Anders Knudby, CC BY 4.0.

1.2	1.1	0.9
1.3	1.2	1.1
1.1	0.9	0.8

The optics (system of lenses) in the camera ensure that each individual detecting element is exposed to light coming from a unique pre-defined direction relative to the camera. Knowing where the satellite is and how it is oriented relative to the Earth, each of these directions can be converted to a set of geographic coordinates (latitude/longitude) on the Earth surface, from where the light registered by each detecting element must have been reflected. In other words, the electrical current generated by each of the 3 x 3 detecting elements can be mapped back to a location on the Earth's surface.

Table 2: Geographic coordinates from which the sunlight producing the electrical current in each element from Table 1 has been reflected. For this example, ignore the fact that the pixel size implicit in the geographic coordinates is not square and very large. By Anders Knudby, CC BY 4.0.

45.1 N / 80.5 W	45.1 N / 80.4 W	45.1 N / 80.3 W
45.0 N / 80.5 W	45.0 N / 80.4 W	45.0 N / 80.3 W
44.9 N / 80.5 W	44.9 N / 80.4 W	44.9 N / 80.3 W

Because the voltage of the electrical current generated by each detecting element is proportional to the radiant flux density of the incoming radiation, with proper calibration the voltage produced by each detecting element can be converted to a measure of incoming radiant flux density at the camera. Furthermore, with knowledge of the distance between the satellite and the Earth's surface, this can in turn be converted to a measure of the radiance coming from the observed area. And with that we have one of the fundamental measurements in remote sensing – the radiance coming from a well-defined area on the Earth's surface.

Calibration

In the previous section, it was mentioned that the voltage generated by a detecting element can be converted to radiant flux density or radiance *with proper calibration*. Such calibration is done for all sensors prior to launch by exposing them to light of varying levels of known radiance, noting down the voltage created by each exposure. From this, a simple conversion equation, typically linear or nearly linear, can be created to convert between voltage and radiance. An example of the kind of setup needed to perform such calibration is shown in Figure 18.

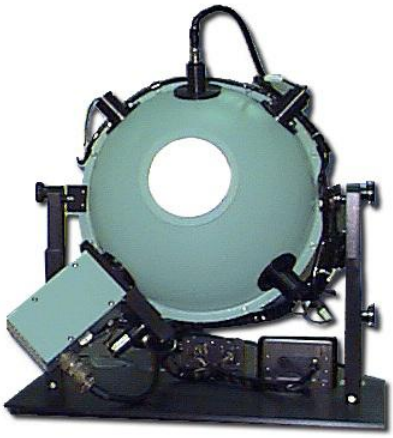

18: Integrating sphere used for sensor calibration. Commercial Integrating Sphere by Electro Optical Industries, Wikimedia Commons, CC BY-SA 3.0.

Multispectral sensing

We have so far ignored the fact that most satellite sensors are designed to detect radiance at multiple distinct ranges of wavelengths, each producing what is known as a 'Band'. The incoming light field reaching the sensor is in the case of reflected sunlight always a mix of radiation with many different wavelengths, so to separate these out and measure the radiance of distinct ranges of wavelengths individually required an additional step beyond what has been outlined above. Different kinds of technologies are used to perform this step, the simplest of which relies on a beam splitter. A beam splitter can be thought of as a kind of fancy prism that redirects electromagnetic radiation within a defined range of wavelengths in one direction, while letting radiation with other wavelengths move on unimpeded. Multiple individual sensors can then be put in the proper locations in the imaging sensor, with the CCD destined to produce data for the 'red' band being located where the beam splitter redirects radiation with wavelengths between 600 and 700 nm (typically considered 'red' light), and so on for each individual CCD. Other technologies than beam splitters are in use, each accomplishing the same basic task of allowing radiation with different wavelengths to be registered separately.

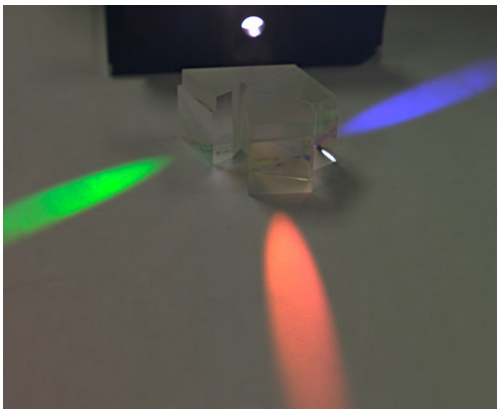

19: The idea of a beam splitter. Radiation enters from the source, and photons are redirected according to their wavelength. Three CCDs can then be positioned to measure the intensity of red, green, and blue light separately. Color Separation Prism by Dick Lyon, Wikimedia Commons, public domain.

Note: For the sake of simplicity, we have used a traditional 'camera' as the model to explain how passive optical sensors detect and measure incoming electromagnetic radiation. While traditional frame cameras are still the instrument of choice in the aerial photography industry, and in the growing field of drone-based remote sensing, passive optical sensors onboard satellites take one of two other forms: push-broom or whisk-broom scanners. The details of these detecting systems are beyond the scope of these notes, but basic information is provided here:

- A push-broom scanner is made of a series of one-dimensional CCDs, each typically consisting of several thousands of individual detectors in one row. Each CCD is positioned carefully in the instrument, 'after' a beam splitter, to record radiation in what becomes one band, more or less as shown in Figure 19. The instrument does not have a shutter, rather the CCDs record incoming radiation from a very wide and narrow line on the Earth, reaching from one side of the sub-orbital swath to the other. These measurements become one line (row) in a satellite image. As the satellite moves along its orbit, another line is recorded, then another, and another, until thousands of lines have been measured. These are then all put together to form an image.
- A whisk-broom scanner relies on a rotating scan mirror to direct radiation from different parts of the Earth's surface onto a smaller number of CCDs. The position of the mirror at any given time can be used to calculate the direction from which the measured radiation arrived at the sensor. Such measurements, with corresponding geolocation information, are taken in rapid succession, and ultimately put together as an image.

This page titled 1.2: Radiometric measurements is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.3: Calculations of TOA radiance and TOA reflectance

Sensor calibration typically results in a simple linear (or almost linear) conversion between the radiance that the sensor is exposed to and the voltage created by the detecting elements. This means that if we record the voltage created by each detecting element, along with the location on the Earth's surface from which the radiation that created the current came, and we arrange all the voltages measured by each of the detecting elements in a spatial pattern according to their locations on the Earth's surface, we can create a map-like image in which the values in each pixel are the voltages measured from that area. In essence that is how satellite images are created.

Quantization

However, the Digital Numbers that we see when we load a satellite image into a software do not equal the measured voltages. There are two related reasons for this – reducing the amount of computer memory needed to store the measurements, while still providing data with adequate precision. Measured voltages take the form of decimal numbers, like 3.560194 mV, which by default require 32 bits of computer memory to store (or 64 bits if a large number of decimals is required). But in reality the measurements of voltage in satellite sensors are not accurate to six decimals (as in the example above) or some other large number of decimals, so there is no need to use 32 bits of memory to store the voltage for each detecting element. What is typically done instead is that during, or even before, sensor development, the necessary precision of the measurements is determined, and the instrument is built to make measurements with that precision (which is then tested later when calibrating the instrument). With the precision of the measurements known, and adding information on the range of radiances/voltages that the instrument should be sensitive to, the number of bits necessary to contain the measurements can then be determined, and the voltage can be converted to a Digital Number that can be stored as an integer, which requires much less computer memory to store.

For example, imagine that you want to measure how tall people are, and you say you want to measure it to the nearest cm. Using units of cm, and knowing that any adult measured in the history of such measurements has been somewhere between 54cm and 272cm tall, you realize that you really only need to be able to measure 219 different values (both extremes included): 54cm, 55cm, 56cm, …. 272cm. To store 219 different values, you would only need 8 bits, because $2^8 = 256$, which even leaves you some extra values you don't need. You then decide to encode a height of 54cm with a value of '0', 55cm with a value of '1', and so on (code values are simply 54 less than the corresponding heights in cm). Now, when you measure someone (with your super-accurate measuring tape) to be 173.234cm tall, you tell yourself that you only need cm precision so you round it off to 173cm, and you encode it as 173 – 54 = '119'. This is the equivalent of the Digital Number that you see in a pixel in a given band in an image. When you know how the code ('119') relates back to the actual measurement (173cm), the value becomes more meaningful. This process of encoding is called *quantization*.

Digital Numbers to TOA Radiance

When you open a satellite image in a software package, the initial value you see in each pixel is called a Digital Number (DN), which is equivalent to the code ('119') in the example above. While the DNs are sufficient to produce a nice visualization of the image, if you want to treat the image as a series of radiometric measurements you need to convert it into a physical quantity (like the 173cm that was the height of the person measured in that example). In remote sensing, the translation of a Digital Number value into a radiance measurement is really a two-step process. In passive optical remote sensing, the calibration of the sensor has resulted in a known relationship between the radiance that the sensor is exposed to, and the voltage generated by the radiation, and quantization has then been used to encode the measure voltage as a Digital Number. To convert this Digital Number (DN) to radiance thus in theory requires first a conversion to voltage and then a conversion to radiance. Luckily it is easy to do these two conversions at the same time. We will look at how this is done for Landsat data as an example, but the same principle is relevant to data from all other passive optical satellite sensors.

Landsat metadata contain radiometric calibration coefficients (Figure 20) that allow you to directly convert DN values to Top-of-Atmosphere (TOA) radiance. The coefficients contain a 'multiplication' term and an 'addition' term, each of which is unique to each band, and they are used like this: $L_\lambda = mult_\lambda * DN + add_\lambda$, where L is the radiance and λ, the common notation for wavelength, indicates a given band.

```
GROUP = RADIOMETRIC_RESCALING
  RADIANCE_MULT_BAND_1 = 1.2168E-02
  RADIANCE_MULT_BAND_2 = 1.2460E-02
  RADIANCE_MULT_BAND_3 = 1.1482E-02
  RADIANCE_MULT_BAND_4 = 9.6820E-03
  RADIANCE_MULT_BAND_5 = 5.9249E-03
  RADIANCE_MULT_BAND_6 = 1.4735E-03
  RADIANCE_MULT_BAND_7 = 4.9664E-04
  RADIANCE_MULT_BAND_8 = 1.0957E-02
  RADIANCE_MULT_BAND_9 = 2.3156E-03
  RADIANCE_MULT_BAND_10 = 3.3420E-04
  RADIANCE_MULT_BAND_11 = 3.3420E-04
  RADIANCE_ADD_BAND_1 = -60.83886
  RADIANCE_ADD_BAND_2 = -62.29974
  RADIANCE_ADD_BAND_3 = -57.40868
  RADIANCE_ADD_BAND_4 = -48.41023
  RADIANCE_ADD_BAND_5 = -29.62465
  RADIANCE_ADD_BAND_6 = -7.36738
  RADIANCE_ADD_BAND_7 = -2.48320
  RADIANCE_ADD_BAND_8 = -54.78706
  RADIANCE_ADD_BAND_9 = -11.57798
  RADIANCE_ADD_BAND_10 = 0.10000
  RADIANCE_ADD_BAND_11 = 0.10000
```

20: Calibration coefficients used to convert DN values to TOA radiance values for Landsat 8. By Anders Knudby, CC BY 4.0.

In order to produce a raster data layer in which each cell value is a measure of TOA radiance, you can therefore use the 'Raster Calculator' available in all GIS / remote sensing software, and some software packages even have dedicated 'Landsat radiometric calibration' tools that find the relevant information in the metadata and apply it for you.

Digital Numbers to TOA Reflectance

Converting DN values to a measure of radiance is an important step toward using remote sensing data in a quantitative way – in which the values in each cell have an actual physical meaning. However, there is one principal problem with radiance measurements: The radiance coming from an area depends both on the physical characteristics of that area, such as what is on the surface, but it also depends on the amount of radiation incident on that area. It is easy to confirm this with a single experiment: Move to a light switch near you, make sure it is on, and focus your gaze on a wall. Note how bright it is (a rough human approximation of the amount of radiance coming from it). Now turn off the light, and note whether the wall got darker or lighter. It got darker, of course, which means that there is now less radiance coming from it. And this happens although the wall itself did not change at all! In the remote sensing context, radiance measurements are therefore rarely used directly, but are typically used as a step toward a radiometric quantity that is more closely linked to the surface it is coming from: reflectance. For isotropic radiation, the reflectance (R) of a given area is the ratio of outgoing vs. incoming radiance, also often expressed on the basis of exitance (E) and irradiance (I):

$$R = \frac{L_{up}}{L_{down}} = \frac{E_{up}}{I_{down}}$$

In the wall example, when you turn off the light both the incoming and the outgoing radiance is reduced, and the reflectance of the wall is unchanged. That's because the reflectance is a physical property of the wall that is independent of illumination. This makes it uniquely suitable for use in remote sensing, because multiple observations with varying illumination can be compared, for uses such as tracking change of an area through time, or for comparing different areas on Earth.

If we convert TOA radiance values to reflectance values, we get what's called TOA reflectance, which is the reflectance of the entire Earth-Atmosphere system. To think about what that means, imagine that you are an astronaut and you put a hula-hoop ring horizontally right at the level of the top of the atmosphere (if the top of the atmosphere were possible to define and find!), and you then measure the amount of radiation passing through the ring in upward and downward directions. Dividing one by the other gives you TOA reflectance.

Calculating TOA reflectance is done a bit differently for different data types. For Landsat 8, there is another set of calibration coefficients that you can apply exactly like you can apply those for TOA radiance (Figure 21).

```
REFLECTANCE_MULT_BAND_1 = 2.0000E-05
REFLECTANCE_MULT_BAND_2 = 2.0000E-05
REFLECTANCE_MULT_BAND_3 = 2.0000E-05
REFLECTANCE_MULT_BAND_4 = 2.0000E-05
REFLECTANCE_MULT_BAND_5 = 2.0000E-05
REFLECTANCE_MULT_BAND_6 = 2.0000E-05
REFLECTANCE_MULT_BAND_7 = 2.0000E-05
REFLECTANCE_MULT_BAND_8 = 2.0000E-05
REFLECTANCE_MULT_BAND_9 = 2.0000E-05
REFLECTANCE_ADD_BAND_1 = -0.100000
REFLECTANCE_ADD_BAND_2 = -0.100000
REFLECTANCE_ADD_BAND_3 = -0.100000
REFLECTANCE_ADD_BAND_4 = -0.100000
REFLECTANCE_ADD_BAND_5 = -0.100000
REFLECTANCE_ADD_BAND_6 = -0.100000
REFLECTANCE_ADD_BAND_7 = -0.100000
REFLECTANCE_ADD_BAND_8 = -0.100000
REFLECTANCE_ADD_BAND_9 = -0.100000
```

21: Calibration coefficients used to convert DN values to TOA reflectance values for Landsat 8. Note that values for bands 10 and 11 are not provided, because these bands do not measure reflected solar radiation (except to a negligible degree). Also note that coefficients are the same for all bands – for Landsat 8 the DN values have been specifically generated for each scene to ensure constant TOA reflectance calibration coefficients. By Anders Knudby, CC BY 4.0.

For older Landsat data, as well as for several other types of sensors, it is necessary to explicitly calculate the amount of incoming solar radiation at the top of the atmosphere, and divide the radiance by this value while correcting for the difference between radiance and irradiance and the angle of incoming solar radiation. This is done generically with the following equation (rearranged for pedagogical reasons from equation 4, page 79, in the Landsat 7 handbook), explained in more detail below:

$$R_{TOA} = \frac{\pi * L}{\frac{1}{d^2} * I_{Sun} * \cos\theta_{Sun}}$$

where R_{TOA} is the top-of-atmosphere reflectance, L is (upward) radiance, d is the Earth-Sun distance in astronomical units, I_{Sun} is the mean extraterrestrial solar irradiance, and θ_{Sun} is the solar zenith angle – the angle between the direction toward the Sun and the normal of the Earth's surface. The equation is a division between the outgoing radiation in all upward directions (i.e. the exitance, in the numerator) and the incoming radiation in all downward directions (i.e. the irradiance, in the denominator).

The numerator is relatively simple: the upward radiance, L, which is quantified as the radiation per exposed unit surface area moving in directions within a specified solid angle, is multiplied by π to convert it to a value that quantifies radiation moving in all upward directions instead, to make it comparable to the incoming solar irradiance in the denominator.

In the denominator, we have I_{Sun}, the average incoming solar irradiance at the top of the atmosphere. Because the distance between the Sun and the Earth varies throughout the year and the amount of solar irradiance varies with it, we need to correct this by the factor $1/d^2$, where d is the actual distance between the Sun and the Earth when the image was acquired. If the Sun were at zenith (directly overhead), this would quantify the incoming irradiance per unit surface area. In cases when the Sun is not at zenith, we must multiply this value by the cosine of the solar zenith angle to arrive at the irradiance per exposed unit surface area.

Note: In the above two equations, the λ subscript indicating that these calculations must be performed on a band-by-band basis has been omitted for clarity, but keep this in mind. L varies by band, as does I_{Sun}, so both must be calculated specifically for each band to be correct.

With these calculations, you are able to convert the DN values in each pixel to a measure of TOA reflectance, which is a physical property of the Earth-Atmosphere system. This fact, that TOA reflectance is not an arbitrary number dependent on sensor calibration and illumination, but rather a physical property of the parts of the Earth seen in the image, allows us to use images captured by different sensors, at different times, and for different areas, and directly compare the values in each pixel. The one remaining problem in this approach is that most researchers are more interested in the reflectance of the Earth's surface than in the reflectance at the top of the atmosphere, but the state of the atmosphere and the state of the surface both influence TOA reflectance. The next step to further process the data from each pixel is thus to remove the influence that the atmosphere has on these values, to arrive at a measure of surface reflectance!

This page titled 1.3: Calculations of TOA radiance and TOA reflectance is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.4: EMR interactions with the Earth's atmosphere and surface

After electromagnetic radiation has been created by the Sun, the part of it that has found its way through the vacuum of space to the top of the Earth's atmosphere must pass through the atmosphere, be reflected by the Earth's surface, pass through the atmosphere again on its way back to space, and then arrive at the sensor in order to be recorded. While nothing happens to the radiation field as it passes through empty space, several things happen as it interacts with the Earth's atmosphere and surface. It is due to these interactions that the measured radiation ends up containing information about the Earth environment, so it is important to take a closer look at exactly what happens in these interactions, and how it affects the radiation field.

Interactions with the atmosphere

The interaction between electromagnetic radiation and the Earth's atmosphere can be considered to have three components: *refraction* that changes the direction of propagation of the radiation field due to density differences between outer space and the atmosphere, *scattering* that changes the direction of propagation of individual photons as they are absorbed and re-emitted by gasses or aerosols or other atmospheric constituents without changing wavelength, and *absorption* that convert photons into vibrations in a molecule, energy which is (later) re-emitted as one or more photons with longer wavelength(s). Each will be considered in more detail below.

Refraction

Refraction is the bending (and slowing down) of the direction of propagation of electromagnetic radiation as it moves between two media with different densities. This happens as radiation arrives from outer space (density ≈0) and enters the atmosphere (density >0). The angle at which the direction of propagation changes is determined by the refractive indices of the two media. The refractive index of a medium (n) is determined as the ratio of the speed of electromagnetic radiation in a vacuum (c) to the similar speed in the medium (c_n): $n=c/c_n$. The refractive index of a standard atmosphere is 1.0003, while the refractive index of water is 1.33. Using the refractive indices of the two media, the amount of refraction can be determined with Snell's Law: $n_1 * \sin\Theta_1 = n_2 * \sin\Theta_2$.

where n are the refractive indices of the two media and Θ are the angles at which the direction of propagation intersects the normal of the surface separating the two media (Figure 22). Refraction is rarely a relevant factor in the practical use of remote sensing data. Its only important influence concerns the georeferencing of imagery collected when the Sun is close to the horizon, and this is a problem that is nearly always dealt with by the image provider. One important situation in which refraction is important and must be considered is when an image analyst needs to precisely geolocate underwater objects (such as features on the seafloor in coastal areas).

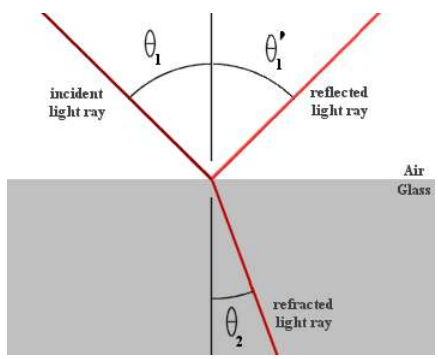

22: Refraction according to Snell's Law. While this image illustrates refraction of light passing from air to glass, the same principle applies to the air-water interface. Example snells law by RJHall, Wikimedia Commons, CC BY-SA 2.0 AT.

Scattering

One of the two remaining processes that influence electromagnetic radiation as it passes through the atmosphere is scattering. Scattering happens when a photon interacts with something in the atmosphere that causes it to change direction. Depending on the size of the object that the photon interacts with, two distinct types of scattering are recognized. *Rayleigh* scattering happens when the object is much smaller than the wavelength of the radiation. In the case of sunlight and the Earth's atmosphere this means that

Rayleigh scattering is cause by atmospheric gasses like N_2, O_2, CO_2 etc. *Mie* scattering happens when the object is similar in size to the wavelength of the radiation, which means that it is caused by aerosols like smoke and dust particles. Additional scattering can happen if radiation interacts with particles larger in size than its wavelength, like water droplets or sand particles.

While refraction is predictable and can be determined by Snell's Law, scattering is an inherently stochastic process: what happens to an individual photon as it passes through the atmosphere is entirely unpredictable, including whether or not it experiences any scattering, and if so which direction it is reemitted in. However, the magnitude and direction of scattering that happens on average to the many photons in a radiation field is predictable.

Rayleigh scattering

A fact that has great importance for remote sensing of the Earth is that the magnitude of Rayleigh scattering is inversely related to the 4th power of the wavelength of the radiation. In other words, radiation with shorter wavelengths is scattered much more by Rayleigh scattering than radiation at longer wavelengths. In the visible wavelengths, this means that blue light is scattered more than green light, which in turn is scattered more than red light. This is the process that makes the Earth's oceans look blue when viewed from space. What happens is that over very dark Earth surfaces, such as the oceans, the majority of radiation reaching the Earth surface is absorbed rather than reflected by it. What is visible from space is thus not radiation reflected by the surface, but rather radiation scattering from within the atmosphere. Because blue wavelengths are those most strongly scattered through Rayleigh scattering, this scattered radiation as a whole looks blue (Figure 23). Another effect of Rayleigh scattering is that regardless of what is on the Earth's surface, a space-based sensor will detect a substantial amount of blue light coming from the Earth-Atmosphere system. This can be a problem because the 'blue signal' form the atmosphere overwhelms variations in 'blue reflectance' on the surface. But it can also be an advantage, because measurements in the blue wavelengths can help assess the strength of Rayleigh scattering across the visible and infrared spectrum, which can in turn be corrected for. This is the basis for the 'aerosol' band that was included on Landsat 8 OLI (but was not found on its predecessor instruments), on Sentinel-2, and on the WorldView-2 and -3 sensors.

