



Tech deep dive:
drone image quality



PERCEPTO

Table of Content

A note from the authors	3
Drone image quality explained - background	4
Introduction	5
Essentials for any drone inspection solution	5
When does image quality come into play?	6
Unpacking image quality	7
Debunking myths, demystifying buzzwords	7
GSD - Ground Sampling Distance	8
MTF - Modulation Transfer Function	8
Photo resolution in pixel count	9
Perceived spatial resolution for photo quality	9
Breaking down superior image quality	9
Superior detail separation	9
White balance	10
Best lens for the job	10
Dedicated landing camera	11
Optimized software parameters	11
Getting the right drone for your business	12
Contact us to set up a demo	13

Authors

We are Percepto's image quality experts, and subject matter experts in visual data analysis for inspection, mapping and modeling purposes. Our goal is to pool together our knowledge, experience and insights on drone image quality for commercial uses, so you can make the best decision of what camera best fits your needs. In so, we have explained some key industry terms, parameters and test results. We hope this helps in making the right drone decisions at your organization.

Happy learning,

- Udi Zohar -
VP Product
- Eilon Hoter -
Product Manager
- Evgeni Iosefovski -
Payload Technical Lead



Drone image quality explained - background

Intelligent automation, advances in machine vision, and innovations in edge computing are bringing about fresh opportunities to reduce costs, improve safety, and get inspections done quickly using autonomous drones. Today, drones are being used for visual inspection in fast-growing areas such as utilities, oil and gas, solar farms, mining, bridges, towers, and more.

The benefits of adopting drones for visual inspection are clear. The ability to easily maneuver, hover, and hold allows for safer, more efficient and more thorough inspections. Advances in drone technology are also delivering more accurate results, especially when it comes to large scale sites or environments that require the precise monitoring of critical assets and infrastructure. In short, drones are making visual inspections faster, less costly, safer, and more precise.

The big advantage of autonomous drone inspections over standard ones that require piloting, is the savings in manpower and expertise. When all you have to do is take your drone out and push a button, it's easier to do more inspections, more often. This translates directly into detecting problems earlier on—before they become expensive to fix or start causing additional damage. What's more, smart autonomous drones will know how to optimize the distance and angle from the target, to maximize the amount of detail captured by the cameras.

When it comes to image quality for drones, everyone wants to know how many megapixels a camera packs or what's the GSD (Ground Sampling Distance). But image quality isn't just a matter of these parameters. As you may have already noticed, there are loads of parameters including, MTF (Modulation Transfer Function), absolute accuracy, relative accuracy, FPS, focal points, image correction, FOV, sharpness, distortion, software, AI algorithms, and more. How do you make sense of these different characteristics and understand which system is right for your needs? After all, security surveillance that requires tracking is different from asset inspection, which requires highly detailed imaging. **This white paper is geared to help you: understand what image quality is right for your use cases, become familiar with the different terms for image quality, and make an educated decision when it comes to the best fit for your operations.**

Without data quality, your inspection system can't reliably deliver insights for your business. There is a clear line between the quality of data collected and the caliber of decision-making. Good quality data imaging provides better insight and enables businesses to take action early on. We like to think of it as 'gold in means gold out'.

Percepto, a leader in visual inspection automation, has been customizing innovative camera hardware and software to deliver cutting-edge imagery, inspection, mapping, and detection. We provide top notch machine vision solutions for clients in the areas of utilities, energy, oil and gas, solar, mining, and more.

This white paper is geared to help you: understand what image quality is right for your use cases, become familiar with the different terms for image quality, and make an educated decision when it comes to the best fit for your operations.

Introduction

With so many options, buzzwords, and parameters out there, the goal of this white paper is to demystify decision-making for autonomous visual inspection solutions. The following sections will take you through some essential background information, introduce the ins and outs of accuracy and quality for visual inspection, and then zero in on image quality and why it's important.

*The information we present refers to RGB cameras and examples using data from the new Percepto Air Max 24MP camera.

Essentials for any drone inspection solution

Let's take a look at two essential components of drone inspection: data capture and data analysis.

- **Data capture** – Any data collected should be topnotch data that contains sharp, high quality images, taken at the correct angle and the optimal time of day, including the most advantageous light, best weather conditions, least wind, and so forth.
- **Data analysis** – Once the data is collected, it should be effortless to quickly upload it to the cloud, as well as accurate, and easy to analyze, share, and compare to historical data.



