ISO/IEC JTC1/SC29/WG1 **N86012**

86th Meeting – Sydney, AU – 18-24 January 2020

INTERNATIONAL ORGANISATION FOR STANDARDISATION
ORGANISATION INTERNATIONALE DE NORMALISATION

ISO/IEC JTC 1/SC 29/WG1
(ITU-T SG16)

Coding of Still Pictures

JBIG	JPEG
Joint Bi-level Image Experts Group	Joint Photographic Experts Group

TITLE: JPEG Pleno Point Cloud – Use Cases and Requirements v1.3

SOURCE: Stuart Perry (Editor)

PROJECT: JPEG Pleno

STATUS: Approved

REQUESTED ACTION:

DISTRIBUTION: Public

Contact:
ISO/IEC JTC 1/SC 29/WG1 Convener – Prof. Touradj Ebrahimi
EPFL/STI/IEL/GR-EB, Station 11, CH-1015 Lausanne, Switzerland
Tel: +41 21 693 2606, Fax: +41 21 693 7600, E-mail: Touradj.Ebrahimi@epfl.ch

Editorial Comments

This is a living document that goes through iterations. Proposals for revisions of the text can be delivered to the editor Stuart Perry, by downloading this document, editing it using track changes and sending it to Stuart.Perry@uts.edu.au.

If you have interest in JPEG Pleno Point Cloud, please subscribe to the email reflector, via the following link: http://jpeg-pointcloud-list.jpeg.org

ISO/IEC JTC1/SC29/WG1 N86012
86th Meeting – Sydney, AU – 18-24 January 2020

JPEG Pleno Point Cloud - Use Cases and Requirements v1.3

January 24th, 2020

1 General background information

1.1 Rationale

This document is intended to describe a number of use cases for the coding of (3D) point clouds. For each use case, a description is given and requirements outlined. From the requirements specific to each use case, a general set of requirements is derived and divided into systems and coding requirements. These requirements can then drive coding proposals and subjective testing.

1.2 Basics of point cloud data

There are a variety of point cloud representations relevant to 2D visualisation, virtual, augmented and mixed reality. One of the key considerations is whether the point cloud geometry is quantised, either by allowing the xyz coordinates of the point cloud to only exist on a regular grid, or by quantising the coordinate values. When the original image grid and camera parameters are present, point clouds can be defined on the image grid and be referred to as an *organized* point cloud (8-bit per colour component and up to 12-bit for each grid entry) [1]. When each point is defined by explicit xyz coordinates without any image grid structure inherent to capture (i.e. all data has been fused in the reconstruction stage to a full 3D reconstruction), the point cloud is considered as *unorganized*. In this case, the number of points (resolution) is arbitrary and depends both on the inputs and the segmentation/reconstruction algorithm. Figure 1 shows examples of organized and unorganized point clouds. xyz coordinates can then be represented by 32-bit floating point values. However, often these are converted to 9-20 bit integers by quantizing the coordinate values. Point clouds with quantized coordinates may be used to represent occupancy grids, which are common in robotics but are also used in virtual and augmented reality. Occupancy grids may be thought of as voxelized point clouds whose quantized coordinates represent occupied volumetric elements or voxels. A floating point representation of the coordinates is often used as input format for graphics APIs, but an integer representation is more suitable for voxelized point clouds as well as scanned 3D point cloud data, as the 3D scanners often have a fixed depth granularity/resolution and accuracy.

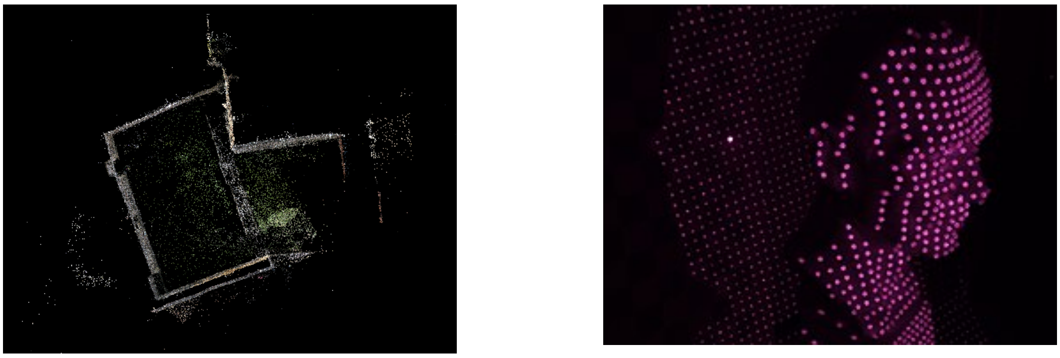

(a) (b)

Figure 1: Examples of unorganized and organized point clouds. (a) An unorganized point cloud captured by the use of photogrammetry to merge images into a 3D point cloud (By John Cummings [CC BY-SA 3.0 (https://creativecommons.org/licenses/by-sa/3.0)], from Wikimedia Commons). (b) An organized point cloud captured using Laser diffraction monofotogrammetry By Jorma Palmén [CC BY-SA 4.0 (https://creativecommons.org/licenses/by-sa/4.0)], from Wikimedia Commons. Notice the grid structure inherent to capture is still visible.

