

A 3D SCANNING primer

Marco Callieri

3D Digital Models



The models seen so far has been generated from real-world objects.

They are not something that "looks like" the original object, but are a faithful, metric digital representation of the surface geometry of their physical counterpart

They have been generated using **3D scanning**



3D SCANNING

3D scanning is not a technology, but a family of technologies (and a quite large one)... but, In all its incarnation, it is an

automatic measurement of geometric properties of objects.

The produced digital model is formed by geometric information that have been measured and have a metric quality.

It may be imprecise and incomplete, but still has all the characteristics of a measurement result

Basics

Geometric properties may be measured in many ways...

What are normally called "3D Scanners", generally, are:


- **DISTANCE**: the instrument does not touch the object
- **OPTICAL**: the measurement relies on some physical law of optics, using visible or near-visible light
- **ACTIVE**: the instrument does project some form of lighting on the object to measure the geometry

Alternatives

Just for completeness, there are alternatives


CONTACT

Geometry is measured by touching it, with articulated arms or industrial touch sensors



NON-OPTICAL

Medical devices measure volumetric data using high-energy radiations: Xray, CT-scan, RM



PASSIVE

Geometry is calculated just by looking at the scene, without adding illumination. 3D from stereo, 3D from images, Photogrammetry, Multi-view stereo






Measurement

Triangulation

A light is projected on the surface and its reflection is read back by a sensor...

Using trigonometry, it is easy to recover the 3D position of the illuminated points. The geometric principle is the simplest possible! The real problems are to be fast and precise.

Knowing the emission and reception angles (α and β), and the distance (d) between the emitter and the sensor, the distance of each sampled point is calculated



Minolta Vivid

Commercial 3D scanner, one of the oldest used in the CH world...


A laser line is swept over the object: 300K points are measured in 2.5 seconds.

No longer in commerce, sadly ☹



NextEngine

Entry-level 3D scanner, simple and cheap. Good quality/price ratio. Ideal to investigate the possible use of this technology in a laboratory / museum



Pro:

- cheap (around 2k Euros)
- good resolution and result coherence
- highly portable (small and lightweight)

Cons:

- fixed working distance,
- sloooooow,
- need parameter tweaking
- does not work well on some materials (dark & shiny)



A dying breed

Laser-line scanners are no longer a popular choice, shadowed by other technologies

Still available some 2nd-hand devices; cheap scanners for small simple pieces; DIY kits.


Mostly, they are now used as industrial sensors.



Structured Light

It is still **TRIANGULATION**, but but different patterns are projected on the object.
Can be more precise than laser-line triangulation, and more resilient to some material-related problems, but require additional hardware and calibration

Many different kind of patterns, many different arrangements of multiple cameras + multiple projectors



Breuckmann GmbH

Industrial sensor, designed for optical metrology.
High precision and high cost.


Accuracy: 0.1 mm (or less)
Cost: 70-80k Euro



GOM

Another high-end metrology scanner.
The device we are currently use for 3D scanning

Accuracy: 0.1 mm (or less)
Cost: 70-80k Euro




Microsoft KINECT (old version)

It is basically a very fast (30fps) structured light scanner. Resolution is not great, very low precision, z-depth is compressed.

However, its cost and performances have shaken the community of 3D hobbyists but also of professionals

Not really usable for 3D scanning, especially in CH



Hand-held

The new "general purpose" scanners. Focused at being easy-to-use, versatile and portable.

They are hybrid scanners, as they use one or a combination of different measurement methods:

- Structured light
- Laser lines
- Time of flight camera (more info later)
- Phase shift (more info later)
- Stereo



Hand-held

This class of devices is rapidly gaining market shares, due to their versatility.

Many different brands, difficult to keep track.
Not every product is suitable for CH applications.

The cost varies a lot, from 200-500 euros (mostly toys) to the same price range of high-end metrology (80k)



ARTEC

Probably the most used in the Cultural Heritage world
Reasonably cheap, quick, portable, versatile.
Very good results on CH objects.

They capture color (not well, though)

Accuracy: 0.5-0.1 mm
Cost: 15-20k Euro



Kinect-derived

Using the same basic sensors of the Kinect (very fast structured-light scanners) work in tandem with tablets.

Super-portable and flexible, but extreme variability in quality/usefulness

Some (like the Structure) are more or less toys

Some (like the XYZ) are good for quick surveys in small environments.

Accuracy: 1-2 mm
Cost: 300-1k Euro



And what about large objects?

This is a very common question... The answer is, you do need a different instrument

Triangulation cannot work on very large objects, it would require an extremely large baseline...


Always remaining in the kingdom of light signals and optical properties, a different strategy is used

These scanners are often labelled as TERRESTRIAL LASER SCANNERS...