When it comes to autonomous inspections, data capture is a pivotal factor. Once you have the best possible data, you need sophisticated on-board and offline services that can analyze the data for an efficient feedback loop.

Automating high quality data capture is tough. At Percepto, we understand the challenges involved in gathering those images, so you capture better quality data. We put this expertise into the logic of the drone's flight plan so it can plan the optimal mission for capturing your data, while putting in the least amount of effort. Even with the best possible lens and algorithms, you need the right parameter settings to get the best possible images. We're focused on delivering a gold standard when it comes to automated flow for the inspection.

Beneath the surface of visual inspection lies a tangled web of tradeoffs associated with different parameters. For example, what's the best way to balance the choice between field of view and spatial resolution (or the smallest objects that can be detected by the camera)? Camera sensors that have a wider FoV lens will provide lower spatial resolution. So, you'll see a larger area in each frame, but fewer details. By having these decisions automatically made for a mapping mission, we are able to optimize the drone area and the flight path. (Mapping is discussed in more detail later on in this white paper.)

We believe that autonomous inspection should match, and even surpass, the level that can be seen by the human eye. With the right level of detail, the ability to easily conduct more frequent inspections, and the computer's sensitivity to changes, we can offer a more accurate understanding of any situation.

Added frequency of inspections can be another game changer for many industries. The value add in cost-per-inspection is clear when the entire process is automated and you don't need a pilot. That's why we made sure to design a solution where you simply deploy the drone on your site and let it run.

When does image quality come into play?

In general, higher quality imaging means you can get accurate measurements and detect tiny objects or faults. Catching a minute fault early on allows you to prevent downtime, improve safety, and minimize the risk of environmental disasters. When a fault is still small, you can:

- **Fix the problem** while it is still relatively easy, less expensive, and faster to repair
- **Protect your employees** by catching a fault early on
- **Prevent a tiny failure** from growing into a larger one

Asset inspection

A lot of organizations rely on automated drones for asset inspection. This is where new solutions can go a long way towards increasing efficiency and reducing the growing costs of asset management. The unique characteristics that make certain asset inspection solutions superior are the drone's ability to collect the maximum amount of detail. It's important to note that having more detail doesn't necessarily mean an image that is more aesthetically pleasing to the human eye. After all, we aren't inspecting facilities to post images on social media.

Here our goal is to capture an image that is easy to understand for the 'machine eye' - not the human eye. It's all about capturing an image from which powerful AI algorithms can extract more insight—whether anomaly detection or minute changes in shapes, lines, or color.

For inspection, most clients are looking for a solution that can see more details, such as cracks in roads, pipes, or insulators showing wear and tear.


For inspection, most clients are looking for a solution that can **see more details**, such as cracks in roads, pipes, or insulators showing wear and tear. **If this is what you need, your solution should be able to detect a crack anywhere in size from 0.4 cm in width to something that would be visible to the human eye.**

Surveying and mapping

Using drones for surveying or mapping requires a solution that can offer real-time processing and high-accuracy. Applications such as mapping, 3D modeling, disaster investigation, or volumetric measurements all need solutions that can monitor a site, continuously collect data, and be able to consistently provide the same camera angle to make comparisons easier.

The answer to all of this lies in technologies that can enhance image quality and stitch together different pieces of data to create accurate maps and 3D models.



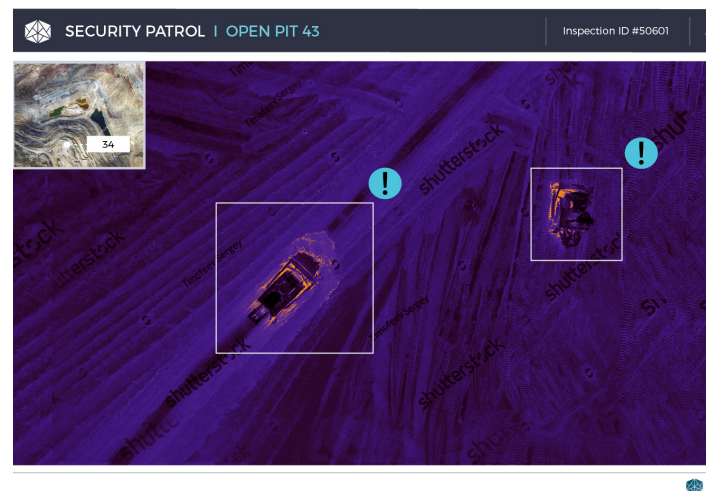


Thermal inspection

Combining drone inspection with thermal cameras offers major advantages when it comes to inspections and emergency response. Thermal imaging can easily detect overheating parts, poor pipe insulation, hot spots and more. The right kind of thermal imaging requires high resolution and algorithms that can compare frames to find tiny differences. In emergency response, it can be used to detect fire, even behind walls or inside structures. This allows response teams to see through smoke, investigate disasters from a safe distance, and search for people or animals when conditions are harsh.