In addition, the colours are represented by 8-12 bits for each component. In cases where reflectance and gloss information is crucial, this may be stored as a separate attribute with up to 16 bits allocated per point depending on the required level of accuracy. For highly realistic rendering of a 3D world scene, support for additional attribute information such as normals, material, radiance and transparency properties, may be desired. For example, normals can be represented by 3 floating point components (e.g. in a spherical coordinate system), while material reflectivity can be represented by material properties (e.g. on a 0-1 floating point scale) and transparency/opacity by point specific attributes.

2 Use cases

This section collects use cases where point cloud representations may play an important role, notably a description and associated requirements. The use cases in this section refer primarily to static point clouds, however some use cases may also relate to dynamic point clouds. This report is not intended to be restricted to only static point clouds.

2.1 Virtual, augmented and mixed reality, computer graphics and gaming

Description: Point clouds are instrumental to support the display of 3D content in virtual, augmented and mixed reality (VR/AR/MR) environments and 2D or 3D computer graphics and gaming scenarios.

The content may be generated artificially by computer-aided design (CAD) design software or created through the scanning of physical objects with 3D scanning equipment. The resolution and number of points of the content will be highly dependent on the type of object and industry. Geometrically, point clouds created using CAD software are likely to be arranged on regular grids and patterns, while data collected using 3D scanning equipment is likely to be arranged in an irregular geometric pattern unless re-meshed in some fashion. Attributes of points in this data is likely to include colour, gloss, bump maps or texture maps as well as bidirectional reflectance distribution function (BRDF) or other material appearance information.

Requirements: 3D point clouds in VR or AR applications may need to be rendered at frame rates of 5 to 30 fps (possibly at a varying frame rate). As such applications introduce stringent requirements on the end-to-end media chain, real-time and low latency are of importance to this use case. In addition to support immersive 3D networking, progressive coding, view selectivity, region of interest and resolution granularity/scalable bit stream capabilities may be also of importance, as is the ability to allow the tuning the quality of the decoded point cloud to suit different display devices. Some point attributes such as intensity, bump maps, texture maps and BRDF information may be encoded in a lossy way; however, support for attributes with lossless encoding may be needed where such attributes support look up or reference to a database such as a database of material properties. Figure 2 provides an example workflow for this use case, illustrating the requirement to render the point cloud from changing viewpoints, scales and resolutions in real time in the course of the gaming or VR experience.

