Quidient Reality® Platform

Conceptual Design Document

Rev 2025-03-11-1

Approvals for this CDD:

____
John Leffingwell, for PDev, Custodian

____
~~*TBD*~~Brett Green, for Engine Group

____
Brevin Tilmon, for Algorithm Group

____
Brett Green, for ~~Product Development~~PDev

____
Scott Ackerson, CEO and President

Table of Contents

1 About This Document

1.1 Description

This Conceptual Design Document (**CDD**) describes key design elements of Quidient's flagship AI software product, the Quidient Reality® Software Platform (**QR**, **Platform**, **QR Platform**). QR's conceptual design draws from Quidient's cumulative body of knowledge and is meant to guide, but not specify in detail, the concrete design and implementation of QR modules toward satisfying PRD requirements. Sections are arranged roughly in order of conceptual importance and therefore impact to concrete design and implementation.

1.2 Nomenclature

Refer to this Glossary of Terms for definitions (like "Scene" and "Platform") that Quidient uses to describe scene reconstruction and processing.

1.3 Custodian

Please consult with the custodian of this document to coordinate substantive changes and approvals.

2 Goal-Driven Processing

QR processing is always driven toward a goal that includes performance criteria.

2.1 Performance criteria

Typical performance criteria that can drive QR processing:

- Perceptual quality of novel viewpoint rendering without relighting
    - A virtualized (sub)scene will be rendered from novel viewpoints under the captured (measured) light field.
    - The range of novel viewpoints can be specified.
    - Examples:
        - Average PSNR calculated over a range of desired novel viewpoints
        - Worst-case PSNR calculated over a range of desired novel viewpoints
- Perceptual quality of relighting
    - A virtualized subscene (e.g., OOI) will be placed into a novel[1] superscene for rendering from a range of viewpoints.
    - The range of novel superscenes can be specified.
    - The range of (novel) viewpoints can be specified.
    - Examples:
        - (Average and/or worst-case) PSNR calculated (on the subscene) over a range of desired viewpoints rendered in a novel hotel lobby (superscene)
        - PSNR calculated over a range of desired viewpoints rendered in a novel city street, forest, and shopping mall interior
- Geometric uncertainty
    - Examples:
        - 1-sigma uncertainty in displacement[2] from a reference plane, calculated over some extent of a virtualized surface
        - 2-sigma uncertainty in surface ~~orientation~~direction (normal vector) calculated over an OOI surface
- Duration (latency)
    - Examples:
        - End-to-end clock time for a reconstruction API call to complete once all captured data has become available to the reconstruction service
        - Time to complete a custom-specified subscene extraction or segmentation (segtree construction)
- Cost (financial)
    - Examples:
        - Dollar cost of GPU compute cycles to complete a reconstruction to a specified target fidelity
- Data size

[1] Novel here means a superscene other than scene in which capture was performed.
[2] Calculated, for example, by interpolating over the origins of persistent elsies that represent a surface

    - Total bytes of cloud storage consumed by a scene model without the later stages of Smart Fidelity processing having been run to increase the parsimony of scene entity representations (and therefore reduce data size)
    - Total bytes of cloud storage consumed by a scene model after the later stages of Smart Fidelity have been run
    - Peak host RAM or GPU memory consumed during a given processing operation
- Disentanglement beyond relightability (e.g., decomposability, analyzability)

3 Elsies

3.1 Overview

In a scene model, QR uses elsies (scene elements) to represent characteristics of matter, light, space, and/or time:

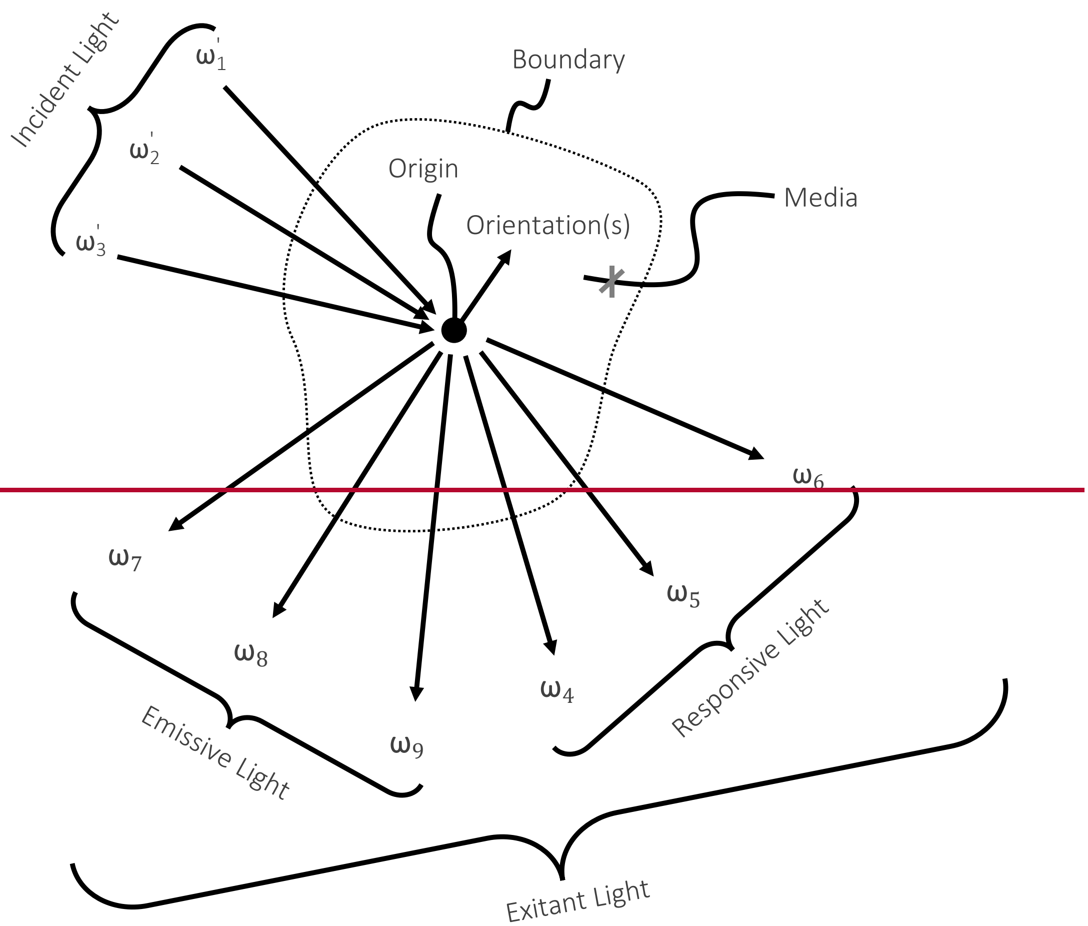

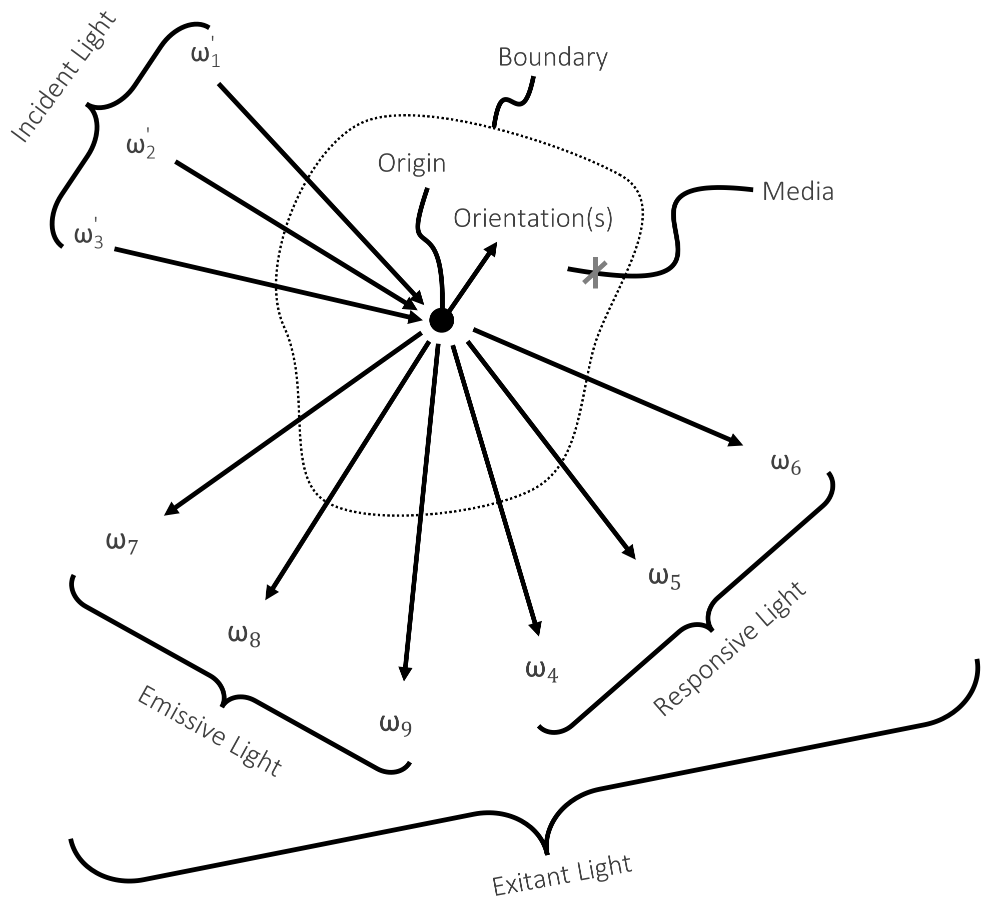

Figure 3.1-1: Elsie – Represents characteristics of matter, light, space, and/or time. A single elsie typically has limited complexity.

3.2 Key attributes

Key attributes of elsies:

- Elsies can efficiently represent natural and manmade matter fields (spatial configurations of media), including prismatic structures[3] and thin structures.
- Elsies can efficiently represent natural and manmade light fields, including radiance distributions ranging from isotropic to highly directional.
- QR stores elsies in a hierarchical spatial (or spatiotemporal) index[4].
- Elsies can represent (temporal) dynamism in matter, in light, and in spatial coordinate frames. The (represented) dynamism can potentially have discontinuities.
- All elsies have an origin (3D point).[5] All other elsie parameters are optional.

[3] "Structure" in this section = configuration of media in the matter field
[4] An acceleration structure in GPU ray tracking, for example
[5] An elsie origin by itself defines a 3D positional coordinate frame, leaving orientation unspecified.

- In general, an elsie with more parameters has greater representational power (aka just "power"). Elsies with higher power can represent richer, more complex structure and behavior in a scene region.[6] Variable (tunable) elsie power is of central importance in Smart Fidelity processing.[7]
- Detail in a plenoptic field (matter and/or light) is representable using a sufficient number / density[8] of elsies of sufficient power.
- An elsie can represent, to variable granularity, the uncertainty in values assigned to its parameters.[9]
- An elsie can represent generalized interaction behaviors between light and matter, including reflection, transmission, refraction, scattering, and absorption.
- An elsie can represent emissivity (light-emitting behavior of matter) alone or in combination with light interaction behavior.
- When interpolating (and extrapolating) elsie parameters over elsies in a scene region, QR evaluates parameter continuity postulates (e.g., C^0, C^1, C^2) in order of decreasing likelihood based on the local configuration of elsies.[10]
- A single elsie typically has limited complexity (limited number of parameters).
- An elsie origin plus a single 3-DOF orientation (2-DOF [pitch, yaw] direction + roll angle) establishes a 6-DOF coordinate frame.[11] An elsie can validly exist for this purpose, having no other parameters beyond its origin and (single) orientation.

3.3 Internal composition

An elsie can be decomposed as follows:

- Elsie = 1 origin + 0:N_{ori}[12] orientation + 0:N_{bnd} boundary + 0:1 mediel + 0:1 incident luminel + 0:1 responsive luminel + 0:1 emissive luminel
- Luminel = 1:N_{rad} radiel

Notes on elsie composition:

- In most QR scene models, mediel data[13] will tend to be more persistent and luminel data will tend to be more ephemeral.[14] Incident luminels will tend to be more ephemeral than responsive and emissive luminels.

---

[6] QR processing doesn't always evolve scene regions toward higher power. Power in a scene region can decrease in order to speed up subsequent operations, reduce data size, etc.
[7] In a multistage reconstruction driven by Smart Fidelity, QR will typically use elsies with lower power in the earlier stages.
[8] A "tunnel of resolution" can be realized as a scene region where the size (extent of influence) of each elsie is substantially smaller (and density of elsies is potentially greater) than in adjacent regions.
[9] An "island of accuracy" can be realized as a scene region with lower uncertainty than adjacent regions.
[10] For example, an intentionally C^2-discontinuous ornamental ridge on a car hood that is otherwise C^2-continuous
[11] Scale (7-DOF frame = 6-DOF frame + scale) is typically represented at the subscene (or occasionally segment) level.
[12] "A:B" here means the range of integers from A to B, inclusive.
[13] "Mediel data" here notably includes macrogeometry, BLIF, and emissivity parameters.
[14] Modern GPU hardware enables fast light transport operations to generate luminels on demand (i.e., to calculate the radiance values of their constituent radiels).

- Any of an elsie’s 3 luminels can be made persistent / cached, for example to represent the following:
    - Light flowing inward or outward across a fenestral boundary
    - Responsive light when the BLIF is computationally expensive[15]
    - Emissive light in fast-access form, e.g., discrete rays in direction space rather than basis functions over direction space

3.4 Example implementation

This section describes one conceptual way to implement elsies with the key attributes listed in §3.2.

3.4.1 Description and key parameters

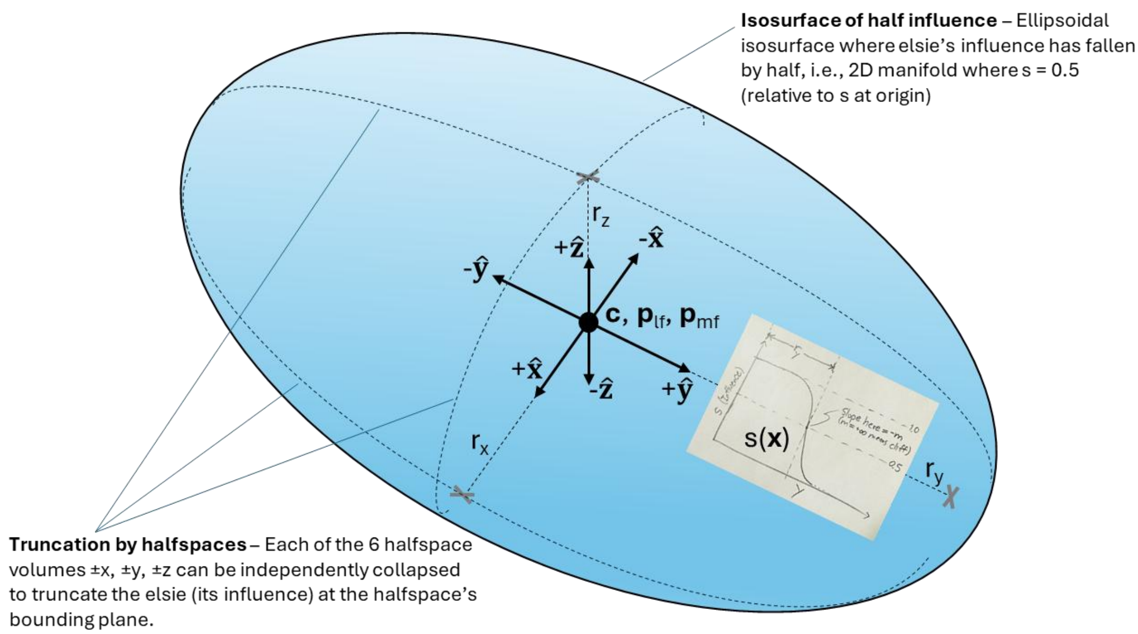

Figure 3.4.1-1: One way to implement elsies

Key parameters in this example implementation:

- **c** [DOF = 3 floats[16]] – Origin (**0** in elsie’s local frame)
    - Important note – The origin position **c**[17] is the only quantity required to be present in an elsie. All other quantities are optional. Even the axis directions are optional, i.e., the elsie’s (coordinate) frame could be 3-DOF vs 6-DOF.
- ±$\hat{\mathbf{x}}$, ±$\hat{\mathbf{y}}$, ±$\hat{\mathbf{z}}$ [DOF = 3 floats (total)] – Semiaxis directions (unit vectors)
    - Defined by 3 angles relative to the parent frame.

[15] Responsive light at an elsie can be cached to speed up queries for incident light at other elsies that are “downstream” in the sense of temporal light flow (photon flow).
[16] “Float” in this section means a floating-point or fixed-point representation of a real number. It doesn’t mean a single-precision float in some particular processor architecture.
[17] Vector / array quantities are **bolded**.

- Together with **c**, these establish a 6-DOF frame. (The scale DOF is typically expressed at subscene or segment level.)

- r_x, r_y, r_z [DOF = 3 floats] – Semiaxis lengths ("radii") from the origin to the s = 0.5 isosurface
- s(**x**) [DOF ≥ 1 float] – Strength of influence of the elsie, as a function of displacement **x** ([x, y, z] in the elsie frame) from its origin.[18, 19]
    - AKA "influence function" or just "influence"
- $\mathbf{p}_{lf}$ [DOF ≥ 1 float] – Parameters of the incident, responsive, & emissive light fields and their optional spatial variation ("SV") function (function of 3D position **x**)
- $\mathbf{p}_{mf}$ [DOF ≥ 1 float] – Parameters of the matter field, including (macro)geometry, emissivity, & BLIF (including microgeometry) and their optional SV function
- **σ** [DOF ≥ 1 float] – Statistical ~~1σ~~1-sigma uncertainty in each parameter in a chosen subset of all elsie parameters
- **trunc** [DOF = 6 Boolean] – Truncation state of each halfspace, e.g., ordered [+x,-x, +y,-y, +z,-z]

3.4.2 Interpolation of parameters

Notes on elsie parameter interpolation in this example implementation:

- QR uses non-gridded interpolation[20] to calculate elsie parameter values at "query points" (3D positions) away from elsie origins and/or under the influence of multiple elsies.
- In addition to a sharp cutoff in influence s(**x**) at truncation faces, the influence function can also have a sharp cutoff at one or more 2D manifolds other than the truncation faces.[21]
- Elsie parameter interpolation honors a spatial variation (SV) function that can optionally be provided for any parameter (besides the origin position). Ways to express spatial variation include the following:
    - Directional derivatives of order 1 to N
    - Orthonormal basis functions
    - Raster-like "map", e.g., material map, normal vector map, displacement map
    - Weights in a trained neural network
- A parameter's SV function is orthogonal to the influence function.
- A parameter's SV function can have discontinuities.
- The interpolated value of a parameter at a query point depends on the following:
    - Value of that parameter at the origin of each "contributing elsie"[22]
    - Particular type of interpolant used[23]
    - Each contributing elsie's influence function evaluated at the query point

---

[18] One example influence function is a downward sigmoid parameterized by the 3 semiaxis lengths ("radii") r_x, r_y, r_z and a sharpness parameter "m" (so named because m tells the negative slope of influence at s = 0.5). https://math.stackexchange.com/a/4798714 with its "k" set to our slope parameter m, and its "t" set to our s = 0.5 isosurface distance in each direction of interest.
[19] A 1-parameter Gaussian falloff might suffice to represent many matter and/or light configuration of interest, especially if expedient in earlier implementation work.
[20] "Interpolation" in this section also includes extrapolation to points beyond the envelope / hull of elsies (their centers) being interpolated over.
[21] See the example in §1.4.3 of a thin wire with sharp cutoff in elsie influence.
[22] An elsie whose influence function evaluated at the query point is above some floor threshold
[23] For example: nearest neighbor, linear, natural neighbor, cubic, (poly)harmonic spline

- From each contributing elsie, that parameter's optional SV function evaluated at the query point
- QR associates surface parameters[24] with 2D manifolds representing media discontinuities. Such manifolds include the following:
  - Elsie truncation faces
  - Influence function cutoffs
  - Discontinuity boundaries in the SV function
  - Discontinuity boundaries in the interpolant itself[25]
- In addition to or instead of SV functions, macroscopic parameter gradients can be represented by multiple elsies in a region and the interpolation over them.
- Where useful[26], QR replaces (interpolated) parameter discontinuities (aka "cliffs") with a narrow sigmoid or similar "smoothed cliff".