23: Earth from space. The Blue Marble (globe west) by Reto Stöckli, Visible Earth (NASA), Visible Earth Image Use Policy.

While any scattering in the atmosphere is a source of noise (for those interested in using satellite imagery to characterize the Earth's surface), Rayleigh scattering is a relatively benign source of noise because its wavelength dependence makes it largely predictable, and because the gasses responsible for it tend to have stable concentrations across space and time. Rayleigh scattering is therefore not a source of great uncertainty for most remote sensing applications.

Mie scattering

Mie scattering, because its strength and wavelength dependence depends on the type and density of the particulates that cause it to happen, varies substantially through time and space. As a result it is one of the most important causes of uncertainty in remote sensing, especially when using satellite data to study dark parts of the Earth's surface from which the amount of reflected radiation is small relative to the total signal from atmospheric scattering. For the same reason it is hard to generalize its importance, but broadly speaking the strength of Mie scattering exceeds that of Rayleigh scattering, and while it still diminishes with increasing

wavelength its influence extends further into the infrared spectrum. Because Mie scattering is caused by atmospheric particulates, it is often dramatically increased during dust storms, forest fires, or other events that caused the atmospheric aerosol load to increase. One such example is seen in Figure 24.

24: Greatly increased Mie scattering caused by forest fires in Australia, 2009. Smoke clouds by Warren, Flickr, CC BY-NC-SA 2.0.

Absorption

The last important thing that happens to electromagnetic radiation as it passes through the atmosphere is that it is partially absorbed by atmospheric gasses (mostly H_2O, CO_2 and O_3). While the energy absorbed is ultimately re-emitted by these gas molecules, the re-emission happens at wavelengths typically outside the spectrum considered in optical remote sensing (but which may be important for thermal remote sensing), so for practical purposes the absorbed photons can be considered gone when absorbed. The strength of absorption is highly dependent on wavelength because it happens most easily when the radiation has a wavelength (frequency) that is similar to a resonant frequency of the gas doing the absorption, which in turn depends on its atomic or molecular structure. For example, due to its molecular structure, O_2 is particularly good at absorbing electromagnetic radiation with wavelengths right around 760 nm, but not at 750 or 770 nm. Similar wavelengths exist at which other gasses are effective or not at absorbing EMR, and in combination the atmospheric gasses let some wavelengths pass through the atmosphere with almost no absorption, while other wavelengths are almost entirely absorbed before they reach the surface of the Earth (Figure 25 and Figure 26). As is especially clear in Figure 26, water vapour is responsible for much of the total gaseous absorption of EMR in the atmosphere, including in the visible spectrum (not clearly shown on that figure). This is an important challenge for remote sensing because while the concentrations of the other gasses are relatively stable through time and space, water vapour concentrations vary greatly through time (humid vs. dry days) and through space (dry arctic vs. humid tropical).

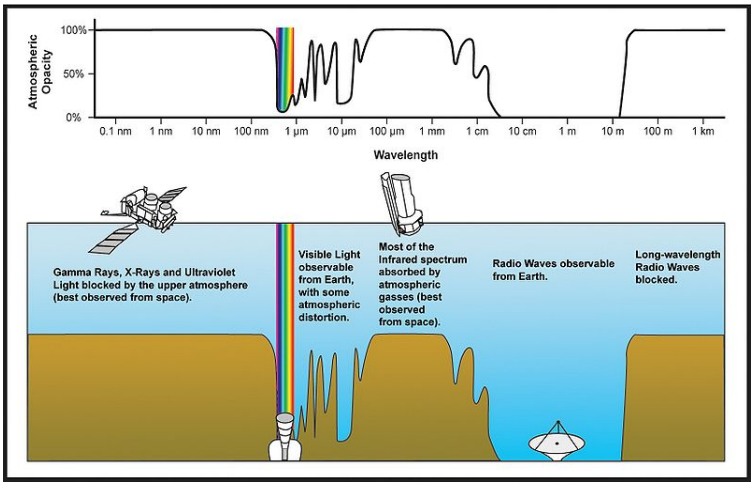

25: Atmospheric opacity by wavelength across the electromagnetic spectrum. Atmospheric electromagnetic transmittance or opacity by Edwtie, created by NASA, Wikimedia Commons, public domain.

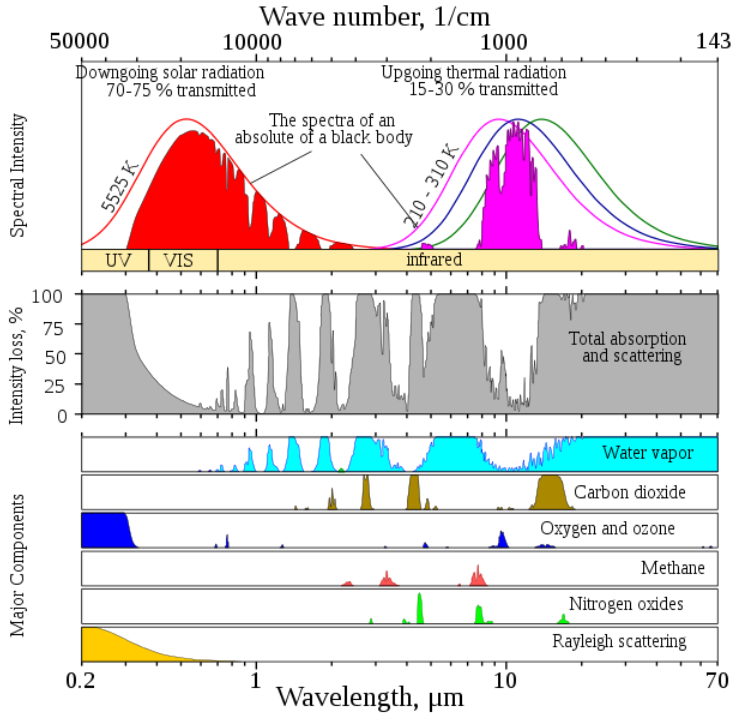

Figure 25, but covering a wider range of wavelengths and showing which gasses are principally responsible for absorption at which wavelength. Atmospheric Transmission by Cepheiden, Wikimedia Commons, CC0 1.0.

Interactions with the surface

The part of the radiation field that has made it through the atmosphere without being absorbed or scattered back toward space now reaches the Earth's surface. For any wavelength that is of relevance to remote sensing, only one of two things can now happen to each individual photon – it can be absorbed by the Earth's surface, or it can be reflected back toward space. The probability of reflection rather than absorption happening is termed the reflectance of the surface, and it depends on the material on the surface as well as the wavelength of the incoming radiation. Each surface material has a unique 'signature' that defines what proportion of radiation is reflected for each wavelength. For example, water reflects a small amount of blue and green wavelengths (typically around 5% – 10% depending on turbidity), less of the red wavelengths, and almost nothing in the infrared wavelengths. Vegetation, on the other hand, reflected around half of all incoming infrared radiation, except for specific wavelengths that are effectively absorbed by liquid water in the leaves. These spectral signatures are commonly portrayed as graphs, with wavelengths along the x-axis and reflectance along the y-axis (as in Figure 27).

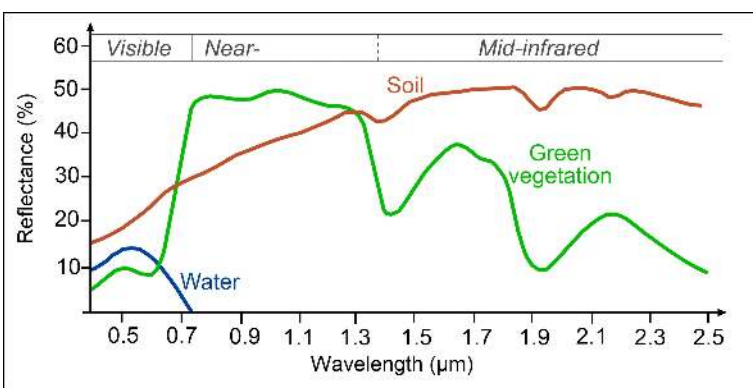

27: Spectral signatures of broadly defined common types of Earth surface materials. Reflectance of water, soil and vegetation at different wavelengths by Science Education through Earth Observation for High Schools (SEOS), CC BY-NC-SA 2.0.

Spectral signatures are what enables us to differentiate between different materials on the Earth's surface when we look at a satellite image. As shown in Figure 27, water has near-zero reflectance at wavelengths longer than 0.7 μm (700 nm), while both soil and green vegetation has reflectances around 40% at 1.3 μm. Measuring the amount of radiation reflected off the Earth-Atmosphere system at 1.3 μm will thus be particularly helpful at differentiating water from the two terrestrial surface types. Similarly, measurements at wavelengths around 1.4 μm (where liquid water in vegetation is a strong absorber) or 1.9 μm (same) can be effective to differentiate between soil and green vegetation.

As a more detailed example, spectral signatures have been effective for large-scale geological surveying/prospecting because different minerals (that may be characteristic of different sub-surface conditions) can be identified through their unique spectral signatures (Figure 28).

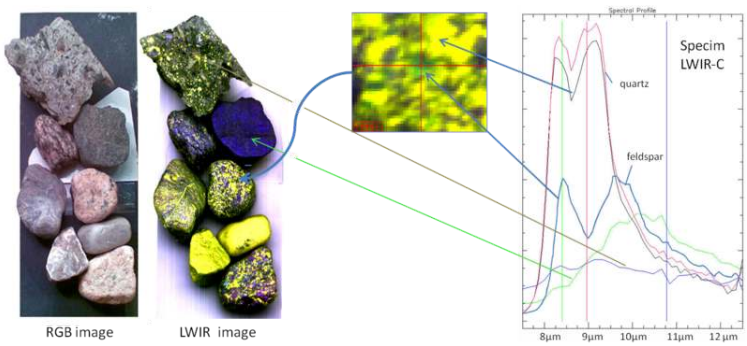

28: Examples of spectral signatures of minerals. HIS LWIR stones by Aappo Roos, Wikimedia Commons, CC BY-SA 3.0.

The part of the radiation field that is reflected by the Earth's surface must naturally make its way back up through the atmosphere, with the attendant refraction, scattering and absorption, before it can be measured by any space-based sensor. While there are many relative advantages and disadvantages to air-borne vs. space-borne sensors, the ability of air-borne sensors to measure the reflected EMR field before it has had to pass through the atmosphere a second time is one distinct advantage.

Atmospheric correction

Unless you are an atmospheric scientist, the interactions between the electromagnetic radiation field and the Earth's surface is the one important thing that is of interest to you in remote sensing, because that is what lets you use the measured radiation to infer something about the state of the Earth's surface – whether it is covered in water or land, what minerals the surface is composed of, how healthy the vegetation is, or how much sediment is in the water. The fact that interactions also happen between the EMR field and the atmosphere, and that these interactions influence the intensity of radiation measured by the sensor at different wavelengths, is an important source of noise. One of the great and difficult challenges in remote sensing thus is to correct for the effect that the atmosphere has on the measured radiation, thus converting our measurements from a description of the Earth-Atmosphere system to a description only of the Earth's surface. This process is called atmospheric correction. In practical terms, atmospheric correction is typically used to convert an estimate of TOA radiance or TOA reflectance to an estimate of surface reflectance. This conversion is important because it lets you make direct comparisons between the reflectances observed in different bands for a pixel and the spectral signatures measured on Earth such as those in Figure 27 and Figure 28.

There are different approaches to atmospheric correction, ranging from the very simple to the very complicated. The simplest approach is to employ the assumption that somewhere in the image you want to perform atmospheric correction for there is a pixel that is completely, perfectly, black, across the bands available in the image. This is called the 'dark pixel assumption'. Strictly speaking this does not have to be a single pixel, one could also assume that at least one pixel has zero reflectance in band 1, one pixel (maybe the same pixel, maybe another one) has zero reflectance in band 2, and so on. Assuming that this is true, for each band, the pixel in the image with the least amount of radiance measured coming from it (i.e. with the lowest DN value) must be the zero-reflectance pixel, and whatever radiance is measured as coming from it must have originated from scattering in the atmosphere (because if it had reached the Earth's surface it would have been absorbed by it). The hypothetical zero-reflectance pixel can thus give a first-order estimate of how much radiance originates from scattering in the atmosphere. Additionally assuming that this amount is constant throughout the image, the 'atmospheric contribution' can be subtracted from all pixels, leaving only what was reflected by the surface, called the 'surface contribution'. Of course, the assumption of the existence of a perfectly black

pixel in the image may not be true, and the atmospheric contribution may also not be the same everywhere, but as a method that can be applied without having any actual knowledge of the state of the atmosphere when the image was acquired, it is not a bad method.

The ‘dark-pixel’ method is an example of what is called image-driven methods of atmospheric correction. Other such methods make other assumptions, such as the presence of dense dark vegetation with typical spectral characteristics, or the presence of a range of surface types with known spectral signatures. These methods work reasonably well when their assumptions are met, but can fail spectacularly when they are not.

More sophisticated methods rely on some form of information about the state of the atmosphere at the time of image acquisition, typically the amount and type of aerosols and the amount of water vapour. With this information, numerical models that quantify the scattering and absorption and thus the transfer of radiation of different wavelengths through the atmosphere can be used to model the surface reflectance that must have existed, in combination with the known atmosphere, to produce the observed TOA reflectance. Some such methods do not require specific information on aerosols and water vapour, but rather need you to tell them what the visibility in the area was at the time of image acquisitions – since visibility is reduced with increasing amounts of water vapour in the atmosphere the methods can work out the water vapour and go from there. Atmospheric correction is a very active area of research in optical remote sensing, and modern satellite sensors are routinely designed to include measurement bands that help in an image-driven estimation of water vapour and aerosol load.

Almost all atmospheric correction methods make the fundamental assumption that the atmosphere is uniform within the satellite image considered. While this is obviously not perfectly true, it is often but not always a reasonable assumption. Hyperspectral sensors, which record radiation in many, often hundreds, of contiguous bands, have the ability to estimate water vapour and aerosols loads on a per-pixel basis, thus effectively mapping the atmospheric components contributing the greatest uncertainty, after which a per-pixel atmospheric correction can be applied.

The importance of atmospheric correction and the estimation of surface reflectance for each pixel is not just to enable comparison with field-measured spectral signatures, or to identify Earth surface materials. More fundamentally, its importance is due to the fact that surface reflectance is a fundamental physical property of the surface, and that it is not influenced by illumination conditions, sensor design, atmospheric state, or any of the other factors that differ between different images of the same area. For a pixel, an estimate of surface reflectance is thus directly comparable between images taken at different times, with different sensors, at different solar zenith angles, and under dry or humid, clear or smoky atmospheres. What this means is that an area can be monitored and in the absence of change the estimated surface reflectance (by band, i.e. the spectral signature) should not change – which in turn means that if the spectral signature is indeed seen to change through time, actual change is happening on the surface. Proper atmospheric correction is thus (with a few exceptions) the basis for a large field of applied remote sensing – detection of change.

This page titled 1.4: EMR interactions with the Earth’s atmosphere and surface is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.5: Visualization and manipulation of satellite images

Most of this course will deal with the typical kind of ‘satellite image’ you might see on the news or on Google Earth, in other words the kind of image produced by passive optical sensors. These images can in many ways be considered to be similar to the photos you take with your phone, so you can do similar things to them to make them better. Except in the context of remote sensing ‘better’ doesn’t mean photoshopping them to remove red eyes and create more vibrant colours, it means showing them in ways that help interpreting them better for specific purposes. This is a form of *scientific visualization*, a process in which displaying data helps explore and understand them. The reason exploring data visually can be more useful than running them through a computer algorithm is that humans are exceptionally adept at interpreting spatial patterns. As an example, just imagine you’re walking through a field in Eastern Canada and see the scene below. You are immediately able to realize that the dark part of the image seems ‘out of place’ and thus might be important, the shape suggests to you that it is a large animal, and you can bring all kinds of context (the time of year, the environment, the fact that bears are known to live in the area) to help in your interpretation of what is going on.

29: People are exceptionally good at recognizing structures in images, and while computer vision is in rapid development, the human brain still often outperforms artificial intelligence in this field. A Bear In The Woods by Linda Tanner, Flickr, CC BY-NC-ND 2.0.

Un-enhanced greyscale display

A lot goes on behind the scenes when you display an image on a computer screen. To look at the details we will take the simplest possible example first, an 8-bit greyscale image, but keep in mind that the same principles apply to data of different kinds. An 8-bit greyscale image (commonly, but misleadingly, also called a black-and-white image) corresponds to a single two-dimensional data layer with values ranging from 0 to 255. Digitally, the image file (e.g. a jpg, png, tif, or other image file) contains a bunch of numbers, as well as some information about the dimensions of the image. In the example below, the image contains 16 pixels arranged in 4 rows and 4 columns, and a value between 0 and 255 is associated with each cell. This is the Digital Number (DN) value. Most computer screens are able to display 256 different intensities of light, and use three colours (red, green and blue) to do so, which enables the production of 256 * 256 * 256 = 16.8 million colours. Think of there being three little ‘coloured light guns’ inside each pixel in the computer screen, which determine how much light to produce, of each colour, in each pixel. When displaying greyscale images, the intensity of light produced in red, green, and blue is all the same, allowing 256 shades of grey. The darkest of these (0 light intensity in red, 0 in green, and 0 in blue) produces a black colour on the screen, and the brightest (255 in red, 255 in green and 255 in blue) produces white. The example shown in Figure 30 is thus the simplest visualization (on the right) of the DN values in the file (on the left).

0	50	85	125
50	85	125	170
85	125	170	215
125	170	215	255

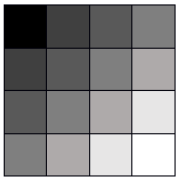

30: Single-layer greyscale image, showing the digital numbers on the left, and their visualization on the right. By Anders Knudby, CC BY 4.0.

Un-enhanced true-colour display

Most images have more than a single data layer, and can thus be used to produce colour images. Phone cameras record image files with three layers, corresponding to red, green, and blue, so visualizing these files is easy with the three 'colour guns' in your computer screen. A true-colour visualization is, not surprisingly, one in which you display the 'red' data layer using the red 'colour gun', and so on. For example:

Red:

237	255	128	128
0	34	34	128
63	63	34	128
0	195	255	0

Blue:

36	0	0	0
0	76	76	0
204	204	76	0
0	195	255	0

Green:

28	242	255	255
0	177	177	255
72	72	177	255
0	195	255	0

True-colour composite:

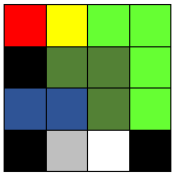

31: DN values in the red, green, and blue bands, and their direct display in a true-colour composite. By Anders Knudby, CC BY 4.0.

Un-enhanced false-colour display

Most passive optical sensors on satellites record light in more than three 'bands', and thus produce image files with more than three data layers. Some of these bands may be different shades of visible light (e.g. violet), while others, such as infrared and ultraviolet, fall outside the spectrum that is visible to humans. To show such data, false-colour visualizations are used. The name false-colour simply states that what is shown does not represent what the images area would look like to a human, it does not imply that there is anything wrong with the image (in fact false-colour displays are very useful). False-colour displays come in many flavours – just imagine the number of different combinations of three bands you can make with data from a sensors that has more than 200 bands… The most common false-colour visualization is often called 'colour-infrared', and it displays a 'near-infrared' band using the red colour gun, a 'red' band using the green colour gun, and a 'green' band using the blue colour gun. Figure 32 provides a side-by-side comparison of true-colour and a false-colour display of a Landsat image containing Zanzibar Town.

32: True-colour and colour-infrared displays of a Landsat image containing Zanzibar Town. By Anders Knudby, CC BY 4.0.

Image enhancement

In the preceding examples, we have imagined that the images always contained values between 0 and 255, neatly providing the 256 different possible values that the colour guns in a typical computer screen can display. There are two situations in which using the DN value found in the image file directly to determine the intensity of light produced by a colour gun is not the best option:

- Not all sensors record data using 8 bits per pixel (per colour). For (older) sensors that use less than 8 bits, or (newer) sensors that use more than 8 bits, a conversion between the DN value and the intensity with which it should be displayed is typically necessary.
- If the sensor has produced an image over a particularly bright or dark area, or at least if such an area is of specific interest, the values found over that area (e.g. between 50 and 80) can be displayed using a wider range of intensities on the screen (e.g. between 20 and 110), thus improving the visual contrast over the area of interest.

To better understand these issues, we need to consider how the DN values are stored in computer memory. We will consider a single pixel and its DN value for a single band. For an 8-bit sensor, this DN value is stored in 8 bits, each of which can either be ON (1) or OFF (0). The value is encoded as illustrated in Figure 33. In that example, the bits representing values of 128, 32, 16, 8 and 2 are all ON, while the others are OFF. The DN value of the pixel is therefore calculated as 128 + 32 + 16 + 8 + 2 = 186. The range of possible values with 8 bits goes from 0 (all bits are OFF) to 255 (all bits are ON), for a total of 256 possible values. In fact, the number of possible values, given the number of bits, is easily calculated as $2^{n,}$ where n is the number of bits.

Value if ON	128	64	32	16	8	4	2	1
ON / OFF	1	0	0	1	1	0	1	0

33: This example illustrates an 8-bit pixel with the value 186. By Anders Knudby, CC BY 4.0.

This is the reason why sensors that use fewer than 8 bits to store the DN for each pixel cannot produce 256 different values. However, just for the sake of argument, if we were to display 4-bit DN values (DN values ranging between 0 and 15) directly without transformation, we would only see very dark shades of grey (no actual sensors record information using only 4 bits, but the first Landsat sensors used only 6 bits, and the issue is the same).

On the other hand, sensors that use 12 bits to store each DN value (as many modern sensors do), contain values ranging from 0 to 2^{12}-1 = 4095. To display this range of values, we need to transform them to all fit between 0 and 255, which is most easily done by simply dividing by 16 and rounding down. Note that in this case, while the computer screen is not capable of displaying the data in all its detail, image processing algorithms will still be able to make use of it.