Time-of-flight (TOF)

The distance of the sampled point is obtained by measuring the time between the laser impulse and the sensor reading the ray coming back, and dividing it by two times the speed of light...



This is a linear distance measurement, to scan the object, the sensor uses a regular polar grid to cover the area of interest




Pure TOF scanners

Older models of terrestrial laser scanners only used TOF
They were extremely slow (20-30 minutes per scan) and expensive (>100k euro)

Accuracy was in the order of 5mm-2cm




Phase Interference / Phase Shift




The distance is measured by analyzing the interference between the emitted ray and the returning ray.

This calculation is combined with the TOF measurement.



Phase Interference / Phase Shift



The same interference principle is used on THREE different scales:


- **CONOSCOPY**: coins, paintings, small relief. Extremely specialized devices, often custom-made
- **INTERFERENCE + TOF**: buildings, terrains
- **FLASH INTERFERENCE**: interference cameras for human-size objects, using fast, synchronized “flashes”. Again, in the hand-held class.

Interference +TOF scanners

The use of interference made scanners much faster (5-10 min per scan) and precise (1mm or less). Most interference+tof scanners have, however, a shorter range.

Hardware has become smaller and cheaper (20k € and up).

This is the most used family of products in CH.



Big Names

Difficult to mention all possible terrestrial laser scanners, here are the main brands:

- **Leica (Cyrax)**: most diffused, produces all possible tools for survey
- **FARO**: affordable and most portable, also produces small-scale 3D scanners
- **RIEGL**: long range scanners, and inertial platforms
- **Z+F**: produces sensor hardware, sometimes re-branded by other companies
- **TOPCON**: extremely popular in US (and east Asia) for engineering and construction works



Data and problems

Limitations

3D scanners can cover a variety of objects, but there are still some limitations. Some of them can be overcome, others are intrinsic:

- **Visibility** (direct, cone of visibility)
- **Non-rigid** objects
- **Color** (black, pure color)
- **Material** (reflective, transparent and semi-transparent, peculiar BRDFs)
- **Non-solid** stuff (hair, beard, feathers, vegetation)
- Acquisition **environment** (temperature, illumination, crowded places)
- **Size vs. Single map** acquired (accumulation of alignment error)

Only Points

Regardless of the technology, 3D scanners only measure the spatial position of POINTS.

All that is returned from a single “shot” is just a series of points in the 3D space.

The characteristics of the points generated by the scanner do depends on the kind of scanner used.


Range map

Almost all optical scanners uses a camera as input device.



What is recovered after a single shot is a depth value for each pixel in its sensor which is converted in a 3D point.

So, from the point of view of the scanner, all the 3D points are on a REGULAR GRID, that is promptly triangulated using this intrinsic regularity.

This is possible (without introducing much error) because of the limited Z-span.



The result of a single scan is generally called a RANGE MAP



Polar Range map

Terrestrial laser scanners measure **one point at a time**. This distance measurement is iterated rotating along two axis.

Each scan, thus, creates a polar grid.

There is still a regular grid, but as the Z-span is too large, it is generally not advisable to triangulate them. TLS scans are normally kept and processed as pointclouds.



Aggregated clouds

Some scanners (mostly handhelds) do produce aggregated clouds, where the grid/radial structure is lost.

This is because some processing (alignment, as we will see later) has already been done.

This restricts the kind of filtering, cleaning and processing you may do on the raw data.

There is not much you can do about it, save that to use it as a whole.

The "ERROR"

Everyone asks “how precise is this scanner / 3D model?”. But this is a very tricky question...

Scanner data sheets are laboratory condition, determined with metrology tests. They are significant as the tech specs of your car (i.e. not that much)

On-the-field conditions do affect the data quite a lot, so do the material of the object, so do the scanner distance/angle. So, it is not even possible to give a single number for the accuracy of a single shot of the scanner, as the value changes point by point.

X-Y error is different from Z error:

- X-Y position is determined by the scanning grid (low error)
- Z (depth) is calculated, and here is most of the error

The "ERROR"

It has been proven error in a single scan is not “white noise”, but still, it can be lessened by redundancy.

There are systematic and recurring errors, sometimes local (specular highlights, black-to-white), sometimes global (vibrations, moiré patterns).

Determination of the error is often a matter of “thumbing it”

Error is bound by the greatest of:

- Resolution (how far are two measured points). Actually, should be half of the resolution for the sampling theorem.
- Scanner sampling error (at least the value in the data sheets, but normally higher)



Data Processing

All that remains

Single scans are already a 3D model. But they generally do not cover the whole object.

Multiple shots are needed. How many? Which one to choose?