3.4.3 Prismatic structures

This example implementation can represent prismatic structures in the matter field, like this slab of solid glass:

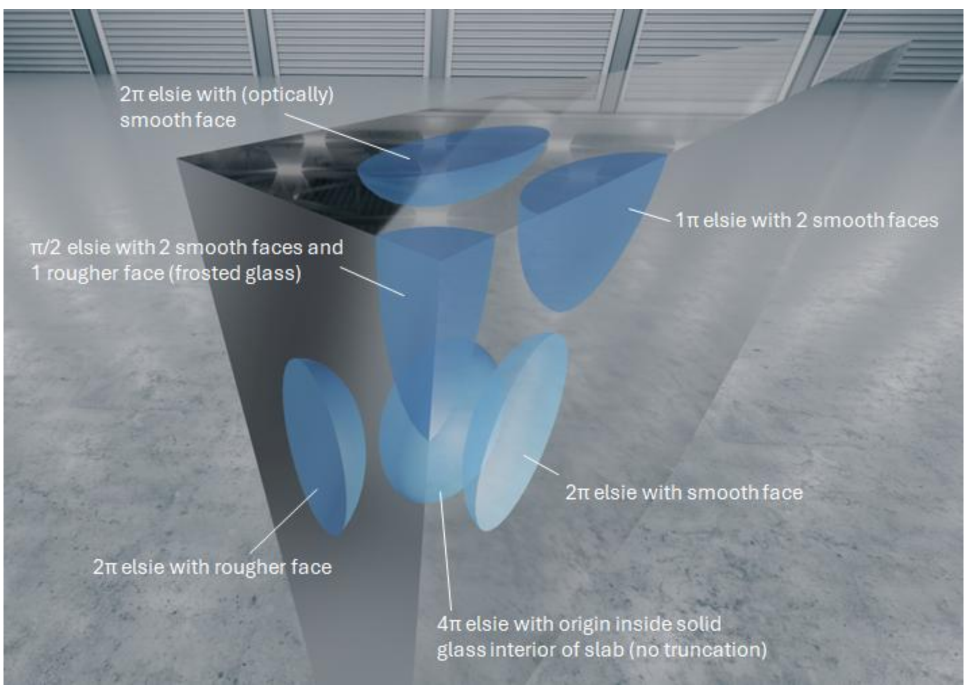

Figure 3.4.3-1: Prismatic slab of glass

[24] Microfacets and/or other microgeometry, for example, where volumes of different media meet
[25] C^0 discontinuity in nearest neighbor interpolation, for example
[26] For example, to avoid invalid derivatives in derivative-based optimization

3.4.4 Thin structures

This example implementation can represent thin structures in the matter field, like this electrical wire:

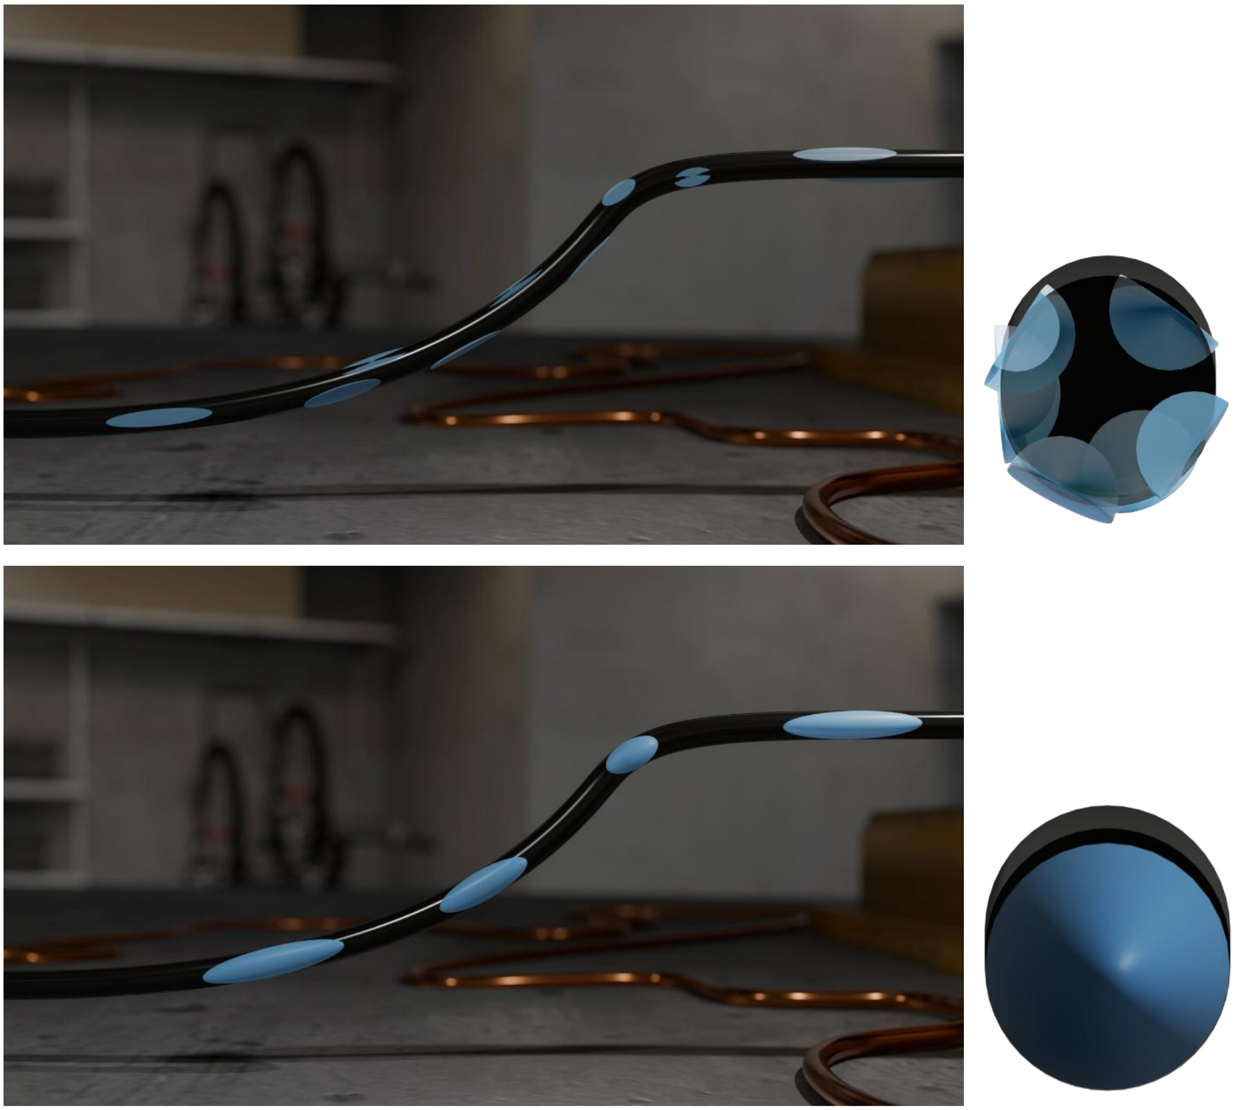

Figure 3.4.4-1: Thin wire represented in 2 ways – Top left: Side view of a wire represented as 2π elsies centered along the wire’s surface. Top right: Axial view of top left scene. Bottom left: Side view of the same wire represented as 4π elsies centered along its main axis. Bottom right: Axial view of bottom left scene.

3.4.5 Other matter fields

This example implementation can represent other matter fields that occur often enough in real scenes to merit representability in QR:

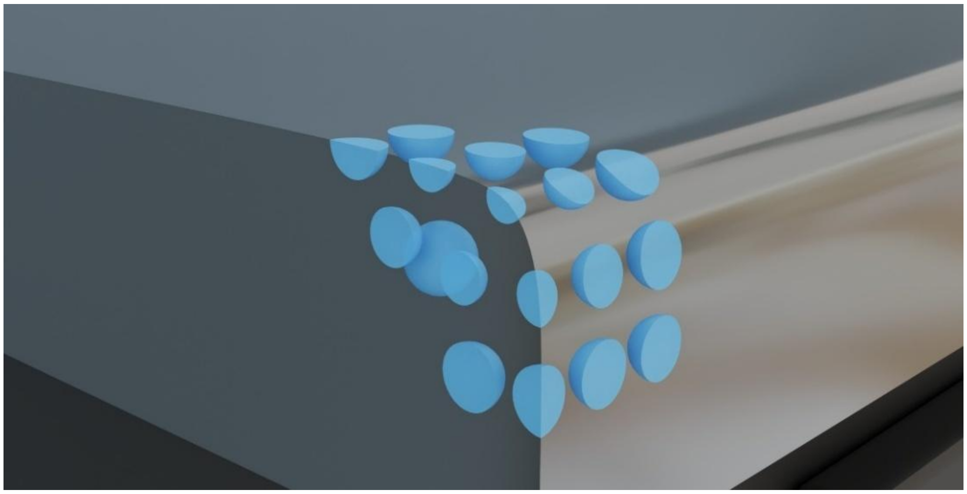

Figure 3.4.5-1: Corner of countertop with one rounded edge

Figure 3.4.5-2: 4 elsies in a treetop ~~with direct~~ – Direct sunlight and indirect other daylight ~~filtering~~filter through the treetop in various proportions.

3.4.6 Light fields

This example implementation can also represent light fields (using luminels) independent of media:

Figure 3.4.6-1: 4 small regions of exitant light ~~each~~ – Each region is representable as ~~an~~a single elsie (4 elsies total~~) –~~). Top row: With radial falloff in exitant radiance (e.g., using spatial variation function for light field parameters). Bottom row: Constant radiance with sharp / abrupt transition to darkness (e.g., using SV function and/or influence function).

4 Scene Model

4.1 Overview

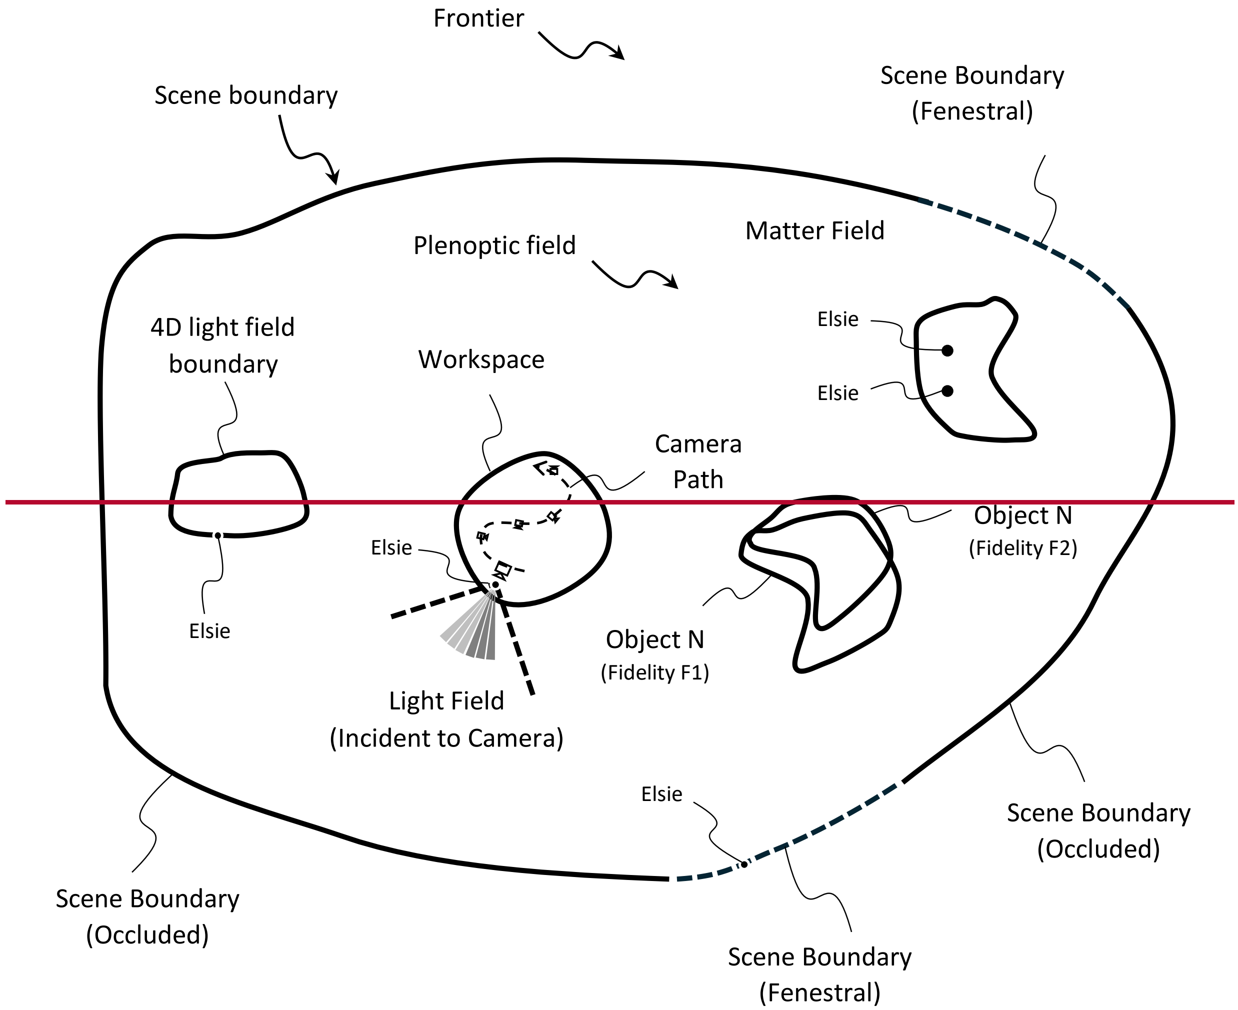

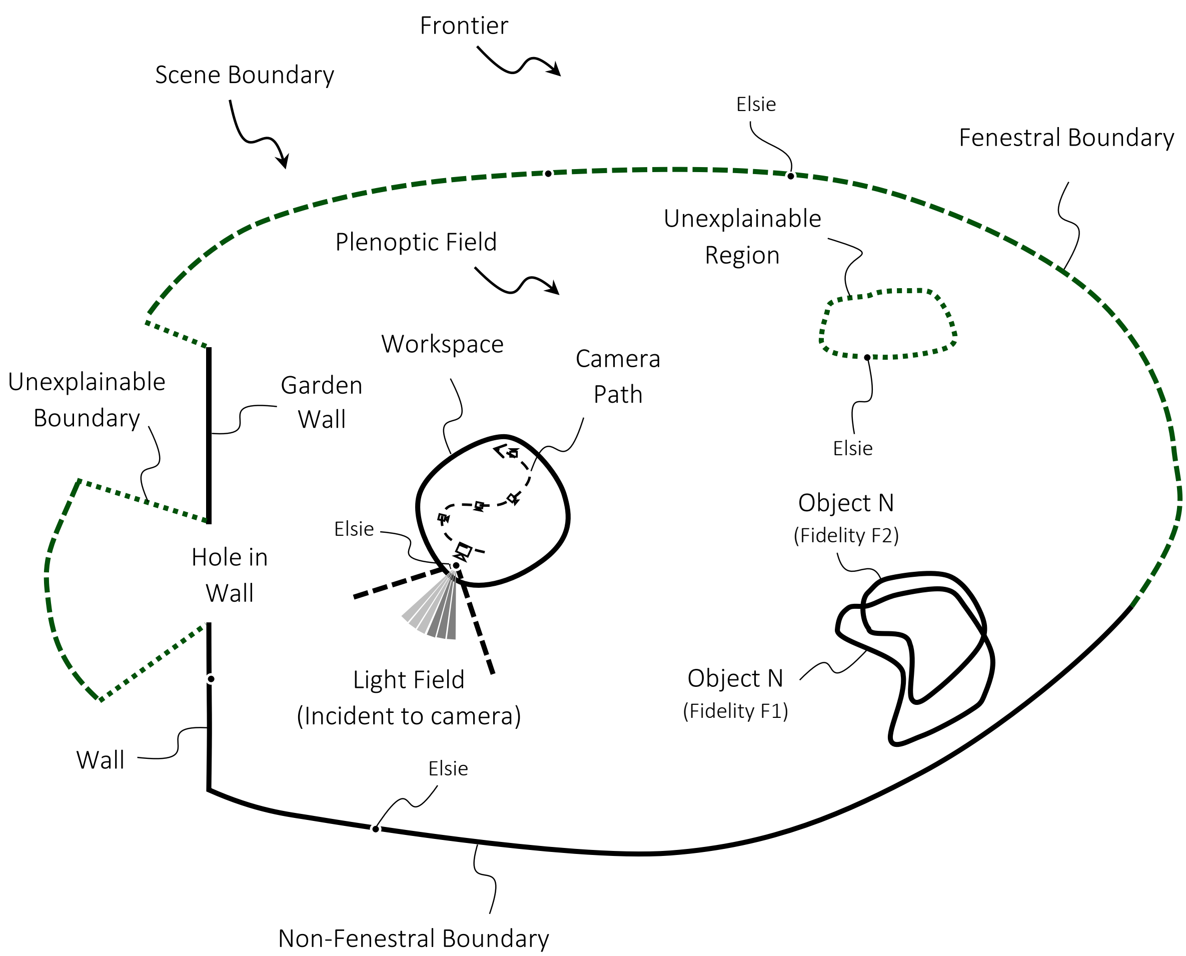

Figure 4.1-1: Scene model – Overhead view of an example scene model containing its (always single) plenoptic field along with other scene model entities that are not part of the plenoptic field.

4.2 Key attributes

Key attributes of scene models:

- The scene model is enclosed by an outer scene boundary and may also contain inner boundaries.
- The outer scene boundary is a closed 2D manifold separating the plenoptic field from the frontier.
- The scene model's plenoptic field lies between the outer scene boundary and any inner scene boundaries.
- A scene model contains exactly one plenoptic field, which comprises elsies.
- Neither matter nor light are represented outside the outer boundary nor inside any closed inner boundaries.

4.3 Scene boundaries

~~A~~In every QR scene, a closed scene boundary lies at the farthest extent of elsies~~ in a scene~~. This section presents an example set of scene boundary configurations encountered in real-world scenes.

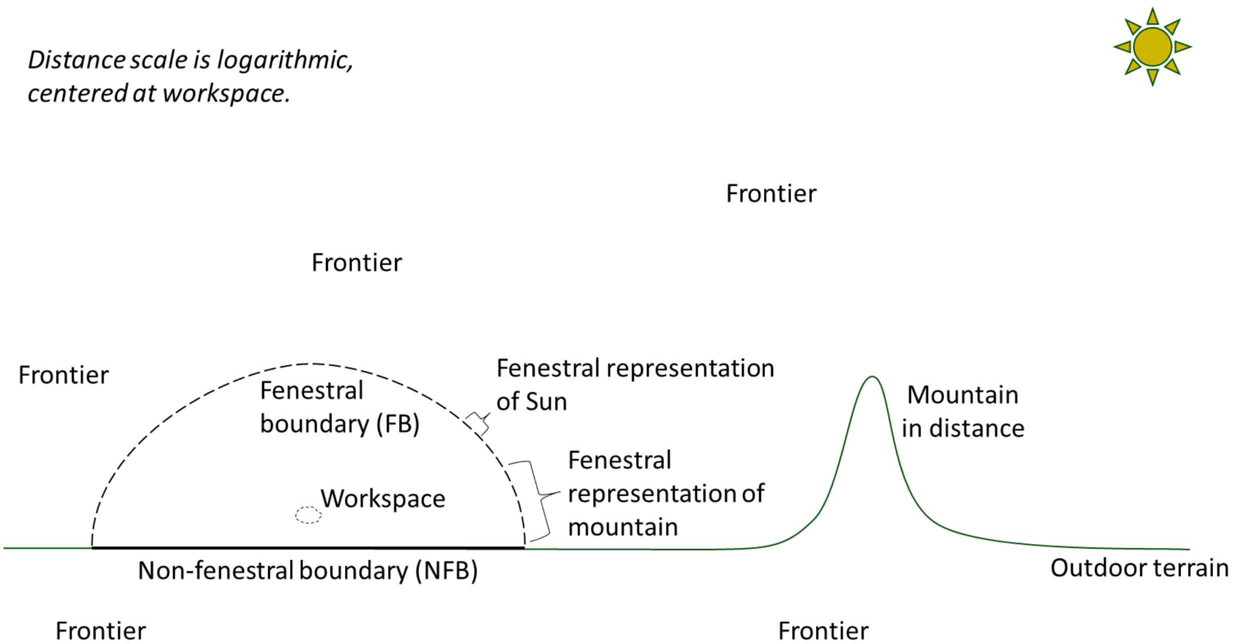

Figure 4.3-1: Outdoor scene model (mature reconstruction) with observations in all directions[27] – In this example scene model at a mature stage of reconstruction, a fenestral boundary (FB) spans approximately the upper hemisphere (2π steradians) of directions, relative to the workspace, in which light and matter can flow through the air (the default medium in this scene). A non-fenestral boundary (NFB) lies along the ground because if blocks the flow of light and matter.