Things get more interesting when we have an image in which a specific range of values are of greatest interest. A common example is the kind of image seen in Figure 32, which contains some water (very dark), land (relatively dark and largely green), and clouds (very bright). Unless we are specifically interested in studying the clouds in this image, we don't really need to be able to distinguish between pixels with DN values of 245 vs. 250 vs. 255. What we can do, then, is to display all 'bright' pixels that we don't really care about with the same colour (white), and thus 'reserve' more of the different brightnesses with which we can

display pixels for the DN values we care more about. Mathematically, this means creating a table that translates each DN value, for each band, into an equivalent brightness value with which it is going to be displayed on screen. Graphically, we can show this as done below in Figure 34, where the DN values are shown on the x axis, and the brightness with which they are displayed is shown on the y axis. The blue line represents the direct use of the DN value, without contrast enhancement. The orange line represents a linear contrast enhancement, in which all DN values less than 50 are displayed with a brightness of 0, and all DN values greater than 240 are displayed with a brightness of 255. This allows us to use the remaining 254 possible brightness values to better display the smaller differences between pixels with DN values between 50 and 240. The grey line is a square root transformation that keeps some contrast along the entire range of DN values, but focuses on displaying differences between pixels with low DN values, at the expense of those with high DN values. Examples of what such contrast enhancement can look like in practice are shown in Figures 35-37:

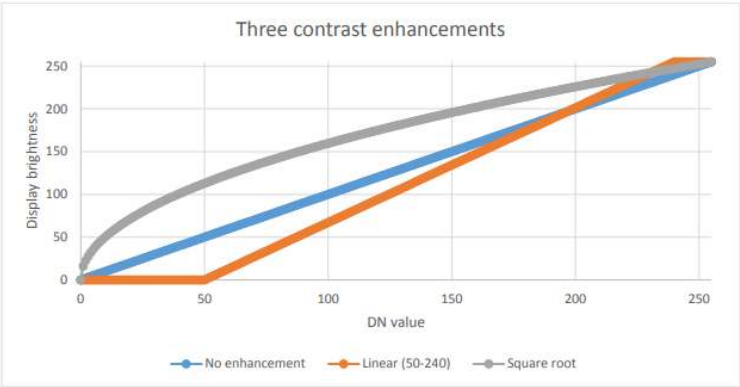

34: Three options for displaying 8-bit data. By Anders Knudby, CC BY 4.0.

35: No contrast enhancement. By Anders Knudby, CC BY 4.0.

36: Linear contrast enhancement using the minimum and maximum values of the image in place of 50 and 240 in the above example. By Anders Knudby, CC BY 4.0.

37: Contrast enhancement known as histogram equalization. This is mathematically more complex than the examples shown above, but often provides good results. By Anders Knudby, CC BY 4.0.

Pan-sharpening

When you watch Hollywood movies and the CIA analyst or the surprisingly good-looking police officer clicks a button to 'enhance' footage from a surveillance camera, thus revealing facial features from what was previously a blurry blob, well, that's Hollywood. The thing about images is that they are just data, and you can't make detail (new data) appear where none exists in the original data. There's a funny explanation of this general idea here. However, with some satellite imagery there is one kind of 'enhancement' that actually kind of works (although it wreaks havoc in subtle ways on the image radiometry). The technique in question is called pan-sharpening, which refers to the fact that most passive optical satellite sensors have a 'pan-chromatic' band that records data on brightness across all the visible bands and at higher spatial resolution than the individual colour bands. For example, the sensors on Landsat satellites 4-7 have blue, green and red bands, all with 30 meter spatial resolution, and a pan-chromatic band that records blue and green as well as red light, with 15 meter spatial resolution. For modern high-resolution sensors, the difference in spatial resolution is even greater, with WorldView-2 having colour bands at 2 meter resolution and a pan-chromatic band at 50 cm resolution. Broadly speaking, we can therefore combine the colour information from the individual colour bands with the spatial detail from the pan-chromatic band, to get something like what you see in Figure 38.

There are several algorithms that can be used to perform this operation, and each employs a different set of assumptions and rules for how to generate colour information for each of the pixels in the pan-sharpened image. While context determines which algorithm is superior for a given image, the results of most algorithms are impressive and can greatly aid in visual interpretation of

satellite imagery. Most commercial image processing software packages have one or more algorithms for pan-sharpening, so while the process itself is fairly complex, actually applying it to an image is done with the click of a button.

Figure 7) by Fu et al, CC BY 4.0.

Band ratios and other band math

The previous sections have mostly dealt with different ways of displaying the DN values found in each band – in greyscale or colour, directly or transformed through contrast enhancement, pan-sharpened or not. An alternative is to create new bands using a process commonly called 'Band Math'.

For example, we might want to produce a band that quantifies how 'red' something is. The problem with the original 'red' band in a typical satellite image is seen by comparing the 'red' DN values from the top left part of Figure 31 with the colours shown in the bottom right part of the same figure. Arguably, the pixel that is in the top left corner is what most people would consider 'red', but the two pixels with the highest DN values in the 'red' band are the yellow and white pixels (both have 'red' DN values of 255). Despite their high 'red' DN value, these pixels do not appear 'red' because they **also** have high DN values in one or more of the other colours. To better quantify the 'red' appearance of each pixel, we can calculate something called the 'red chromatic coordinate' (RCC). For each pixel, this is calculated as RCC = Red / (Red + Green + Blue), where Red, Green and Blue represent the respective DN values in each pixel. As you can see from the calculations in Figure 39, the red pixel in the top left corner has a higher RCC value than any of the other pixels, followed by yellow pixel next to it (because yellow is 'close to' red). Note that band math operates on each pixel individually, and completes the same mathematical operation on all pixels in an image. In doing so, it produces a new band that is a function of the original bands. Other common uses of band math is to calculate 'band ratios', such as the vegetation index called the 'Simple Ratio' that is calculated as SR = NIR / Red, where NIR is a value (ideally the TOA reflectance) in the near-infrared band, and Red is the same value in the red band. Because healthy plants reflect near-infrared light very effectively, while they absorb most red light (and thus reflect very little red light for the sensor to detect), areas with vegetation tend to have high reflectance and DN values in the near-infrared band and low values in the red band. This leads areas with vegetation to have very high Simple Ratio values, and this band ratio can therefore be used to quickly and easily produce a visualization that helps people distinguish between vegetated and un-vegetated areas. An even more popular vegetation index is the Normalized Difference Vegetation Index (NDVI), which is calculated as NDVI = (NIR – Red) / (NIR + Red).

Red:

237	255	128	128
0	34	34	128
63	63	34	128
0	195	255	0

Blue:

36	0	0	0
0	76	76	0
204	204	76	0
0	195	255	0

Green:

28	242	255	255
0	177	177	255
72	72	177	255
0	195	255	0

RCC:

0.8	0.5	0.3	0.3
NA	0.1	0.1	0.3
0.2	0.2	0.1	0.3
NA	0.3	0.3	NA

Figure 31. The Red Chromatic Coordinate (RCC) is calculated as RCC = Red / (Red + Green + Blue). Pixels with DN values of 0 in all bands (i.e. black pixels) have been given NoData (NA) values. By Anders Knudby, CC BY 4.0.

Band math can be used to calculate new bands with any combination of input bands you can think of, and different kinds of band math are used, among many other applications, to monitor global deforestation, track seasonal changes in snow cover, map water depth, and monitor urban sprawl.

Tasseled cap transformation (TCT)

A particular application of band math called the tasseled cap transformation has been developed for Landsat data, and is now commonly used to produce a standardized and readily interpretable set of bands from the original data. It was originally developed for an instrument called Landsat MSS (which does not operate anymore), and has since been updated for Landsat TM (the sensors on Landsat 4 and 5), ETM+ (the sensor on Landsat 7), and OLI (the passive optical sensor on Landsat 8). The idea behind the tasseled cap transformation is to produce three new bands, 'greenness', 'brightness' and 'wetness', that describe the actual variations in those parameters across the image. In a typical satellite image of a rural landscape, areas covered by water (blue dots) and bare soil (red dots) will form a "line" in two-dimensional (NIR / RED) space, while vegetation (green dots) will lie away from the line and various stages of plant growth together will form a shape roughly like a tasseled cap (Figure 40). The three new bands are created by effectively rotating the axes of the coordinate system, so 'brightness' is determined along the 'line of soils' in Figure 40, 'greenness' is determined along an axis that departs from the line of soils and increases in value as it moves toward the top of the 'tasseled cap', and 'wetness' is determined along a third axis (not shown in Figure 40) perpendicular to the first two. An example of the results, based on an IKONOS image from Idaho, USA, is shown in Figure 41. In practical terms, each component of the tasseled cap is calculated as a weighted sum of the values in the original bands. For example, for Landsat TM data, the 'brightness' is calculated as BR = 0.3037 * BLUE + 0.2793 * GREEN + 0.4343 * RED + 0.5585 * NIR + 0.5082 * SWIR1 + 0.1863 * SWIR2 (SWIR1 and SWIR2 are two shortwave infrared bands also present on Landsat TM sensors).

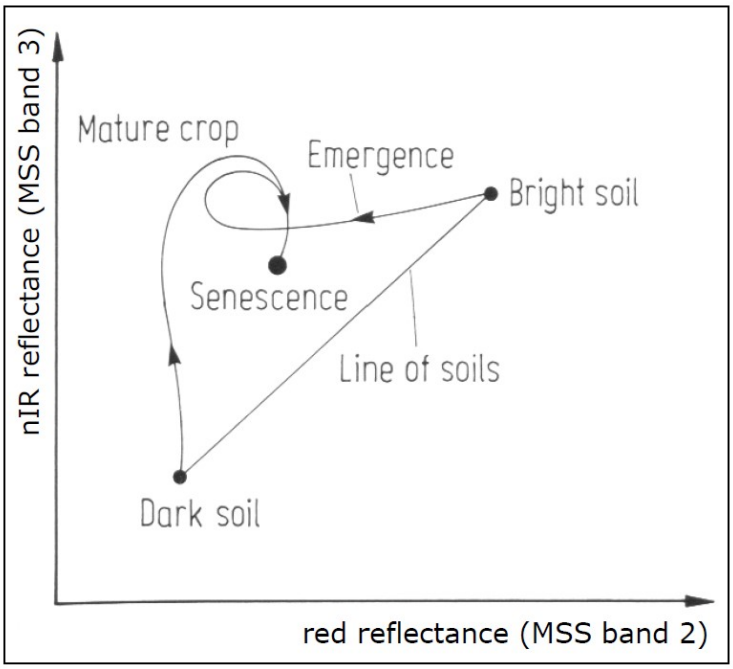

40: The principle of the tasseled cap transformation. Bare soils of varying brightness fall along a line when plotted with their values in the RED and NIR bands, while vegetation of increasing density is found within the 'cap', up and left from the line of soils. Vegetation at various growth stages, growing on soils of different brightness, thus form a shape that looks like a tasseled cap. The original formulation of the tasseled cap included a third dimension, using the green band. Tasseled Cap Transformation for bright and dark soil by Gurkengräber (adapted from Richards and Jia), Wikimedia Commons, CC BY-SA 4.0.

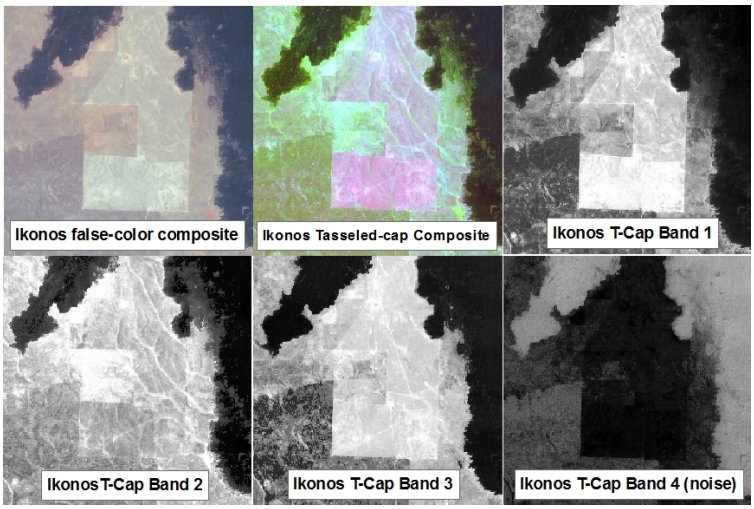

41: An example of the tasseled cap transformation applied to an IKONOS image from Idaho, USA. tcap_example by Jason Karl, The Landscape Toolbox, CC0 1.0.

Principal Components Analysis (PCA)

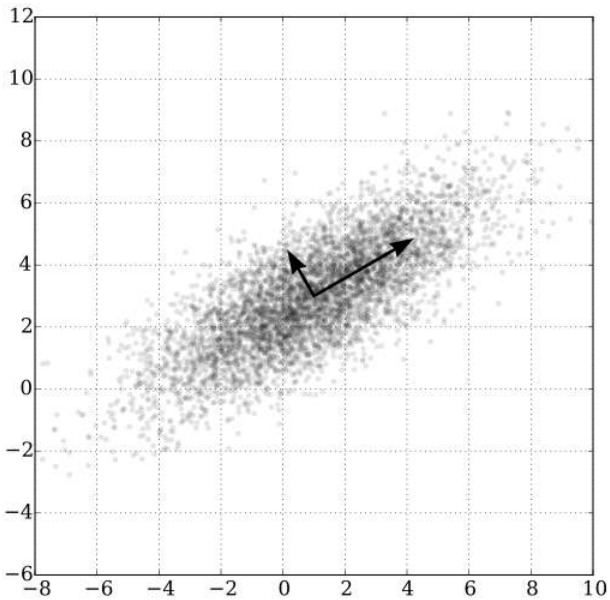

42: A two-dimensional example of Principal Components Analysis. GaussianScatterPCA by Nicoguaro, Wikimedia Commons, CC BY 4.0.

The last kind of band math we will look at in this chapter is similar to the tasseled cap transformation in that the calculation involves a weighted sum of the original bands, but it is different in the sense that the weights are determined from the properties of the image itself, and as such they vary from image to image. Figure 42 provides a good illustration of how it works. We can imagine that each points in the scatter plot is a pixel from our image, and that it is plotted according to its value in one band on the x-axis and another band on the y-axis. It is clear from the distribution of points in Figure 42 that points with high values in the one band tend to also have high values in the other band, and that most of the distribution of the points runs along an imaginary line from the bottom left to the top right of the plot. The orientation of this imaginary line is called the first principal component (PC_1), and is illustrated with the longest arrow. But there is also some variation orthogonally (perpendicular) to PC_1. This is the

orientation of the second principal component, PC_2, illustrated with the shorter arrow. Depending on how many bands are in an image, one could continue to find new such lines in the third, fourth etc. dimensions, each oriented along the greatest distribution of points in any orientation perpendicular to all previous principal components. In Figure 42 the original data set had only two dimensions, so two principal components were created. The length of each arrow corresponds to the amount of total distribution captured by each principal component, and it is clear that PC_2 has captured rather little of the distribution in the original points. This can be interpreted as PC_2 containing very little information, and often the higher-numbered principal components can be removed entirely with little loss of information. With this knowledge, new bands can be created, in which the value of each pixel is no longer its DN value in the original bands, but rather its value along each of the principal component axes. This is the main strength of PCA: that it allows you to transform the data into a smaller number of bands (thus taking up less hard disk space, and allowing quicker reading and processing) without losing much information. Research has even shown that for some kinds of images (e.g. a cloud-free Landsat image of an agricultural area), specific principal components tend to correspond to specific environmental variables, such as soil humidity, vegetation cover, and so on. However, a note of caution is important here. Because the orientation of the principal components (the arrows in Figure 42) depend on the distribution of the totality of pixel values in each individual image, the principal component values of a pixel are not directly interpretable, and may vary dramatically between images, even for identical pixels. PCA can thus help you to reduce the data volume of an image, but it destroys (or at least complicates) interpretation in the process. Another really good explanation of PCA, including an animated scatterplot that I think explains things really well, is found here.

This page titled 1.5: Visualization and manipulation of satellite images is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.6: Classification

While visual interpretation of a satellite image may suffice for some purposes, it also has significant shortcomings. For example, if you want to find out where all the urban areas are in an image, a) it will take a long time for even an experienced image analyst to look through an entire image and determine which pixels are ‘urban’ and which aren’t, and b) the resulting map of urban areas will necessarily be somewhat subjective, as it is based on the individual analyst’s interpretation of what ‘urban’ means and how that is likely to appear in the image. A very common alternative is to use a classification algorithm to translate the colour observed in each pixel into a thematic class that describes its dominant land cover, thus turning the image into a land cover map. This process is called image classification.

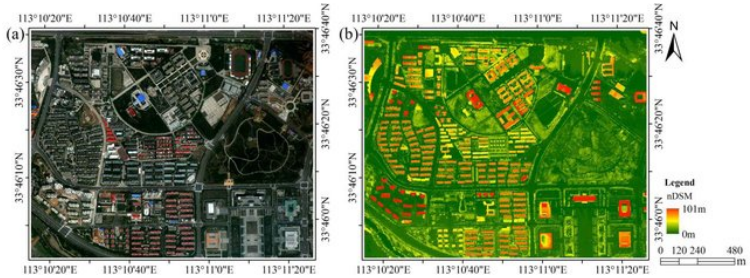

Figure 1) by Wu et al., CC BY 3.0.

You can identify two categories of approaches to image classification. The traditional and easiest way is to look at each pixel individually, and determine which thematic class corresponds to its colour. This is typically called per-pixel classification, and it’s what we will look at first. A newer and increasingly popular method is to first split the image into homogeneous segments, and then determine which thematic class corresponds to the attributes of each segment. These attributes may be the colour of the segment as well as other things like shape, size, texture, and location. This is typically called object-based image analysis, and we’ll look at that in the second half of this chapter.

Even in the category of per-pixel classification, two different approaches are available. One is called ‘supervised’ classification, because the image analyst ‘supervises’ the classification by providing some additional information in its early stages. The other is called ‘unsupervised’ classification, because an algorithm does most of the work (almost) unaided, and the image analyst only has to step in at the end and finish things up. Each has its advantages and disadvantages, which will be outlined in the following.

Supervised per-pixel classification

The idea behind supervised classification is that the image analyst provides the computer with some information that allows calibration of a classification algorithm. This algorithm is then applied to every pixel in the image to produce the required map. How this works is best explained with an example. The image shown in Figure 44 is from California, and we want to translate this image into a classification with the following three (broad) classes: “Urban”, “Vegetation”, and “Water”. The number of classes, and the definition of each class, can have great impact on the success of the classification – in our example we are ignoring the fact that substantial parts of the area seem to consist of bare soil (and so does not really fall into any of our three classes). And we may realize that the water in the image has very different colour depending on how turbid it is, and that some of the urban areas are very bright while others are a darker shade of grey, but we are ignoring those issues for now.

Figure 44: True-colour composite showing part of a Landsat image from California. An urban area with some turbid water is seen in the top right corner, bordered by a mix of vegetated and barren areas and a lake with clear water, and a largely forested mountainous area is seen in the bottom left part of the image, with some bright cleared areas near the centre left. By Anders Knudby, CC BY 4.0, based on Landsat 5 data (USGS).

The "supervision" in supervised classification almost always comes in the form of a calibration data set, which consists of a set of points and/or polygons that are known (or believed) to belong to each class. In Figure 45, such a data set has been provided in the form of three polygons. The red polygon outlines an area known to be "Urban", and similarly the blue polygon is "Water" and the green polygon is "Vegetation". Note that the example in Figure 45 is **not** an example of best practice – it is preferable to have more and smaller polygons for each class spread out throughout the image, because it helps the polygons to only cover pixels of the intended class, and also to incorporate spatial variations in e.g. vegetation density, water quality etc.

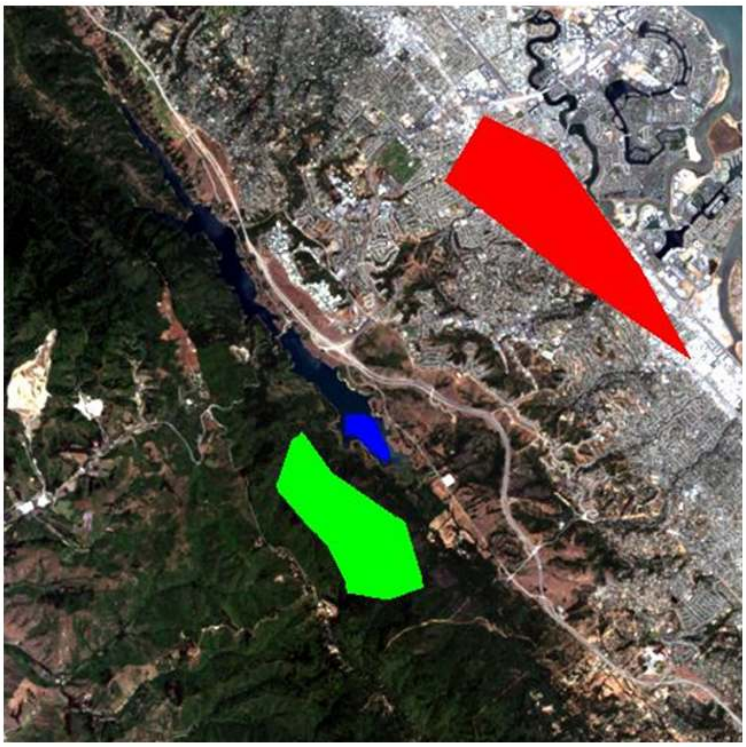

Figure 44, with three polygons overlaid. The red polygon represents an area known by the image analyst to be “Urban”, the blue polygon is “Water”, and the green polygon is “Vegetation”. By Anders Knudby, CC BY 4.0.

Now let’s look at how those polygons help us turn the image into a map of the three classes. Basically, the polygons tell the computer “look at the pixels under the red polygon – that’s what ‘Urban’ pixels look like”, and the computer can then go and find all the other pixels in the image that also look like that, and label them “Urban”. And so on for the other classes. However, some pixels may look a bit “Urban” and a bit like “Vegetation”, so we need a mathematical way to figure out which class each pixel resembles the most. We need a classification algorithm.

If we take all the values of all the pixels in Landsat bands 3 and 4 and show them on a scatterplot, we get something like Figure 46. This image has an 8-bit radiometric resolution, so values in each band theoretically range from 0 to 255, although in reality we see that the smallest values in the image are greater than 0. Values from band 3 are shown on the x-axis, and values from band 4 on the y-axis.

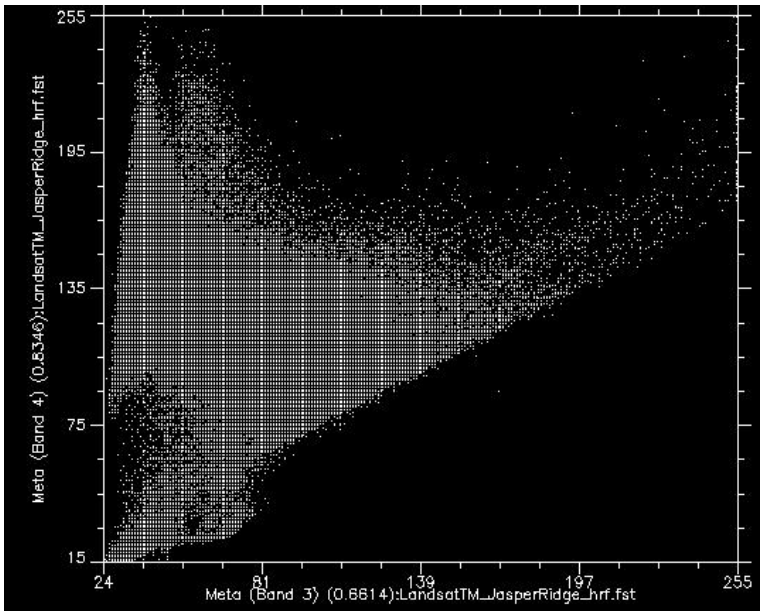

Figure 44. Scatterplot created using ENVI software. By Anders Knudby, CC BY 4.0.