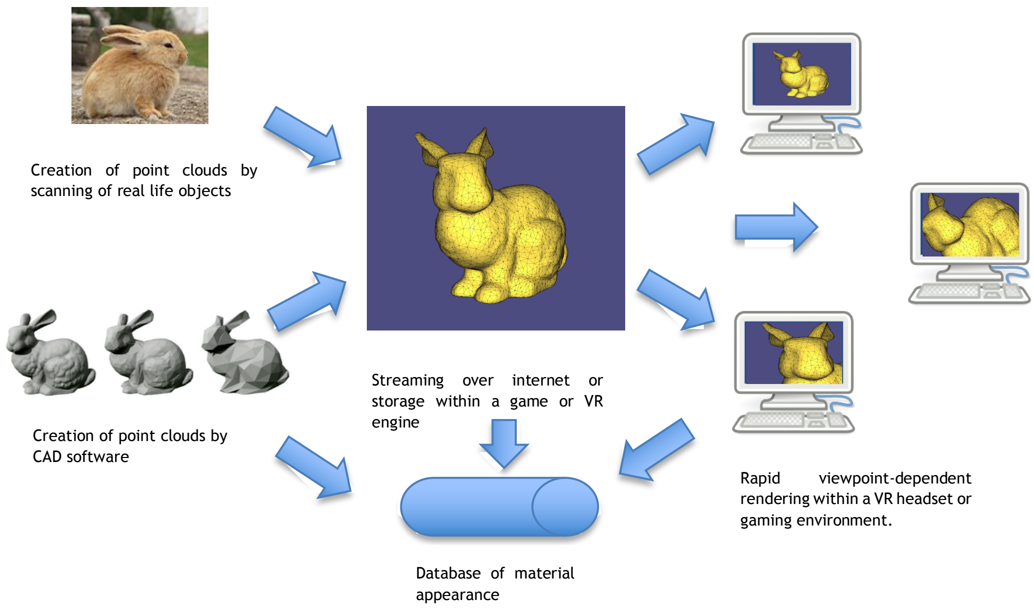

Figure 2: An example of use case 2.1. A 3D model is created by either the scanning of real objects or using Computer-Aided Design (CAD) software. This point cloud is streamed over the internet, a local network or stored within a game or VR engine to be displayed to the user. Depending on the viewpoint and scale of the object in the virtual environment, selective view selectivity or scalable bit stream capabilities may be needed for low latency display. In the case of computer generated content, a database of material appearance may be referenced at various points in the use case workflow. Image of live rabbit by Anna Frodesiak [CC BY-SA 3.0 (https://creativecommons.org/licenses/by-sa/3.0)] and central image of rabbit mesh courtesy of G1malitm [CC BY-SA 4.0 (https://creativecommons.org/licenses/by-sa/4.0)]

2.2 3D content creation

2.2.1 Design, manufacturing and 3D printing

Description: This use case involves artificially created point clouds to support the manufacture of 3D content for 3D printing or traditional manufacture.

The resolution and number of points in the point clouds will be highly dependent on type of object and industry; however, geometrically, data is likely to be arranged on regular grids and patterns associated with CAD software. Attributes of points in this data are likely to include colour as well as other material appearance information, potentially linked to external databases containing a set number of materials.

Requirements: This use case requires accurate representation of point clouds so lossless representation will be important. In addition, support for attributes with lossless coding may be needed where such attributes support lookup or reference to a database such as a database of material properties. To protect intellectual property there is a requirement to provide means to guarantee the privacy and security needs associated with content. Figure 3 shows an example workflow for this use case. In this example, an object is created and then transmitted over an insecure channel for later 3D printing. It is essential that a means for ensuring the privacy and security of the content are available to protect the intellectual property of the creator.

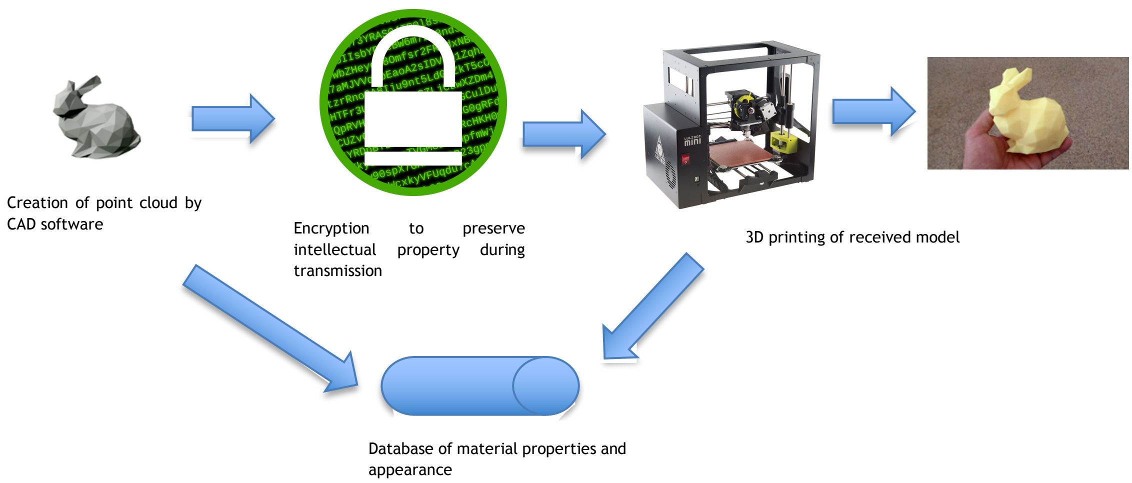

Figure 3: An example of use case 2.2.1. CAD software is used to create a 3D point cloud intended for printing. The point cloud requires security mechanisms to ensure protection of the creator's intellectual property during transmission. This use case may also make use of a database of material properties and appearance to aid in selection of the correct printing material during the 3D printing process. Image of 3D printed rabbit courtesy of Creative Tools [CC BY 2.0 (https://creativecommons.org/licenses/by/2.0/)]. Image of 3D printer courtesy of SparkFun Electronics (https://www.flickr.com/photos/sparkfun/16135614727) [CC BY 2.0 (https://creativecommons.org/licenses/by/2.0)].