[27] Combined FOV of the workspace (union of FOV-per-image over all viewpoints) is omnidirectional, i.e., covers 4π steradians.

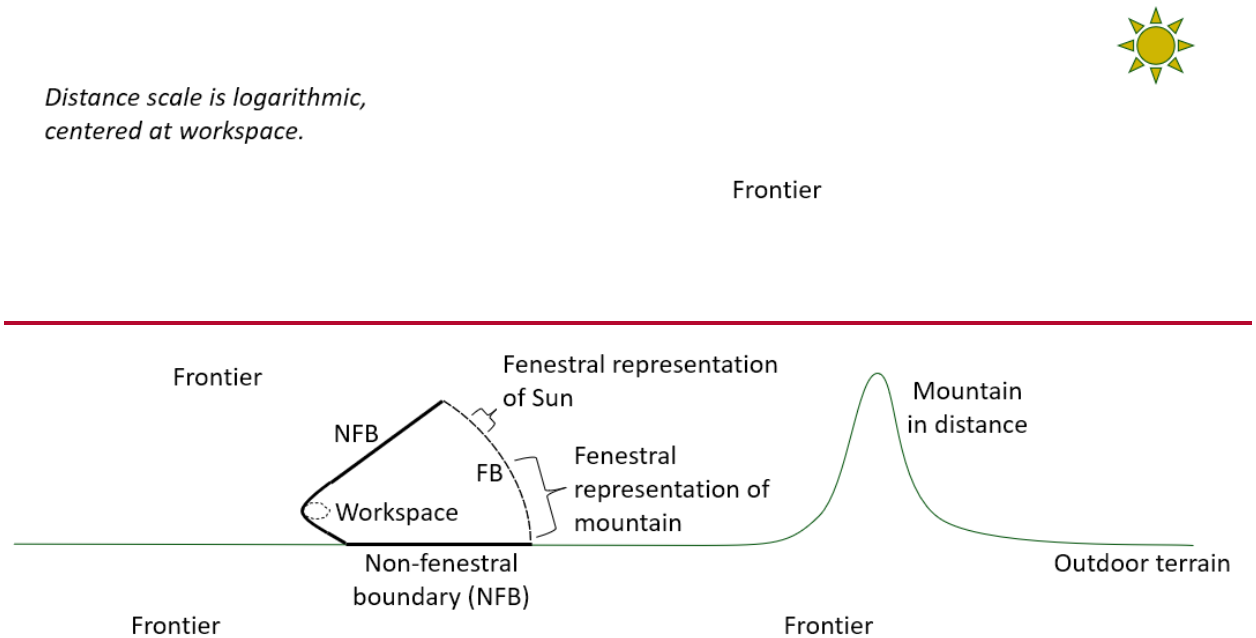

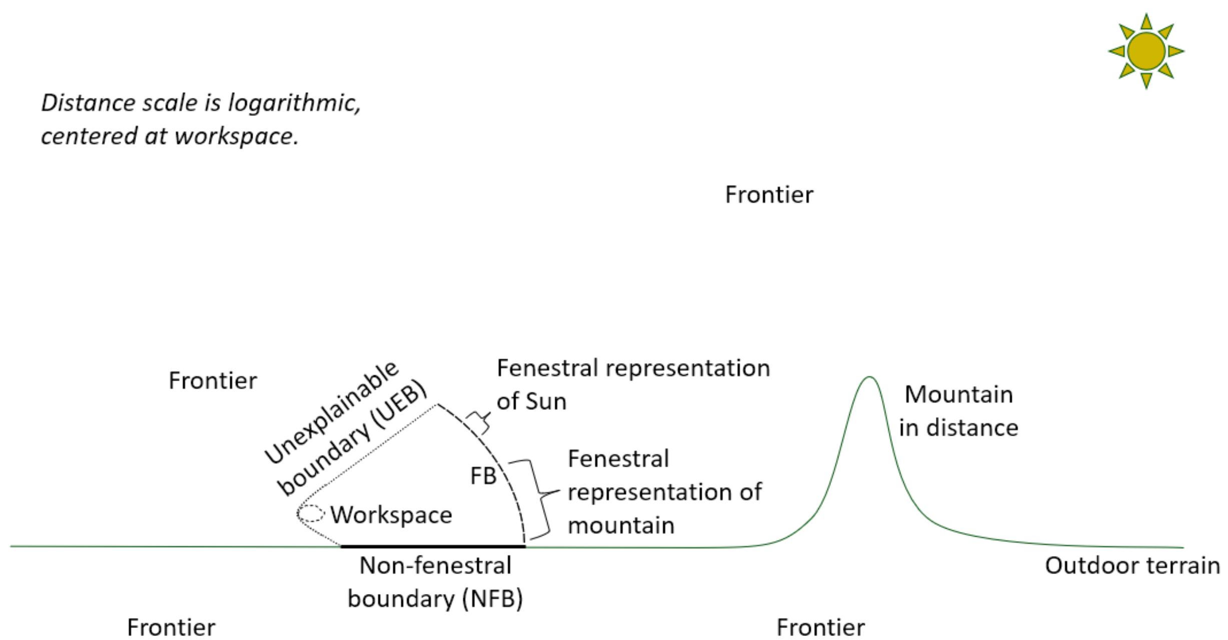

Figure 4.3-2: Outdoor scene model (mature reconstruction) with observations in a limited range of directions[28] – In this example scene model at a mature stage of reconstruction, the boundary of the workspace FOV is ~~non-fenestral~~unexplainable because QR ~~(~~in this ~~particular~~ case~~) doesn't~~ hasn't made a (sufficiently strong) determination about whether to represent ~~any~~the flow of light ~~or~~& matter vs blockage to such flow~~ across it~~. See figure 4.3-4 for a case where ~~an NFB~~a UEB is punctured by a corridor leading to a second FB.

[28] Combined FOV of the workspace isn't omnidirectional, and in this example is << 4π steradians.

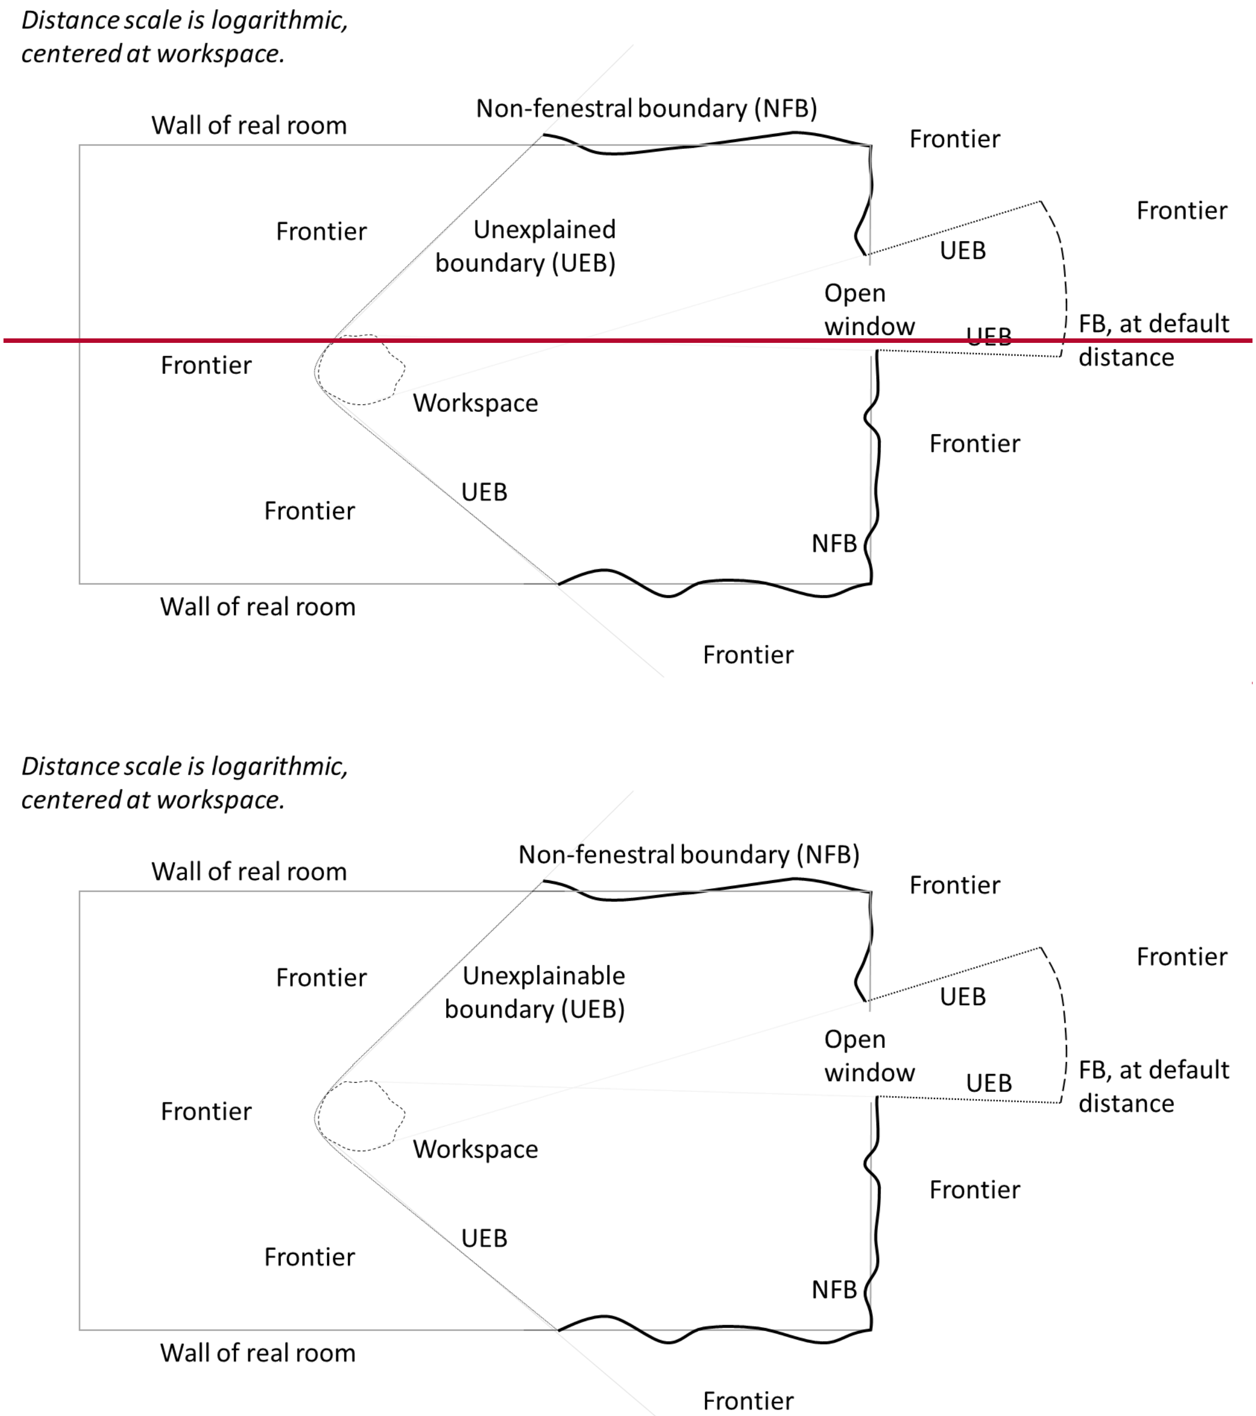

Figure 4.3-3: Indoor/outdoor scene model early in reconstruction processing (and where workspace FOV << 4π steradians) – In this example scene model at an early stage of reconstruction, one part of the scene boundary coincides with ("follows") the workspace FOV boundary and remains "~~unexplained~~unexplainable" (UEB) because QR hasn't yet determined, beyond a specified fidelity floor, whether to represent the flow of light (especially incident light) and/or matter through the boundary. The single FB is at a default middle-of-the-road distance

away from the workspace. NFBs approximately follow the observed parts of the room walls, which are represented to lower fidelity at this early stage.

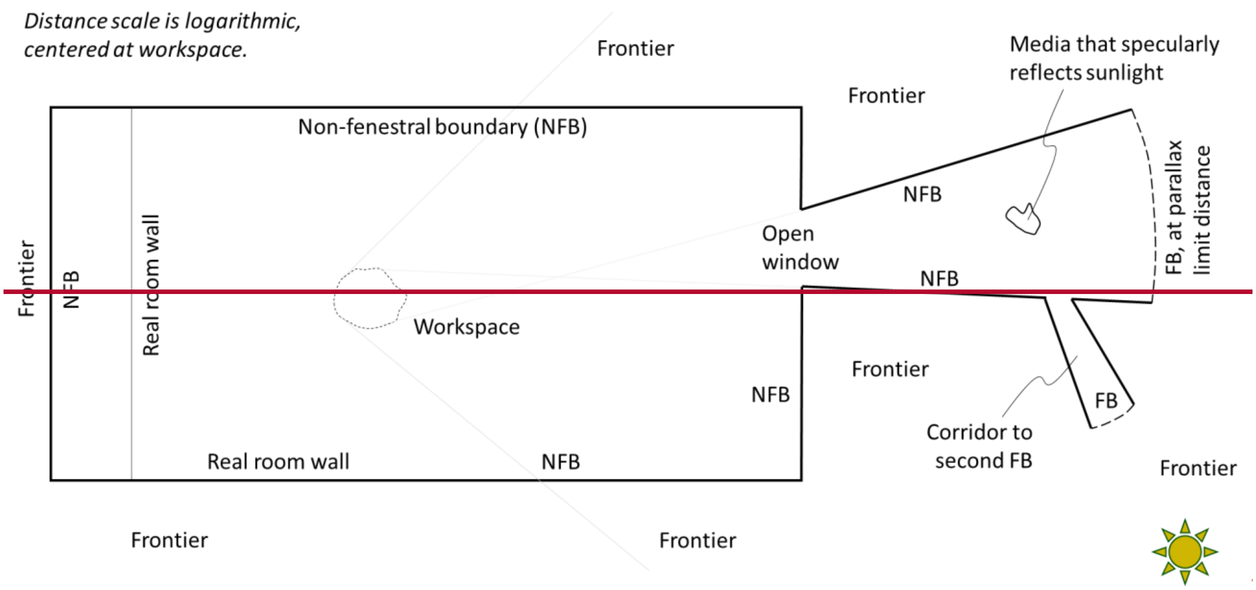

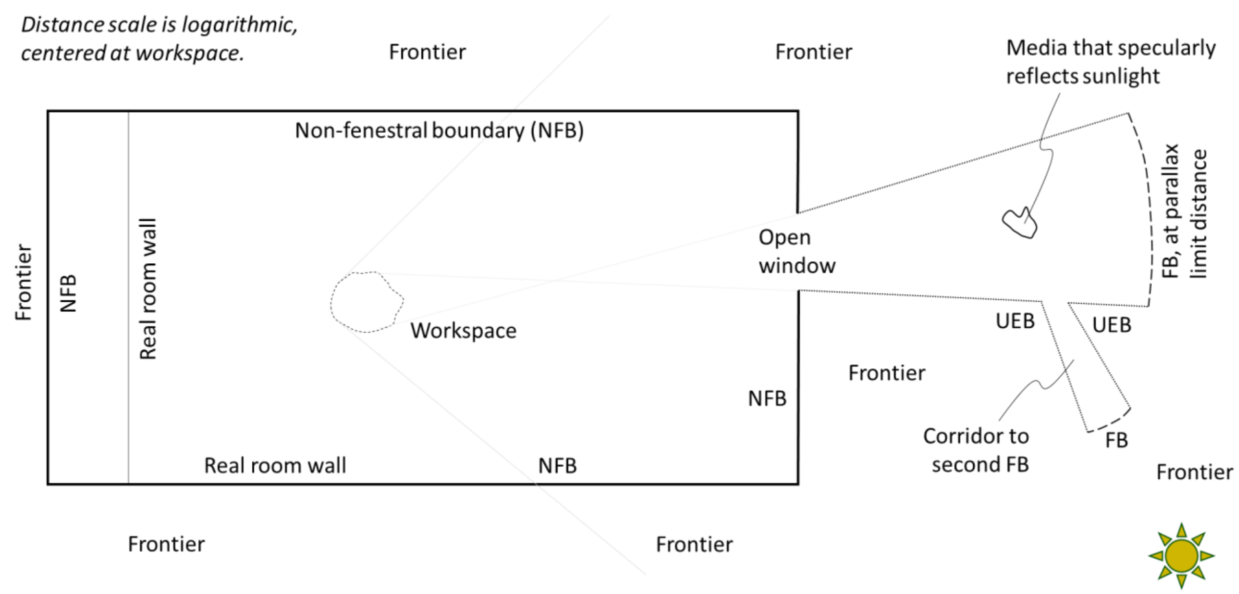

Figure 4.3-4: Indoor/outdoor scene model late in reconstruction processing (reconstruction is mature, and where workspace FOV << 4π steradians) – This figure shows a later stage of processing the working scene model of figure 4.3-3. The NFB that follows the room walls has been refined to high fidelity in the observed part of the room. In the unobserved[29] part of the room, QR has extrapolated the NFB to follow the 2 longer walls of the room using the apriori assumption that the room has a prismatic shape with 90° angles. Without direct observation of

[29] Not directly observed from the workspace

the rear shorter wall, however, QR places the corresponding part of the room NFB beyond the real wall.[30] Through the open window, a region of specular media in the meso-field has been reconstructed to a medium target fidelity that includes specular BLIF modeling. Captured workspace images contain observations of sunlight reflected off the specular media. QR has access to an (apriori) model that predicts solar radiance given the capture location, date, time of day, atmospheric conditions, etc. By combining the reflected sunlight observations with the solar radiance model, QR reaches suitably low uncertainty in the Sun's unoccluded presence and radiance parameters. A new corridor of default media is then created that punctures one of the ~~NFBs~~UEBs and ends at a second FB where the relevant sunlight is represented in elsies with low power.