Now, if we colour all the points that come from pixels under the red polygon (i.e. pixels we "know" to be "Urban"), and do the same with pixels under the blue and green polygons, we get something like Figure 47. There are a few important things to note in Figure 47. All the blue points ("Water") are found in the bottom left corner of the figure, below the yellow circle, with low values in band 3 and low values in band 4. This is indeed typical of water, because water absorbs incoming radiation in the red (band 3) and near-infrared (band 4) wavelengths very effectively, so very little is reflected to be detected by the sensor. The green points ("Vegetation") form a long area along the left hand side of the figure, with low values in band 3 and moderate to high values in band 4. Again, this seems reasonable, as vegetation absorbs incoming radiation in the red band effectively (using it for photosynthesis) while reflecting incoming radiation in the near-infrared band. The red points ("Urban") form a larger area near the centre of the figure, and cover a much wider range of values than either of the two other classes. While their values are similar to those of "Vegetation" in band 4, they are generally higher in band 3.

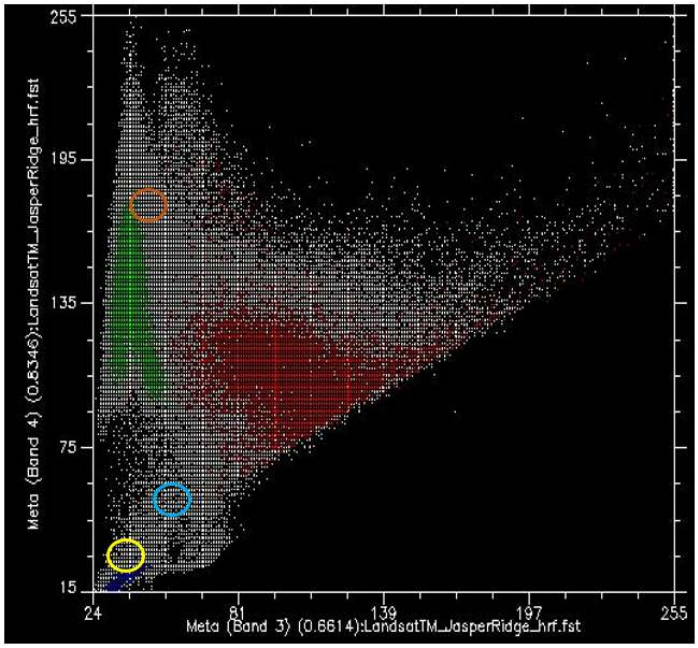

Figure 46, but with each point coloured by the polygon it is under. Scatterplot created using ENVI software. By Anders Knudby, CC BY 4.0.

What we want the supervised classification algorithm to do now, is to take all the other pixels in the image (i.e. all the white points on the scatter plot) and assign them to one of the three classes based on their colour. For example, which class do you think the white dots in the yellow circle in Figure 47 should be assigned to? Water, probably. And what about those in the light brown circle? Vegetation, probably. But what about those in the light blue circle? Not quite so easy to determine.

Note: The classification algorithm can make use of all the bands in the Landsat image, as well as any other information we provide for the entire image (like a digital elevation model), but because it is easiest to continue to show this in two dimensions using bands 3 and 4 only we will continue to do so. Just keep in mind that the scatter plot is in reality an n-dimensional plot, where n is equal to the number of bands (and other data layers) we want to use in the classification.

Minimum distance classifier

One way to estimate which class each pixel belongs to is to calculate the "distance" between the pixel and the centre of all the pixels known to belong to each class, and then assign it to the closest one. By "distance", what we mean here is the distance in "feature space", in which dimensions are defined by each of the variables we are considering (in our case bands 3 and 4), as opposed to physical distance. Our feature space, therefore, is two-dimensional, and distances can be calculated using a standard Euclidian distance.

As an example, for the points in Figure 48 we have calculated the mean values of all green, red, and blue pixels for bands 3 and 4, and indicated them with large dots. Let's say they have the following values:

Table 3: Mean values in bands 3 and 4 for the "Urban", "Vegetation" and "Water" classes show in Figure 49.

Mean values	Red points ("Urban")	Green points ("Vegetation")	Blue points ("Water")
Band 3	100	40	35
Band 4	105	135	20

Then let's say a pixel indicated by the yellow dot in Figure 48 has a value of 55 in band 3, and 61 in band 4. We can then calculate the Euclidian distance between the point and the mean value of each class:

Distance to red mean: (100-55)2+(105-61)2 = 62.9

Distance to green mean: (40-55)2+(135-61)2 = 75.5

Distance to blue mean: (35-55)2+(20-61)2 = 45.6

Because the Euclidian distance is shortest to the centre of the blue points, the minimum distance classifier will assign this particular point to the "Blue" class. While the minimum distance classifier is very simple and fast and often performs well, this example illustrates one important weakness: In our example, the distribution of values for the "Water" class is very small – water is basically always dark and blue-green, and even turbid water or water with lots of algae in it basically still looks dark and blue-green. The distribution of values for the "Vegetation" class is much greater, especially in band 4, because some vegetation is dense and some not, some vegetation is healthy and some not, some vegetation may be mixed in with dark soil, bright soil, or even some urban features like a road. The same goes for the "Urban" class, which has a wide distribution of values both in bands 3 and 4. In reality, the yellow dot in Figure 48 is not very likely to be water, because water that has such high values in both bands 3 and 4 basically doesn't exist. It is much more likely to be an unusual kind of vegetation, or an unusual urban area, or (even more likely) a mix between those two classes. The next classifier we will look at explicitly takes into account the distribution of values in each class, to remedy this problem.

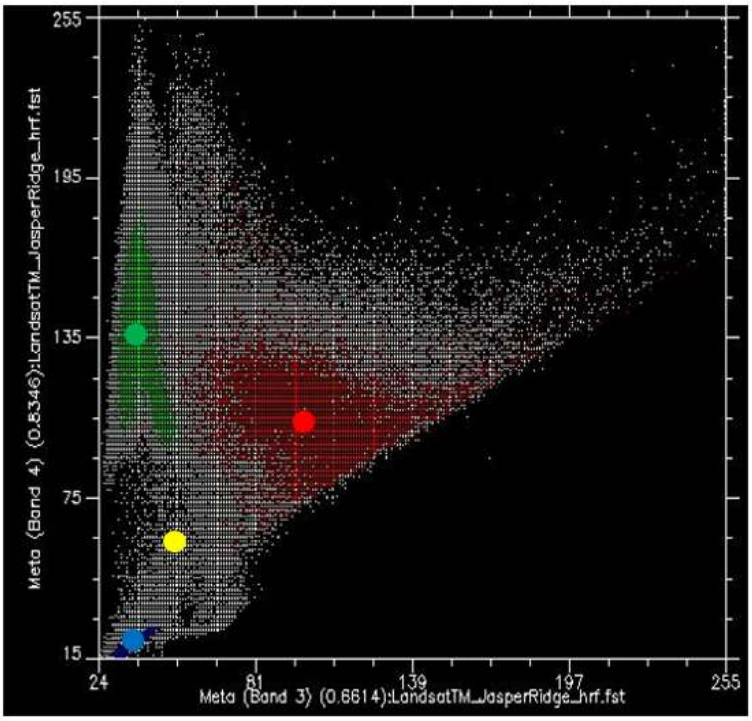

48: The minimum distance classifier assigns the class whose centre is nearest (in feature space) to each pixel. The mean value of all the red points, in bands 3 and 4, is indicated by the large red dot, and similarly for the green and blue points. The yellow dot indicates a pixel that we wish to assign to one of the three classes. Scatterplot created using ENVI software. By Anders Knudby, CC BY 4.0.

Maximum likelihood classifier

Until around 10 years ago, the maximum likelihood classifier was the go-to algorithm for image classification, and it is still popular, implemented in all serious remote sensing software, and typically among the best-performing algorithms for a given task. Mathematical descriptions of how it works can seem complicated because they rely on Bayesian statistics applied in multiple dimensions, but the principle is relatively simple: Instead of calculating the distance to the centre of each class (in feature space) and thus finding the closest class, we will calculate the probability that the pixel belongs to each class, and thus find the most probable class. What we need to do to make the math work is to make a few assumptions.

- We will assume that before we know the colour of the pixel, the probability of it belonging to one class is the same as the probability of it belonging to any other class. This seems reasonable enough (although in our image there is clearly much more "Vegetation" than there is "Water" so one could argue a pixel with unknown colour is more likely to be vegetation than water… this can be incorporated into the classifier, but rarely is, and we will ignore it for now)
- We will assume that the distribution of values in each band and for each class is Gaussian, i.e. follows a normal distribution (a bell curve).

To start with a one-dimensional example, our situation could look like this if we had only two classes:

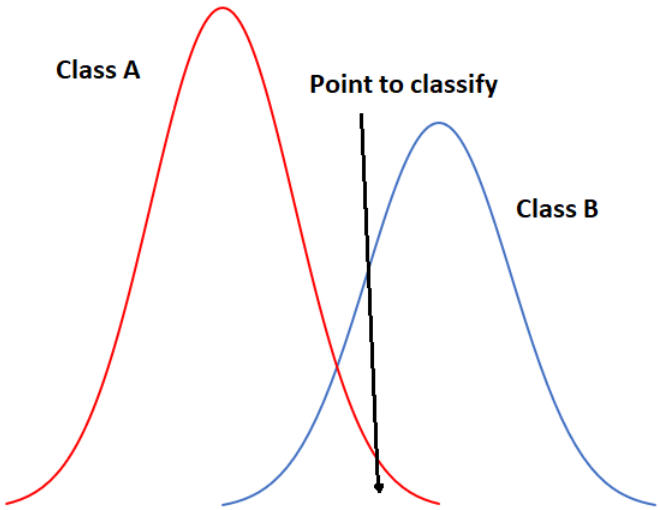

49: One-dimensional example of maximum likelihood classification with two classes. By Anders Knudby, CC BY 4.0.

In Figure 49, the x-axis represents values in an image band, and the y-axis shows the number of pixels in each class that has a given value in that band. Clearly, Class A typically has low values, and Class B typically has high values, but the distribution of values in each band is significant enough that there is some overlap between the two. Because both distributions are Gaussian, we can calculate both the mean and the standard deviation for each class, and we can then calculate the z-score (how many standard deviations we are away from the mean). In Figure 49, the two classes have the same standard deviation (the 'bells' have the same 'width'), and because the point is located a little closer to the mean of Class B than it is to Class A, its z-score would be lowest for Class B and it would be assigned to that class. A slightly more realistic example is provided below in Figure 50, where we have two dimensions and three classes. The standard deviations in band 4 (x-axis) and band 3 (y-axis) are shown as equiprobability contours. The challenge for the maximum likelihood classifier in this case is to find the class for which the point lies within the equiprobability contour closest to the class centre. See for example that the contours for Class A and Class B overlap, and that the standard deviations of Class A are larger than those of Class B. As a result, the red dot is closer (in feature space) to the centre of Class B than it is to the centre of Class A, but it is on the third equiprobability contour of Class B and on the second of Class A. The minimum distance classifier would classify this point as "Class B" based on the shorter Euclidian distance, while the maximum likelihood classifier would classify it as "Class A" because of its greater probability of belonging to that class (as per the assumptions employed). Which is more likely to be correct? Most comparisons between these classifiers suggest that the maximum likelihood classifier tends to produce more accurate results, but that is not a guarantee that it is always superior.

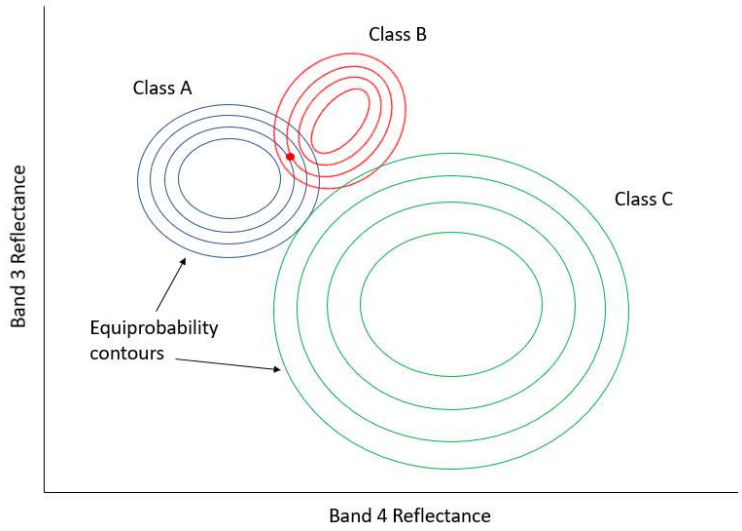

50: Two-dimensional example of the maximum likelihood classification situation, with six classes that have unequal standard distributions. By Anders Knudby, CC BY 4.0.

Non-parametric classifiers

Within the last decade or so, remote sensing scientists have increasingly looked to the field of machine learning to adopt new classification techniques. The idea of classification is fundamentally very generic – you have some data about something (in our case values in bands for a pixel) and you want to find out what it is (in our case what the land cover is). A problem could hardly be more generic, so versions of it are found everywhere: A bank has some information about a customer (age, sex, address, income, loan repayment history) and wants to find out whether (s)he should be considered a “low risk”, “medium risk” or “high risk” for a new $100,000 loan. A meteorologist has information about current weather (“rainy, 5 °C”) and atmospheric variables (“1003 mb, 10 m/s wind from NW”), and needs to determine whether it will rain or not in three hours. A computer has some information about fingerprints found at a crime scene (length, curvature, relative position, of each line) and needs to find out if they are yours or somebody else’s). Because the task is generic, and because users outside the realm of remote sensing have vast amounts of money and can use classification algorithms to generate profit, computer scientists have developed many techniques to solve this generic task, and some of those techniques have been adopted in remote sensing. We will look at a single example here, but keep in mind that there are many other generic classification algorithms that can be used in remote sensing.

The one we will look at is called a decision tree classifier. As with the other classifiers, the decision tree classifier works in a two-step process: 1) Calibrate the classification algorithm, and 2) apply it to all pixels in the image. A decision tree classifier is calibrated by recursively splitting the entire data set (all the pixels under the polygons in Figure 45) to maximize the homogeneity of the two parts (called nodes). A small illustration: let’s say we have 7 data points (you should **never** have only seven data points when calibrating a classifier, this small number is used for illustration purposes only!):

Table 4: Seven data points used to develop a decision-tree classifier.

Data point	Band 1 value	Band 2 value	Actual class
1	10	30	A
2	20	40	A
3	30	40	A
4	15	55	B
5	35	40	B
6	40	35	B
7	45	35	B

The first task is to find a value, either in band 1 or in band 2, that can be used to split the data set into two nodes such that, to the extent possible, all the points of Class A are in one node and all the points in Class B are in the other. Algorithmically, the way this is done is by testing all possible values and quantifying the homogeneity of the resulting classes. So, we see that the smallest value in Band 1 is 10, and the largest is 45. If we split the data set according to the rule that “All points with Band 1 < 11 go to node X, and all others go to node Y”, we will end up with points split like this:

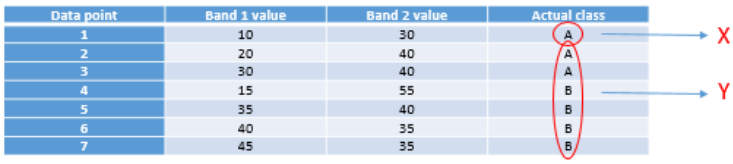

51: Points split according to threshold value 11 in band 1. By Anders Knudby, CC BY 4.0.

As we can see, this leaves us with a single “A” point in the one node (X), and two “A” points and four “B” points in the other node (Y). To find out if we can do better, we try to use the value 12 instead of 11 (which gives us the same result), 13 (still the same), and so on, and when we have tested all the values in band 1 we keep going with all the values in band 2. Eventually, we will find that using the value 31 in band 1 gives us the following result:

Data point	Band 1 value	Band 2 value	Actual class
1	10	30	A
2	20	40	A
3	30	40	A
4	15	55	B
5	35	40	B
6	40	35	B
7	45	35	B

X

Y

52: Points spit according to threshold value 31 in band 1. By Anders Knudby, CC BY 4.0.

This is almost perfect, except that we have a single “B” in node X. But ok, pretty good for a first split. We can depict this in a “tree” form like this:

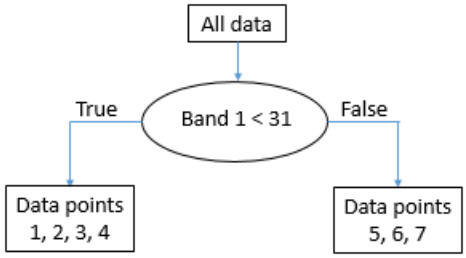

53: The emerging “tree” structure from splitting the data on threshold value 31 in band 1. Each collection of data points is called a node. The ‘root node’ contains all the data points. ‘Leaf nodes’ are also called ‘terminal nodes’, they are the end points. By Anders Knudby, CC BY 4.0.

The node with data points 5, 6, and 7 (all have band 1 values above 31) is now what is called a “pure node” – it consists of data points of only one class, so we don’t need to split it anymore. Nodes that are end points are also called “leafs”. The node with data points 1, 2, 3 and 4 is not “pure”, because it contains a mix of Class A and Class B points. So we go through again and test all the different possible values we can use as a threshold to split that node (and only that node), in both bands. It so happens that the point from Class B in that node has a value in band 2 that is higher than all the other points, so a split value of 45 works well, and we can update the tree like this:

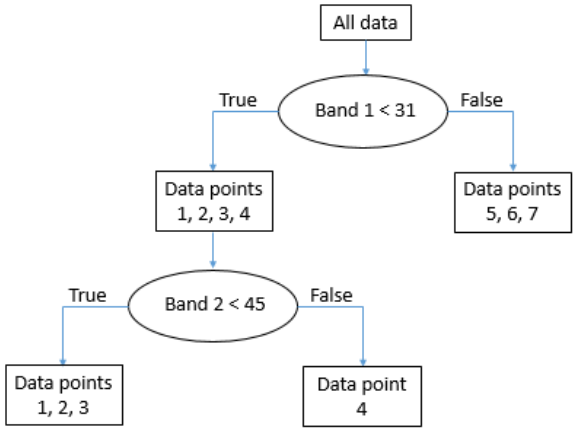

54: The final “tree” structure. All nodes (final parts of the data set) are now pure. By Anders Knudby, CC BY 4.0.

With the “tree” in place, we can now take every other pixel in the image and “drop it down” the tree to see which leaf it lands in. For example, a pixel with values of 35 in band 1 and 25 in band two will “go right” at the first test and thus land in the leaf that contains data points 5, 6 and 7. As all these points were from Class B, this pixel will be classified as Class B. And so on.

Note that leafs do not have to be “pure”, some trees stop splitting nodes when they are smaller than a certain size, or using some other criterion. In that case, a pixel landing in such a leaf will be assigned the class that has the most points in that leaf (and the information that it was not a pure leaf may even be used to indicate that the classification of this particular pixel is subject to some uncertainty).

The decision tree classifier is just one out of many possible examples of non-parametric classifiers. It is rarely used directly in the form shown above, but forms the basis for some of the most successful classification algorithms in use today. Other popular non-parametric classification algorithms include neural networks and support vector machines, both of which are implemented in much remote sensing software.

Unsupervised per-pixel classification

What if we don’t have the data necessary to calibrate a classification algorithm? If we don’t have the polygons shown in Figure 45, or the data points shown in Table 4? What do we do then? We use an unsupervised classification instead!

Unsupervised classification proceeds by letting an algorithm split pixels in an image into “natural clusters” – combinations of band values that are commonly found in the image. Once these natural clusters have been identified, the image analyst can then label them, typically based on a visual analysis of where these clusters are found in the image. The clustering is largely automatic, although the analyst provides a few initial parameters. One of the most common algorithms used to find natural clusters in an image is the K-Means algorithm, which works like this:

1) The analyst determines the desired number of classes. Basically, if you want a map with high thematic detail, you can set a large number of classes. Note also that classes can be combined later on, so it is often a good idea to set the number of desired classes to be slightly higher than what you think you will want in the end. A number of “seed” points equal to the desired number of classes is then randomly placed in feature space.

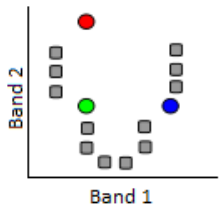

55: K-Means classification step 1. A number of “seed” points (coloured dots) are randomly distributed in feature space. Grey points here represent pixels to be clustered. Modified from K Means Example Step 1 by Weston.pace, Wikimedia Commons, CC BY-SA 3.0.

2) Clusters are then generated around the “seed” points by allocating all other points to the nearest seed.

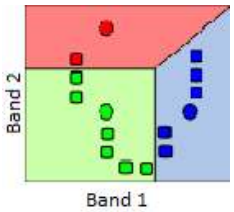

56: A cluster is formed around each seed by allocating all points to the nearest seed. Modified from K Means Example Step 2 by Weston.pace, Wikimedia Commons, CC BY-SA 3.0.

3) The centroid (geographic centre) of the points in each cluster becomes the new “seed”.

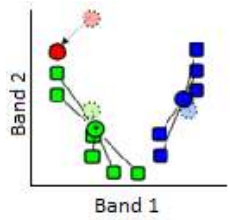

57: The seeds are moved to the centroid of each cluster. The centroid is calculated as the geographic centre of each cluster, i.e. it becomes located at the mean x value of all points in the cluster, and the mean y value of all points in the cluster. Modified from K Means Example Step 3 by Weston.pace, Wikimedia Commons, CC BY-SA 3.0.

4) Repeat steps 2 and 3 until stopping criterion. The stopping criterion can be that no point moves to another cluster, or that the centroid of each cluster moves less than a pre-specified distance, or that a certain number of iterations have been completed.

Other unsupervised classification algorithms do the clustering slightly differently. For example, a popular algorithm called ISODATA also allows for splitting of large clusters during the clustering process, and similarly for merging of small nearby clusters. Nevertheless, the result of the clustering algorithm is that each pixel in the entire image is part of a cluster. The hope, then, is that each cluster represents a land cover type that can be identified by the image analyst, e.g. by overlaying the location of all pixels in the cluster on the original image to visually identify what that cluster corresponds to. That is the final step in unsupervised classification – labeling of each of the clusters that have been produced. This is the step where it can be convenient to merge clusters, if e.g. you have one cluster that corresponds to turbid water and another that corresponds to clear water. Unless you are specifically interested in water quality, differentiating the two is probably not important, and merging them will provide a clearer map product. Also, you may simply have two clusters that both seem to correspond to healthy deciduous forest. Even if you work for a forest service, unless you can confidently figure out what the difference is between these two clusters, you can merge them into one and call them “deciduous forest”.

As an example, the image below shows the original image in the background, and the central pixels coloured according to the product of an unsupervised classification. It is clear that the “blue” area corresponds to pixels covered by water, and the green area corresponds largely to vegetation. More detailed analysis of the image would be necessary to label each area, especially the red and grey ones, appropriately.