Security and tracking

When it comes to security and tracking, real-time decisions need to be made on the spot, without waiting for post-processing. Our approach to security gives top priority to two aspects: the solution's ability to detect and its ability to track an intruder. The computer vision algorithms, the camera rate in frames per second (FPS), and the video quality must all provide the needed information for prompt and accurate object detection and tracking. Resolution will impact the algorithms' ability to detect an intruder. The video's FPS will also impact the ability to understand what is being seen in the video so the algorithms can track an item autonomously. By comparing frames to each other, the algorithm can measure an object's movement in relation to other features in the frames. If the FPS is too low, the distance an intruder moves between each frame will be larger and the algorithm will be more likely to fail in its ability to track the item. For security applications, you'll need a system that can provide at least 20 FPS.



Unpacking image quality

Debunking myths, demystifying buzzwords

There are a number of important metrics that describe drone image quality, but there is no single answer or parameter that sets the benchmark. At Percepto, we understand that good quality is not just a matter of camera features. It ultimately depends on what you need the image for. Although a camera's resolution is expressed in megapixels, the resolution of an imaging system has to include a spatial metric. And the more details the better. That's why we count spatial resolution (i.e. spatial resolution /GSD) and MTF among the most important metrics when calibrating our equipment for superior image quality. Let's take a closer look at these terms and some common misconceptions.

GSD - Ground Sampling Distance

Myth: Lower GSD means better inspection

Truth: Although GSD can precisely calculate the level of detail that can be detected, you can still have the same GSD for a very crisp image and a very blurry one.

Ground sampling distance (GSD), which represents the difference between the center of two pixels, is the most common statistic noted in discussions of resolution. For example, a 1 cm GSD means that every pixel reflects one centimeter down on the ground. In short, the same camera flying at a low altitude will have a smaller (i.e., better) GSD than it would at a high altitude.

When choosing a drone system, you'll want to take this tradeoff into account and get the best GSD for your needs. For example, most mapping is done at a distance of 60 to 80 meters from the ground surface. If mapping is your thing, your solution's GSD should be optimized for this range. When it comes to inspections, the GSD should be optimized to provide maximum detail at about 20 meters from the object.

GSD doesn't take into account diffractions, defocus, aberrations, distortion, lens curvature, motion blur, and other artifacts

The industry standard GSD for mapping and modeling is a value less than 0.5 cm at almost any altitude you are flying. Although GSD can precisely calculate the level of detail that can be detected, you can still have the same GSD for a very crisp image and a very blurry one. This occurs because **GSD doesn't take into account diffractions, defocus, aberrations, distortion, field curvature, motion blur, and other artifacts**. This gives us a compelling reason to look at MTF as an essential parameter.

MTF - Modulation Transfer Function

Myth: MTF is a precise measurement of contrast detail

Truth: There is no standard way to measure MTF. Each camera sensor has its own MTF chart as does each lens. The lens MTF should match or exceed that of the sensor.

Modulation transfer function (MTF) measures the ability to differentiate between pixels, and is an indication of image sharpness. An MTF can be presented as either a table of values or a chart, and it provides the contrast and the resolution of some part of an optical system, for a particular set of parameters, which includes the distance to an object, lens focusing distance, and aperture. This means, for example, that an MTF chart is relevant for a specific distance. If you want to look at an object from a different distance, a new MTF chart for the new distance is required.






Figure 1: MTF curve for high quality lens

In the case of a lens, MTF plots the contrast and resolution of the image created by the lens from its center to the edges of the image. A theoretical perfect lens, which can create a perfectly sharp image, would have an MTF of "1" from the center of the image up to its edges. Of course, this perfect lens doesn't exist due to the imperfection of the material, tolerances in lens making and more.

The high and flat curves plotted on the MTF graph below indicate superior performance. The higher the line on the Y axis, the more contrast we see in the image details. In the MTF chart for a Nikon lens below, a perfectly flat

curve across the X axis indicates that the pixel at the edge of the image is of identical quality to the pixel in the center.