4.4 Fenestral boundaries

Further prominent attributes of fenestral boundaries:

- A light field represented at a fenestral boundary is "2D" in nature.
- Fenestral boundaries are often ephemeral. During maintenance of a hosted[31] scene model (LSM or XSM), earlier FBs will tend to be garbage-collected out of existence. And then new FBs will be created when new subscenes are extracted for dispatched refinement or transmission to application software.[32]
- If the scene is dynamic, the fenestral boundary will generally consist of a plurality of layers (or equivalent) that allow light exiting from objects in the frontier (e.g., an aircraft, the sun, the moon) to move.
- Notes on fenestral elsies[33]:
  - In a standalone / isolated subscene, fenestral elsies represent light flowing from the frontier inward into the subscene's meso-field. That inward-flowing false emissive light conveys information into the subscene that represents what's happening outside the subscene.
  - In a standalone / isolated subscene, fenestral elsies also represent light flowing from the subscene's meso-field out into the frontier. That outward-flowing light is part of the energy balance of the scene that enables net true emission of light from a scene to be calculated.
  - If the subscene later gets merged into (not merely reconciled against) its enclosing superscene, a disentangling force during refinement will minimize that false emissive light by discovering that it can be explained as incident light that exited other elsies in the superscene.
  - When a subscene's camera workspace(s) expand, elsies that were previously fenestral can be "pulled into" the subscene's plenoptic field and come to represent low-powered 4D elements (e.g., have some occlusion and some false emission).
- In a hosted scene, fenestral LFs have a timestamp and thereafter might or might not be used often.

[30] In this example scene, the room is empty and the walls have a matte finish. If there were a mirror or other highly specular reflectors in the observed part of the room, then (depending on target fidelities)
[31] Stored and manipulated at the host / CPU layer.
[32] Fidelity varies regionally across each FB according to Smart Fidelity logic.
[33] These notes apply when the FB implementation uses elsies.

4.5 Further attributes of scene model entities

Prominent further attributes of entities in scene models:

- In a QR scene model, a feature is typically realized as a set of one or more scene entity[34] parameters, along with each parameter's value, that have some meaning or utility to users of QR processing output or internally within QR. Example features include a dot, a scratch, a glint of light, and a building. Features come in various types depending on the parameters involved, including parameters of light field radiels. For example, an identifiable SIFT feature in an image (in the sense of conventional 2D computer vision) can be realized as a pattern of radiels at a focal plane (pixels in an image).
- In a QR scene model, a "seg" (segment) is typically realized as a set of elsies that share a sufficient degree of similarity in the value of one or more (groups of) parameters (e.g., macrogeometry, BLIF). A given elsie can fall into (be associated with) multiple segments. Segments exist in hierarchical "segtrees" (segment trees), and a scene model can have multiple segtrees. Each segtree arises from a segmentation process that operates on one or more features of interest.[35] Example features of interest for segmenting a car scene: geometric continuity / curvature, "base color" of BLIF, degree of transmissivity (glass vs opaque media), 2D (optionally multiview) NN-based classification by car panel / part type.
- In a QR scene model, an object is typically realized as a set of elsies, segments, and/or analytic entities in a spatial (and optionally temporal) arrangement that humans would recognize as an "object". Example objects include a kitchen table, a glass window, and a tree.
- At positions along a camera path, observed radiels may be grouped together to represent images, typically at a focal plane of pixels, where each pixel stores the radiance value(s) of the corresponding radiel.
- One or more scene graphs may point to entities in the plenoptic field. A scene graph may also point to analytic entities not currently manifested in the plenoptic field. A scene graph is arranged into a hierarchy of nodes defining the relationships, spatial or otherwise, between the referenced entities.

[34] Including elsies, segments, and analytic entities
[35] Each feature of interest is represented as one or more parameters with its assigned value.

4.6 Quidient Reality Model (QRM)

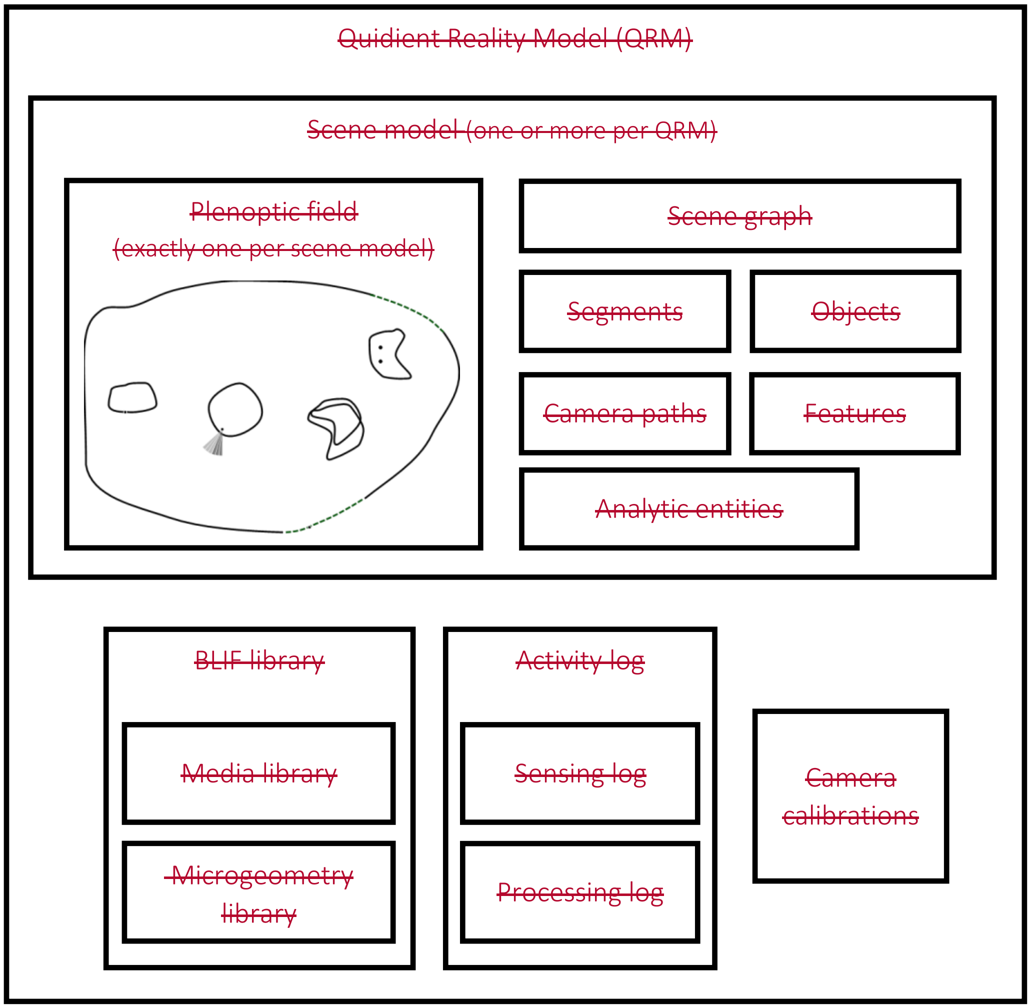

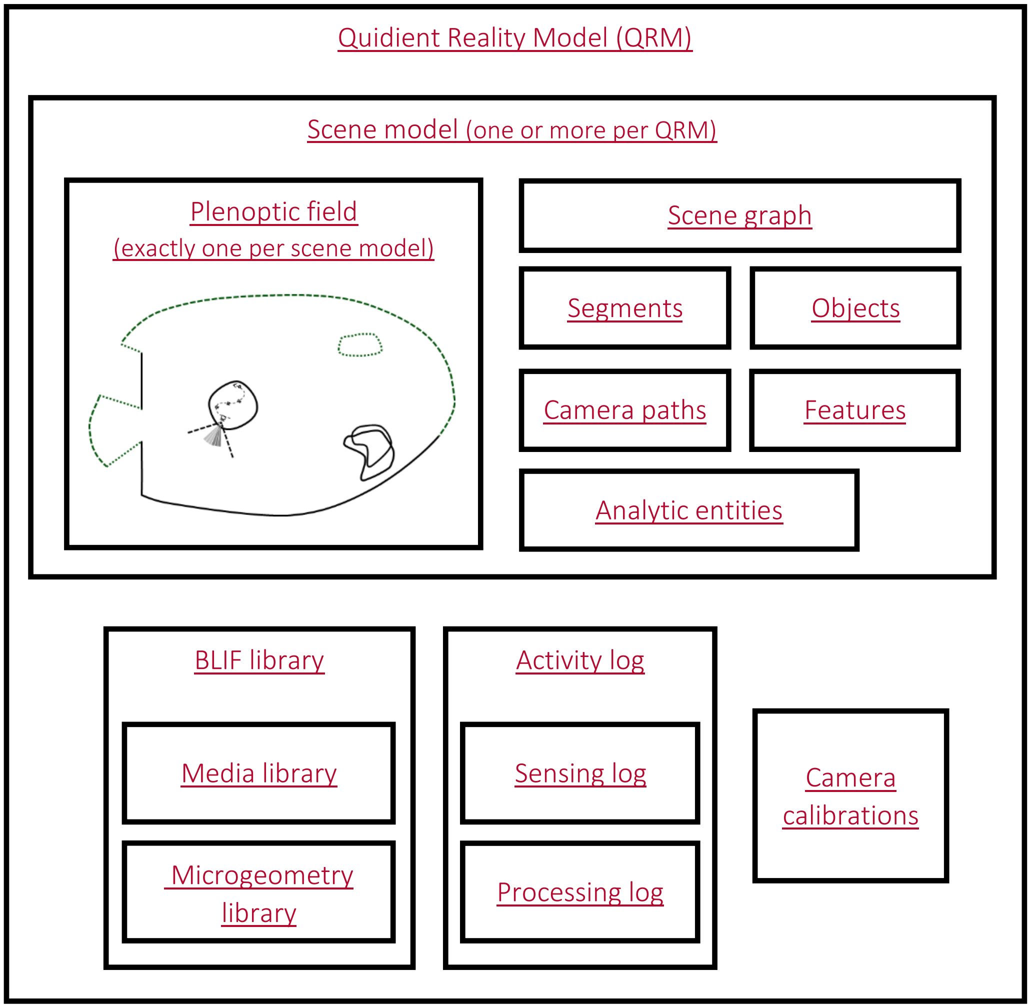

Figure 4.6-1: Quidient Reality Model (QRM) – Block diagram of entities in a Quidient Reality Model

Entities that can exist (be contained, data-wise) in a Quidient Reality Model, but not in any of the scene models the QRM contains, include the following:

- BLIF library – BLIFs, typically pertaining to one or more of the QRM's scene models, are stored in a library of BLIFs. Within the BLIF library are a library of media (types) and a library of microgeometry configurations.
    - Each entry in the media library represents a medium (type of media) that can interact with light. Example media library entries include air, water, fog, chrome, red automotive paint, and soda-lime glass.

- Each entry in the microgeometry library represents the microgeometry[36] configuration into which one or more media may be arranged. A prominent example of a microgeometry library entry is the distribution of microfacets of a surfel, described for example by the GGX[37] model.
- Each BLIF (library entry) in the BLIF library is defined by one or more media from the media library and one or more microgeometries from the microgeometry library.

- Activity log – Scene capture actions are logged in a sensing log. Scene processing actions, including reconstruction operations, are logged in a processing log. Among other purposes, these logs enable "forensic" tracing of various levels of processed scene information, including all the way back to the original sensor observations, depending on the specified level of rigor (verbosity) of the log. This log information also enables reinvocation of scene processing at various stages using alternative goals, settings, observations, prior models, etc.
- Camera calibrations – A scene database stores camera calibrations, including compensation parameters and other data related to calibration of physical or virtual cameras used in imaging, display, and other analysis operations on a scene model. Lens distortion parameters, flat field, and optional polarimetric factors are examples of such calibration information.

4.7 Example implementation

This section describes one conceptual way to implement selected aspects of scene models with the key attributes listed in §4.2.

4.7.1 Segmentation

Here we show examples of plenoptic field segmentation. (Note that in these particular examples, elsies are shown on a regularly spaced grid, which is a valid plenoptic field configuration but not a general requirement of QR scene models.)

In the following 2 figures, a faceted glass tower has been segmented on features of geometry (sensitive to the tower's faceting) and features of materiality (sensitive to types of media):

[36] "Microgeometry" here means (a description of) geometry present at a spatial scale much finer than the matter field (elsie) resolution being considered.
[37] https://cs.cornell.edu/~srm/publications/EGSR07-btdf.pdf

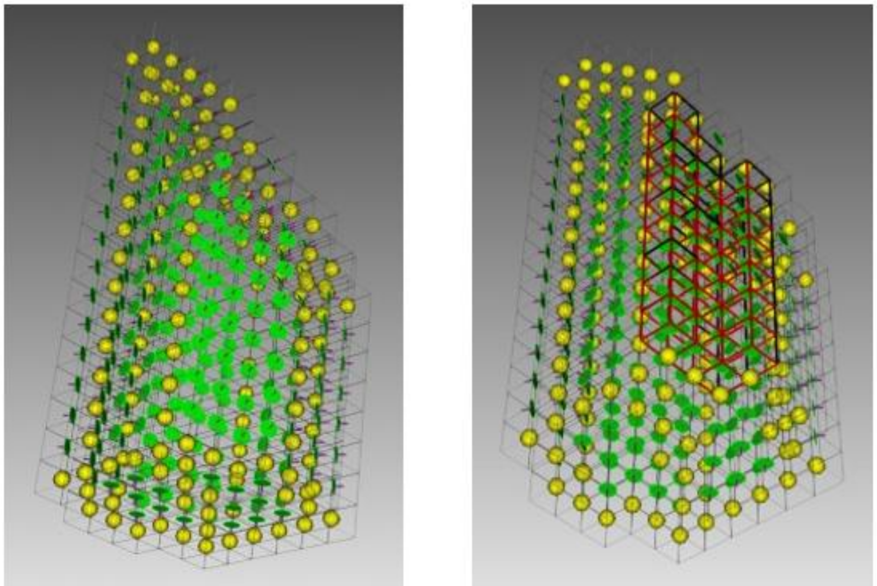

Figure 4.7.1-1: Analytic views of segmentation of a faceted glass tower ~~virtualized by QR – (~~– Left~~)~~: Surface elsies. Green discs symbolize smooth surface elsies (“simple surfels”). Yellow spheres symbolize faceted edge & corner elsies. ~~(~~Right~~)~~: One segment of surface elsies on a vertical face of the tower is highlighted.

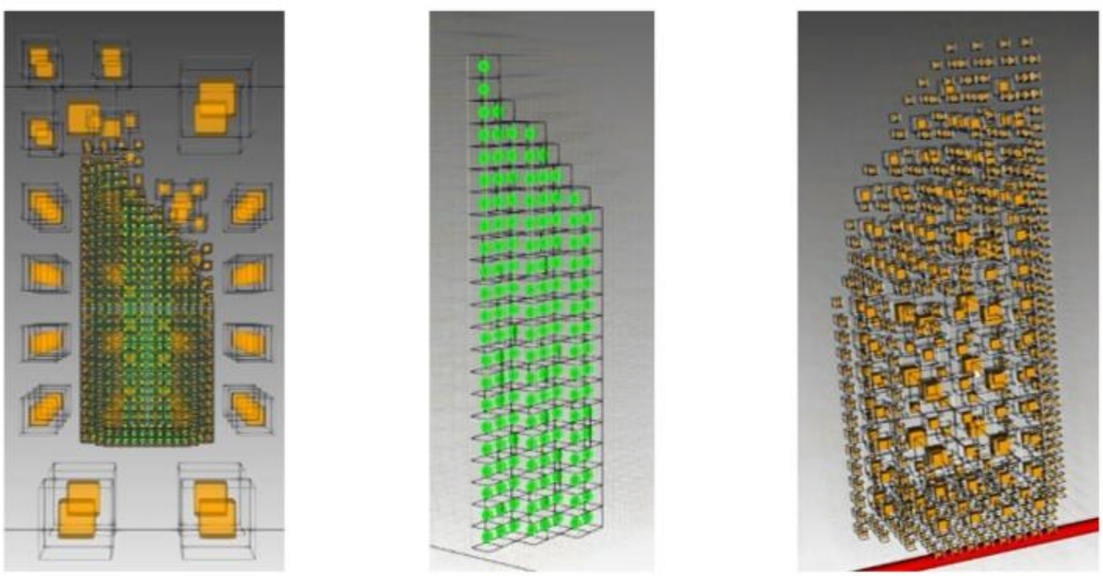

Figure 4.7.1-2: Further analytic views of segmentation of a faceted glass tower ~~virtualized by QR – (~~– Left~~)~~: Elsies of air (homogeneous default medium) symbolized as orange cubes. (Middle) One vertical face segment shown in isolation. ~~(~~Right~~)~~: Interior glass elsies (homogeneous non-default medium) shown in isolation.

In the following figure, a small dent on a car hood has been segmented apart from the surrounding surface by virtue of violating the expected “Class A” geometric smoothness / curvature condition of the hood:

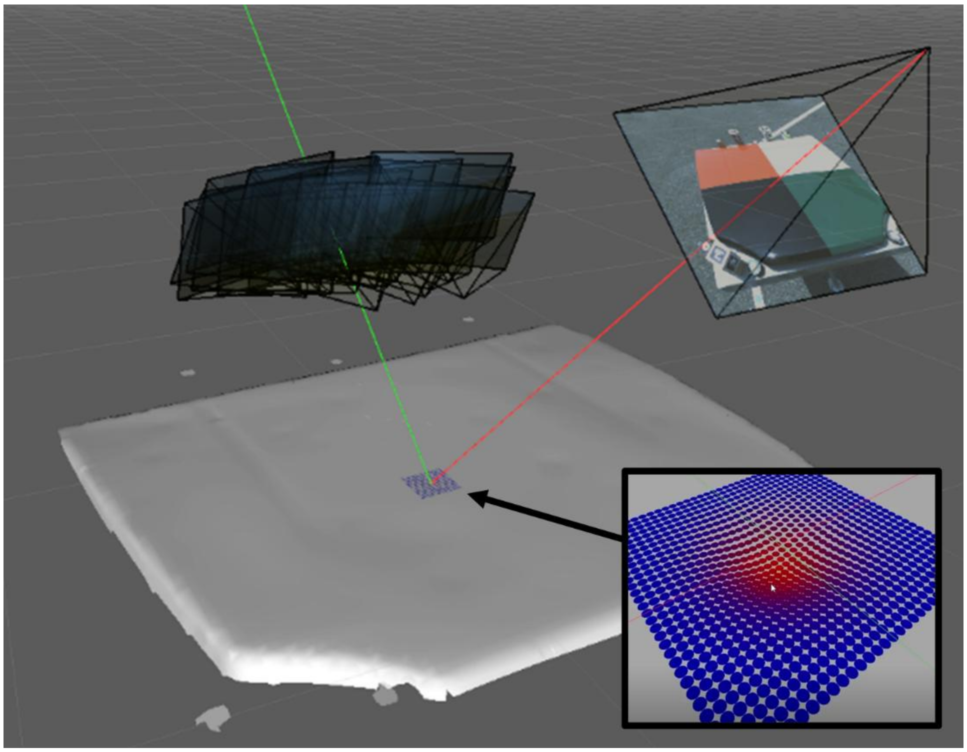

Figure 4.7.1-3: Analytic view of curvature-driven segmentation of a dented car hood

4.7.2 Analytic entities

A scene model can contain analytic entities that coexist with the plenoptic field but are not contained within it. Example of analytic entities are CAD-like spheres, cones, etc. that might have been fit to (or otherwise associated with) local groups of elsies. An analytic entity needn’t be associated with elsies nor regard interaction between light and matter. Another example of an analytic entity is a CV-style SIFT feature formed by a pattern of pixels in an image, without any necessary relation to radiels in the plenoptic field.