2.2.2 Motion capture

Description: Motion capture with or without markers often results in point clouds where points correspond to each marker (either real or virtual).

These types of motion capture data are relevant for motion analysis and synthesis for various applications such as sports analysis, video game production, choreography, and movie production. Point clouds are a useful representation for data resulting from multiple sensors or from marker-based/markerless motion reconstruction.

Requirements: Efficient and interoperable coding and storage methods are a necessity in such systems. This use case shares requirements with the "Design, manufacturing and 3D printing use case described in Section 2.2.1 above.

2.3 Medical applications

2.3.1 3D medical imaging

Description: This use case refers to point clouds created by the 3D scanning of internal or external anatomy of human and animal subjects for the purpose of research, medical record keeping, diagnosis and the preparation of treatment plans.

The point cloud data may be tagged with important additional information such as tissue parameters or type, but traditional colour attributes are not likely to play a strong role.

Requirements: Point clouds collected for medical purposes need to be encoded in a way that preserves the ability for subsequent analysis. This requires accurate representation of point clouds so lossless or very low loss coding will be important. May often require multiple scans to scan anatomical structures and internal detail, with potentially the need to keep point clouds from individual scans distinct in the coded format for record keeping or future reconstruction purposes. In addition to support examination by medical professionals, view selectivity, region of interest and resolution granularity/scalable bit stream capabilities may be of importance. Some point attributes may be encoded by a lossy method; however, support for attributes with lossless coding will be needed where such attributes are crucial to diagnosis and analysis. To protect patient confidentiality, there is a requirement to provide means to guarantee the privacy and security needs associated with content.

2.3.2 Prosthesis and body parts design and manufacture

Description: This use case involves artificially created point clouds to support the manufacture of 3D prosthetics for medical purposes.

The resolution and number of points in the point clouds will be highly dependent on the type of object. Data is likely to be arranged on regular grids and patterns associated with CAD software; however, the data may also be arranged in an irregular geometric pattern associated with point clouds captured from scanned body parts. Attributes of points in this data is likely to include colour as well as other material appearance information, potentially linked to external databases containing a set number of materials. Highly accurate colour and material appearance representations may be needed to match the prosthetic against skin.

Requirements: This use case shares requirements with the "Design, manufacturing and 3D printing use case described in Section 2.2.1 above, although the need for security and privacy is to protect patient confidentiality and not intellectual property.

2.4 Construction and manufacturing

2.4.1 Analysis of 3D structure for defect detection

Description: This use case concerns point clouds resulting from the 3D scanning of construction material and industrial parts with the goal of detecting defects, damage, shape tolerance or cracks.

Requirements: This use case requires accurate representation of point clouds so lossless or very low loss coding will be important. Coding that preserves sharp edges and fine detail is essential. In addition, support for attributes with lossless coding may be needed where such attributes support lookup or reference to a database such as a database of material properties. May often require multiple scans to handle object internal structure and occlusions, with potentially the need to keep point clouds from individual scans distinct in the coded format for record keeping or future reconstruction purposes.

2.4.2 3D scanning to support project management

Description: Point clouds are used to support the display of 3D content for visualisation of architecture and infrastructure in context of a building site or surrounding area.

An example would be to visualise planned construction elements against already completed construction elements or visualise a planned construction in the context of existing surroundings. The content may be generated artificially by CAD design software or created through the scanning of physical objects with 3D scanning equipment or may be a merger of both sources of data. The resolution and number of points of the content will highly dependent on the type of object and industry. Geometrically, point clouds created using CAD software are likely to be arranged on regular grids and patterns, while data collected using 3D scanning equipment is likely to be arranged in an irregular geometric pattern unless geometrically re-positioned in some fashion. Attributes of points in this data is likely to include colour, gloss, bump maps or texture maps as well as BRDF or other material appearance information.

Requirements: 3D point clouds created for this use case are generally intended for rendering and not detailed analysis. This use case hence shares the requirements of the “Virtual, augmented and mixed reality, computer graphics and gaming” use case described in Section 2.1 above.

2.4.3 Repair and analysis of confined or dangerous places

Description: This use case refers to 3D scanning of cavities, confined, or dangerous places (for example inside pipes, grain silos or infrastructure) with the goal of detecting defects, damage, shape distortions or cracks.

For some diagnostic purposes, colour information may be very important. In addition, capture may be under ambient or uncontrolled lighting, so high-dynamic range colour may be required.