Photo resolution in pixel count

 Myth:	More megapixels means better quality images
 Truth:	Too many megapixels can harm a camera's performance. The sensor size is more important when it comes to getting sharper details.

The photo resolution reflects the tradeoff that can be chosen between FOV and GSD. A higher resolution means that we can either increase FOV without compromising GSD, or increase GSD without compromising FOV—or even increase both.

A higher GSD does not necessarily mean higher spatial resolution. In fact, the higher the photo resolution, the harder it usually is to properly utilize each pixel. Thus, by itself, the photo resolution neither reflects the amount of details nor the quality of the image. That said, if high photo resolution is paired with a compatible optical system, it's possible to achieve both a high level of details and a large field of view. In the commercial drone industry, the standard photo resolution ranges from about 12MP—including the retired Percepto Sparrow, Skydio 2+ for Enterprise, and DJI mini 2—to 24MP for the Percepto Air Max, and can go up to about 48MP as it does for the DJI Mavic 2 Enterprise Advanced.

Perceived spatial resolution for photo quality

 Myth:	The more lines per centimeter, the better the system
 Truth:	Higher spatial resolution means a smaller area is photographed, and you'll need more flight time to cover the same area.

Spatial resolution, measured in line pairs per centimeter (lp/cm), describes the amount of distinguishable details captured on a physical surface. This is generally the first and the most common measure for the quality of a dataset. But more lines per centimeter also means you are getting a smaller area photographed, which translates directly into more flight time or more footage needed to cover the same area. To balance this tradeoff, we calibrated our solution to get the best possible detail for inspection at close range so that any cracks or anomalies are visible—while ensuring that the same configuration works well for mapping applications, which don't require as much detail.

Breaking down superior image quality

Although we already looked at four important and well-known parameters for image quality, there are other less-known factors that, together, impact the viability of different inspection solutions.

Superior detail separation

Inspection and anomaly detection requires the best image for differentiating between pixels. This is not about the best looking image, rather the one in which the smallest faults can be detected at an early stage, and from a distance. By offering a more robust algorithm with high detail separation, Percepto allows issues to be detected earlier on. Our sophisticated machine vision algorithms also make autonomous inspection feasible even for sites with vast differences in elevation.

By offering a more robust algorithm with high detail separation, Percepto allows issues to be detected earlier

Figure 2 below compares examples of detail separation for different drone inspection solutions. The




Figure 2: Differences in detail separation at 20m distance

images show a test pad that each drone videoed from a height of 20 m.

White balance

Correct coloring is vital for inspections and when detecting problems depends on the image being represented accurately. White balance is the process of adjusting the color temperature of your data captures so they match what our eyes would see. This makes white balance one of the more important post-processing features for fine detail inspections. By developing a sophisticated software solution for white balance, we make problem detection easier and lower the rate of false alarms.




Figure 3. White balance helps ensure colors appear as seen by the human eye

Best lens for the job

The optimal lens is a function of how you combine spatial resolution, MTF, FoV, and distortion to best meet your application's needs. For example, as we discussed earlier, choosing a different FoV will influence the GSD and distortion. After calibration, each lens has specific distances at which it performs best, based on its MTF, aperture, focal distance, and some other parameters. For any calibrated distance, the lens will have a specific depth of field. Images taken beyond the range established for the lens, will have lower image quality.

At Percepto, we put a lot of effort into getting the right lens and lens calibration, based on our understanding of the inspection process and the distances at which our targets are situated. Our decision to use a lens with 50 degrees FoV gives us 1) higher spatial resolution than a wider FoV 2) the ability to use a low distortion lens 3) and the choice of a quality manufacturer who uses a superior polishing technique to boost the MTF.

This same lens has a very high focus coefficient with a resolution multiplier limit that makes it a surefire choice for the distances at which inspection drones are flying. For example, with most customer scenarios, our drone won't need to go closer than 15 meters from an object we are inspecting and will not stay more than 75 meters away.

Combining the best lens and payload, even at a distance gives us an advantage over other solutions,

which need to get close up to provide the same level of detail, and ultimately need to take many more images to cover the same area. Being able to create a high quality model while flying further away means we save time and computer costs by processing less data to get the same 3D model.