Here we show examples of analytic geometric entities (“analytic geometries”) that happen to be associated with elsies in the plenoptic field of their scene model:

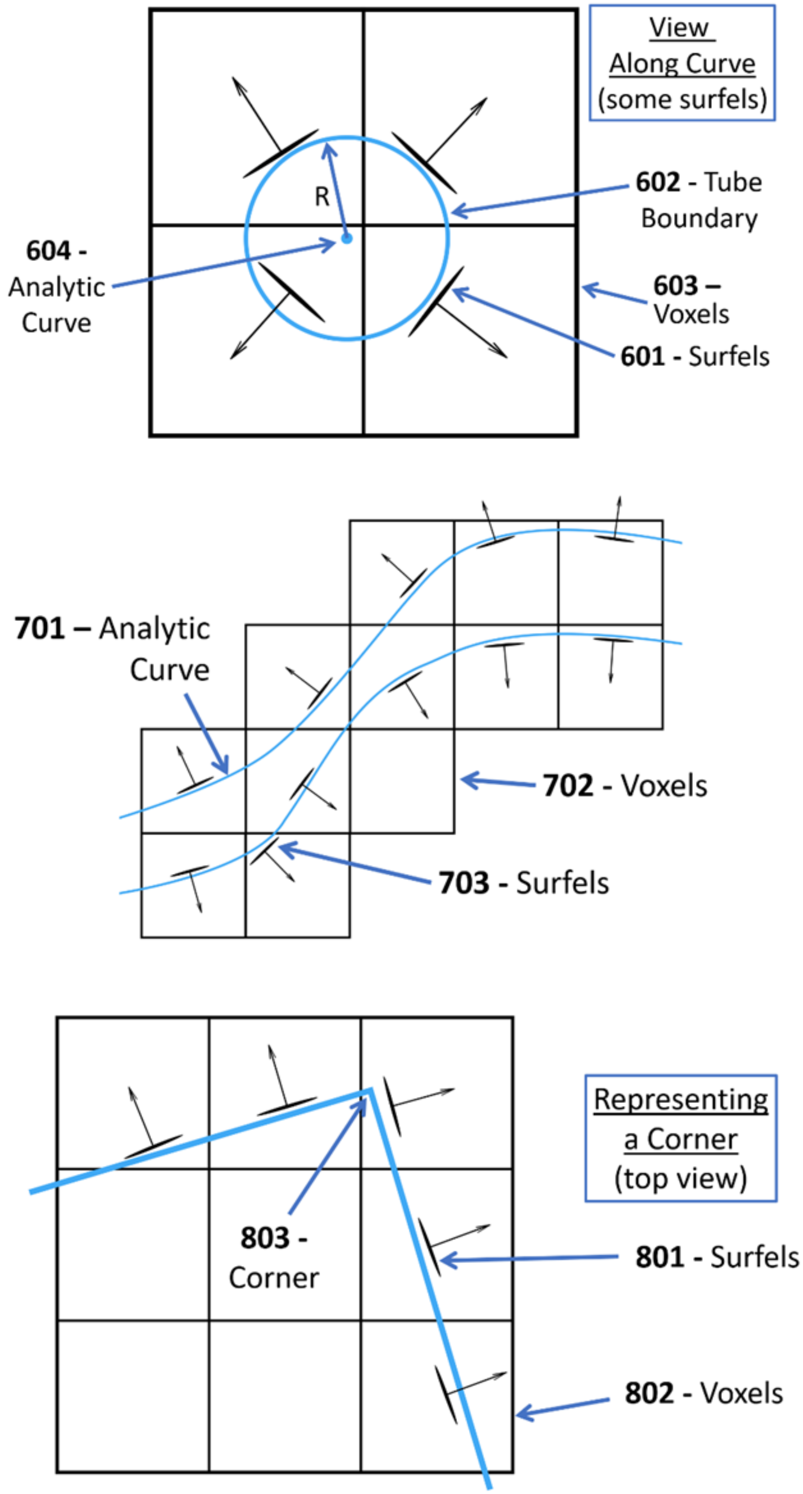

Figure 4.7.2-1: Examples of analytic geometries (geometric entities) – The voxel grids shown are of secondary importance.

5 Smart Fidelity

5.1 Overview

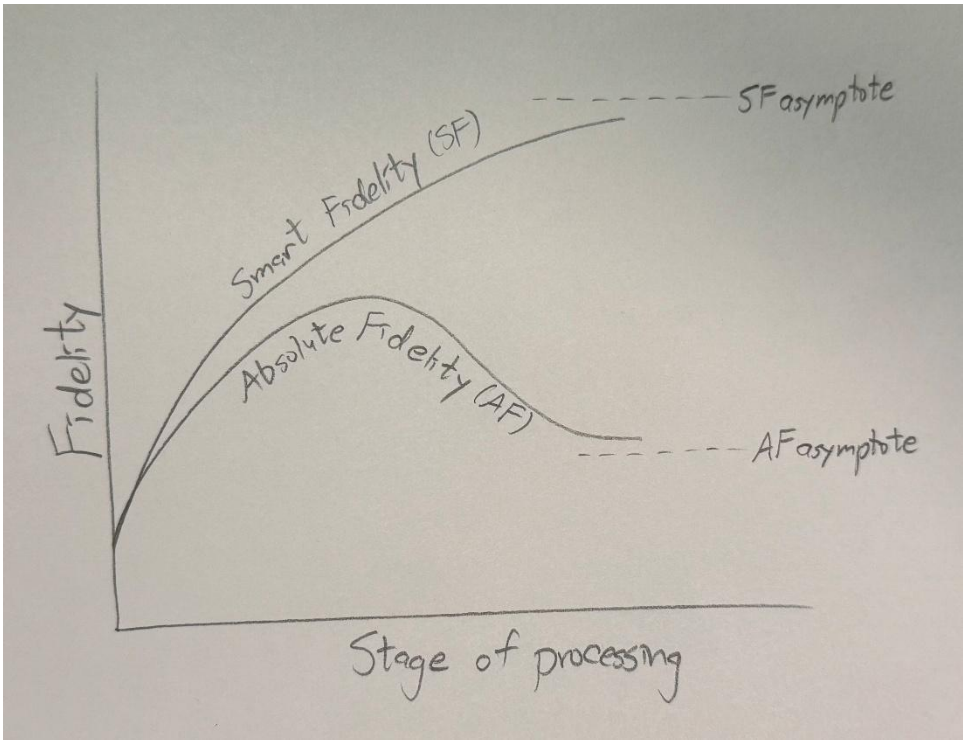

Figure 5.1-1: Smart Fidelity & ~~absolute fidelity~~Absolute Fidelity as processing proceeds through stages

5.2 Key attributes

Key attributes of Smart Fidelity processing:

- "Fidelity" has 3 orthogonal (independent) dimensions:
    - Resolution
    - Uncertainty[38, 39]
    - Power (representational power)
        - Power closely determines a scene region's degree of relightability, its utility for geometric measurement ("dimensioning"), and its utility for quantifying materiality.
        - For elsies, power pertains to the elsies' representation of (any or all of) matter, light, space, and time characteristics.

[38] When no ground truth is available, uncertainty here can be calculated by propagating uncertainty values from parts of the scene model whose parameters connect (usually via light transport) to the given region's parameters.
[39] In engineering-level communication within Quidient, we generally use the term "uncertainty" rather than "accuracy". The more loosely defined term "accuracy" is sometimes used in communications to external stakeholders. (Internally, "accuracy" could be appropriate to use when we have a trusted "10x" GT in the parameters of interest.)

- QR processing (reconstruction is one very important kind of processing) proceeds in a heterogeneous "waterfall" of stages where different postulates[40] are evaluated at each stage.
    - The waterfall of stages will look different in different scene regions and depending on current processing goals.
- Early stages will typically use primitives with lower power than later stages.
- Later stages of processing will in general simplify scene regions to be represented more parsimoniously[41] (i.e., trade decrease in ~~absolute fidelity~~Absolute Fidelity for increase in Smart Fidelity).
    - QR customers (developers) can implement simplification to customized representations for their use cases.
    - The final processing stages of non-customized core QR will end at elsies and other primitive entities documented in this CDD.
- Processing typically proceeds in these 2 broad (super-)stages:
    1. Add more ~~absolute fidelity~~Absolute Fidelity (AF) until reaching the per-region[42] target fidelity specified in the current processing goal.
    2. Simplify[43] scene regions to lower AF while SF keeps rising.
- Target Smart Fidelity is different than target ~~absolute fidelity~~Absolute Fidelity and tends to be more useful, especially in “granddad” cases.
- Typical drivers of target fidelity include the following:
    - Desired geometric accuracy
    - Desired perceptual quality[44]
    - Desired relightability
- In Smart Fidelity processing, automated QR logic sets a target resolution, uncertainty, & representational power of each region of a scene before & during processing. The assigned fidelity is tailored to the processing goal at hand.
    - For example, in the “Grandma's Mustang” case, use of QR in a car insurance app, plus heavy capture attention[45] paid to the dent ROI, trigger the lowering of target uncertainty in surface ~~orientation~~direction (normal vectors) vs other regions of the car surface.
- Absolute ~~fidelity~~Fidelity typically decreases outward from the center of a scene’s near-field(s).
- ROIs[46] can and will exist in the meso-field.
    - For example, a car that is several camera workspace diameters away from the current workspace receives fine pixel resolution & substantial parallax using a telephoto zoom setting during video capture.
- ROIs can and will exist in the far-field.

---

[40] Postulated values for sets of parameters of elsies in the plenoptic field, as well as for any other scene entities outside the plenoptic field

[41] Using fewer primitive entities, with each entity usually having higher representational power

[42] Objects (including OOIs) and matter field segments are particular types of scene region.

[43] Includes merging (coalescing) elsies. In this broad stage, QR typically lowers (absolute) resolution by using elsies that have higher power but are fewer in number. In all cases, QR honors any explicitly specified target resolution, target uncertainty, and target power.

[44] Note that resolution tends to have outsize importance in the perception of most humans, as compared to rigorous quantitative radiometric uncertainty.

[45] “Attention” here includes intention(ality), not unwitting / perchance lingering of a camera FOV on a particular scene region.

[46] Regions of Interest have higher fidelity than adjacent scene regions.

- One common type of far-field ROI comprises an array of neighboring luminels[47] with some combination of finer directional resolution and lower radiometric uncertainty.
- For example, the Moon having received heavy capture attention vs other regions of the sky.

- The longstanding Quidient terms “tunnel of resolution”[48] and “island of accuracy” are each a (sub)dimension of the composite higher fidelity at an ROI compared to adjacent regions.

5.3 Fidelity variation over scene regions

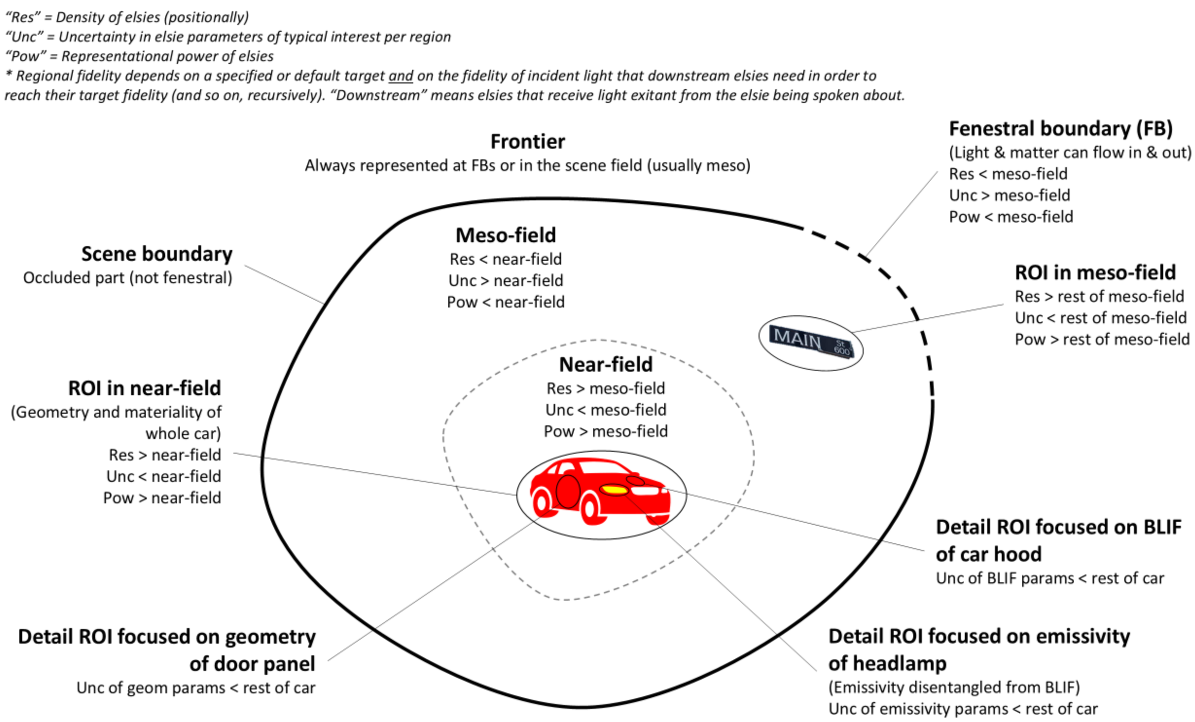

Figure 5.3-1: Typical fidelity in different regions of a mature scene model – Note that this concentrically layered “onion” configuration is one among many scene model configurations that occur in real-world QR cases (e.g., “ant colony” of chamber-like subscenes connected by elongated corridors).

5.4 Tunnels of resolution

Prominent attributes of the “tunnel of resolution” dimension of Smart Fidelity:

- A tunnel of resolution is a region of a plenoptic field where elsie resolution is spatially and/or temporally much finer than in adjacent regions. Spatial resolution here pertains to position (e.g., elsie origin) and/or ~~direction~~orientation (e.g., elsie orientation, normal vector, radiels in luminels).

[47] In this case, each luminel will be the predominant component of an elsie.
[48] For an especially deep tunnel of resolution (e.g., beyond 4 OOM), QR can extract a separate subscene centered at the tunnel region.

- A tunnel of resolution can extend arbitrarily far in both the coarser and finer directions of resolution.
- Tunnels of resolution are created and modified on demand as determined by the particular scene processing operations at hand.
- Example of a tunnel of resolution in a matter field – In the middle of a large wall of uniform color, QR represents a small dot as elsies that are subdivided much more finely than elsies of the surrounding wall where material variation is much lower frequency.
- Example of a tunnel of resolution in a light field – The disc of direct sunlight is surrounded by a large region of cloudless sky. A chrome motorcycle helmet receives incident light from that sky region. Radiels directed toward the boundary (transition region) between the Sun disc and its immediately adjacent sky are subdivided to quite fine (directionally narrow) resolution. Radiels in directions other than the Sun are much coarser (directionally wider).

5.5 Example implementation

This section describes one conceptual way to implement selected aspects of Smart Fidelity processing with the key attributes listed in §5.2.

5.5.1 Stages of processing

Notes on elsie parameter interpolation in this example implementation:

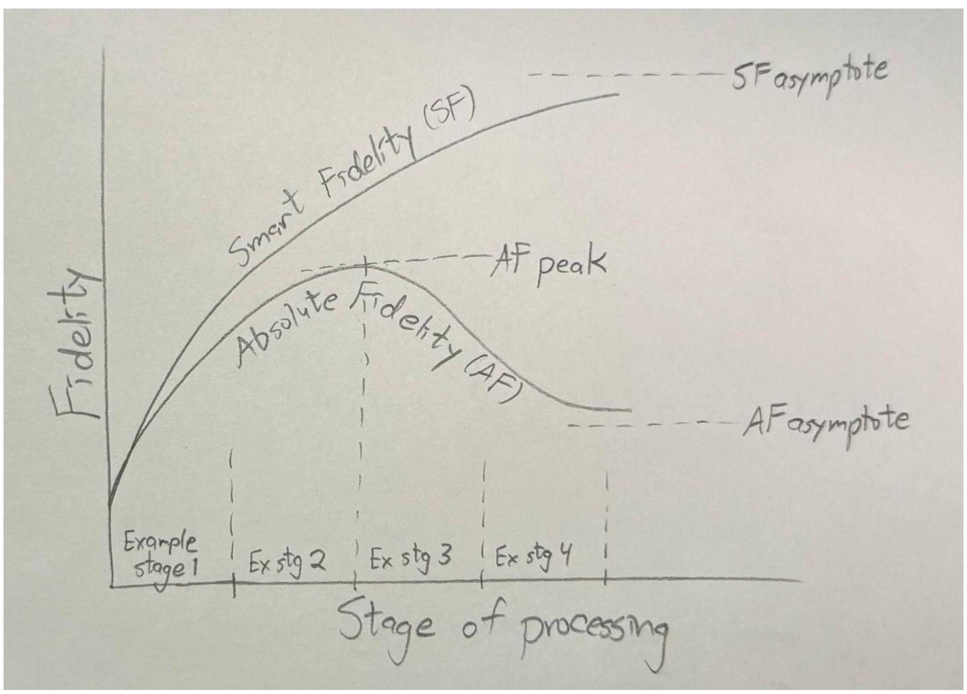

Figure 5.5.1-1: Example stages of Smart Fidelity processing

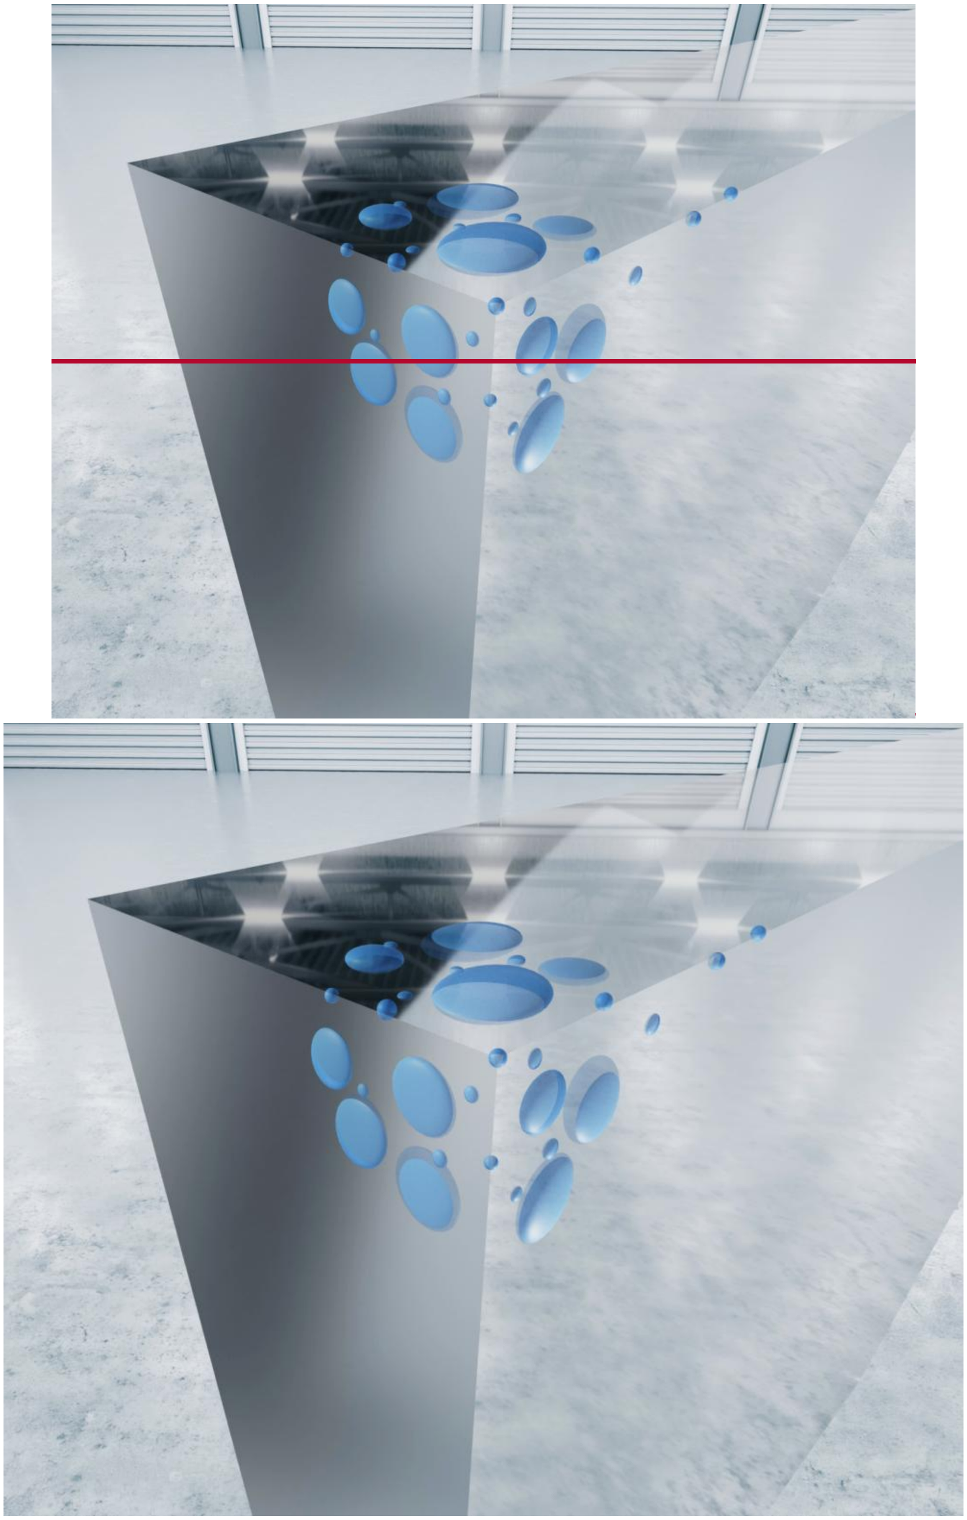

Figure 5.5.1-2: End of stage 1 – Example working scene model of a slab of glass (~~showing~~ only ~~the~~ nearest corner is shown)