Requirements: This use case shares the requirements of the “Analysis of 3D structure for defect detection” use case described in Section 2.4.1 above, with the additional requirement for the storage of high accuracy colour information.

2.4.4 Point clouds for urban planning and urban analysis

Description: Point clouds are used to record 3D shape and details of objects greater than 20 metres in extent, especially in large urban areas.

The data is intended for analysis and rendering, and capture is likely to be under ambient or uncontrolled lighting.

Requirements: Data from sources such as this may need to be rendered for visualisation as well as processed for analysis purposes. This use case hence shares the requirements of the "Virtual, augmented and mixed reality, computer graphics and gaming" use case described in Section 2.1 above. In addition, this use case may often involve multiple scans to handle object internal structure and occlusions and it may be desirable to keep each scan separate in the coding format.

2.4.5 Crack and damage detection for roads, bridges and other infrastructure

Description: Point clouds are used to record 3D shape and details on large items of infrastructure.

The goal is to spot defects, damage or cracks, and capture is likely to be under natural lighting or circumstances of limited control and consist of a number of separate scans. Accurate colour may be required to spot various types of damage.

Requirements: This use case shares the requirements of the "Analysis of 3D structure for defect detection" use case described in Section 2.4.1 above.

2.5 Consumer and retail

2.5.1 Small objects: jewelry, decorations, shoes, etc.

Description: Point clouds are used to record 3D shape and details of small objects, e.g. for on-line shopping.

Point clouds for this use case are likely to be complicated in structure, high resolution with complex material appearance.

Requirements: Data from sources such as this will often be rendered on websites and mobile devices for retail. This use case hence shares the requirements of the "Virtual, augmented and mixed reality, computer graphics and gaming" use case described in Section 2.1 above with additional requirements for the rendering to be perceptually lossless. To protect intellectual property there is a requirement to provide means to guarantee the privacy and security needs associated with content.

2.5.2 Mid-size objects: (cars, motorcycles, furniture, etc.) and large objects (houses, apartments, etc.)

Description: Point clouds are used to record 3D shape and details of mid-size objects (up to around 10 metres in extent) and large objects (beyond 10 metres in extent), e.g. for online shopping, 3D content creation, semi-professional level manufacture, 3D printing, real estate, advertising and VR experiences.

May often require multiple scans to handle object internal structure and occlusions. Point clouds for this use case are likely to be complicated in structure, and have high resolution with complex material appearance. Mixtures of individually scanned objects merged into a whole object is possible with resulting different resolutions and point cloud properties within the single point cloud entity.

Requirements: Data from sources such as this will often be rendered on websites and mobile devices for retail or rendered for display (either 2D or 3D/VR/AR). This use case hence shares the requirements of the “Virtual, augmented and mixed reality, computer graphics and gaming” use case described in Section 2.1 above with additional requirements for the rendering to be perceptually lossless. Support in the coding format for separate point clouds which together form a whole object may be needed to allow for object interactivity or disassembly. To protect intellectual property there is a requirement to provide means to guarantee the privacy and security needs associated with content.

2.6 Cultural heritage

2.6.1 Jewellery, pottery, bones, fossils and other small artefacts

Description: Point clouds are used to record 3D shape and details of small objects.

In this use case the area scanned is small, but at high resolution. Point clouds for this use case are likely to be complicated in structure, with high geometric resolution and complex material appearance. Point clouds for cultural heritage often have a dual purpose of enabling detailed analysis and being visualised by the public on websites or mobile devices.

Requirements: This use case shares the requirements of the “Virtual, augmented and mixed reality, computer graphics and gaming” use case described in Section 2.1 above with additional requirements for the rendering to be lossless or have very low loss to support accurate analysis. Capturing an entire object may often require multiple scans to handle object internal structure and occlusions, with potentially the need to keep point clouds from individual scans distinct in the coded format for record keeping or future reconstruction purposes. Some point attributes such as intensity, bump maps, BRDF and texture maps may be encoded by a lossy method for the purpose of rendering with changes in viewpoint rendered with low latency; however, support for attributes with lossless encoding may be needed where such attributes are essential to analysis (such as colour and gloss maps). To protect cultural property there is a requirement to provide means to guarantee the privacy and security needs associated with content.

2.6.2 Paintings

Description: Point clouds are used to record 3D shape and details of paintings and murals.

The goal of this use case is potentially the display of the artwork, and analysis of the artwork to recognise the artist, understand how the artwork was created, or detect forgeries. Multiple scans may be required to capture the entire object and point clouds for this use case are likely to be complicated in structure, with high geometrical resolution and complex material appearance. Highly accurate colour and material appearance is likely to be needed for this use case. There is a possibility that point clouds for this use case may be represented as a height map above a surface, reflecting the flat geometry of paintings and murals.