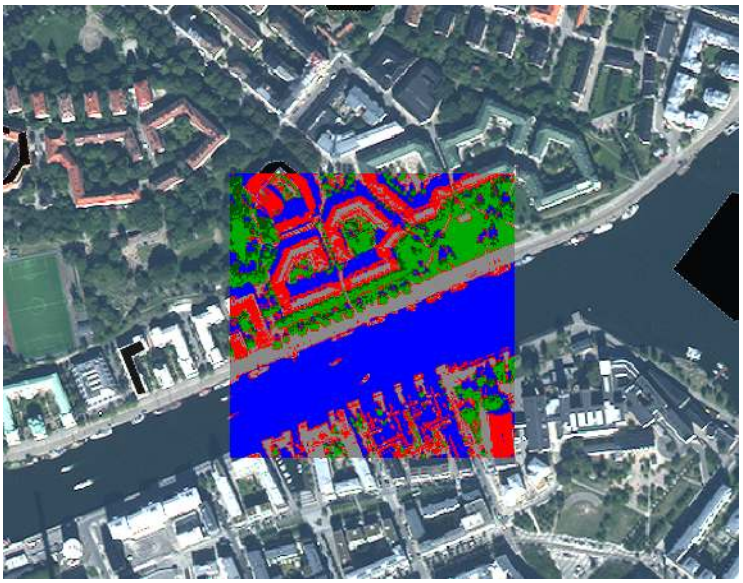

58: Example of correspondence between the original image and the clusters formed in an unsupervised classification process. By Anders Knudby, CC BY 4.0.

Land cover classification is one of the oldest uses of remote sensing, and is something many national governments do for their territory on a regular basis. For example, in Canada the Canada Centre for Remote Sensing works with partners from the US and

Mexico to create a North American Land Cover map. Global land cover maps are also produced by various institutions, like the USGS, the University of Maryland, ESA, and China, to name just a few.

One of the typical shortcomings of image classification schemes that operate on a pixel-by-pixel level is that images are noisy, and the land cover maps made from images inherit that noise. Another more important shortcoming is that there is information in an image beyond what is found in the individual pixels. An image is the perfect illustration of the saying that "the whole is greater than the sum of its parts", because images have structure, and structure is not taken into account when looking at each pixel independent from the context provided by all of the neighbouring pixels. For example, even without knowing the colour of a pixel, if I know that all of its neighbouring pixels are classified as "water", I can say with great confidence that the pixel in question is also "water". I will be wrong occasionally, but right most of the time. We will now look at a technique called object-based image analysis that takes context into account when generating image classifications. This advantage often allows it to outperform the more traditional pixel-by-pixel classification methods.

Object-based Image Analysis (OBIA)

Many new advances in the remote sensing world come from the hardware side of things. A new sensor is launched, and it has better spatial or spectral resolution than previous sensors, or it produces less noisy images, or it is made freely available when the alternatives were at cost. Drones are another example: the kind of imagery they produce is not substantially different from what was previously available from cameras on manned aircraft – in fact it is often inferior in quality – but the low cost of drones and the ease with which they can be deployed by non-experts has created a revolution in terms of the amount of available imagery from low altitudes, and the cost of obtaining high-resolution imagery for a small site.

One of the few substantial advances that have come from the software side is the development of object-based image analysis (OBIA). The fundamental principle in OBIA is to consider an image to be composed of objects rather than pixels. One advantage of this is that people tend to see the world as composed of objects, not pixels, so an image analysis that adopts the same view produces results that are more easily interpreted by people. For example, when you look at Figure 59, you probably see the face of a man, (if you know him you will also recognize who the man is).

59: Robert De Niro. Or, if you're a pixel-based image analysis software, a three-band raster with 1556 columns and 2247 lines, each pixel and band displaying a varying level of brightness. Robert De Niro KVIFF portrait by Petr Novák (che), Wikimedia Commons, CC BY-SA 2.5.

Because this is a digital image, we know that it is actually composed of a number of pixels arranged neatly in columns and rows, and that the brightness (i.e. the intensity of the red, green and blue colour in each pixel) can be represented by three numbers. So we could classify the bright parts of the image as "skin" and the less bright parts as "other", a mixed class comprising eyes, hair, shadows, and the background. That's not a particularly useful or meaningful classification, though! What would be more meaningful would be to classify the image into classes like "eye", "hand", "hair", "nose", etc.

An example somewhat more relevant to remote sensing is seen below in Figure 60, in which an urban area has been classified into objects, including an easy-recognizable stadium, streets, individual buildings, vegetation etc.

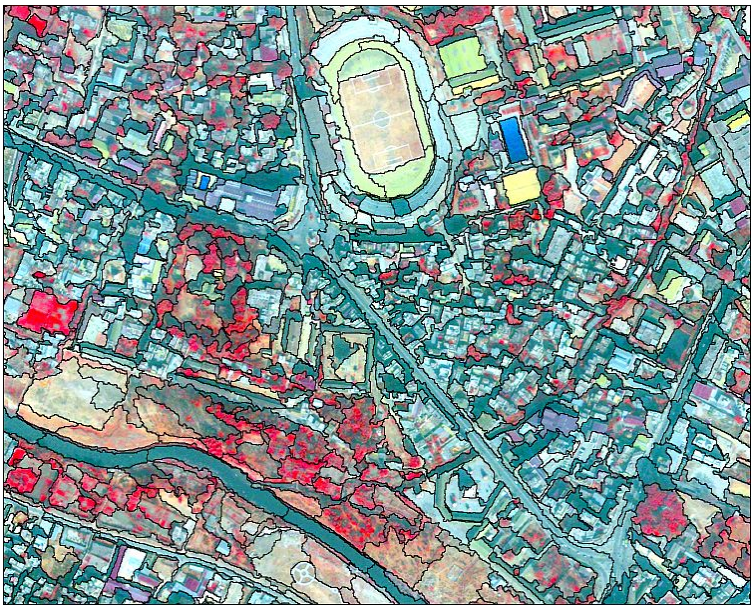

60: Classification of an urban area using object-based image analysis. Object based image analysis by Uddinkabir, Wikimedia Commons, CC BY-SA 4.0.

Image segmentation

The purpose of image segmentation is to take all the pixels in the image and divide them into segments – contiguous parts of the image that have similar colour. Image segmentation is useful because it takes us away from analyzing the image on a pixel-by-pixel basis and instead allows us to analyze the individual segments. There are a couple of very important advantages to this. First of all, we can look at the "average colour" of a segment and use that to classify that segment, rather than using the colour of each individual pixel to classify that pixel. If we are dealing with noisy imagery (and we always are!), using segment averages reduces the influence the noise has on the classification. Secondly, segments have a number of attributes that may be meaningful and can be used to help classify them – attributes pixels do not have. For example, an image segment is made up of a number of pixels, and we can thus quantify its size. Since not all segments will have the same size, there is information in the attribute 'segment size' that may be used to help classify the segments. For example, in Figure 60, notice that the segment covering the lake is quite large compared to all of the segments on land. That is because the lake is a very homogeneous part of the image. The same difference is seen between different parts of the land area – the segments just west of the lake are generally larger than those southwest of the lake, again this is because they are more homogeneous, and that homogeneity may tell us something important about what kind of land cover is found there. Apart from size, segments have a large number of other attributes that may or may not be helpful in a classification. Each segment has a specific number of neighbouring segments, and each segment is also darker or lighter or somewhere in the middle compared to its neighbours. See this page for a list of the attributes you can calculate for segments in ENVI's Feature Extraction Module (which is **not** the most comprehensive such module, by far). The ability to get all this information about segments can help to classify them, and this information is not all available for pixels (e.g. all pixels have exactly four neighbours, unless they are located on an image edge, so the number of neighbours is not helpful for classifying a pixel).

Different kinds of segmentation algorithms exist, and all are computationally complicated. Some are open source, while others are proprietary, so we don't even really know how they work. We will therefore not go into detail with the specifics of the segmentation algorithms, but they do have a few things in common that we can consider.

- Scale: All segmentation algorithms require a 'scale factor', which the user sets to determine how large (s)he wants the resulting segments to be. The scale factor does not necessarily equal a certain number of pixels, or a certain number of segments, but is usually to be considered a relative number. What it actually means in terms of the on-the-ground size of the resulting segments is typically found by a process of trial and error.
- Colour vs. shape: All segmentation algorithms need to make choices about where to draw the boundaries of each segment. Because it is usually desirable to have segments that are not too oddly shaped, this often involves a trade-off between whether

to add a pixel in an existing segment if a) that pixel makes the colour of the segment more homogeneous but also results in a weirder shape, or b) if that pixel makes the colour of the segment less homogeneous but results in a more compact shape. One or more parameters usually control this trade-off, and as for the 'scale' parameter finding the best setting is a matter of trial and error.

In an ideal world, the image segmentation would produce segments that each correspond to one, and exactly one, real-world object. For example, if you have an image of an urban area and you want to map out all the buildings, the image segmentation step would ideally result in each building being its own segment. In practice that is usually impossible because the segmentation algorithm doesn't know that you are looking for buildings… if you were looking to map roof tiles, having segments correspond to an entire roof would be useless, as it would be if you were looking to map city blocks. "But I could set the scale factor accordingly" you might say, and that is true to some extent. But not all roofs are the same size! If you are mapping buildings in an image that contains both your own house as well as the Pentagon, you are unlikely to find a scale factor that gives you exactly one segment covering each building… The solution to this problem is typically some manual intervention, in which the segments produced by the initial segmentation are modified according to specific rules. For example you might select a scale factor that works for your house and leave the Pentagon divided into 1000 segments, and then subsequently merge all neighbouring segments that have very similar colours. Assuming that the Pentagon's roof is fairly homogeneous, that would merge all of those segments, and assuming that your own house is surrounded by something different-looking, like a street, a backyard, or a driveway, your own roof would not be merged with its neighbouring segments. This ability to manually fiddle with the process to achieve the desired results is both a great strength and an important weakness of object-based image analysis. It is a strength because it allows an image analyst to produce results that are extremely accurate, but a weakness because it takes even expert analysts a long time to do this for each image. An example of a good segmentation is shown in Figure 61, which also illustrates the influence of changes in the segmentation parameters (compare the bottom left and the bottom right images).

It is worth noting here that object-based image analysis was initially developed and widely used in the medical imaging field, to analyse x-ray images, cells seen under microscopes, and so on. Two things are quite different between medical imaging and remote sensing: In medical imaging, people's health are very directly affected by the image analysis, so the fact that it takes extra time to get an accurate result is less of a constraint than it is in remote sensing, where the human impacts of poor image analysis are very difficult to assess (although they may ultimately be just as important). The other issue is that the images studied by doctors are created in highly controlled settings, with practically no background noise, the target of the image always in focus, and can be redone if difficult to analyze. In remote sensing, if haze, poor illumination, smoke, or other environmental factors combine to create a noisy image, our only option is to wait until the next time the satellite passes over the area again.

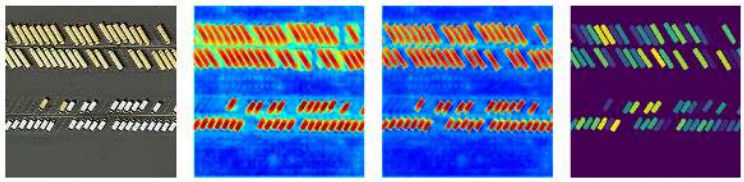

Figure 6) by Chen et al., CC BY 4.0.

Segment classification

Segment classification could in theory follow the supervised and unsupervised approaches outlined in the chapter on per-pixel classification. That is, areas with known land cover could be used to calibrate a classifier based on a pre-determined set of segment attributes (supervised approach), or a set of pre-determined attributes could be used in a clustering algorithm to define natural clusters of segments in the image, which could then be labeled by the image analyst (unsupervised approach). However, in practice object-based image analysis often proceeds in a more interactive fashion. One common approach is that an analyst develops a rule-set that is structured as a decision tree, sees what result that rule-set produces, modifies or adds to it, and checks again, all in a very user-intensive iterative process. This is one of those areas where remote sensing seems more like an art than a science, because image analysts gain experience with this process and become better and better at it, basically developing their own 'style' of developing the rule sets needed to achieve a good classification. For example, after having segmented an image, you may want to separate man-made surfaces (roads, sidewalks, parking lots, roofs) from natural surface, which you can typically because the

former tend to be grey and the latter tend not to be. So you develop a variable that quantifies how grey a segment is (for example using the HSV "saturation" value, more info here), and manually find a threshold value that effectively separates natural from man-made surfaces in your image. If you are doing urban studies, you may want to also differentiate between the different kinds of man-made surfaces. They are all grey, so with a pixel-based classifier you would now be out of luck. However, you are dealing with segments, not pixels, so you quantify how elongated each segment is, and define a threshold value that allows you to separate roads and sidewalks from roofs and parking lots. Then you use the width of the elongated segments to differentiate roads from sidewalks, and finally you use the fact that parking lots have neighbouring segments that are roads, while roofs don't, to separate those two. Finding all the right variables to use (HSV saturation value, elongation, width, neighbouring segment class) is an inherently subjective exercise that involves trial-and-error and that you get better at with experience. And once you have created a first decision-tree (or similar) structure for your rule set, it is very likely that your eyes will be drawn to one or two segments that, despite your best efforts, were still misclassified (maybe there is the occasional roof that extends so far over the house that its neighbouring segment is indeed a road, so it was misclassified as a parking lot). You can now go through and create additionally sophisticated rules that remedy those specific problems, in a process that only ends when your classification is perfect. This is tempting if you take pride in your work, but time-consuming, and it also ends up resulting in a rule-set that is so specific as to only ever work for the single image it was created for. The alternative is to keep the more generic rule set that was imperfect but worked fairly well, and then manually change the classification for those segments you know ended up misclassified.

In summary, object-based image classification is a relatively new methodology that relies on two steps: 1) dividing the image into contiguous and homogeneous segments, followed by 2) a classification of those segments. It can be used to produce image classifications that nearly always outperform per-pixel classifications in terms of accuracy, but also require substantially more time for an image analyst to conduct. The first, best, and most-used software for OBIA in the remote sensing field is eCognition, but other commercial software packages such as ENVI and Geomatica have also developed their own OBIA modules. GIS software packages such as ArcGIS and QGIS provide image segmentation tools, as well as tools that can be combined to perform the segment classification, but generally provide less streamlined OBIA workflows. In addition, at least one open source software package has been specifically designed for OBIA, and image segmentation tools are available in some image processing libraries, such as OTB.

This page titled 1.6: Classification is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.7: Accuracy assessment

Once we have produced a land cover (or other) classification from a remote sensing image, an obvious questions is "how accurate is that map?" It is important to answer this question because we want users of the map to have an appropriate amount of confidence in it. If the map is perfect, we want people to know this so they can get the maximum amount of use out of it. And if the map is no more accurate that a random assignment of classes to pixels would have been, we also want people to know that, so they don't use it for anything (except maybe hanging it on the wall or showing it to students as an example of what **not** to do…).

The subject of accuracy assessment also goes beyond classifications to maps of continuous variables, such as the Earth's surface temperature, near-surface CO_2 concentration, vegetation health, or other variables come in the form of continuous rather than discrete variables. Regardless of what your map shows, you'll want people to know how good it is, how much they can trust it. While there are similarities between assessing maps of categorical and continuous variables, the specific measures used to quantify accuracy are different between the two, so in this chapter we will treat each in turn.

Accuracy assessment for classifications

The basic principle for all accuracy assessment is to compare estimates with reality, and to quantify the difference between the two. In the context of remote sensing-based land cover classifications, the 'estimates' are the classes mapped for each pixel, and 'reality' is the actual land cover in the areas corresponding to each pixel. Given that the classification algorithm has already provided us with the 'estimates', the first challenge in accuracy assessment is to find data on 'reality'. Such data are often called 'ground-truth' data, and typically consist of georeferenced field observations of land cover. A technique often used is to physically go into the study area with a GPS and a camera, and take georeferenced photos that in turn allow the land cover to be determined visually from each photo. Because people can visually distinguish between different kinds of land cover with great accuracy, such data can reasonably be considered to represent 'reality'. In many cases, though, the term 'ground truth' oversells the accuracy of this kind of information. People may be good at distinguishing between 'desert' and 'forest' in a photo, but they are clearly less good at distinguishing between 'high-density forest' and 'medium-density forest'. Especially if the difference between two classes is based on percentage cover (e.g. the difference between medium-density and high-density forest may be whether trees cover more or less than 50% of the surface area) field observations may not always lead to a perfect description of reality. Many remote sensing scientists therefore prefer the term 'validation data', suggesting that these data are appropriate as the basis for comparison with remote-sensing based classifications, while at the same time acknowledging the potential that they do not correspond perfectly to the 'truth'.

Creating validation data

If you want to produce an honest and unbiased assessment of the accuracy of your land cover map (and I assume you do!), there are a couple of things to consider as you create your validation dataset:

- You should have validation data covering all the different land cover classes in your map. If you don't you will really only be able to assess the accuracy of the parts of the map covered by the classes you have data for.
- You should also ideally have validation data that are distributed randomly (or more or less evenly) throughout your study area. To produce a set of validation data that both covers all classes and has a good spatial distribution in your study area, a stratified random selection of validation points is often used (i.e. including a number of points from each class, with those points belonging to each class being randomly distributed within the area covered by that class).
- The number of data points used for each class should either be the same, or should reflect the relative extent of each class on your map. The former approach is most suitable if you want to compare classes and find out which ones are mapped better than others. The latter approach is most suitable if you want to produce a single accuracy estimate to the entire map.
- The more validation data, the better. However, creating validation data can take time and money, so getting 'enough' data is often a reasonable objective. Rules of thumb exist about what constitutes 'enough' data (e.g. 50 per class), but there are many exceptions to those rules.

If you use field observations to create your validation data, it is important to remember that the validation data should be comparable to the classes derived from your image, in several ways:

- The definitions used for each class should be the same between the classification and the validation data. For example, if in your classification you considered a 'water body' to be at least 0.1 km^2 in size, you need to keep this in mind as you create your validation data, so when you go in the field and one of your data points are in a puddle you do not consider it a 'water body' but instead figure out what the land cover around the puddle is.
- Related to the above point, keep in mind the spatial resolution of the image used to produce your classification. If you have based your classification on Landsat (TM, ETM+, OLI) imagery without pan-sharpening, each pixel corresponds to an area of approximately 30 x 30 meters on the ground. So when you go in the field, you should be documenting the dominant land cover in 30 x 30 meter areas, rather than the land cover at the exact coordinates of the data point.

An alternative approach to creating validation data, useful when going to the study area and collecting field observations is too costly, is visual inspection of high-resolution remote sensing imagery. If you choose this approach, you have to be confident that you can visually distinguish all the different classes, from the image, with high accuracy. People sometimes use imagery from Google Earth for validation, or they use visual interpretation of the same image used for the classification. The latter option seems a bit circular – as in 'why use a classifier in the first place, if you can confidently assign classes based on visual interpretation of the image?' However, visual interpretation may be entirely appropriate for accurately defining land cover for a number of validation data points, while doing visual interpretation of an entire image could be an enormously labour-intensive task. The same considerations outlined in bullet points above apply whether the validation data are created using field observations of visual interpretation of imagery.

An interesting new approach to creating validation data is to use publicly available geotagged photos, such as those available through Flickr or other sites where people share their photos. Especially for cities and popular tourist sites, the Internet contains a vast repository of geotagged photos that may be used by anyone as field observations. Some quality control is needed though, as not all photos available online are geotagged automatically with GPS (some are manually 'geotagged' when posted online), and most photos show land cover conditions at a time that is different from when the remote sensing image was acquired (e.g. winter vs. summer).

The confusion matrix

Once you have created a set of validation data that you trust, you can use their georeference to pair them up with the corresponding land cover mapped in the classification. You can think of the resulting comparison as a table that looks something like this:

Table 5: Comparison of estimates derived from a classification algorithm (left) with validation data (right).

Mapped land cover (estimate)	Validation data (reality)
Forest	Forest
Water	Water
Forest	Grassland
Grassland	Grassland
Grassland	Bare soil
Bare soil	Bare soil
…	…

With many validation data points, a method is required to summarize all this information, and in remote sensing the method that is used universally, and has been for decades, is called the confusion matrix (also called 'error matrix' or 'contingency table'). Using the four classes listed in the example above, the frame of the confusion matrix would look like this:

Table 6: Example of a confusion matrix, before values have been entered.

		Validation data				
	Class	Forest	Water	Grassland	Bare soil	Total
Classification	Forest					
	Water					
	Grassland					
	Bare soil					
	Total					

Read along the rows, each line tells you what the pixels classified into a given class are in reality according to the validation data. Read along the columns, each column tells you what the validation data known to be a given class were classified as. For example:

Table 7: Confusion matrix with values entered.

		Validation data				
	Class	Forest	Water	Grassland	Bare soil	Total
Classification	Forest	56	0	4	2	62
	Water	1	67	1	0	69
	Grassland	5	0	34	7	46
	Bare soil	2	0	9	42	53
	Total	64	67	48	51	230

Reading along the rows, the table above tells you that 56 pixels classified as ‘forest’ were also considered ‘forest’ in the validation data, that 0 pixels classified as ‘forest’ were considered ‘water’ in the validation data, 4 pixels classified as ‘forest’ were considered ‘grassland’ in the validation data, and 2 pixels classified as ‘forest’ were considered ‘bare soil’ in the validation data, for a total of 62 pixels classified as forest. And so on.

User, producer, and overall accuracy

Using the information in the confusion matrix, we can find answers to reasonable questions concerning the accuracy of the land cover map produced with the classification. There are three kinds of questions typically asked and answered with the confusion matrix.

The **user accuracy** answers a question of the following type: ‘If I have your map, and I go to a pixel that your map shows as class ‘x’, how likely am I to actually find class ‘x’ there?’ Using the example of ‘grassland’ from the table above, we can see that a total of 46 pixels classified as ‘grassland’ were checked against validation data. Of those 46 pixels, 34 were considered to be ‘grassland’ in the validation data. In other words, 34 pixels, out of the 46 pixels classified as ‘grassland’ are actually ‘grassland’. 34 out of 46 is 74%, so the user accuracy of the classification, for the ‘grassland’ class, is 74%. User accuracies vary between classes, as some classes are easier to distinguish from the rest than other classes. Water features tend to be easy to map because they are dark and blueish and not many features found on land look like them. In the example above, the user accuracy for the ‘water’ class is 67 out of 69, or 97%.

The **producer accuracy** answers a question of the following type: ‘If an area is actually class ‘x’, how likely is it to also have been mapped as such?’ Again using the example of ‘grassland’, we see that a total of 48 validation data points were considered to be ‘grassland’, and 34 of those were also classified as such. 34 out of 48 is 71%, so the producer accuracy for the ‘grassland’ class is 71%.