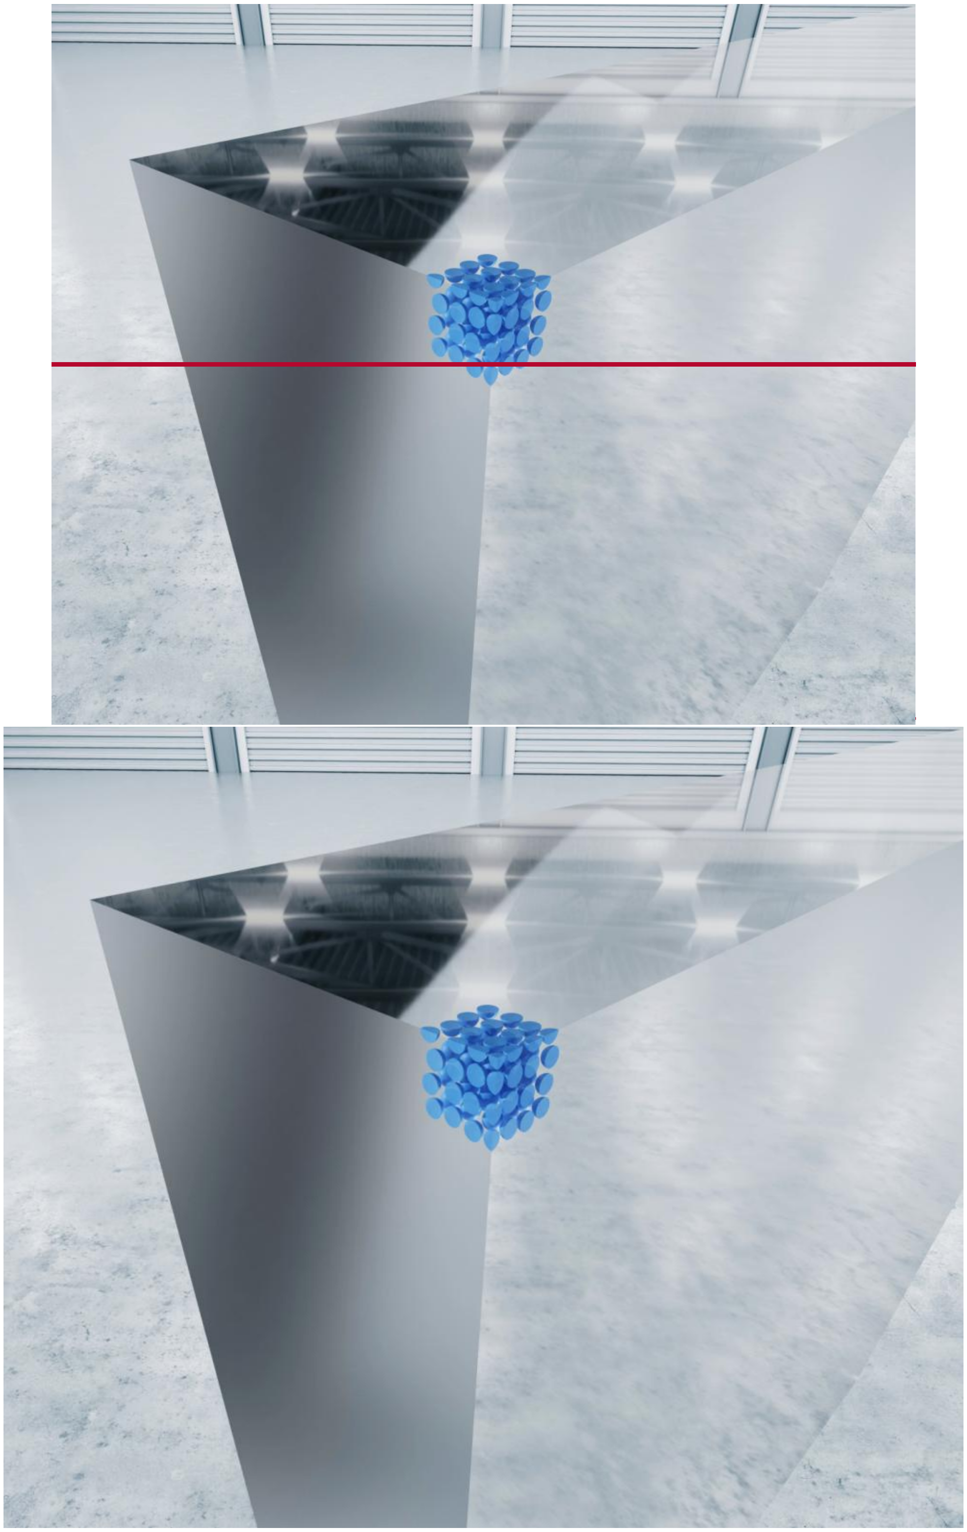

Figure 5.5.1-3: End of stage 2 – Example working scene model of ~~a~~the slab of glass ~~(showing~~ (only ~~the~~ nearest corner is shown)

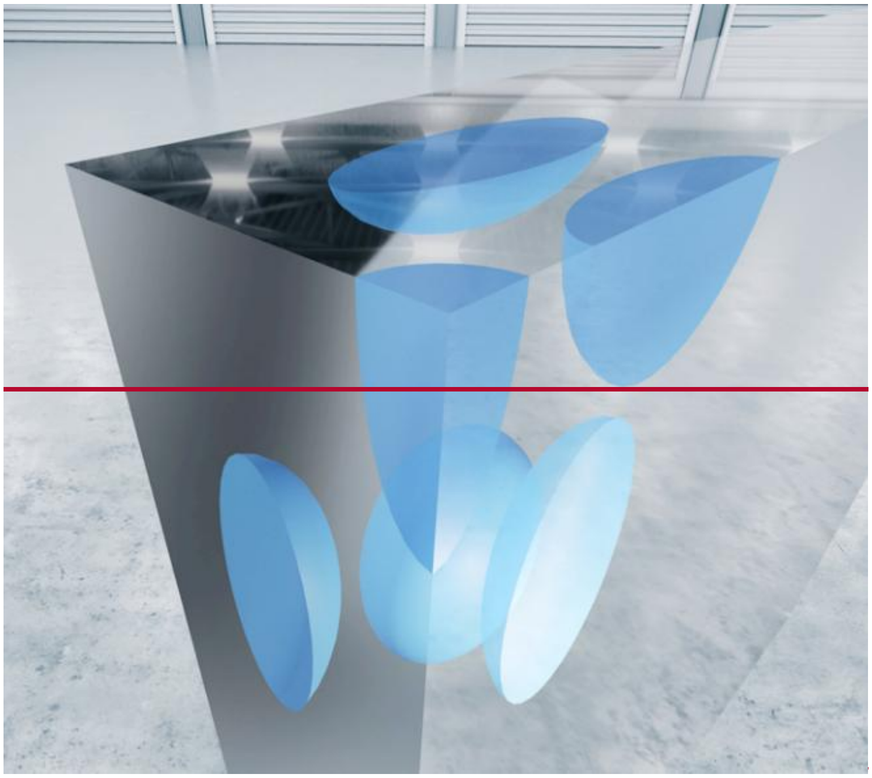

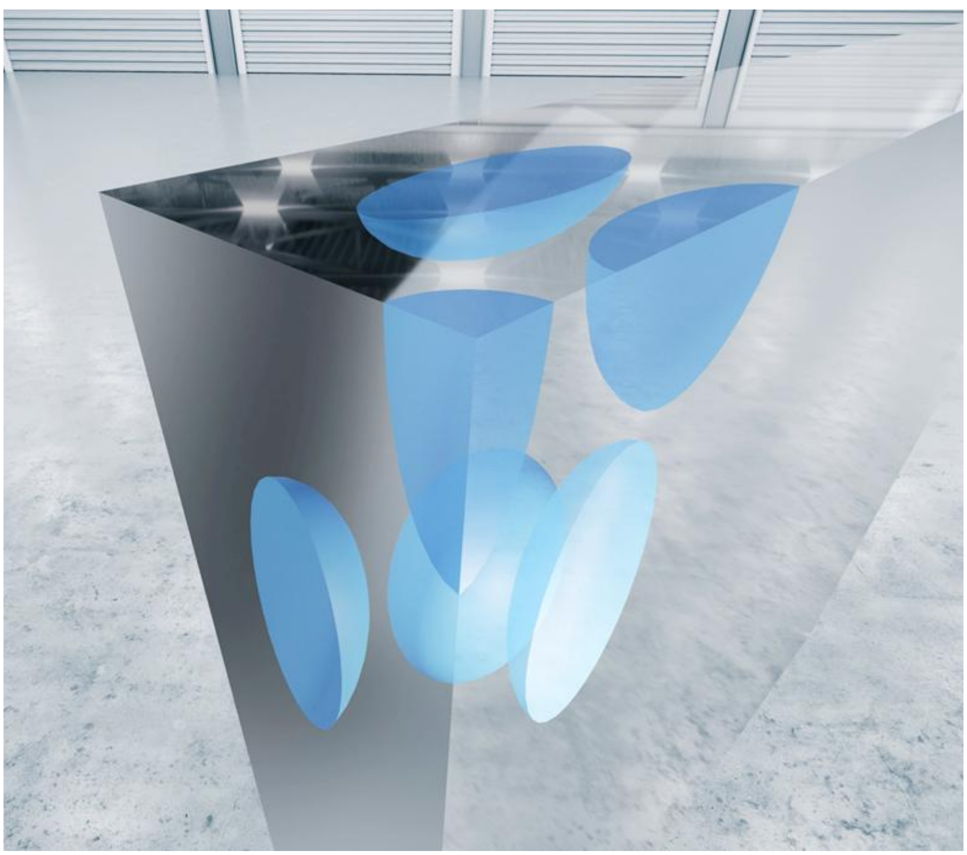

Figure 5.5.1-4: End of stage 3 – Example working scene model of ~~a~~the slab of glass ~~(showing~~ (only ~~the~~ nearest corner is shown)

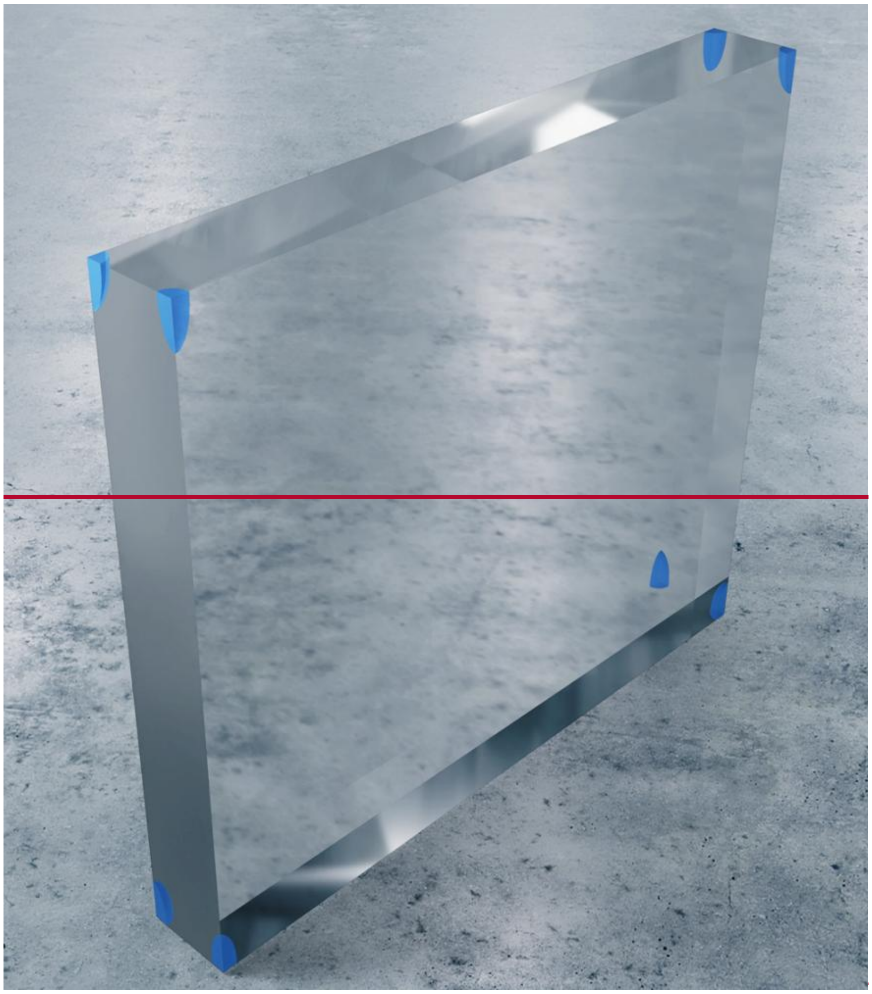

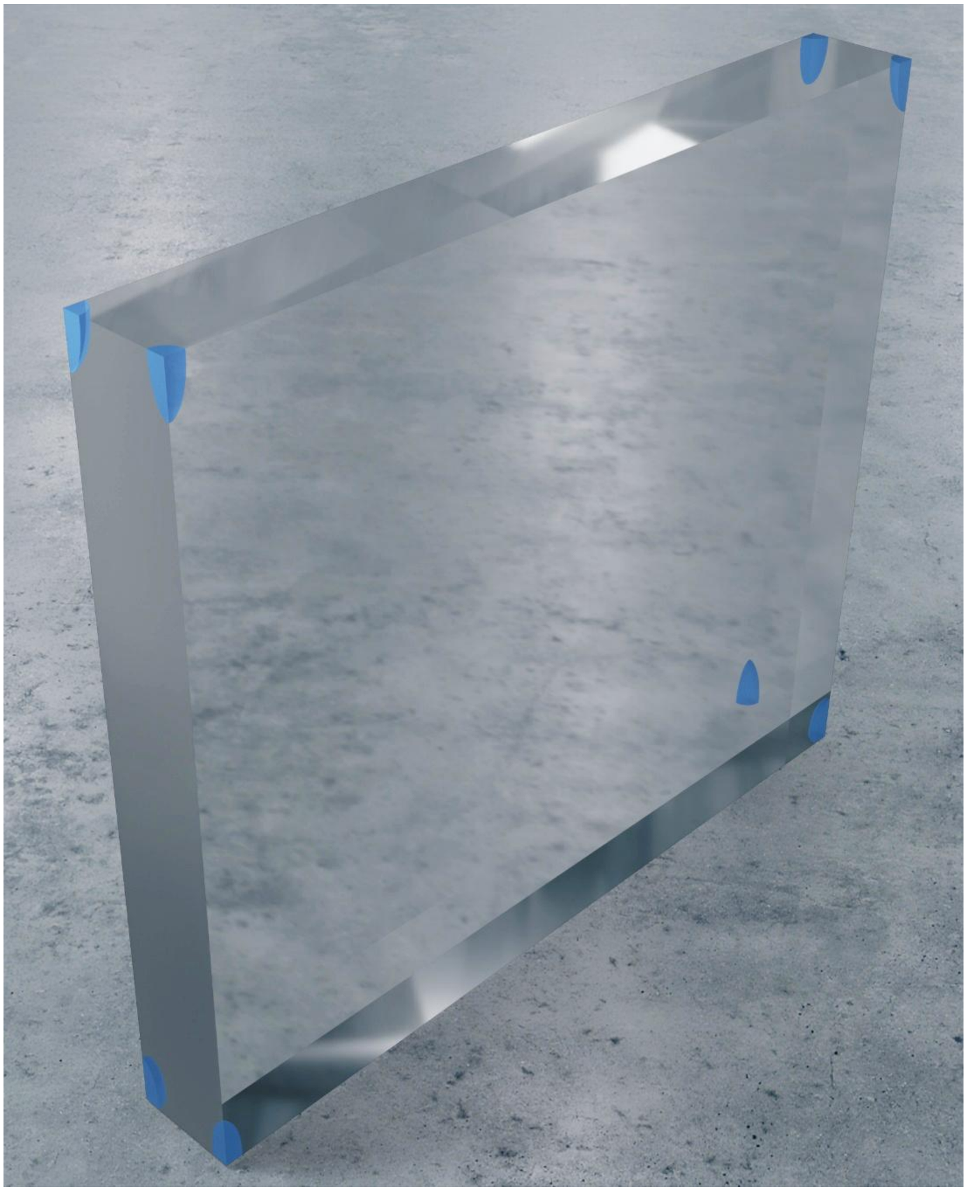

Figure 5.5.1-5: End of stage 4 – Example working scene model of ~~a~~the slab of glass~~ (showing only the nearest corner)~~

5.5.2 Use case: Virtualizing collision damage to higher fidelity

This section lays out a real-world use case where Smart Fidelity and intuitive capture let a grandmother virtualize collision damage on her car using her insurance carrier's QR-powered claims app:

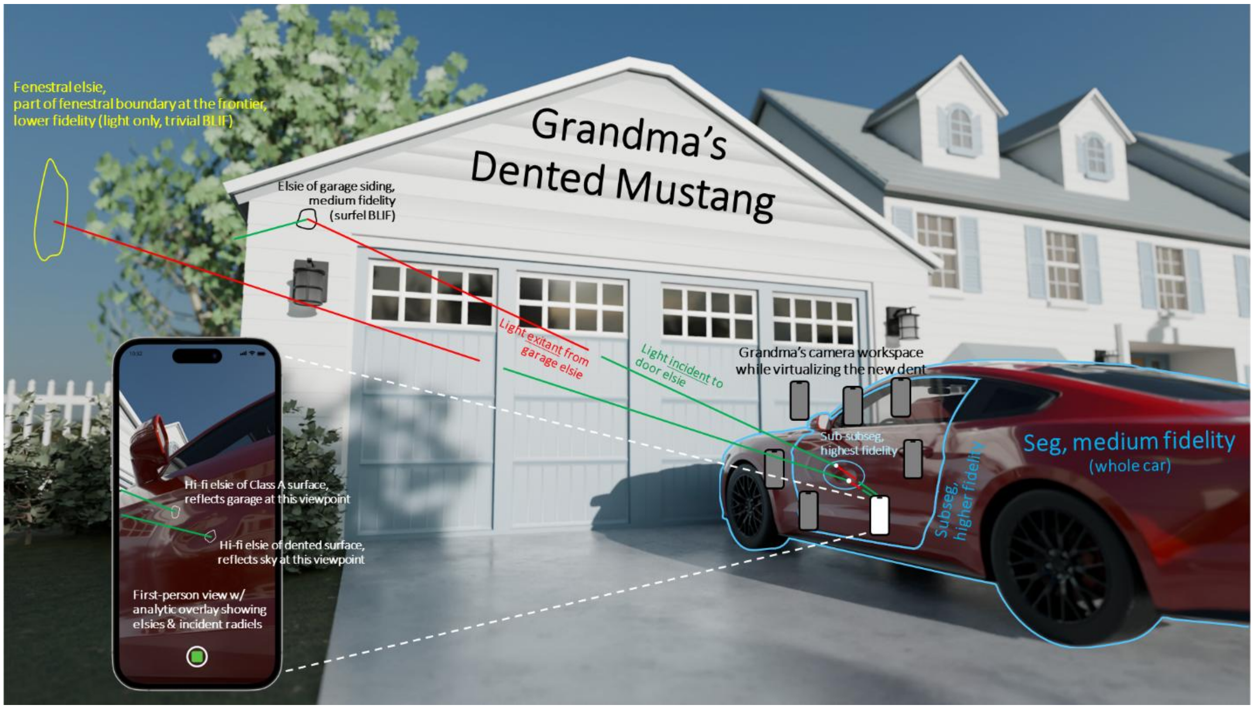

Figure 5.5.2-1: Virtualizing collision damage to higher fidelity

Example case – Grandma's red Mustang gets dented in a collision

- **Pseudo-algorithm step: Example case description (before first virtualization of car)**
  - Grandma's local municipal government maintains a QR very large scene model (XSM) of her whole town, updated frequently from image data captured casually / intuitively by residents and professionally by a small fleet of drones & autonomous cars.
  - Grandma's neighborhood is represented at medium fidelity in the "town XSM"[49].
  - The town XSM has been processed into canonical form[50] and is in medium-temperature cloud storage.
  - A few months ago, grandma's visiting granddaughter installed Q Virtualize on grandma's iPhone and virtualized her new red Mustang inside the garage~~ at night~~:
    - (Here we suppress detail re: virtualizing the car because such detail is given for the later virtualization grandma performs after her minor car accident.)