Requirements: This use case shares the requirements of the “Cultural heritage: Jewellery, pottery, bones, fossils and other small artefacts” use case described in Section 2.6.1 above.

2.6.3 Statues, mid-size artefacts

Description: Point clouds are used to record 3D shape and details of objects up to around 1-4 metres in extent.

Multiple scans may be required to capture the entire object and point clouds for this use case are likely to be complicated in structure, with high geometrical resolution and complex material appearance.

Requirements: This use case shares the requirements of the "Cultural heritage: Jewellery, pottery, bones, fossils and other small artefacts" use case described in Section 2.6.1 above.

2.6.4 Facades, monuments, archaeology

Description: Point clouds are used to record 3D shape and details of objects up to around 4-20 metres in extent.

Multiple scans may be required to capture the entire object and point clouds for this use case are likely to be complicated in structure, high resolution with complex material appearance. Scans are likely to be collected under conditions of uncontrolled lighting.

Requirements: This use case shares the requirements of the "Cultural heritage: Jewellery, pottery, bones, fossils and other small artefacts" use case described in Section 2.6.1 above.

2.7 Remote Sensing and Geographical Information Systems

2.7.1 Wide area scanning

Description: Point clouds are used to record 3D shape and details of objects greater than 20 metres in extent.

Lighting during capture is likely to be uncontrolled and multiple scans may be required to reconstruct the scene. Detailed colour capture is sometimes, but not always, a feature of this use case, but the large-scale nature of the scene being captured often negates the need to capture detailed material appearance attributes such as transparency and gloss. Figure 4 shows an example of a point cloud typical of this use-case. The data has a high resolution but does not encode the material appearance of the scene imaged. In this particular example, color information has not been recorded.

Requirements: This use case shares the requirements of the "Cultural heritage: Jewellery, pottery, bones, fossils and other small artefacts" use case described in Section 2.6.1 above, however, unlike that use case, this use case may not require the storage of very accurate color and material appearance information.