The actual scanning is just ONE of the steps to create a 3D model



The long, winding road

- [planning]
- **Acquisition** of multiple range maps
- Range map **cleaning**
- **Alignment** of range maps
- **Merging** of range maps
- Mesh **cleaning**
- Color mapping
- [model preparation and archival]



Scanning
pipeline

Not always a straight pipeline, sometimes you go to previous steps, fix things, and retry

Working

Every scanner is bundled with a control software. The software is able to do all the processing. Nevertheless, for large or complex project, you may use other, specialized tools

- **GEOMAGIC**: commercial, the most used tool by professionals.
- **MeshLab**: opensource, 3D meshes and pointclouds processing, powerful and versatile. Not really user friendly.
- **CloudCompare**: opensource, for the processing of pointclouds. Very advanced and versatile. Even less user friendly.

Alignment

Each part of the model is in its own coordinate system. There is no spatial relationship between the different parts, as they have been generated in a different shots

Goal: bring all the parts in a common reference system (like a 3D jigsaw puzzle)

Two steps:

1. Rough alignment: user manually positions the various chunks in more or less the correct position
2. Fine alignment: the computer automatically perfects the alignment using the shared area between the range maps



1 Rough: correspondences alignment

First step.

It is necessary to have an overlap region with some common feature.

Common method: picking shared reference points.



Models are roughly positioned according to the selected couples of points. Not a perfect alignment, but enough to start the next phase



2 Fine: geometric alignment

The scans are now close enough. Local geometry of the scans is used to move one surface to the other, until complete matching.

ICP (iterative closest points) is a widely used method.



More than 2

This strategy is for two meshes. What happens when there are more than 2?

Geometric alignment is done on all the overlapping couples

Global optimization, a.k.a. Bundle adjustment is used to combine all individual movements, and evenly distribute the error

This geometric+global strategy is used by all software tools, with a lot of differences in implementation and interface.

Not always necessary



Not all scans need the alignment step or, at least, an explicit alignment step.

- Scanner tracking
- Progressive tracking
- Scans can be aligned using reference markers
- Automatic matching and alignment is possible in some cases

Scanner tracking

If scanner position is known in each shot, alignment phase can be reduced (rough alignment) or completely eliminated

- Rotary stage: PC-controlled, 1 DOF angle rotation. Simple and effective
- Arm positioning system: 2 to 6 axis, complex and costly, but very high precision (active or passive)
- Tracking system: generally wireless, less precise than a physical tracking, but flexible



Progressive tracking

Most hand-held scanners scan continuously (or a few times per second)...
Each «shot» is aligned to the previous one (or to the accumulated pointcloud).
The scanner/software assume a slow motion, and rely on geometry, color information, inertial data (accelerometers) or markers (see later).

Error accumulates... So, in most cases, a global optimization is carried out at the end of the scan session.

Different session has to be aligned as described before.



Markers

Markers are physical objects placed near/onto the surface to be acquired that are recognized by the scanner (known patterns/geometries, color-codes, materials).

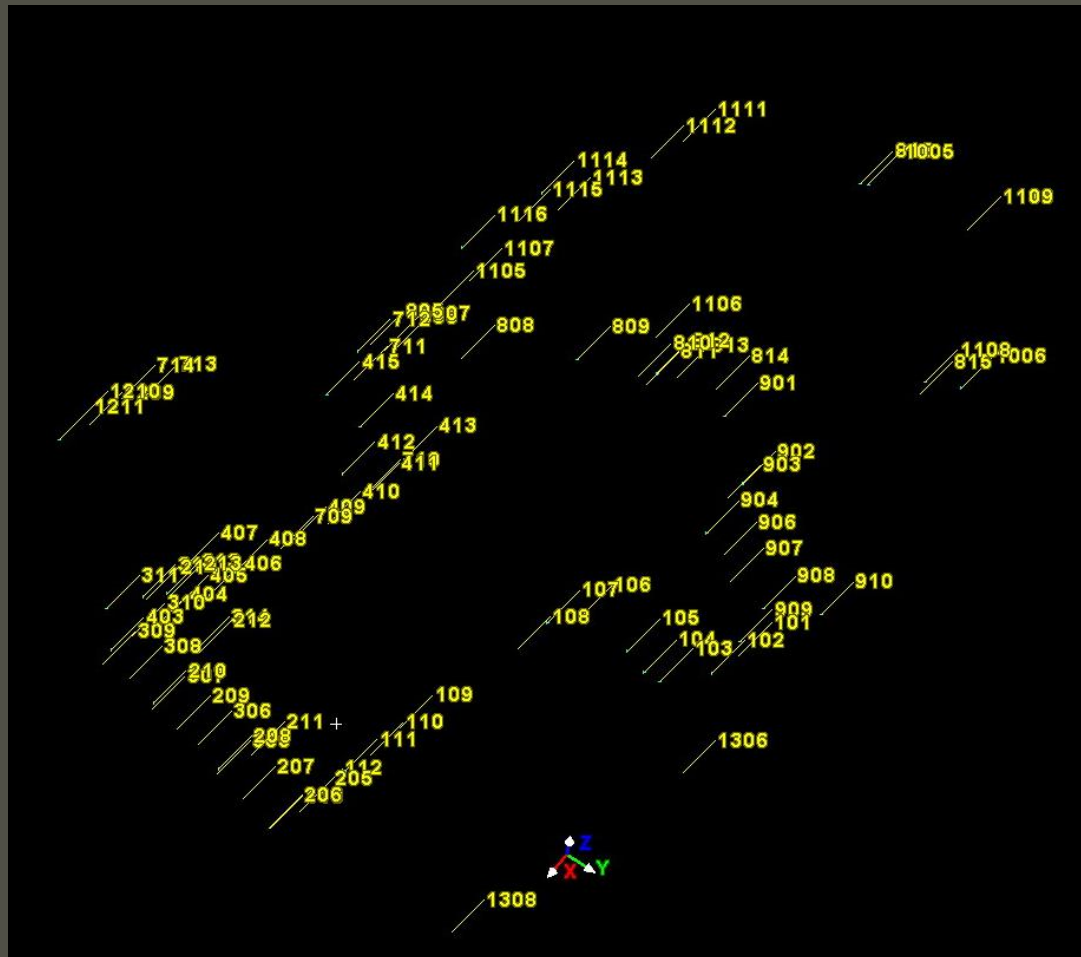
Their position is used as a reference for rough and fine registration