[49] The town XSM could exist in one or more hierarchies of XSMs up through county, state / province, region, nation, whole Earth. Storage temperature (access speed, dollar cost) decreases upward through such a hierarchy. Quidient would tend to license the QR Platform rather than manage such XSMs directly.

[50] "Canonical form" here means that elsies have been disentangled to relatively high power and all false (faux) emissive light has been eliminated, leaving the canonical light field as the only light in the scene model. "Canonical light field" = Incident light entering the scene at fenestral boundaries plus (presumed) true emissive light arising from energy conversion within the ~~scene' s~~scene's plenoptic field.

        - The garage door was shut, and no appreciable outdoor light entered through the garage window(s).[51]

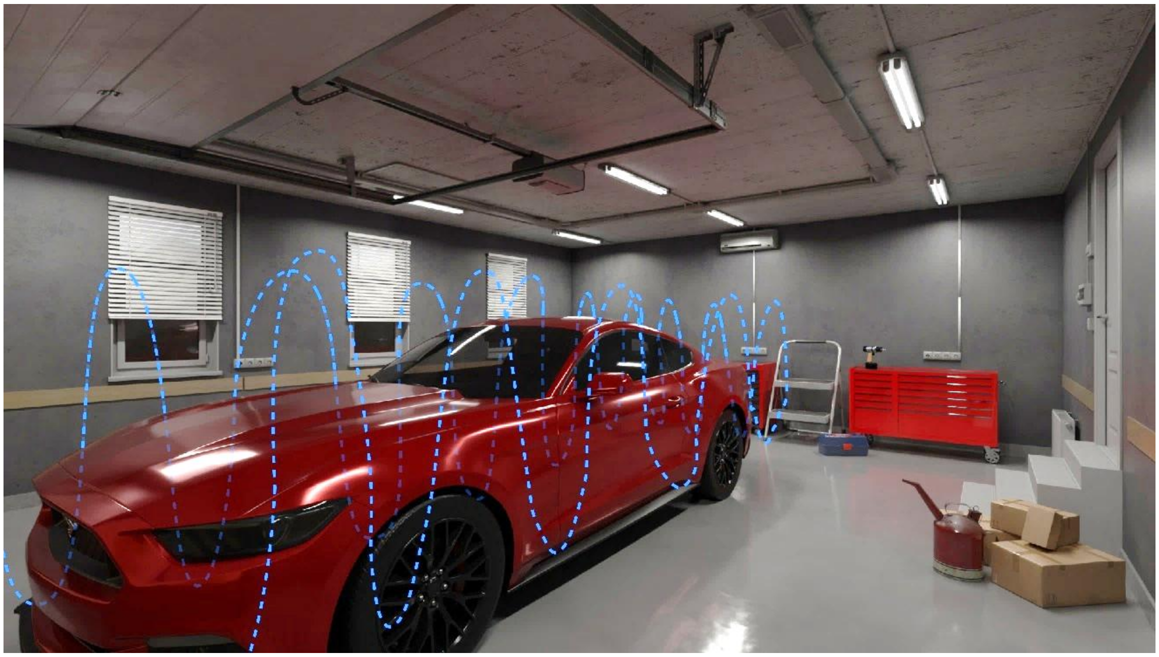

Figure 5.5.2-2: Grandma's Mustang before collision damage – Grandma's granddaughter captures the car along the dashed camera path in the garage at night.

- **Pseudo-algorithm step: Virtualize OOI in reference condition (first virtualization, algo detail suppressed)**
    - **Pseudo-algorithm step: Set target fidelities**
        - QR sets target resolution (floor), uncertainty (ceiling), and power (floor) for grandma's car (or its constituent regions) to meet the default "virtualization profile" that Q Virtualize operates in. The default virtualization profile targets part-in-1,000 dimensional uncertainty (calculated at statistical 1-sigma relative to the OOI's estimated size) and relit rendering PSNR > 30 in a variety of common environments[52].
    - **Pseudo-algorithm step: Initialize working scene (model)**
        - At the start of the virtualization session, QR queries for priors[53] when initializing the working scene model (falling back to a default working scene like "grayworld").
        - Phone GPS readings tell QR that the camera workspace is likely inside grandma's garage. ML-based recognition of a parked car and other typical garage entities bolsters that postulate.
        - **Pseudo-algorithm step: Extract initial working scene from superscene**

[51] In another salient case, light from streetlights, the Moon, and/or other outdoor entities enters the garage ~~window(s).~~.
[52] "Replacement" superscenes (could be virtualized from real scenes, or could be otherwise created) that the virtualized OOI subscene can be dropped into for rendering
[53] Parameters and their values that quantify characteristics of entities in the working scene model

    - Based on the postulated workspace location inside grandma's garage, QR queries "grandma's universe", which is synced[54] to the town XSM, and performs a (copy-type) subscene extraction centered on the anticipated workspace for an intuitive orbit around the car using grandma's type of smartphone.
    - The extracted subscene includes a medium-fidelity model of outdoor entities.
    - The extracted subscene becomes the working scene for the granddaughter's virtualization session.
  - **Pseudo-algorithm step: Initialize segtree(s)**
    - QR postulated earlier in recon processing that the main OOI is a car and therefore expects[55]:
      - Shiny opaque surfels to model the painted surfaces and chrome trim
      - Cleanly transparent surfels[56] to model the windshield & other glass panels
      - G^2 smoothness (Class A[57] shape) across the lion's share of all surfaces[58]
  - **Pseudo-algorithm step: Relight initial working scene**
    - QR sees that no appreciable light is entering through the garage window and therefore doesn't expend processing cycles relighting any outdoor regions of the scene.
      - In a different case, such relighting would be done for the purpose of predicting incident light on the car due to transmission through the window from outdoors.

- **Pseudo-algorithm step: Refine working scene (HWA dispatch happens here)**
  - QR "compiles" the working scene into HWA-friendly[59] form and dispatches it for HWA-based refinement to meet the target fidelity specified for the car OOI.
    - This step takes the scene (model) to its peak ~~absolute fidelity~~Absolute Fidelity (see figure 5.5.1-1).
    - Before the car model has been refined to high maturity, in regions where glass meets opaque surfaces, strips of elsies of non-negligible width are likely to exist in a hybrid state that combines the appearance and behavior of shiny opaque car paint and cleanly transparent glass. As refinement proceeds to maturity, those elsie strips will narrow (to negligible width, depending on target resolution).
  - After the desired rounds of refinement have been performed on the "dispatched working scene", QR manifests the resulting refinements[60] into the "hosted working scene" from which the dispatched working scene was produced.

- **Pseudo-algorithm step: Reduce ~~absolute fidelity~~Absolute Fidelity[61]**

[54] Parts or all of grandma's universe can be declared private, in which case the town XSM receives no information about those private regions beyond a specified fidelity ceiling (which could be 0 in the case of total privacy).
[55] "Expect" includes initialization of a working library of BLIFs.
[56] Or "sandel" pairs of back-to-back surfels
[57] https://en.wikipedia.org/wiki/Class_A_surface#Normalisation proposes a quantitative definition of Class A surfaces that merits consideration.
[58] The G^2 / Class-A smoothness postulate will impose a regularizing force on the origin position and orientation of each surfel relative to its neighbors (nearest and potentially more distant).
[59] Friendly to GPU ERT in current concrete implementations of QR
[60] Includes metrics about the refinements performed, in addition to refined parameter values themselves
[61] This optional step to reduce absolute fidelity of the working scene could be performed on the hosted and/or dispatched working scenes.

- As shown in figure 5.5.1-1, QR can optionally (and iteratively / incrementally) apply further Smart Fidelity processing to reduce ~~absolute fidelity~~Absolute Fidelity (and therefore data size of the working scene) while continuing to raise Smart Fidelity.
- This step typically lowers average elsie resolution and raises average elsie power.

  - **Pseudo-algorithm step: Persist working scene into superscene**
    - After refinement and optional reduction of ~~absolute fidelity~~Absolute Fidelity, the maturely refined working scene exists in host storage format as a hosted working scene.
    - Depending on QR virtualization settings, QR can optionally merge the refined working scene into its originating superscene, e.g., grandma's universe.
    - In the updated superscene, the refined car will typically be a distinct branch in one or more segtrees.[62] It will also be recognized as a human-meaningful "object".

- **Pseudo-algorithm step: Example case description (before second virtualization of car)**
  - Earlier today, a pickup truck lightly sideswiped grandma's red Mustang, creating a large dent in its passenger-side rear door.
  - Upon arriving back home around noon, grandma parks in her driveway and opens the claims app from her car insurance carrier, which offers QR-powered damage / condition inspection.

Figure 5.5.2-3: Grandma's Mustang after collision damage – Grandma parks her car in her driveway to virtualize a dent that has been newly created in the driver-side door.

- **Pseudo-algorithm step: Virtualize new dent in Mustang door (second virtualization)**
  - **Pseudo-algorithm step: Set target fidelities**

[62] There can be multiple segtrees, each one arising from segmentation that operates on one or more features in the scene. See §4.3 for examples of segmentation on various features. In this example case, segmentation might typically be performed on geometric continuity / curvature, "base color" of BLIF, degree of transmissivity (glass vs opaque media), and (optionally multiview) 2D NN-based classification by car panel / part type.

- Set a target fidelity for each region of the working scene. In this second virtualization of the Mustang, target geometric accuracy of the new dent is the main target fidelity that drives other target fidelities throughout the working scene.

- **Pseudo-algorithm step: Initialize working scene that includes previously virtualized OOI (in reference condition without the new dent)**
  - Phone GPS readings tell QR that the anticipated camara workspace is near grandma's driveway. Use of the insurance claims app tells QR that the main OOI is a passenger vehicle.
  - Under suitable privacy controls, QR accesses grandma's universe. Grandma's universe now includes the non-dented Mustang that her granddaughter virtualized a month before the collision (using grandma's phone).
  - **Pseudo-algorithm step: Extract initial working scene from superscene**
    - QR performs a (copy-type) subscene extraction from grandma's universe, centered where the real car is parked in the driveway.
      - The extracted subscene initially extends a default distance[63] outward from the anticipated camera workspace. Parts of its scene boundary can move outward & inward as capture begins and the actual (vs anticipated) workspace develops.
    - The extracted subscene becomes the working scene model.[64]
    - The claims app tells QR to look for grandma's car. With some combination of ML object search and plenoptic search / template-matching, QR recognizes the presence of her car in the real scene and localizes the car model from her QR universe into the working scene, including 7-DOF pose estimation.
  - **Pseudo-algorithm step: Initialize segtree(s)**
    - QR knows there's a car that has incurred likely damage and therefore should expect:
      - Segs of grandma's car that was previously virtualized in an undamaged state
      - Regions (could become segs) of anomalous surfels, explored in "waterfall" order from typical to unusual damage patterns
  - **Pseudo-algorithm step: Relight initial working scene**
    - Different relighting processing is called for here as compared to the previous nighttime capture of the undamaged Mustang in grandma's garage. In the nighttime scene, no appreciable light entered the garage ~~window~~ from outside, and therefore relighting the initialized outdoors wasn't useful. In the current driveway scene at noon, though, plenty of outdoor light is incident on the car and needs to be modeled to suitably high fidelity.
  - In the real scene, the next-door neighbor's house and large tree[65] happen to reflect on the door dent at most of the practical viewpoints to which grandma can comfortably move her iPhone.
  - The working scene model includes the neighbor's tree & house from the town XSM at the medium fidelity they were last virtualized to.

---

[63] Example of default distance: logarithmically half of the limit distance at which parallax can no longer be resolved on scene features, given target fidelities and camera capabilities.
[64] We could alternatively say that the subscene copy-extracted from grandma's universe gets merged into a default "grayworld" initial working scene.
[65] In this example, let's say the wind speed is negligible and we therefore needn't model tree motion. In cases where the wind speed is non-negligible, QR must account for the tree's motion by either recognizing and ignoring observations that are significantly influenced by the tree, or by modeling the tree's motion.

- The working scene model includes the undamaged Mustang to high fidelity — at least the fidelity that will be needed by the claims app to estimate the repair cost.

Figure 5.5.2-4: Camera workspace to virtualize the collision dent

- **Pseudo-algorithm step: Plenoptic diff to measure matter field anomalies to lower fidelity**
    - Grandma's memory & eyes tell her where the dent is, so she starts performing an intuitive capture focused on the dent.
        - Note that QR might have enough information from the town XSM + daytime sky model for [lat/lon, date, time] to predict medium-fidelity incident light at the dent w/o grandma needing to capture video of any scene region besides the dent itself from convenient viewpoints.
    - QR continues grandma's intuitive capture session, using "exaggerated deviation" visuals and potentially more poignant cues for missing viewpoints.
    - With the light field information grandma has captured so far, QR is able to do a medium-fidelity "plenoptic diff" on the dent ROI, meaning a comparison between these 2 light fields:
        - Exitant light that is currently observed by grandma's iPhone
        - Exitant light predicted from the car door in its non-dented former state (using the medium-fidelity incident light field estimated a few bullets above)
    - Such a plenoptic diff is computationally cheap because it involves mainly forward modeling (prediction of exitant light).

Figure 5.5.2-5: Dent ROI surfels before surface refinement – Analytic visualization of surfels in the dent ROI before QR performs surface refinement

- **Pseudo-algorithm step: Check attained fidelities against target fidelities**
  - In some cases, the above plenoptic diff alone could be enough to reveal the lateral extent and severity / depth of surface deformation needed for the insurance carrier to calculate an informed estimate. But let's say that the plenoptic diff doesn't suffice in this case.
  - Recognizing that the plenoptic diff won't quantify the dent damage to sufficient fidelity when translated into claim dollars (a generalized "k-factor"), QR proceeds to do a medium-fidelity refinement of the dent ROI, starting at initial state = undamaged door region from last month.

- **Pseudo-algorithm step: First refinement to measure matter field anomalies to medium fidelity**
  - The refinement proceeds as described in step "Refine working scene" of the non-dented car virtualization above.
  - Dent surfels start off initialized to the undamaged Class A state that grandma's granddaughter virtualized last month.
  - QR assesses the severity of damage using ML inference and/or plenoptic diff, and then tailors the waterfall of elsie refinement postulates, including:
    - Geometry – G^2 smoothness is likely to be violated (especially at the dent perimeter), so quickly fall through to G^1 and even G^0.
    - BLIF:
      - Media – Check for changes in media due to abrasion & cracking.
      - Microgeometry – Check for changes in microfacet roughness in regions of light abrasion.
    - The tailoring can be done according to a model of statistical likelihood of various surface anomalies due to the present kind of sideswipe impact of the offending pickup truck on a Mustang like grandma's.

- QR finds that in their initially postulated Class A state, nearly all surfels in the dent ROI violate an "uncertainty ceiling" that has been set based on the k-factor relating matter field fidelity to insurance claim dollars.
  - The uncertainty ceiling gets violated because iPhone-captured exitant light heavily disagrees with exitant light that's been predicted based on Class A surfels. This is a per-elsie version of plenoptic diff, aka a tightly local "rendering loss" for a single elsie calculated over a sparse sample of exitant directions.
- The ceiling-violating uncertainties across dent surfels impel them to search, in waterfall order, a range of alternative parameter values in an effort to minimize disagreement between predicted and captured exitant light.
  - In this case, since QR thinks / knows it's dealing with a dented Class A surface, it could be computationally natural and efficient to trial-displace[66] each surfel along its Class A normal vector direction.

- **Pseudo-algorithm step: Check attained fidelities against target fidelities**
  - The medium-fidelity refined matter field of the dent still fails to translate to sufficient claim dollar accuracy (in this example case). QR recognizes that it needs to refine the dent to higher fidelity, expressed as some combination of geometry and material fidelity.

Figure 5.5.2-6: Dent surfel refinement – Left: Analytic visualization of dent ROI surfels before QR performs surface refinement. Right: Analytic visualization of dent ROI surfels after QR performs surface refinement.

- **Pseudo-algorithm step: Further refinement to measure matter field anomalies to higher fidelity**
  - QR recognizes that the light incident on the dent from the neighbor's tree & house isn't of sufficient fidelity (directional resolution and/or radiometric uncertainty) to reach the target fidelity calculated / specified for the dent matter field.

[66] "Trial-displace" means to explore postulated positions for a given surfel within some range of displacements. The displacement is 1D along the initial (Class A) normal vector direction in this case. There's limited utility in letting dent surfels displace laterally (tangent to Class A surface) when the indentation has limited severity.

- One way for QR to obtain the needed high-fidelity incident light is to tell grandma to place her phone right in front of the dent and do a brief pan/tilt sequence to directly record the needed incident light.
- But in this case, QR decides (according to some operating settings re: UX and various computation budgets) that it's best to further refine the neighbor's house and tree elsies higher fidelity, to then predict to sufficiently high fidelity the light that exits those elsies and fares toward surfels of the dent.
- In an alternative case, QR could instead decide[67] that grandma should turn around and capture new video of the tree and/or house to supplement the light field information that already exists in the town XSM. This could happen, for example, if QR finds that the town XSM's existing elsies & fenestral light field can't produce[68] a light field of sufficient fidelity to support further refinement of house and tree elsies to the target fidelity.
    - In this and a great many related cases, QR performs bookkeeping on data capture timestamps. This is important because a very different quasistatic light field can exist in conjunction with an unchanging matter field over even seemingly brief time intervals, e.g., a cloud that moves to cover the Sun can drastically alter the outdoor light field within a few seconds of clock time.
    - If grandma ends up capturing supplemental video of the neighbor's property, a privacy question arises of whether to merge that info back into the town XSM.
    - Dynamic sky (moving clouds or even dawn / dusk rapid sun change) is a prominent factor that could prompt a new grandma capture of the environment.
- The refinement proceeds as described in step "Refine working scene" of the non-dented car virtualization above.
- QR is now satisfied with the fidelity of predicted incident light at dent surfels coming from a wide range of directions, and proceeds to refine the dent to a higher fidelity that meets the target geometric accuracy requested by the insurance claims app.

---

[67] Reasons QR could decide this include:
- QR tried the refinement in the preceding bullet and found the result unsatisfactory.
- Computation budget(s) favor new data capture over further refinement of existing information.
- QR estimates that the current light field differs sufficiently from the stored light field.