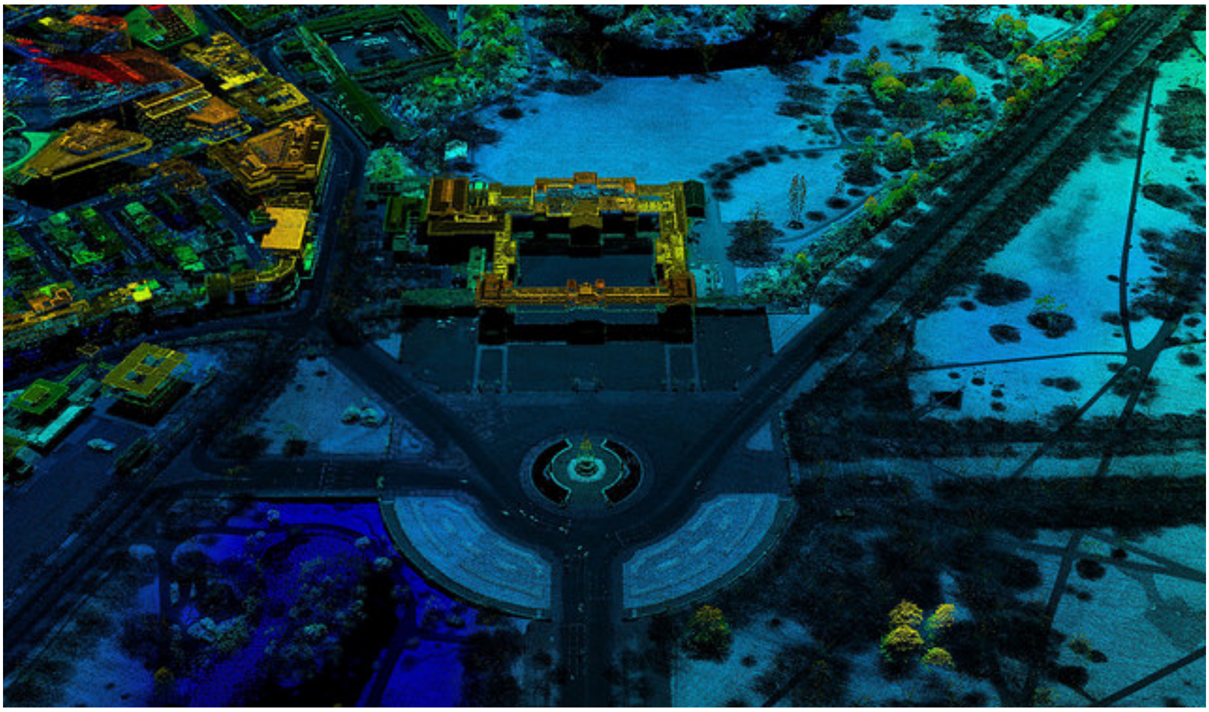

Figure 4: An example of a point cloud that would fall into use case 2.7.1. In keeping with this use case, the data is of high resolution, but without accurate material appearance information or color. The color in this figure represents height and is not related to the color of the imaged scene. The image is a LIDAR scan of Buckingham Palace, UK and is courtesy of Environmental Agency (https://www.flickr.com/photos/environment-agency/27489358013) [CC BY 2.0 (https://creativecommons.org/licenses/by/2.0/)]

2.8 Autonomous vehicles, drones

2.8.1 Autonomous vehicles

Description: Point clouds are used to record 3D shape and details of environment surrounding an autonomous vehicle.

The goal of this use case is to analyse the point cloud to spot objects of interest such as pedestrians and obstructions. Resolution of the scanned area may vary with finer resolution in particular directions of interest, such as the direction of the motion of the vehicle. Capture is likely to be under uncontrolled lighting. Point clouds created for this use case may not have or require accurate colour and material appearance.

Requirements: Analysis requires accurate representation of point clouds so lossless or very low loss coding will be important. Coding that preserves sharp edges and fine detail is essential. Requirements for real-time and low latency are of importance to this use case due to the need to continuously collect and analyse the incoming data. In addition, view selectivity, region of interest and resolution granularity/scalable bit stream capabilities may be of importance.

2.9 Surveillance

2.9.1 Point clouds for search and rescue in confined spaces

Description: Point clouds are used to record 3D shape and details of objects (in particular cavities up to around 1-10 metres in extent).

The goal of this use case is to spot objects of interest such as trapped survivors. Point cloud capture is likely to be under semi-controlled lighting, but in potentially chaotic environments and point clouds from this use case are not likely to require accurate colour or material appearance information, although colour may be important to aid object identification.

Requirements: This use case shares the requirements of the “Analysis of 3D structure for defect detection” use case described in Section 2.4.1 above.

3 Royalty-free goal

The royalty-free patent licensing commitments made by contributors to previous standards, e.g. JPEG 2000 Part 1, have arguably been instrumental to their success. JPEG expects that similar commitments would be helpful for the adoption of a JPEG point cloud coding standard.

4 Requirements

This section presents the overall set of requirements that have been extracted from the above described use cases. Requirements are split between “core requirements” which are essential for the standard and “complementary requirements” which are nice to have and will be decided depending on their cost. Depending on their type, the requirements will be classified as coding and systems requirements.
When considering geometry and attribute scalability, it is important to be aware that there is a relationship between the scalability of attributes and the scalability of geometry. Although it is possible to have use cases that require attribute scalability without corresponding geometry scalability, it is difficult to visualise use-cases where geometry scalability is needed without corresponding attribute scalability.

In addition, in the description below:

- An object is set of points within the point cloud representing a semantically meaningful entity, e.g. a chair or the floor within a room setting. Object based scalability is expected to be addressed at the systems level.
- A region is understood to be a set of points within the point cloud selectable at the decoding stage. For example, a particular region might be the points associated with the face in a point cloud representing a human being. This set of points may be contained within an object. This type of scalability is expected to be addressed at the coding level.

4.1 Core requirements

This section lists the set of core requirements derived from the full set of use cases.

4.1.1 Systems requirements

1. ***JPEG Systems Compatibility:*** The systems elements of the standard shall comply with the relevant JPEG Systems specification.

2. ***Multiple Point Clouds:*** The standard shall support the storage of multiple independent point clouds in a single stream or file. The standard shall allow selective access to an independent point cloud in the case that multiple point clouds are stored in the same stream or file.

3. ***Large Point Clouds***: The standard shall support the coding of point of clouds of any size, eventually using appropriate segmentation. The standard shall support the means for the appropriate decoding and recomposition of point clouds coded as multiple segments.

4. ***Metadata:*** The standard shall provide appropriate description tools for efficient search, retrieval, filtering and calibration of content. Strategies and technical solutions shall fit within the JPSearch and JPEG Systems frameworks.