While the user and producer accuracies focus on individual classes, the **overall accuracy** answers the following question: 'What proportion of the map is correctly classified?', which can often be interpreted simply as 'how accurate is the map?'. Looking at the values in the diagonal of the confusion matrix in the above example, we see that 56 pixels were considered 'forest' in the validation data and had also been classified as 'forest', and we see similar numbers of 67 for 'water', 34 for 'grassland', and 42 for 'bare soil'. These sum up to 56+67+34+42=199, out of a total 230 pixels in the validation data set. 199 out of 230 is 87%, so based on the validation data we estimate that 87% of the map is correctly classified.

The overall accuracy needs to be reported with care, as the following example will illustrate. Imagine that the image you used for the classification covered a coastal zone, and the sub-orbital track of the satellite had been a bit off-shore, so 80% of the image was covered by 'water'. The remaining 20% of the image was covered by 'bare soil' or 'vegetation'. If you reflected this uneven distribution in the creation of your validation data, 80% of your validation data would be over water, and since water is relatively easy to distinguish from the other surface types, your confusion matrix might look something like this:

Table 8: Another example of a confusion matrix. Note the very different number of validation points used for each class.

		Validation data			
	Class	Water	Vegetation	Bare soil	Total
Classification	Water	82	1	0	83
	Vegetation	0	12	2	14
	Bare soil	0	2	9	11
	Total	82	15	11	108

While the user and producer accuracies for 'vegetation' and 'bare soil' are not impressive in this scenario, as expected 'water' has been classified almost perfectly. The dominance of 'water' pixels influences the calculation of the overall accuracy, which ends up as 82+12+9=103 out of 108, an overall accuracy of 95%. If the purpose of the map is to find out where the coastline is, or something else that only truly requires separating water from land, this might be acceptable as an estimate of how good the map is. But if you have made the map for a local government agency tasked with monitoring coastal vegetation, the overall accuracy of 95% may falsely provide the idea that the map should be used with confidence for that purpose, which largely will require separating 'vegetation' from 'bare soil'.

In general, as long as you report a) how you produced the map, b) how you produced the validation data, and c) the entire confusion matrix along with any additional accuracy measure derived from it, an intelligent reader will be able to judge whether the map is appropriate for a given purpose, or not.

Accuracy assessment for classifications when you are only trying to map one thing

A special case of accuracy assessment presents itself when you are making a map of one type of object, like houses, swimming pools, and so on. While this is still rare in remote sensing, it is becoming increasingly necessary with object-based image analysis, which is an effective means of mapping specific object types. We'll use swimming pools as an example. Imagine that you have created an object-based image analysis workflow that takes a high-resolution satellite image and attempts to detect all swimming pools in the area covered by the image. The product of that workflow is a set of polygons that outlines all swimming pools identified in the image. Similarly, your validation data consist of a set of polygons that outline all swimming pools manually identified in the image, for a small part of the image you are using for validation. So you now have two sets of polygons to compare, one being your 'estimate', the other being 'reality'. Your confusion matrix can be set up to look like this (explanation follows below):

Table 9: Confusion matrix used to calculate the F1 score when accuracy is assessed for detection of only one kind of objects.

		Validation data		
	Class	Presence	Absence	Total
Classification	Presence	TP	FP	Precision = TP / (TP + FP)
	Absence	FN		
	Total	Recall = TP / (TP + FN)		

In this table, 'presence' indicates the presence of a swimming pool (in either data set) and absence indicates the absence of a swimming pool (also in either data set). TP is the number of True Positives – swimming pools that exist in the validation data, and that were correctly identified in your map as being swimming pools. FP is the number of False Presences – object identified in your map as being swimming pools, but which are in reality something else. FN is the number of False Negatives – swimming pools that exist in reality, but which your map failed to detect. Note that in this table, there are no True Negatives (objects that are in reality not swimming pools, and were also not identified in the image as swimming pools). This has been omitted because, in the case of an image analysis that aims to find only one thing, no other objects are identified in the image, nor in the validation data.

The goal of a good image analysis is, of course, to have a large number of True Presences, and a small number of False Presences and a small number of False Negatives. To quantify how well the image analysis succeeded in this, the value typically calculated is called the F1 score, which is calculated as: F1 = (2*Precision*Recall) / (Precision+Recall). The F1 score has the nice property of having values that range from 0 (worst) to 1 (best), which makes it easy to interpret.

Accuracy assessment for continuous variables

When dealing with continuous variables, comparing 'estimates' and 'reality' is no longer a case of checking whether they are identical or not, because when measured with enough detail they never are. For example, you may have mapped a pixel as having a surface temperature of 31.546 °C while your corresponding field observation says that it is in reality 31.543 °C. Despite how the two values are not identical, you would probably not want that to simply be considered 'no match'. Instead, what we need to do is to provide users of the map with an idea of what the **typical difference** is between the mapped estimate and reality.

Creating validation data

As when assessing accuracy of classification, you need a set of validation data that are considered to represent reality. These almost universally come from field measurements, and it is important to remember that, as when assessing accuracy of classifications, the validation data should be comparable to the measures derived from your image. Especially the issue of spatial resolution can be problematic here, because it is difficult to make accurate measurements over large areas with most field equipment. Consider the case of surface temperature, which is typically measured with a handheld infrared thermometer (Figure 62).

An infrared thermometer (like the ear thermometers used to check if you have a fever or not) measures radiation coming from a small circular area of the Earth's surface, wherever the thermometer is pointed at. Satellites essentially measure the same radiation and estimate temperature in the same way as a handheld infrared thermometer, except that they integrate the radiation measured over a larger area (100 x 100 meters, in the case of TIRS on Landsat 8) while the handheld version measures a very small area (e.g. 0.5 x 0.5 meters). Because surface temperature varies with soil moisture, vegetation cover, and surface material, among other things, taking a measurement in the field that is representative of the average surface temperature in a 100 x 100 m area is an almost insurmountable challenge in heterogeneous environments. One way around this is to create all your validation data in areas that are as homogeneous as possible, another is to take multiple measurements and use average values at each site to produce a single validation data point. Whatever your strategy, matching the validation data to the spatial resolution of the image is essential to produce a meaningful accuracy assessment.

62: Handheld infrared thermometer of the kind used to measure surface temperature for a small area. Isometric Medical Digital Non-Contact Infrared Thermometer Temperature Measurement Device blue white by Jernej Furman, Flickr, CC BY 2.0.

Once you have a good set of validation data, the table that is the basis for accuracy assessment is one that compares numerical values, as in the below example of woody biomass per ha in a forest:

Table 10: Example of table used to compare mapped estimates of a continuous variable with validation data from field observations:

Mapped biomass, tons per ha (estimate)	Observed biomass, tons per ha (reality)
10.1	9.2
5.7	4.8
3.2	4.0
6.7	6.6
7.8	7.1
9.3	9.1
…	…

One of the first things we can quantify with this kind of data is whether the mapped estimates are typically underestimates or overestimates, when compared against the validation data. If this is the case, the workflow (algorithm) that produced the estimates is considered biased, and revising it may be in order. The bias can be calculated as the mean error, i.e. the mean value of the estimate minus the validation. We can calculate that by updating Table 10 to get to Table 11:

Table 11: Bias is calculated as the mean error of estimates.

Mapped biomass, tons per ha (estimate)	Observed biomass, tons per ha (reality)	Error (estimate– reality)
10.1	9.2	0.9
5.7	4.8	0.9
3.2	4.0	-0.8
6.7	6.6	0.1
7.8	7.1	0.7
9.3	9.1	0.2
…	…	
Mean Error:		0.33

In Table 11, we can see that the estimates have a mean error of 0.33. In other words, on average the biomass estimates in our map overestimate reality by 0.33 tons per ha. This may or may not warrant refining the algorithm used to produce the estimates, but at least it is something worth reporting to potential users of the biomass map that contains the overestimates.

Another crucial thing to quantify is how wrong the estimates typically are. This is important because you can imagine having an unbiased map that severely overestimates biomass in one part and severely underestimates it elsewhere. Such a map, unbiased, would still be wrong everywhere! The simplest way to report this is using the Mean Absolute Error, as shown in Table 12:

Table 12: The Mean Absolute Error (MAE) is typically reported to show how 'accurate' the map is.

Mapped biomass, tons per ha (estimate)	Observed biomass, tons per ha (reality)	Error (estimate– reality)	Absolute Error
10.1	9.2	0.9	0.9
5.7	4.8	0.9	0.9
3.2	4.0	-0.8	0.8
6.7	6.6	0.1	0.1
7.8	7.1	0.7	0.7
9.3	9.1	0.2	0.2
…	…		
Mean Absolute Error:			0.6

The MAE tells you what difference you should expect between the estimate and reality, but not whether this difference takes the form of an overestimate or an underestimate. As most algorithms used to map continuous variables are very effective at minimizing bias, the MAE is often used to compare different approaches to mapping a continuous variable, considering the approach that leads to the lowest MAE as the 'best' one.

An alternative measure of 'accuracy' for continuous variables is the Root Mean Squared Error (RMSE), which is calculated as in Table 13:

Table 13: Calculation of the Root Mean Squared Error.

Mapped biomass, tons per ha (estimate)	Observed biomass, tons per ha (reality)	Error (estimate– reality)	Squared Error
10.1	9.2	0.9	0.81
5.7	4.8	0.9	0.81
3.2	4.0	-0.8	0.64
6.7	6.6	0.1	0.01
7.8	7.1	0.7	0.49
9.3	9.1	0.2	0.04
…	…		
Mean Squared Error:			0.47
Root Mean Squared Error:			0.68

The RMSE is larger than the MAE (or equal to it, in extremely rare cases) as it weighs large errors more heavily in its calculation. Algorithms that avoid large errors thus benefit when compared using the RMSE, while algorithms that often have very low errors but occasionally have very large errors do not.

Finally, another value that is typically reported when mapping continuous variable is the coefficient of determination, or R^2, which indicates how much of the variance in the validation data values ('reality') is predictable from the estimates.

While it is important to understand why and how to report on the accuracy of any map you produce, in practical terms it is easier than what it might seem in the above sections. Most software used for image processing has built-in functions for accuracy assessment, and often even provides functionality to help you create the validation data. If the software you use doesn't let you calculate a specific measure (e.g. the MAE, or the F1 score), it will probably help you export the necessary data to a file that is easily read in Excel, R, or whatever other software you are familiar with and can use to do the rest of the calculations manually.

This page titled 1.7: Accuracy assessment is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.8: Change detection

Much environmental change occurs at temporal and spatial scales that make it challenging and potentially costly to study. Obvious examples include some of the manifestations of climate change: the increasing temperatures and melting ice in the Arctic and the subtle vegetation changes occurring throughout Canada, to name a few. Monitoring these things exclusively with field measurements would be incredibly costly, or heavily biased towards populated areas and, when historic measurements are not available, impossible. Remote sensing data come to the rescue here, because aerial photography and satellite imagery is routinely stored and catalogued and they can thus function as data from both the past and the present – as long as we can extract the necessary information from them. Just imagine answering these questions with confidence without having access to remote sensing data:

- Is the annual area burned by forest fires increasing or decreasing in Canada?
- Which part of Canada is warming the fastest?
- Where is there still multi-year ice in the Arctic, and how fast is that area shrinking?

Before we start to use remote sensing data to detect environmental change, it is helpful to consider what we mean by that – what is *environmental change*? The environment changes all the time – it is snowing outside as I write this and it didn't snow two days ago – that is a kind of environmental change, but not the kind most people want to detect and map with remote sensing data. Nevertheless, it is exactly that kind of change that is visible in images of the Earth from space, so whether we like it or not we detect a lot of irrelevant change with remote sensing data. To better think about what change we want to detect and what change we don't, one useful distinction to make is between short-term/instantaneous change and long-term/gradual change. Sometimes we are interested in finding the short-term change, while at other times, when more concerned with longer gradual trends, we consider it noise.

Instantaneous change

Remote sensing can be effective at detecting instantaneous change, i.e. change that happens between one image acquisition and the next. This change may not necessarily be instantaneous in the strict sense of the word, but for practical remote sensing purposes a change that happens in the period between two image acquisitions can be considered 'instantaneous'. Detecting such change is often done quite easily with visual interpretation of two images, comparing them side by side to easily identify areas of change. An example of this kind of comparison is shown in Figure 63, which shows changes in an urban area in Taiwan between February (left) and December (right) 2002. The vegetation (red areas) has clearly changed in some places, and so has a few of the other areas. Many image processing algorithms have also been developed for this purpose, to automate the process of finding the areas that have changed; we will look at those later in this chapter.

63: Example of side-by-side comparison of two images showing Zanzibar Town in 1995 (left) and 2020 (right). Urban expansion that has occurred in the period between the two images were acquired can be identified visually or algorithmically. By Anders Knudby, CC BY 4.0.

Gradual change

Visual or other image-image comparison can be effective in highlighting obvious abrupt changes such as those caused by urban development, landslides, forest fires and so on. However, more subtle change, such as annual or decadal variation in soil moisture or vegetation health, or change in the depth of the winter snow pack, often requires a different approach to change detection, one that relies on quantitative analysis of changes observed over many images. This is especially the case when individual image-to-image changes may obscure the more gradual long-term trend because of weather-related or seasonal patterns. For example, to study how climate change is affecting vegetation patterns in the Canadian Arctic, a comparison of two images provides, at best, a snapshot of the years those images are from, and at worst it simply tells us that there is more vegetation in 2002 (August) than there was in 1994 (February), which really has nothing to do with a change between the two years but rather is a function of the time of year each image was acquired. Either way, such image-to-image comparisons do not tell us very much about the long-term trend. Instead, a measure of the variable of interest (e.g. vegetation density or health) must be quantified at regular intervals over a long time period to detect such trends. An example is provided in Figure 64, in which the extent of Arctic Sea Ice has been quantified on a near-daily basis for the period 1979 – 2018.

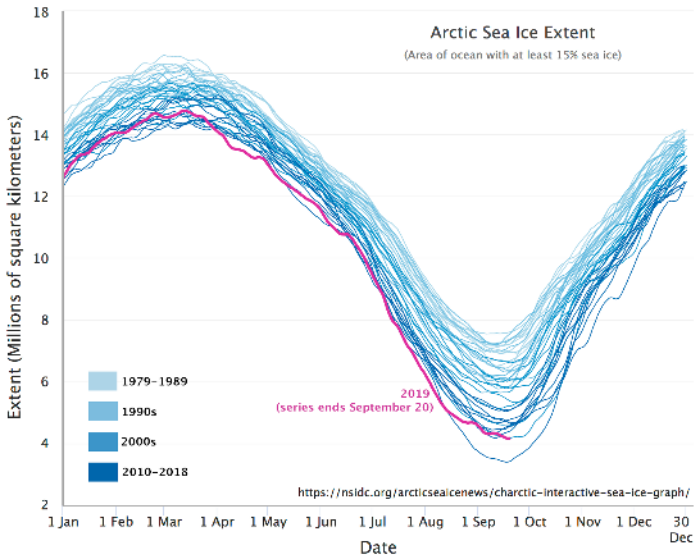

64: Near-daily estimates of Arctic Sea Ice extent. Arctic Sea Ice Extent by M. Scott, National Snow & Ice Data Center (NSIDC), NSIDC Use and Copyright.

The specific methodology used to detect long-term gradual change depends entirely on the kind of change in question (e.g. vegetation vs. sea ice change), so it is difficult to provide details about how to go about detecting change in this case. However, one core concern for any such kind of 'trend detection' is that the estimate (e.g. of sea ice extent) should be relatively unbiased throughout the period of observation. In other words, sea ice extent should neither be under- or over-estimated in the 1979-1989 period, and should also neither be under- or over-estimated in 2010s, or at any point in between. This is important to avoid 'detecting' trends that are caused by bias in data acquisition or processing, but do not exist in reality.

Separating noise from actual instantaneous change

When comparing one image to another with the intent to detect changes between them, the principal challenge is to detect real environmental change while not detecting change that hasn't actually happened. As is always the case in binary classification (which this is an example of), there are four possible combinations of reality (change or no change) and estimate (change detected vs. no change detected, Table 14):

Table 14: A binary classification (one that involves two categorical options) can be described with a table like this one. For any image processing system, the challenge in binary classification is to optimize the number/rate of true positives and true negatives.

	Change detected	**No change detected**
Change	True positive (TP)	False negative (FN)
No change	False positive (FP)	True negative (TN)

The goal, then, is to optimize the rate at which true positives and true negatives are found. Most change detection algorithms operate on a pixel-by-pixel basis, so this means correctly detecting pixels that have actually changed without incorrectly 'detecting' change in pixels that have not changed. To do this, we need a way to separate three different situations, each of which can occur in a given pixel:

- No change: the pixel literally looks exactly the same in each image.
- Noise: The pixel looks different in the two images, but the difference is small enough that it is probably caused by factors unrelated to *real environmental change*. These can include differences in atmospheric conditions between the two images, random noise in the images, imperfect georeferencing, and other issues.
- Actual change: *Real environmental change* has happened in the pixel, and it shows as a substantial difference between how the pixel looks in the first image and in the second image.

The first situation, no change, happens very rarely because there is noise inherent to the image creation process, and this noise is unlikely to be identical between two different images. The real challenge, then, is to separate situations 2) and 3). There are two principles involved in doing this.

Noise reduction

First of all, it is important to remove as many sources of noise as possible before comparing the data in the two images. Some strategies often employed to this end include:

- Use 'anniversary dates'. Choose two images that were captured on the same date (or almost) in different years. This is a good way to eliminate large differences between each image related to seasonal changes in soil moisture, vegetation state, snow cover, and other environmental factors that change with the seasons.
- Use images captured by the same sensor. Given that no sensor is perfectly calibrated, using images from two different sensors could potentially introduce a difference that is sensor-based rather than environment-based. For example, if radiance in the 500-600 nm spectrum is slightly overestimated by one sensor and slightly underestimated by another sensor, comparing images between the two may 'detect change' where none exists.
- Use images captured at similar atmospheric states. This is difficult because we often don't have accurate information on aerosol loads, water vapour, wind speed etc., but definitely avoid comparing images with obvious differences in haze, visibility, and other visible atmospheric factors.
- Compare images based on their surface reflectance rather than TOA radiance or TOA reflectance. This is because surface reflectance is a physical attribute of the surface, and is at least in principle independent of atmospheric state and illumination.
- For tilting sensors, if possible compare images taken at more or less the same geometry.

Identifying a threshold

Once as many sources of noise as possible have been eliminated, and the images have been converted to surface reflectance for comparison, there will still be some noise that causes surface reflectance values in the two images to differ slightly between the images, even for pixels that have not experienced the kind of change we want to detect. To separate situations 2) and 3) above, it is therefore important to consider what 'real' environmental change means and define a threshold below which any 'change' observed in the image comparison is considered 'noise' rather than 'real change'. For example, imagine that you are studying changes in vegetation, and you have two images of the same forest. The real change that happened between these two images is that one leaf fell off one of the trees in the forest. This is an actual observable change in the environment, but in any real sense it does not warrant the label 'deforestation'! One prerequisite for separating situations 2) and 3) above then is to define how much change is

required for you to consider the area 'changed'. Unless you have field observations (and you typically don't because change is detected backwards in time and you cannot go and get field data from the past), this requires defining the threshold based on the images themselves, something that is typically an interactive and subjective process.

Methods for detecting instantaneous change

Some of the simplest approaches to detecting the *magnitude* of change are based on 'image math'.

Band difference

For example, the magnitude of change can be defined as the difference between surface reflectance values in band 1 in the two images, plus the difference in band 2, etc.:

$$\text{Change} = (B1_{I2}-B1_{I1}) + (B2_{I2}-B2_{I1})\ldots$$

Band ratio

Or ratios can be used instead of differences:

$$\text{Change} = (B1_{I2}/B1_{I1}) + (B2_{I2}/B2_{I1})\ldots$$

Euclidian distance

A more commonly used method is to calculate the 'Euclidian Distance', using each band as a dimension and each image as a point. Imagine that you plot the surface reflectances of each image as points in a coordinate system. The Euclidian distance between them would then equal:

$$\text{Change} = (B1_{I2}-B1_{I1})^2 + (B2_{I2}-B2_{I1})^2\ldots$$

For each of these approaches, to separate noise from real environmental change, a threshold value will need to be defined. Pixels that have experienced no change or change smaller than the threshold value will then be considered effectively 'unchanged', while those that have experienced more change will be considered to contain 'real change'. Also note that for each of these approaches, the number of bands involved in the calculation have been limited to two in the equations above, but can be extended to include any number of bands present in the images.

Methods for detecting and attributing instantaneous change

All three above approaches can be effective at *detecting* change, but they also all suffer from the drawback that they tell us very little about the *kind of* change that has happened in a given pixel. While detection of change is a great first step, it is not easy to base management decisions on, so figuring out a bit more about the kind of change that has been detected, called change *attribution*, is useful.

Change vectors

Change vectors can help with this. Change vector analysis is an extension of calculating the Euclidian distance, in which the direction of change is also calculated. An example with only two dimensions is provided in Figure 65. In the 'before' image, the surface reflectance value in band 1 is 0.1, and the NDVI value is 0.5. In the 'after' image, these values have changed to 0.3 and -0.2 respectively. This illustrates two things: 1) you do not need to use actual band values as inputs to change detection analyses, and 2) the direction of change – increase in band 1 reflectance, decrease in NDVI, can be found.

In change vector analysis, the magnitude of change is calculated as the Euclidian distance, as per the equation above. The direction of change can be calculated in degrees (e.g. compass direction) or, to ease interpretability, in major directions: A) up and right, B) down and right, C) down and left, D) up and left. If the 'bands' used can be interpreted in a meaningful way, attributing the change can be easy. Changes that involve a decrease in a vegetation index involve a loss of vegetation, while those that lead to an increase in, say, surface temperature involve, well, an increase in surface temperature (these two often go together because vegetation helps keep a surface cool). While calculating Euclidian distance is easy when more than two dimensions are used in change vector analysis – you just extend the equation to contain more terms – the direction of change becomes more difficult to define and categories are typically developed for a specific application. As with the methods simply used to detect change, a threshold can be applied to the magnitude of the change vector, below which 'no change' is detected.

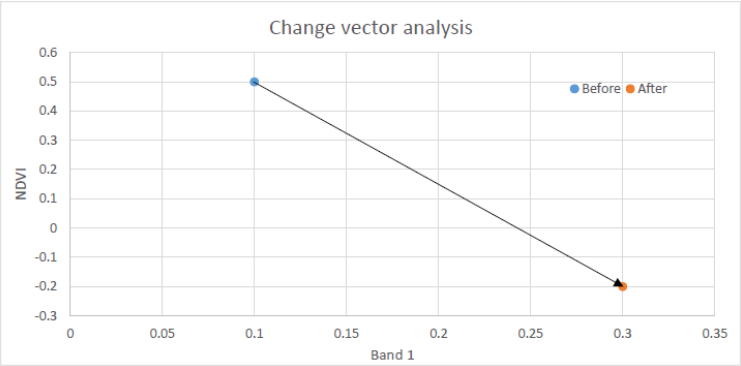

65: Example of change vector analysis with only two bands. NDVI is the Normalized Difference Vegetation Index, described in more detail in the next chapter. Anders Knudby, CC BY 4.0.