“Total Station” is used in surveying and laser 3d scanning of building, a theodolite is used to determine the position of reference points.

This technique is quite slow but really precise and reliable (we have used it in the last 7000 years)



Markers



Markers

- All terrestrial laser scanners uses markers, and this option is natively present in their software
- All terrestrial laser scanners software tools do accept external total station reference points
- Some triangulation scanners support markers (natively or with an add-on)
- It may be possible to mix reference points / markers / geometric alignment, but heavily depends on the software.
- Some hand-held scanners use also markers for progressive tracking


Automatic matching

The software, just by looking at the scans, can determine the relative position, by matching geometrical features.

Assumptions:

- Enough details on the surface
- Reasonably small change in the point of view of the scanner

Automatic matching replace the initial rough alignment, geometric alignment is used afterwards




Surface creation

When all maps have been generated, cleaned and aligned, it is time to generate a single surface

Why? to cover the entire surface we need all scans, but now we have multiple maps covering the same area, intersecting each other. Moreover, the sum of all map has too much data to be useful.

- All points represent the surface we need
- Each single point has been measured, but contains error
- Redundancy does help



Surface creation


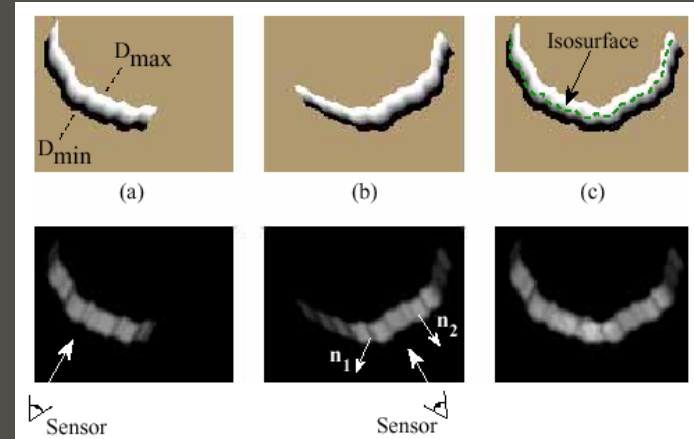

We need to find the “consensus surface” that better approximates the input points.

Generally, this happens in a **volumetric grid**.

All points are used to generate a volume scalar field that approximate the surface.

This is done using distance fields, implicit functions, and other geometrical/mathematical representations.

Then, we triangulate it.



Marching Cube


Most of the merging software use some variant of this algorithm.. It works on a scalar field, defining for each point in space how far is from the surface to be extracted, and if we are outside or inside it.

For each cell, the sign (I/O/OUT) of its vertices is computed. This configuration determines the triangles that will be in that cell. Triangle position is then chosen using the field value (isosurface level 0)



Marching Cubes: A high resolution 3D surface construction algorithm

William E. Lorensen, Harvey E. Cline (siggraph 87)


Up to few years ago, it was covered by a patent ☺



Marching Cube



Reconstruction
resolution



Marching cube

The resolution of the grid controls the final resolution of the model. Smaller cells means more detail, but also much more resources needed to calculate it.

Marching cube is responsible for that “blocky” look of the triangulated surface.