5. ***Privacy and Security:*** The standard shall provide means to guarantee the privacy and security needs associated with content from the various imaging modalities. Strategies and technical solutions shall fit within the JPEG Privacy & Security and JPEG Systems frameworks.

4.1.2 Coding requirements

6. ***Data Types:*** The standard shall be able to code both the geometry data and the (global and local) attributes of a point cloud, where these data types are defined as follows:
    a. ***Geometry:*** Positional information of the points in the point cloud taken as integers or floats.
    b. ***Local/Point Attributes:*** Attributes on a point by point basis that may be integer or float such as colour, reflectance, links to external databases of materials, temperature or any indexes to lists. This includes attributes whose value changes with the angle of view to the associated point.
    c. ***Global Attributes:*** Global point cloud attributes that may be integer or float such as texture maps, bump maps and normal maps.

7. ***Coding Accuracy:*** The standard shall support both lossless and lossy coding of the geometry and attributes of the point cloud.

8. ***Compression Efficiency:*** The standard shall compress both geometry and attributes in an efficient manner, e.g. as measured in terms of bits per point.

9. ***Tuneable Quality:*** The standard shall allow tuning the quality of the decoded point cloud. Quality shall be measured by subjective and objective means as set forth by SC29/WG1.

10. ***Low Complexity:*** The standard shall allow for implementations of feasible encoding and decoding complexity, notably within reasonable expectations of time for the use cases listed above.

11. ***Parallelisation:*** The standard shall support efficient parallelisation of the encoding and decoding processes.

12. ***Geometry Scalability***: The standard shall allow for decoding the point cloud geometry using only a part of the full bitstream wherein the compressed bitstream shall be structured with more than one layer to decode. This shall provide:
    a. A point cloud with increasing number of points as additional layers are decoded.
    b. A point cloud with increasing quality, for a specified bit depth, as additional layers are decoded.
    c. A point cloud with increasing precision, in terms of increasing bit depth, as additional layers are decoded.

    If increased scalability comes at the cost of a reduced compression efficiency, then it should be possible to tune/adjust the scalability granularity, notably to a minimum of scalability to reduce the compression efficiency penalty.

13. ***Attributes Scalability***: The standard shall allow for decoding point cloud attributes using only a part of the full bitstream wherein the compressed bitstream shall be structured with more than one layer to decode. This shall provide point cloud attributes of increasing quality as additional layers are decoded.
    If increased scalability comes at the cost of a reduced compression efficiency, then it should be possible to tune/adjust the scalability granularity, notably to a minimum of scalability to reduce the compression efficiency penalty.

14. ***Random Access:*** The standard shall allow the selective decoding of a portion of the point cloud independently of the rest of the point cloud without decoding the entire bitstream. For example:
    a. Selectively decoding a portion of the point cloud corresponding to a volume in 3D space.
    b. Selectively decoding a portion of the point cloud visible from a given viewpoint.

4.2 Complementary requirements

This section lists the set of complementary requirements derived from the full set of use cases.

4.2.1 Systems requirements

15. ***Metadata:*** The proposals for metadata must meet two concurrent requirements: i) support the JPEG Pleno Point Cloud use cases, and ii) be sufficiently flexible framework that is usable for other JPEG Pleno projects (e.g., Light Field and Holography), i.e., be a metadata framework. In addition, the proposals should be extensible to allow for emerging use cases for JPEG Pleno.

4.2.2 Coding requirements

16. ***Region-based Geometry Scalability***: The standard should allow for decoding the point cloud geometry using only a part of the full bitstream wherein the compressed bitstream should be structured with more than one layer to decode. This should provide a point cloud where an increasing number of regions are decoded as additional layers are decoded.

17. ***Hybrid Region-Based Geometry Scalability***: The standard should allow for decoding the point cloud geometry using only a part of the full bitstream wherein the compressed bitstream should be structured with more than one layer to decode. This should provide:
    a. A point cloud with increasing number of points as additional layers are decoded.
    b. A point cloud with increasing quality, for a specified bit depth, as additional layers are decoded.
    c. A point cloud with increasing precision, in terms of increasing bit depth, as additional layers are decoded.

    wherein the above provisions are controllable on a region by region basis.

18. ***Region-based Attributes Scalability***: The standard should allow for decoding point cloud attributes using only a part of the full bitstream wherein the compressed bitstream should be structured with more than one layer to decode. This should provide point cloud attributes of increasing quality as additional layers are decoded controllable on a region by region basis.

5 References

[1] JPEG Pleno – Scope, use cases and requirements Ver.1.6, Doc. ISO/IEC JTC 1/SC 29/WG1 N73030, Chengdu, China, Oct. 2016.