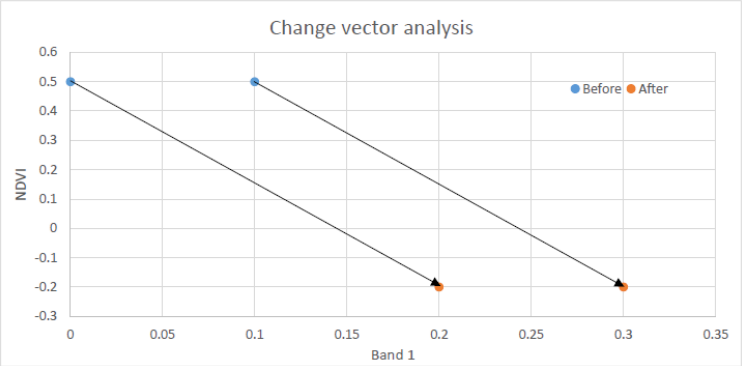

66: One drawback of change vector analysis is that two identical change vectors can represent different kinds of change. Anders Knudby, CC BY 4.0.

One drawback of change vector analysis is that two different changes can have the exact same change vector. For an example, look at the two arrows in Figure 66. While the two vectors are identical – they have the same direction and magnitude – they are likely to represent different kinds of change because they start and end in different places. It is thus difficult, in practice often impossible, to use change vector analysis to specifically figure out what the surface was before and after the change.

Post-classification change detection

Probably the simplest way to both detect and attribute change in an area is to conduct a land cover classification on the 'before' image, conduct another land cover classification on the 'after' image, and then find pixels that got classified differently in the two images. While this is appealingly straight-forward, and can work occasionally, it is subject to a significant drawback: It is very difficult to produce accurate results in this way. The reason is that no classification is perfect, and when comparing two imperfect classifications the errors combine. I therefore strongly suggest avoiding this approach.

Change classification

If you really want to know what the changed area was before and after the change, an alternative to post-classification change detection is to combine all bands from the two images into a single image. For example, if you have two images, each with six bands, you could 'stack' them to get a single 12-band image. This allows you to run a single classification on the 12-band image. With good field data to calibrate this classification each class can be defined according to the combination of before-and-after land covers. The classification can be either supervised or un-supervised, pixel-based or segment-based. This approach allows the two vectors in Figure 66 to end up in two different classes, where maybe one is a change from coniferous forest to marsh while the other is a change from deciduous forest to bog.

Change classification also allows you to remove class combinations that are impossible or highly unlikely, such as a change from 'ocean' to 'coniferous forest', or from 'industry' to 'wetland'. Deciding which before-after combinations of classes are likely and

which are unlikely in a given area requires some expertise, and complex change detection systems are often built using a combination of change vector analysis, change classification, and such expert input.

This page titled 1.8: Change detection is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.9: Vegetation and Fire

The first chapters of these notes have largely covered generic subjects in passive optical remote sensing – issues of general importance and image processing techniques that are not specific to one field of application or another. In this chapter we will look at some techniques developed and widely used specifically for remote sensing of two things: vegetation, and fire. Similar chapters could be written on many other specific applications of remote sensing – vegetation and fire are used here because the use of satellite imagery to map and monitor them is well-developed and used operationally at both national and global scales.

Vegetation

Given how much of Earth is covered by some form of green vegetation, it is not surprising that much research effort has gone into using remote sensing to map and assess various aspects of vegetation. This is important for basic industries that support human livelihoods across the globe, such as forestry and agriculture, as well as for almost any environmental assessment of terrestrial areas. In fact, so much research has been done that entire books have been written just on remote sensing of vegetation.

In the simplest terms, the ability to detect vegetation in passive optical remote sensing is based on its spectral signature (Figure 67), which is quite different from other surface types.

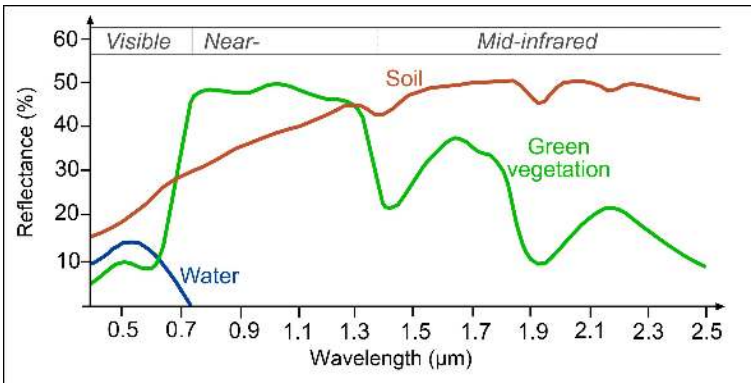

Figure 27. Reflectance of water, soil and vegetation at different wavelengths by SEOS, Science Education through Earth Observation for High Schools (SEOS), CC BY-NC-SA 2.0.

Visible wavelength region

Healthy green vegetation absorbs most of the visible light incident on it, using it to grow through the process of photosynthesis. Photosynthesis is a relatively complex biochemical process that begins with the absorption of a photon by a pigment molecule, which sets off an electronic transition (the movement of an electron from one energy level to a higher one) that starts a chain-reaction leading to production of plant material from the basic chemical components of CO_2 and water. A form of chlorophyll is the dominant pigment in almost all vegetation, and chlorophyll is very good at absorbing visible light, which is why vegetation has low reflectances across the visible wavelengths. Chlorophyll is a little better at absorbing blue and red wavelengths than it is at absorbing green wavelengths, so relatively more of the green radiation is reflected, giving vegetation its green appearance.

Infrared wavelength region

Incoming photons at wavelengths in the near-infrared region individually contain less energy (recall that the energy in a photon is proportional to its frequency, and hence inversely proportional to its wavelength). These photons are therefore unable to cause an electronic transition so pigments used for photosynthesis cannot use them, and in general do not absorb them. Other parts of plants, especially liquid water found in the leaves, do absorb these photons, which rotate and stretch chemical bonds in the water or in plant cells, effectively heating up the absorbing material. In general, plants have no need for heating up, and except for wavelengths around 1.4 μm and 1.9 μm where water is a strong absorber, plants reflect much of the incoming near-infrared radiation.

The combination of low reflectance in the visible and high reflectance in the near-infrared wavelengths is the most characteristic signature of vegetation used in remote sensing, and is used for quick-and-dirty mapping of vegetation on Earth. Consider that most land areas contain a mix of vegetated and unvegetated areas. Sometimes this ‘mix’ is very one-sided, as in the Amazon, the Sahara desert, or the Canadian Arctic, and sometimes it is truly more mixed, as on the African savanna (Figure 68).

Vegetation indices

Using the reflectance in the visible wavelengths alone does not tell us how much vegetation is in an area, because water also has low reflectance in this wavelength region. And using the reflectance in the near-infrared region alone also does not tell us much about vegetation, because other surface types, such as the bright sand and bright snow depicted in Figure 68 also have higher near-infrared reflectance. However, the ratio of the visible and near-infrared reflectances is a useful indicator for the amount of vegetation present in an area because no other surface type has *both* as high near-infrared *and* as low visible reflectance as vegetation. This observation led initially to the development of what is called the Simple Ratio (SR):

$$SR = NIR / Red$$

where NIR is the surface reflectance in the near-infrared wavelength region (typically around 700-1000 nm), and RED is the same for the red wavelength region (600-700 nm).

68: Four areas with rather different vegetation characteristics. Top left: The African savanna contain a mix of trees and grass cover. Trees On The African Savanna by Lynn Greyling, PublicDomainPictures.net, CC0 1.0. Top right: The Sahara desert is in most places void of vegetation. Sahara Desert by Azer Koçulu, Wikimedia Commons, CC0 1.0. Bottom left: The Amazon rainforest is covered in dense vegetation. Amazon Rainforest by CIFOR (Neil Palmer/CIAT), Flickr, CC BY-NC-ND 2.0. Bottom right: The Canadian Arctic has large areas with no vegetation. Patterned ground in Canadian tundra by Raymond M. Coveney, Wikimedia Commons, CC BY-SA 3.0.

While surface reflectance values should in principle always be used to calculate this and any other vegetation index, in practice these calculations are often done instead based on TOA reflectance or even TOA radiance (or even DN values!). The Simple Ratio has one drawback that means it is rarely used – its values approach infinity when red reflectance approaches zero, which makes differences between large values difficult to interpret. A simple improvement called the Normalized Difference Vegetation Index (NDVI), based on the same two measurements, is therefore more commonly used:

$$NDVI = (NIR-Red) / (NIR+Red)$$

NDVI values range from a minimum of -1 to a maximum of +1. Typical values for bare soil are around 0, for water around -0.2, and for vegetation in the range 0.1-0.7 depending on vegetation health and density. Because NDVI relies only on two measurements, in the near-infrared and red wavelength regions, it can be calculated with data from the earliest multispectral sensors in orbit, including the Landsat series (since 1972) and the AVHRR series (since 1978). This allows easy visualizations of vegetation cover, on small or large scales, such as the one shown in Figure 69.

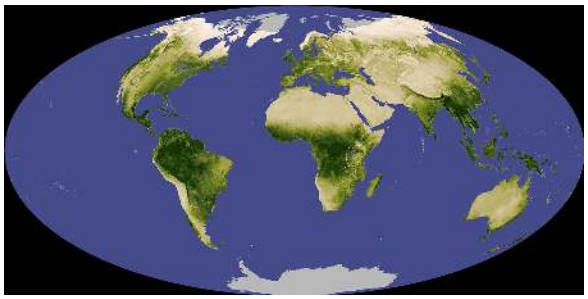

69: NDVI composite of Earth, showing relative differences in vegetation cover. Globalndvi tmo 200711 lrg by Reto Stockli and Jesse Allen (NASA), Wikimedia Commons, public domain.

Vegetation indices can be used for more than producing pretty pictures of global vegetation distribution. As a result of global warming, which is especially pronounced in the Arctic region, the northern parts of Canada are undergoing rapid changes in vegetation distribution and growth patterns. Monitoring such changes over this large and very sparsely populated region is based on long-term NDVI trends in the summer months, which among other things demonstrates that region undergoing faster warming also experience more rapid growth in vegetation density. Another important use of vegetation indices is in crop yield forecasting: by tracking NDVI through an agricultural growing season, the development of the crops can be quantified and their eventual yields predicted statistically. Probably the most important use of such crop prediction models is to forecast droughts, crop failures and famines in parts of the world where people depend on local agriculture for their food supply. This is done systematically for countries that are prone to famine by the Famine Early Warning Systems Network (FEWS), a US-based organization, and it is also done by many national governments and regional/global organizations. Vegetation indices thus literally have the potential to save lives – if the global community responds to the predictions generated by them.

Many refinements have been introduced since the development of NDVI, to create vegetation indices that perform better under a range of conditions. These include the Soil-Adjusted Vegetation Index, the Enhanced Vegetation Index used with MODIS data, and the Global Vegetation Index used with MERIS data. These newer vegetation indices were specifically designed to be more robust to varying environmental conditions, such as variations in soil brightness and atmospheric conditions. Despite such refinements, all vegetation indices suffer from a fundamental drawback: while their interpretation is quite clear in relative terms (higher values indicate more green vegetation), their relationship to real measureable vegetation characteristics used by people outside the field of remote sensing is not clear.

A vegetation index not intended to portray vegetation density, but rather to portray relative differences in vegetation water content, has been developed specifically to take advantage of the water absorption feature around 1.4 μm shown in Figure 67. It is called the Normalized Difference Water Index (NDWI), and its generic equation is:

$$NDWI = (NIR-SWIR) / (NIR+SWIR)$$

where SWIR denotes a measurement in the short-wave infrared region (1.4 μm). Dry vegetation contains less water, and therefore absorbs less of the SWIR radiation, reducing the value of NDWI.

For any vegetation index, the specific bands that are used as the ‘RED’, ‘NIR’ and ‘SWIR’ bands depends on the sensor in question, and not all sensors have bands appropriate for calculating one or another index.

Mapping real vegetation attributes

Vegetation attributes that are more used by non-remote sensing people include among other things above-ground biomass, absorbed photosynthetically active radiation (APAR), and the Leaf Area Index (LAI). Above-ground biomass can be measured (by cutting and weighing it) in kg/ha, and is of fundamental interest in monitoring vegetation trends. APAR is the amount of radiation that is absorbed by vegetation (per areal unit), and is strongly linked to vegetation growth and crop yield. LAI is the (one-sided) area of leaves divided by the area of ground, so it is a unitless measure of leaf density. Because (water and) gas exchange between vegetation and the atmosphere happens through the leaves, LAI can be used as a proxy for this exchange, which is important for vegetation’s function as a carbon sink, among other things. All of these vegetation attributes are thus more directly linked to the function of vegetation in the environment, and certainly more directly interpretable, than is a vegetation index. But they are also much more difficult to map with remote sensing data.

Empirical approach

A simple approach to mapping LAI is illustrated in Figure 70, in which NDVI values have been regressed against field-measured LAI. As is obvious from Figure 70, what we measure in remote sensing data is related, but not very closely, and certainly not linearly, to actual LAI measured on the ground. Nevertheless the regression lines can be used to produce a first approximation of LAI from these remote sensing data. A similar approach can be taken to map other vegetation attributes.

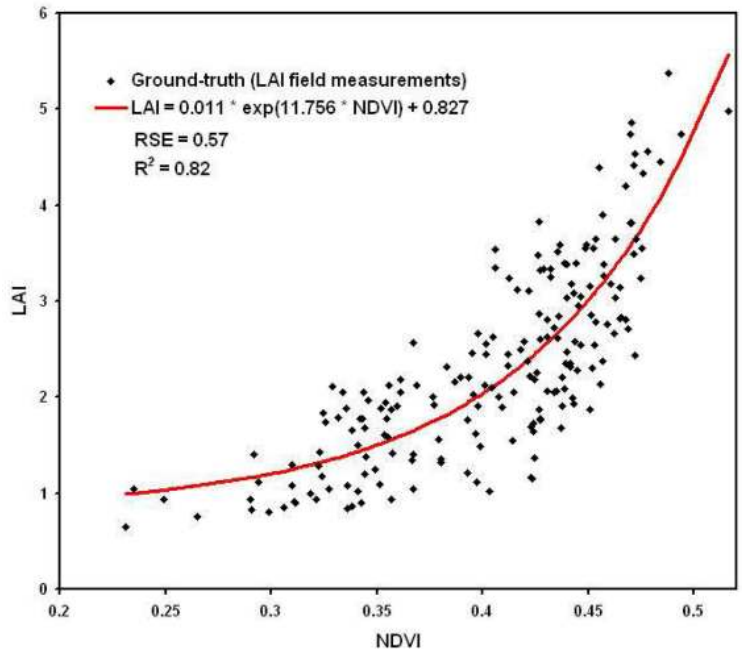

Figure 8) by LeLong et al., CC BY 3.0.

Radiative transfer approach

More sophisticated approaches to mapping vegetation attributes rely on radiative transfer models. Radiative transfer models build up a simple 3D world consisting of a) the ground, b) some vegetation on the ground, c) the atmosphere, and d) the Sun. Each of these elements have specific attributes that define how they interact with radiation: The Sun emits a certain amount of radiation at each modeled wavelength, the atmosphere absorbs and scatters this radiation to some degree, and the vegetation and the ground absorb and reflect the radiation. Once specified, the model can then be 'run' to determine what the TOA radiance (or other measureable radiation parameter) would be for a given combination on element definition. These elements can then be changed and the model 'run' again to see how the change affects the TOA radiance. For example, different densities of vegetation can be included, representing different values of LAI, each producing a unique combination of TOA radiances in different wavelengths. To produce a robust picture of how LAI influences TOA radiance, the other elements must also be varied, so the influence of soil colour and atmospheric constituents also get captured in the process. Even more importantly, non-LAI variation in vegetation must also be considered. Individual leaves may have different spectral signatures depending on their pigment composition, leaves of different shapes and orientation interact with the Sun's radiation differently, and the structure of the canopy also exerts its own influence (e.g. are the leaves distributed randomly, evenly, or in a clumped manner?). Varying all of these parameters, each time running the radiative transfer model to determine what their specific combination would look like at the top of the atmosphere, allows you to build up a look-up table with a) environmental conditions including the parameter of interest (e.g. LAI) and b) the resulting TOA (spectral) radiance. When mapping LAI with a satellite image, the observed TOA radiance of each pixel can then be matched to the entry in the table it matches most closely to determine the probably value of LAI in that pixel.

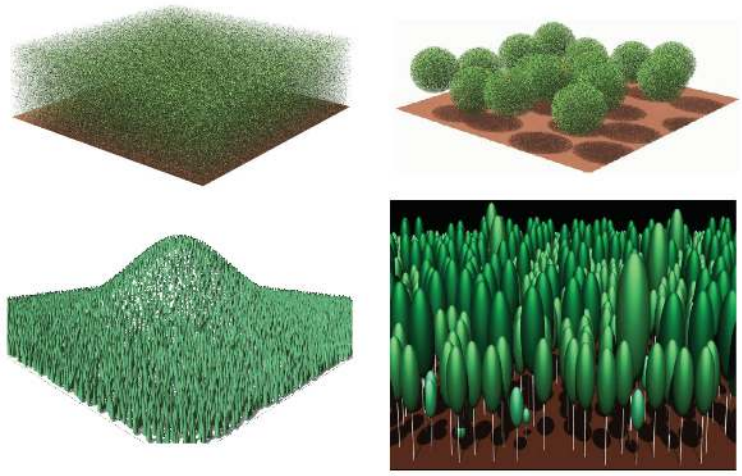

Figure 7) by Huaguo Huang, CC BY 4.0.

Fire

Another sub-field of remote sensing concerns mapping of fire. Active fire is relatively easily detected because it is hot, and heat can be detected with thermal sensors, and because it creates smoke, and smoke is visible in satellite imagery. Detecting active fire with satellite data is so well-developed that a near-instantaneous 'active fire product' is made each time a MODIS or VIIRS sensor passes over an area.

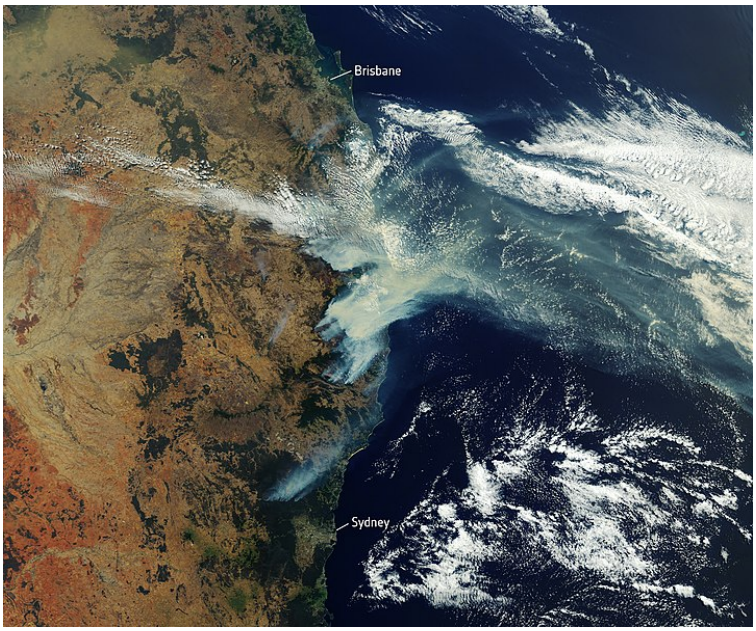

72: Example of active fire detection from space during the Australian wildfires of 2019. Satellite image of bushfire smoke over Eastern Australia by the European Space Agency, Flickr, CC BY-SA 2.0.

Once a fire is over, it may be important to assess how much area was burned, and how intensely it was burned. This has important implications for the recovery of the vegetation and associated biota in the area, and it can also help us answer important questions related to climate change, like how the frequency, duration, size, and intensity of wildfires are changing as the climate warms. While there are different ways to assess these parameters, the most commonly used approach involves quantifying something called the Normalized Burn Ratio (NBR). The NBR is calculated as:

$$NBR = (NIR-SWIR) / (NIR+SWIR)$$

which, you might say, is the same as the definition of NDWI! And it is close, but the term SWIR covers a fairly wide range of wavelengths, and while for NDWI the original definition of 'SWIR' referred to wavelengths around 1.24 µm, for NBR 'SWIR' was specifically designed for band 7 on Landsat 4 and 5, which is centered around 2.2 µm. Nevertheless, NBR is really just a vegetation index, and while it is not the only index one can use for mapping fire impacts, and not even necessarily the best one, it is the most commonly used. To use it for quantification of fire severity, one needs to compare NBR values in two images, one from before and one from after the fire, which leads to the calculation of δNBR:

$$\delta NBR = NBR_{before} - NBR_{after}$$

To minimize the effects of other environmental factors on this calculation, it is ideal to use 'anniversary dates' to calculate δNBR, although this is not necessary if two images are available from a small period covering just before and just after the fire (Figure 73).

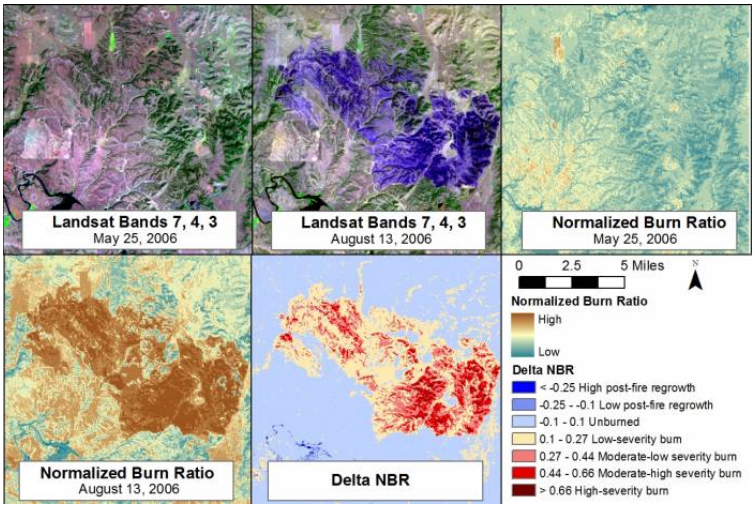

73: Example of a pre-fire image (top left), a post-fire image (top middle), pre-fire NBR (top right), post-fire NBR (bottom left), and δNBR (bottom middle). While the impact of the fire is evident in the original imagery, the use of δNBR helps highlight the difference between the two images. Example of normalized burn ratio from a fire at Camp Gurnsey, Wyoming in 2006 by Jason Karl, The Landscape Toolbox, CC0 1.0.

Because NBR is calculated as a normalized difference, its values range between -1 and +1, so δNBR ranges, in theory, from -2 to +2, with higher values indicating greater fire severity. As with other vegetation indices, extreme values are rare though, so even values of 0.2-0.3 indicate relatively severe fire impacts.