Comparison of 24MP camera with 50 degree FOV vs. 80 degree FOV

50 degree FoV	80 degree FoV
<ul style="list-style-type: none"> Improved spatial resolution Easier to see minute details on the surface being observed Can keep more distance between the drone and electricity poles/client assets during the mission Low distortion images are more aesthetically pleasing and offer reduced lens shading Need to fly more missions to map the same area but maps will have a higher level of detail (assuming all other parameters are equal i.e., altitude, overlap, sensor) 	<ul style="list-style-type: none"> Reduced spatial resolution Harder to see the details of the surface being observed May require manually maneuvering the drone closer to assets being inspected Barrel distortion may occur at the edges of the image and higher level of lens shading Advantageous in mapping / modeling missions because you can cover more area with less distance flown

Dedicated landing camera

Because of the need to calibrate the lens for larger distances, Percepto’s drone solutions are challenged when it comes to vision landing. Our solution relies on our cutting-edge vision landing algorithm, which by nature involves close up distances. We overcame this barrier by introducing a separate camera for landing, giving us the freedom to use the best camera for inspection during the flight, without taking a hit on image quality when it comes to landing. This creative solution frees the primary payload from landing responsibilities, allowing it to focus on the main task of inspection, and offers a superior landing.

Our landing algorithm takes into account multiple factors including the resolution of the camera, the size of the frame, the input, and more. By using computer vision to land, the drone system remains autonomous. It eliminates the need for a specialized pilot and allows our drone to return safely to its starting point every time.

Optimized software parameters

The Percepto payload team was originally founded to push image quality to the forefront of the competitive drone manufacturing landscape. We are continuously working to improve the software correction capabilities via the onboard Tegra GPU. The team is dedicated to providing the best software custom-built for our payload and hardware—so they are working in concert to achieve the same target.



We built a best-in-class Debayer that takes the raw pixels and transforms them to compose an image. This process runs at the edge in real-time on the drone's Tegra GPU. With a powerful feedback loop between the camera and software, we are able to adjust all the image parameters to get the best possible image during flight. The drone continuously adjusts its position and direction to get the best image, while our proprietary three-payload gimbal keeps the cameras steady. What's more, our system automatically chooses the exposure and gain that will produce an image with the lowest amount of noise, lowest blur, and maximum amount of recognizable detail.

Lens:

- Focuses the light entering the lens onto the sensor

Sensor:

- Senses the photons coming from the lens and translated to a digital value passed along to the camera

Camera:

- Received digital value from the sensor and translates to raw pixels
- Unique Debayer processes the raw image on Tegra

Software (Percepto):

- Picture correction and post processing capabilities
- White balance
- Hue gradient and color enhancement algorithm
- Image alignment

Getting the right drone for your business

Although there is no one size fits all solution when it comes to drone inspection, you can still get the best in class image quality available in today's marketplace.

This document is geared to demystify decision-making for autonomous visual inspection solutions. So, if you're shopping around for a drone for your site's inspection, you should now know:

- More megapixels doesn't guarantee better image quality and neither does GSD.
- An inspection system should help you see enough details such that you can detect cracks, leaks, or damage anywhere in size from 0.4 cm in width to something that would be visible to the human eye.
- For inspection, GSD should be optimized to provide maximum detail at about 20 meters from the object. If mapping is your thing, your solution's GSD should be optimized for 60 to 80 meters. Being able to create a high quality model while flying further away means you can save time and computer costs by processing less data to get the same 3D model.
- The MTF of the lens should be on a par, or even higher, than the MTF of the camera—allowing the full benefit of the camera's capabilities. It's not an independent standard.
- The lens and lens calibration should be appropriate for your needs and optimized for the distances at which your targets are situated. For example, at Percepto we use a lens with 50 degrees FoV, as opposed to one with a wider FoV, so we get higher spatial resolution.
- Find out if you are getting a low distortion lens and whether it comes from a quality manufacturer that can get the most MTF from the lens.
- Ideally, the drone should continuously adjust its position and direction to get the best image, while automatically selecting the exposure and gain that will produce an image with the lowest amount of noise, lowest blur, and maximum amount of recognizable detail.

- Keep an eye out for the solution's software features. You need sophisticated on-board processing to make sure that any data collected contains sharp, high-quality images, taken at the correct angle and the optimal time of day, including the most advantageous light, best weather conditions, least wind, and at an optimal distance.
- Find a solution that is easy to deploy and doesn't need a specialized pilot. By using computer vision to land, the Percepto system remains autonomous, allowing the drone to return safely to its starting point every time.

By giving top priority to our lens configuration, image quality, and software algorithms, Percepto allows you to enjoy the ease-of-use, quality, and cost savings you deserve.

Contact us to set up a demo

Put us to the test!

Drop us an email or visit our web site to consult with one of our experts.