74: An example of both vegetation and fire, both of which can be mapped well with remote sensing. Forest Fire, by YIvers, Pixabay, Pixabay License.

This page titled 1.9: Vegetation and Fire is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

1.10: Photogrammetry and Structure-from-Motion

All of the applications of remote sensing data covered in the previous chapters have been based either on individual images (e.g. classification), or on time series of images (e.g. change detection). However, when two (or more) images are acquired of the same area, in rapid succession and from different angles, it is possible to extract geometric information about the mapped area, just as our eyes and brain extract 3D information about the world around us when we look at it with our two eyes. This is the sub-field of remote sensing called photogrammetry.

Traditional (single-image and stereo) photogrammetry

Photogrammetry involves using information about imaging geometry (distance between sensor and object, sensor position and orientation, and internal geometry of the sensor) to derive information about positions, shapes and sizes of objects seen in the imagery. Photogrammetry has been used since the earliest use of aerial photography, and the basic idea in this field is to enable measurement of distances in image. As such, photogrammetry allows you to find out how long a road segment is, how far it is from point A to point B, and how large your local skating rink actually is (assuming you have a well-calibrated image of it, that is). You can think of it in terms of taking an aerial photograph and manipulating it in a way that you can overlay it in a GIS, or on Google Earth. Once geographic coordinates are associated with the image, measurement of distances becomes possible.

Single vertical image, flat terrain

The simplest possible scenario in photogrammetry is that you have an image produced by a sensor pointing directly toward Earth, and that the image contains perfectly level terrain. If we assume that the image was taken by a sensor located either on an aircraft flying at very high altitude, or on a satellite, and that the image covers only a very small area, we can make the simplifying assumption that every part of the image has been imaged exactly from zenith, and that as a result the distance between the sensor and the Earth surface is the same throughout the image (Figure 75).

In this case, the scale of the image is equal to the ratio of the camera's focal length (f, or an equivalent measure for sensor types other than pinhole cameras) and the distance between the sensor and the ground (H):

$$\text{Scale} = f / H$$

If you know the scale of your image, you can then measure distances between real-world objects by measuring their distance in the image and dividing by the scale (note that the scale is a fraction that has 1 in the numerator, so dividing by the scale is the same as multiplying by the denominator).

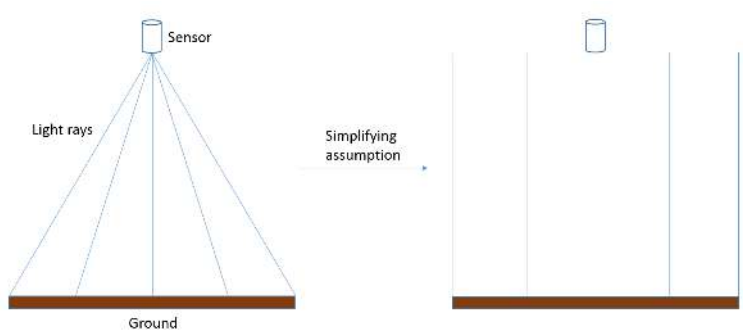

75: In reality, an image is created from light reaching the sensor at different angles from different parts of the Earth's surface. For present purposes, we can assume that these angles are in fact not different, and that as a result the distance between the sensor and the Earth's surface is the same throughout an image. In reality, this assumption is not correct, and scale/spatial resolution actually changes from the image centre toward the edges of the image. Anders Knudby, CC BY 4.0.

This definition of scale is only really relevant for analog aerial photography, where the image product is printed on paper and the relationship between the length of things in the image and the length of those same things in the real world is fixed. For digital aerial photography, or when satellite imagery is used in photogrammetry, the 'scale' of the image depends on the zoom level, which the user can vary, so the more relevant measure of spatial detail is the spatial resolution we have already looked at in previous chapters. The spatial resolution is defined as the on-the-ground length of one side of a pixel, and this does not change depending on

the zoom level (what changes when you zoom in and out is simply the graininess with which the image appears on your screen). We can, however, define a 'native' spatial resolution in a manner similar to what is used for analog imagery, based on the focal length (f), the size of the individual detecting elements on the camera's sensor (d), and the distance between the sensor and the ground (H):

$$\text{Resolution} = (f*d) / H$$

Single vertical image, hilly terrain

The situation gets more complicated when an image is taken over terrain that cannot reasonably be assumed to be perfectly flat, or when the distance between the sensor and the ground is sufficiently small, or the area covered by the image sufficiently large, that the actual distance between the sensor and the ground changes substantially between different parts of the image:

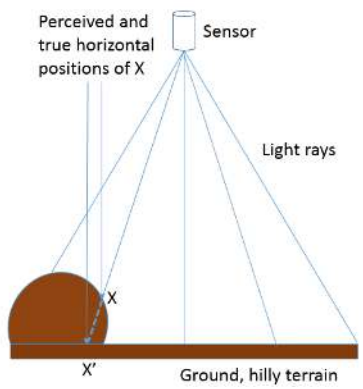

76: When an image is taken closer to the Earth's surface, or over more hilly terrain, the distance between the sensor and the Earth's surface cannot be assumed to be constant throughout the image. If not taken into account, this leads to both vertical and horizontal displacement of the imaged objects. Anders Knudby, CC BY 4.0.

Hilly terrain causes a number of problems in photogrammetry, most importantly the vertical and horizontal displacement of imaged objects. As illustrated in Figure 76, if point X is assumed to lie on the flat surface (and not on the hill where it is actually found), based on the imaging geometry its horizontal position will be assumed to be at X', because that is where the 'light ray' intersects the flat ground. In reality, of course, its horizontal position is not at X'. If the surface terrain of the imaged area is known, for example through availability of a high-quality digital elevation model, it can be taken into account when georeferencing the image and the correct intersection between the 'light ray' and the ground (on the hill) can be identified. However, if no digital elevation model exists, the effect of terrain – known as relief displacement – cannot be explicitly accounted for, even if it is obvious in the resulting imagery.

Stereo imagery

Until a decade ago, probably the most advanced use of photogrammetry was to use a stereo-pair (two images covering the same area from different points of view) to directly derive a digital elevation model from the imagery, which also enabled things like measuring the heights of things like buildings, trees, cars, and so on. The basic principle of stereoscopy, as the technique is called, is a replication of what you do with your two eyes all the time – by viewing an object or a landscape from two different angles it is possible to reconstruct a three-dimensional model of it. While there is a sophisticated art to doing this with high accuracy (an art we cannot explore as part of this course), we will review the basic idea here.

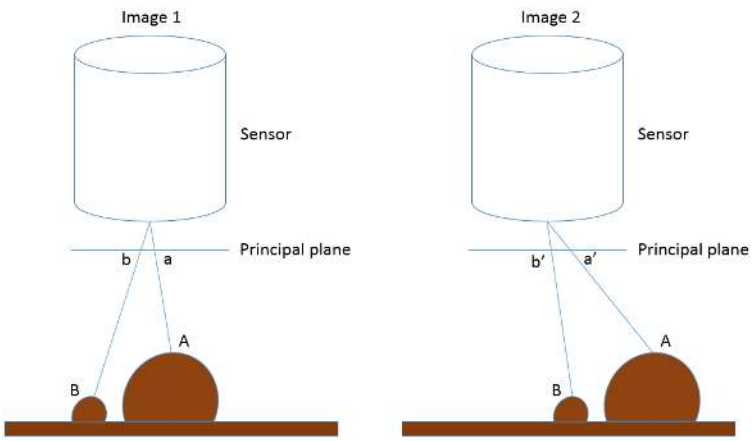

77: Illustration of how the apparent horizontal positions of point A and B change between the two images that form a stereo view of the area. Sensor size has been greatly exaggerated to show the location of each point on the principal plane. Anders Knudby, CC BY 4.0.

In Figure 77, Image 1 and Image 2 represent a stereo view of the landscape that contains some flat terrain and two hills, with points A and B. On the left side of the figure (Image 1), a and b denote the locations each point would be recorded at on the principal plane in that image, which corresponds to where the points would be in the resulting image. On the right side of the figure, the camera has moved to the left (following the airplane's flight path or the satellite's movement in its orbit), and again b' and a' denote the locations each point would be recorded at on the principal plane in that image. The important thing to note is that a is different from a' because the sensor has moved between the two images, and similarly b is different from b'. More specifically, as illustrated in Figure 78, the distance between a and a' is greater than the distance between b and b', because A is closer to the sensor than B is. These distances, known as parallax, are what convey information about the distance between the sensor and each point, and thus about the altitude at which each point is located. Actually deriving the height of A relative to the flat ground level or relative sea level or a vertical datum, or even better deriving a digital elevation model for the entire area, requires careful relative positioning of the two images (typically done by locating lots of surface features that can be recognized in both images) as well as knowing something about the position and orientation of the sensor when the two images were acquired (typically done with GPS-based information), and finally knowing something about the internal geometry of the sensor (typically this information is provided by the manufacturer of the sensor).

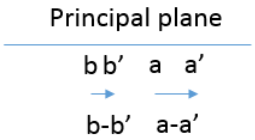

Figure 77. Anders Knudby, CC BY 4.0.

The math applied in stereoscopy is fairly complicated, but in practice even a novice can produce useful digital elevation models from stereo imagery with commercial software. While some specialized photogrammetry software exists, the large commercial remote sensing software packages typically include a 'photogrammetry module' that can help you produce digital elevation models even if you are not an expert photogrammetrist. In the above, we have assumed nadir-looking imagery (i.e. that the sensor was pointed straight toward Earth for each image) because the concepts get more complicated when this is not the case, but mathematical solutions exist for oblique imagery as well. In fact, one of the useful attributes of several modern satellites carrying imaging sensors is their ability to tilt, which allows two images to be captured by the same sensor during the same orbit. As the satellite approaches its target, it can be pointed forward to capture an image, and later in the orbit it can be pointed backward to image the same area again from a different angle.

Structure-from-Motion (multi-image photogrammetry)

As it happens occasionally in the field of remote sensing, techniques developed initially for other purposes (for example medical imaging, machine learning, and computer vision) show potential for processing of remote sensing imagery, and occasionally they

end up revolutionizing part of the discipline. Such is the case with the technique known as Structure-from-Motion (SfM, yes the 'f' is meant to be lower-case), the development of which was influenced by photogrammetry but was largely based in the computer vision/artificial intelligence community. The basic idea in SfM can be explained with another human vision parallel – if stereoscopy mimics your eye-brain system's ability to use two views of an area acquired from different angles to derive 3D information about it, SfM mimics your ability to close one eye and run around an area, deriving 3D information about it as you move (don't try this in the classroom, just imagine how it might actually work!). In other words, SfM is able to reconstruct a 3D model of an area based on multiple overlapping images taken with the same camera from different angles (or with different cameras actually, but that gets more complicated and is rarely relevant to remote sensing applications).

There are a couple of reasons why SfM has quickly revolutionized the field of photogrammetry, to the point where traditional stereoscopy may quickly become a thing of the past. Some of the important differences between the two are listed below:

- Traditional (two-image) stereoscopy requires that the internal geometry of the camera is characterized well. While there are established processes for figuring out what this geometry is, they require a moderate level of expertise and require the user to actually go through and 'calibrate' the camera. SfM does this camera calibration using the very same images that are used to map the area of interest.
- In a similar vein, traditional stereoscopy requires that the position and orientation (pointing) of the camera is known with great precision for each image. SfM benefits from approximate information on some or all of those parameters, but does not strictly need them and is able to derive them very accurately from the images themselves. Because many modern cameras (like the ones in your phone) have internal GPS receivers that provide approximate location information, this allows consumer-grade cameras to be used to derive remarkably accurate digital elevation models.
- For both techniques, if the camera positions are not known, the 3D model reconstructed from the images can be georeferenced using imaged points of known location. In this way the two techniques are similar.
- The rapid evolution of drones has provided average people with a platform to move a consumer-grade camera around in ways that were previously limited to people with access to manned aircraft. Images taken by drone-based consumer-grade cameras are exactly the kind of input necessary in an SfM workflow, while they do not work well in the traditional stereoscopy workflow.

As mentioned above, SfM uses multiple overlapping photos to produce a 3D model of the area imaged in the photos. This is done in a series of steps as outlined below and shown in Figure 79:

Step 1) Identify the same features in multiple images. To be able to identify the same features in photos taken from different angles, from different distances, with the cameras not always being oriented in the same way, possibly even under different illumination conditions, in extreme cases even with different cameras, SfM starts by identifying features using the SIFT algorithm. These features are typically those that stand out from their surroundings, such as corners or edges, or objects that are darker or brighter than their surroundings. While many such features are easy to identify for humans, and others are not, the SIFT algorithm is specifically designed to allow the same features to be identified despite the challenges circumstances outlined above. This allows it to work especially well in the SfM context.

Step 2) Use these features to derive interior and exterior camera information as well as 3D positions of the features. Once multiple features have been identified in multiple overlapping images, a set of equations are solved to minimize the overall relative positional error of the points in all images. For example, consider the following three (or more) images, in which the same multiple points have been identified (shown as coloured dots):

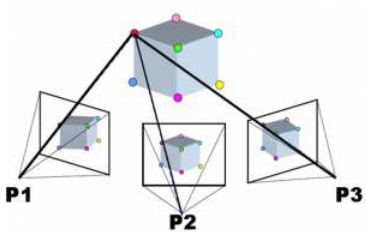

79: Example of bundle block adjustment in SfM. Sfm1 by Maiteng, Wikimedia Commons, CC BY-SA 4.0.

Knowing that the red point in P1 is the same as the red point in P2 and P3 and any other cameras involved, and similarly for the other points, the relative position of both the points and the camera locations and orientations can be derived (as well as the optical parameters involved in the image-forming process of the camera). (Note that the term 'camera' as used here is identical to 'image').

The initial result of this process is a *relative* location of all points and cameras – in other words the algorithm defined the location of each feature and camera in an arbitrary three-dimensional coordinate system. Two pieces of additional information can be used to transform the coordinates from this arbitrary system to real-world coordinates like latitude and longitude.

If the geolocation of each camera is known prior to this process, for example if the camera produced geotagged images, the locations of the cameras in the arbitrary coordinate system can be compared to their real locations in the real world, and a coordinate transformation can be defined and also applied to the locations of the identified features. Alternatively, if the locations of some of the identified features is known (e.g. from GPS measurements in the field), a similar coordinate transformation can be used. For remote sensing purposes, one of these two sources of real-world geolocation information is necessary for the resulting product to be useful.

Step 3) Based on the now geo-located features, identify more features in multiple images, and reconstruct their 3D position. Imagine that there is a little grey dot close to the red point in Figure 79, say ¼ of the way toward the green point, in that direction (there isn't, which is why you need to imagine it!). This little grey dot was too faint to be used as a SIFT feature, and so was not used to derive the initial location of features and orientations of cameras. However, now that we know where to find the red dot in each image, and we know the relative position of each of the other dots, we can estimate where this grey dot should be found in each of the images. Based on this estimate, we might actually be able to find the grey dot in each image, and thus estimate its location in space. Once identified and located in space, we can also assign a colour to the point, based on its colour in each of the images it was found in. Doing this for millions of features found in the complete set of images is called point cloud densification, and is an optional but almost universally adopted step in SfM. In areas with sufficient difference between the colours of neighbouring pixels in the relevant images, this step can often lead to the derivation of the location of each individual pixel in a small area of an image. The result of this step therefore is an incredibly dense 'point cloud', in which millions of point are located in the real world. An example is shown below:

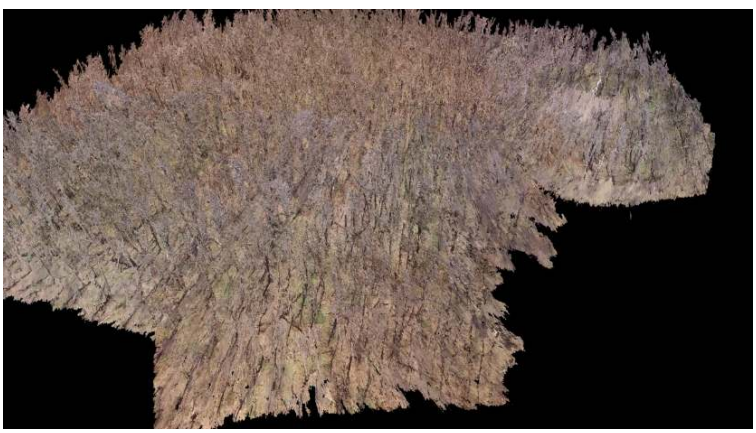

80: Example of dense point cloud, created for a part of Gatineau Park using drone imagery. Anders Knudby, CC BY 4.0.

Step 4) Optionally, use all these 3D-positioned features to produce an orthomosaic, a digital elevation model, or other useful product. The dense point cloud can be used for a number of different purposes – some of the more commonly used is to create an orthomosaic and a digital elevation model. An **orthomosaic** is a type of raster data, in which the area is divided up into cells (pixels) and the colour in each cell is what it would look like if that cell was viewed directly from above. An orthomosaic is thus similar to a map except that it contains information on surface colour, and represents what the area would look like in an image taken from infinitely far away (but with very good spatial resolution). Orthomosaics are often used on their own because they give people a quick visual impression of an area, and they also form the basis for map-making through digitization, classification, and other approaches. In addition, because the elevation is known for each point in the point cloud, the elevation of each cell can be identified in order to create a **digital elevation model**. Different algorithms can be used to do this, for example identifying the lowest point in each cell to create a 'bare-earth' digital terrain model, or by identifying the highest point in each cell to create a

digital surface model. Even more advanced products can be made, for example subtracting the bare-earth model from the surface model to derive estimates of things like tree and building height.

Change detection with photogrammetry

While the creation of digital elevation models or other products using stereoscopy or Structure-from-Motion approaches can be useful on its own, an increasingly common application of these techniques is for surface change detection. If, for example, a digital elevation model is created for a glacier one year, and a new one created two years later, changes in the shape of the glacier's surface can be mapped with great precision. Landslides, thaw slumps, vegetation change, snow depth changes throughout a season – many changes in the physical environment can be detected and quantified by measuring changes in the horizontal or vertical position or the Earth's surface. Given the relatively recent development of SfM technology, and the still-developing capabilities and cost of camera-equipped drones, this is a rapidly developing field likely to be useful in surprising and unsuspected parts of the remote sensing field. Figure 81 illustrates this principle comparing the outline and of the glacier and changes in surface elevation at two different times.

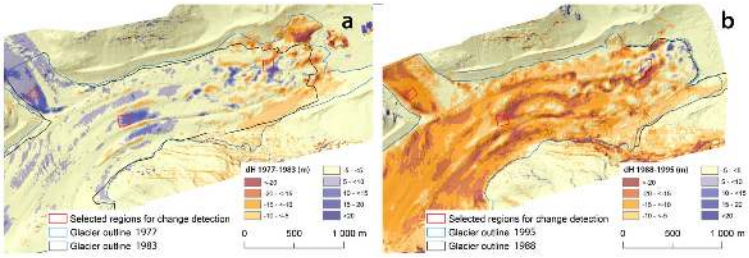

Figure 7) by Mölg and Bolch, CC BY 4.0.

This page titled 1.10: Photogrammetry and Structure-from-Motion is shared under a CC BY 4.0 license and was authored, remixed, and/or curated by Anders Knudby (eCampus Ontario) .

CHAPTER OVERVIEW

Back Matter

Index
Glossary
Detailed Licensing

Index

D

dire

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: Remote Sensing (Knudby)

Webpages: 31

All licenses found:

- Undeclared: 61.3% (19 pages)
- CC BY 4.0: 38.7% (12 pages)

By Page

- Remote Sensing (Knudby) - *CC BY 4.0*
  - Front Matter - *Undeclared*
    - TitlePage - *Undeclared*
    - InfoPage - *Undeclared*
    - Table of Contents - *Undeclared*
    - Licensing - *Undeclared*
    - Welcome to Remote Sensing - *CC BY 4.0*
  - Chapters - *Undeclared*
    - Front Matter - *Undeclared*
      - TitlePage - *Undeclared*
      - InfoPage - *Undeclared*
      - Table of Contents - *Undeclared*
      - Licensing - *Undeclared*
    - 1.1: Introduction - *CC BY 4.0*
    - 1.2: Radiometric measurements - *CC BY 4.0*
    - 1.3: Calculations of TOA radiance and TOA reflectance - *CC BY 4.0*
    - 1.4: EMR interactions with the Earth's atmosphere and surface - *CC BY 4.0*
    - 1.5: Visualization and manipulation of satellite images - *CC BY 4.0*
    - 1.6: Classification - *CC BY 4.0*
    - 1.7: Accuracy assessment - *CC BY 4.0*
    - 1.8: Change detection - *CC BY 4.0*
    - 1.9: Vegetation and Fire - *CC BY 4.0*
    - 1.10: Photogrammetry and Structure-from-Motion - *CC BY 4.0*
    - Back Matter - *Undeclared*
      - Index - *Undeclared*
      - Glossary - *Undeclared*
      - Detailed Licensing - *Undeclared*
  - Back Matter - *Undeclared*
    - Index - *Undeclared*
    - Glossary - *Undeclared*
    - Detailed Licensing - *Undeclared*

Index

D

dire

Glossary

Sample Word 1 | Sample Definition 1

Detailed Licensing

Overview

Title: Remote Sensing (Knudby)

Webpages: 31

All licenses found:

- Undeclared: 61.3% (19 pages)
- CC BY 4.0: 38.7% (12 pages)

By Page

- Remote Sensing (Knudby) - *CC BY 4.0*
  - Front Matter - *Undeclared*
    - TitlePage - *Undeclared*
    - InfoPage - *Undeclared*
    - Table of Contents - *Undeclared*
    - Licensing - *Undeclared*
    - Welcome to Remote Sensing - *CC BY 4.0*
  - Chapters - *Undeclared*
    - Front Matter - *Undeclared*
      - TitlePage - *Undeclared*
      - InfoPage - *Undeclared*
      - Table of Contents - *Undeclared*
      - Licensing - *Undeclared*
    - 1.1: Introduction - *CC BY 4.0*
    - 1.2: Radiometric measurements - *CC BY 4.0*
    - 1.3: Calculations of TOA radiance and TOA reflectance - *CC BY 4.0*
    - 1.4: EMR interactions with the Earth's atmosphere and surface - *CC BY 4.0*
    - 1.5: Visualization and manipulation of satellite images - *CC BY 4.0*
    - 1.6: Classification - *CC BY 4.0*
    - 1.7: Accuracy assessment - *CC BY 4.0*
    - 1.8: Change detection - *CC BY 4.0*
    - 1.9: Vegetation and Fire - *CC BY 4.0*
    - 1.10: Photogrammetry and Structure-from-Motion - *CC BY 4.0*
    - Back Matter - *Undeclared*
      - Index - *Undeclared*
      - Glossary - *Undeclared*
      - Detailed Licensing - *Undeclared*
  - Back Matter - *Undeclared*
    - Index - *Undeclared*
    - Glossary - *Undeclared*
    - Detailed Licensing - *Undeclared*