[68] For example, via propagation of plenoptic field information from the XSM and fenestral light field into the working scene, the propagation occurring in the downstream and/or upstream polarities

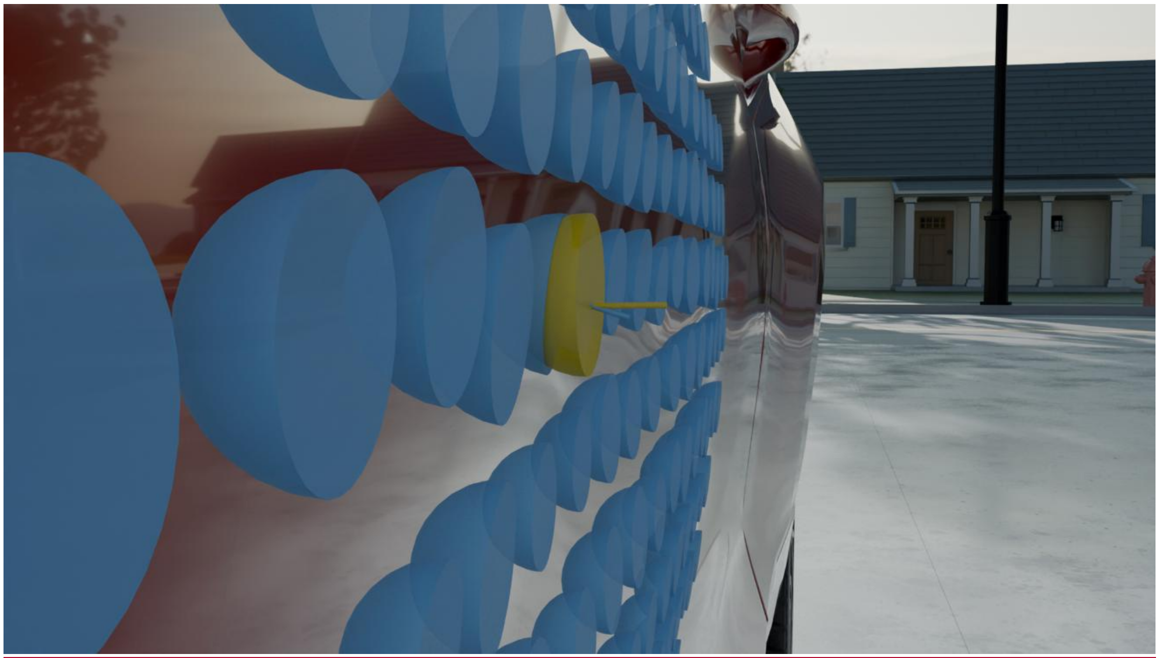

Figure 5.5.2-7: **Closeup of dent surfel refinement** – Oblique view of (analytic visualization of) dent ROI surfels refined to their dented state, with a single surfel shown in yellow in its unrefined initial state for contrast. Not the angle difference between the yellow unrefined normal vector and the corresponding blue refined normal vector.

- **Algo step: Persist working scene into superscene**
  - QR proceeds to persist the output of the refinement as described in step “Persist working scene into superscene” of the non-dented car virtualization above.
  - In addition, the matter field (geometry and materiality) of the highly refined dent ROI makes its way to a processing layer in that quantifies the severity of the dent for repair cost estimation.

6 Subscenes

6.1 Overview

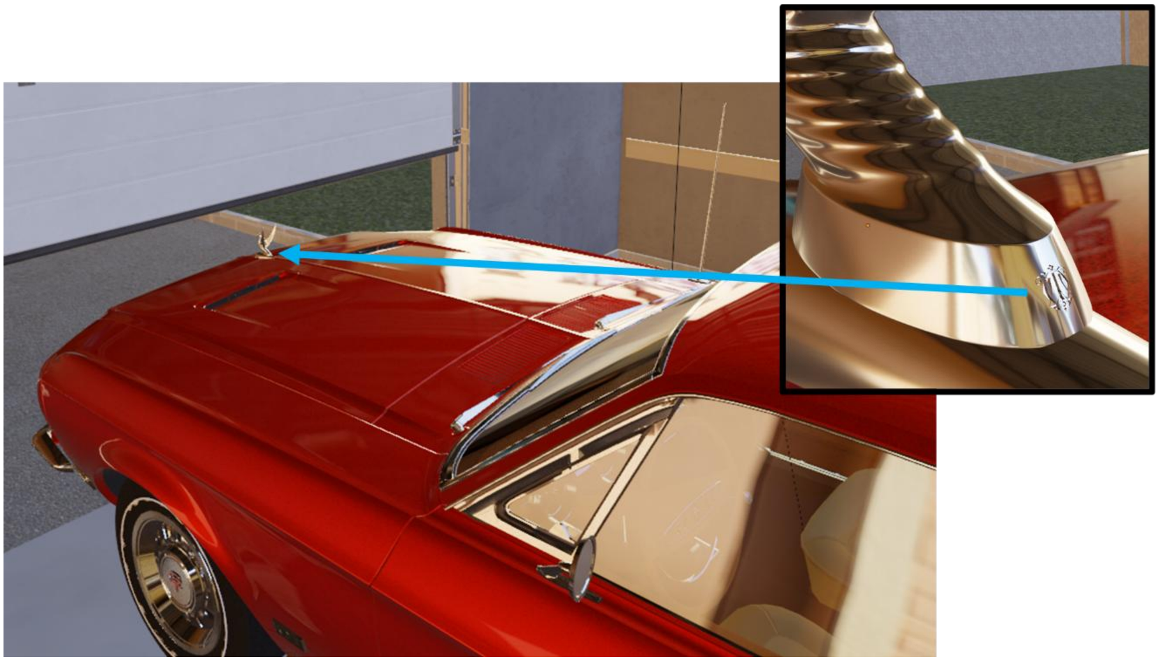

Figure 6.1-1: Subscene centered on a vintage car hood ornament (inset shows part of subscene)

6.2 Purposes

Prominent reasons to create subscenes in QR include:

- **Demarcation**. Demarcate a scene region in a way that's useful to humans (including end-users) or subsequent processing.
    - Examples:
        - A pair of shoes, separate from the (unimportant) ROS in which they were virtualized
        - An open toolbox and the tools inside it
        - A single room in a house
        - A house & its surrounding lot delimited by legal boundaries (property lines)
        - A city or county delimited by its political boundaries (borders)

- **Privacy**. Limit access to information beyond a specified fidelity in a specified region in a scene model.
    - Methods for privacy control in QR include the following:
        - Use coarser elsie resolution on a private regions (similar to 2D image blurring).
        - Filter out private segments.
        - Add opaque barriers, including "reskinning" glass windows as opaque.
    - Examples:

    - A grandmother virtualizes collision damage on her sports car using her insurance carrier's QR-powered claims app. After virtualization, QR extracts the car's relightable matter field as a separate subscene and transmits only that subscene to the carrier, thereby excluding information about the garage, driveway, and adjacent areas.[69]

- **Tunnel of resolution**. Enable local resolution that's much finer than in adjacent regions.
  - This purpose specifically regards the case where QR needs finer local resolution than the parent (super)scene's range of resolutions[70] can accommodate.
  - This kind of subscene can be usefully analogized to capturing scene detail close-up with a macro lens, or zooming in to a very narrow FOV with a telephoto lens.
  - Examples:
    - See the figure in §1.1. An end-user is virtualizing a car 5 meters in length. The overall car (super)scene supports a finest resolution of 500 microns, which is 4 OOM smaller than the car's length. The car's hood ornament has fine geometric detail on the order of 200 microns. QR determines that a resolution of 40 microns (1/5 the scale of the detail) will be needed in order to meet fidelity criteria on the ornament. Since 40 microns is (more than) yet another OOM finer than the car scene's resolution of 500 microns, QR extracts a new subscene centered about the ornament. The new subscene's scale is 1/20th that of the car superscene and has a finest resolution of 25 microns.

- **Island of accuracy (aka "island of lower uncertainty").** Enable lower target uncertainty (in some parameters of interest) in a local region.
  - Examples:
    - Following a collision, an automotive technician removes the engine's cylinder head and uses a QR-powered smartphone app to check the head's bottom surface flatness against a tolerance of 75 microns (0.003") "out of flat" across the surface. QR extracts a subscene with reference scale slightly larger than the cylinder head and assigns a target uncertainty tighter than 15 microns (1/5 of the 75 micron flatness tolerance) for elsie origin positions relative to a suitable local datum.[71]
    - A machine shop technician virtualizes the portion of the shop visible from a camera workspace centered on an empty tabletop that's well lit. The technician proceeds to virtualize many small machine parts, each one being placed on the tabletop. For each part, QR recognizes that the part is small enough not to (non-negligibly) impact the light incident at the tabletop. QR extracts an island-of-accuracy subscene centered on the given machine part and initializes its incident fenestral light field to the previously measured light field.

---

[69] Reflections of people in the car surfaces (e.g., a grandchild playing around the car in a bathing suit) would not be transmitted in this private data.
[70] QR has a guiding range of resolutions (e.g., 4 OOM) that applies to most parameters in most (single) scenes.
[71] An example such datum might consist of 3 non-colinear strong pointlike features on or near the bottom surface of the cylinder head to establish a 6-DOF coordinate frame, plus an additional feature of scale, for example some apriori known dimension of the head. QR needs to resolve the datum-establishing features to an uncertainty lower than the target uncertainty of 15 microns desired for this subscene's elsie positions.

- **Host performance.** Gain improvement in at least one dimension of performance (e.g., fidelity, speed, cost, data size).
    - Examples:
        - For dispatch to (further) host processing, extract a subscene whose spatial extent and elsie density are limited to a specified host memory budget.

- **Hardware acceleration (HWA) performance.** Gain improvement in at least one dimension of HWA performance (e.g., fidelity, speed, cost, data size).
    - Examples:
        - Parallelize over GPU cores to shorten processing time while holding to a specified target fidelity.
        - Parallelize over GPU cores to raise fidelity while holding to a specified target processing time.
        - Extract a goal-directed subscene to limit how many bytes QR needs to transmit to an end-user device.
        - Extract a subscene whose spatial extent and elsie density are limited to a specified GPU memory budget.

6.3 Key attributes

Key attributes of scenes, superscenes, and subscenes (we'll call all three "subscenes" in this section):

- Subscenes need to be independently refinable & jointly reconcilable using identical processing logic regardless of how or why the subscenes were created.
- A subscene extends all the way to its frontier.
- Entities in a scene created at various fidelities, times, etc. can be colocated (or approximately colocated) in the plenoptic field.
- A subscene's plenoptic field can have unrepresented regions.[72, 73]
- A subscene's frontier contains zero or more fenestral boundaries.[74]
- A "universe" is the union of all scenes at all levels in a QRM.
- The "root scene" of a universe is the top-level superscene in a Quidient Reality Model (QRM). Its direct and indirect subscenes can be arbitrarily related or unrelated to each other. And those relationships can change drastically over represented scene time and processing stages (algorithmic processing time).
- Subscenes & superscenes exist at the host (CPU) & HWA (GPU) levels. Hosted subscenes can be "compiled" into HWA-friendly form and dispatched for HWA processing.
- QR handles reconciliation (and potential merging) of multiple, disconnected capture sessions, all of which could be done by "granddad" end-users using intuitive capture.
- The hierarchy of subscenes & superscenes exists at the host / CPU level.
- Any dispatched QR scene (model), not including its subscenes, intentionally contains an amount & complexity of data that can be efficiently processed through many refinement iterations in a GPU or similar HWA.

[72] The unobserved region outside a windowless room is an example of an unrepresented part of a

[73] A subscene's entire 3D field can even be unrepresented.

[74] An example of zero FBs is a windowless room where all incident light arriving at an ROI in the near-field traces back to virtualized surfaces in the ROS

- A GPU (or similar HWA) refines a self-contained dispatched (sub)scene and then reports the results of the refinement to its calling CPU-level process, including performance / fidelity metrics.
- QR intentionally applies identical processing logic to a scene regardless of where it lives in the hierarchy of subscenes and superscenes (i.e., scene processing is conceptually recursive).

6.4 Extraction

Further prominent attributes of subscene extraction:

- Respect for target fidelity is a primary driving factor in subscene extraction.
- For hosted scene maintenance, QR aims to extract subscenes that are minimally coupled to each other in terms of the strength of parameter update influence between them.[75]
- Minimal coupling between dispatched subscenes ~~avoid~~avoids “reconciliation thrashing”[76].
- QR can detect and extract named, human-meaningful regions of a scene.
- A working subscene can “reach out” a significant distance along a corridor to include elsies that heavily influence elsies in ROIs in the working subscene.

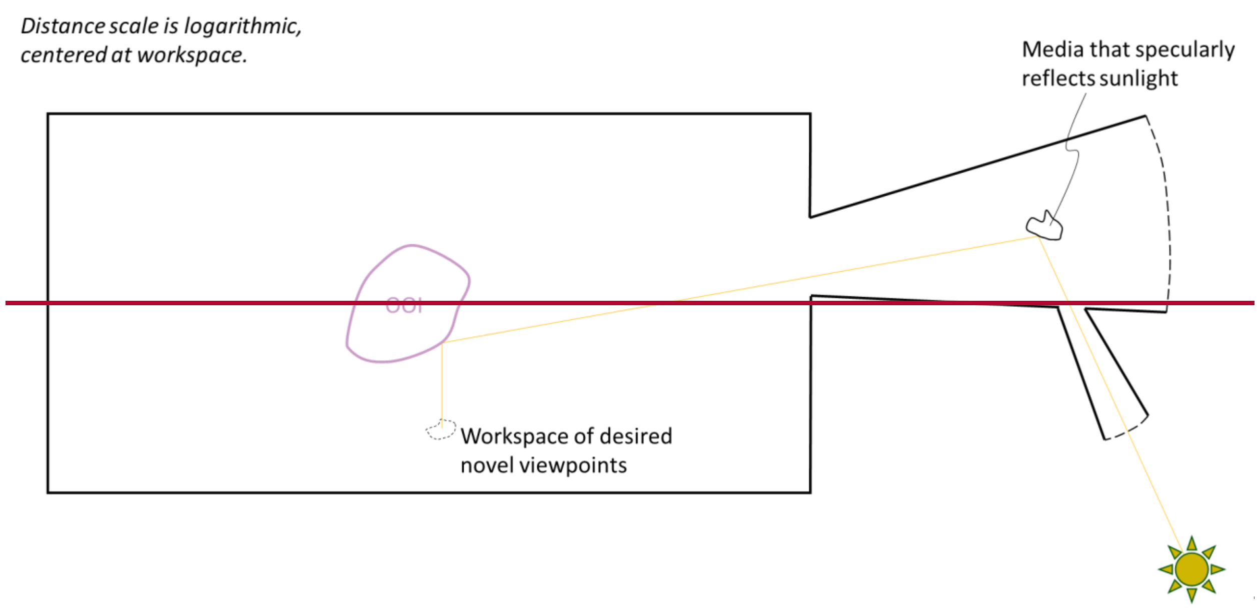

[75] We say that 2 subscenes A & B are significantly coupled if updates to (parameters of elsies of) subscene A significantly change the light incident on (elsies of) subscene B in a way that impels significant updates to (parameters of elsies of) subscene B during subsequent refinement (done by minimizing some radiometric loss). The strength of such coupling will highly depend on the geometry (geometric configuration) and BLIFs of elsies in the 2 subscenes. One potentially useful technique for decoupled extraction is to operate on a graph of elsie nodes and edges weighted by “plenoptic influence”.

[76] "Reconciliation thrashing" here means an undesirably frequent need to reconcile subscenes with each other, usually at the host / CPU processing layer above HWA / GPU.

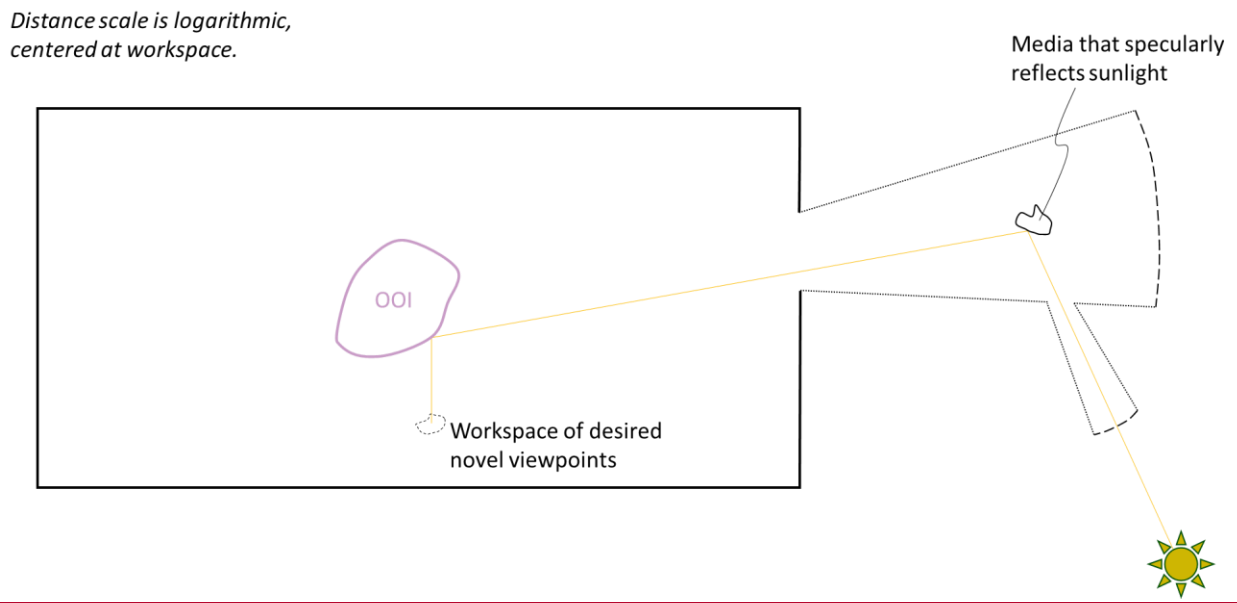

Figure 6.4-1: Before extraction – Full scene with an OOI and workspace of desired novel viewpoints (one example light path shown)

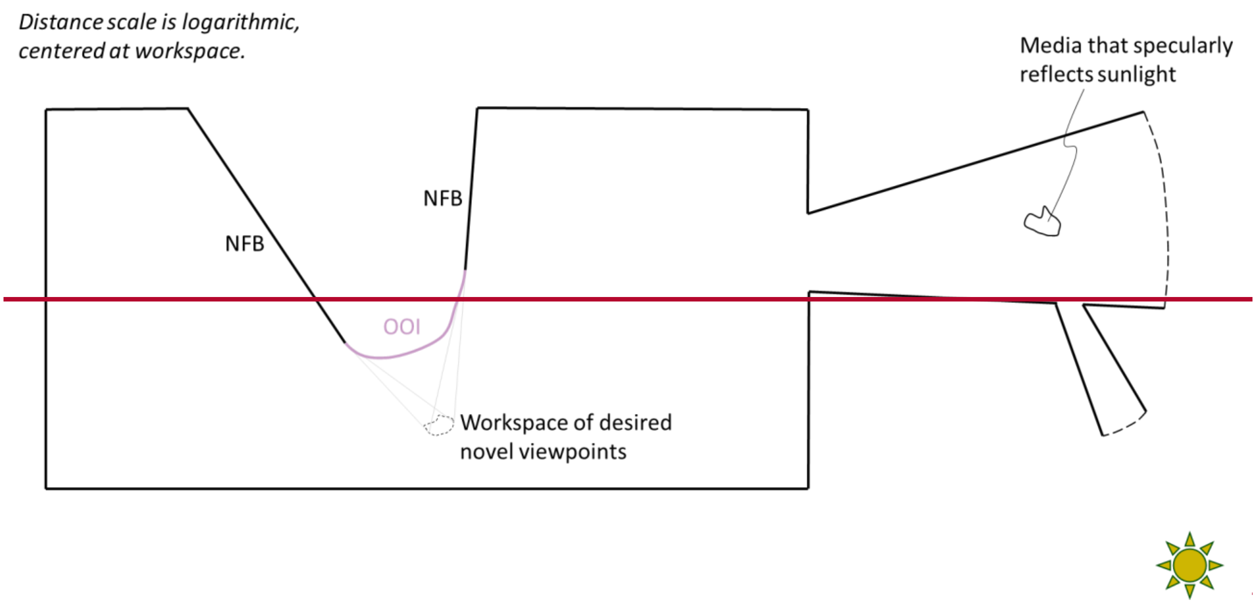

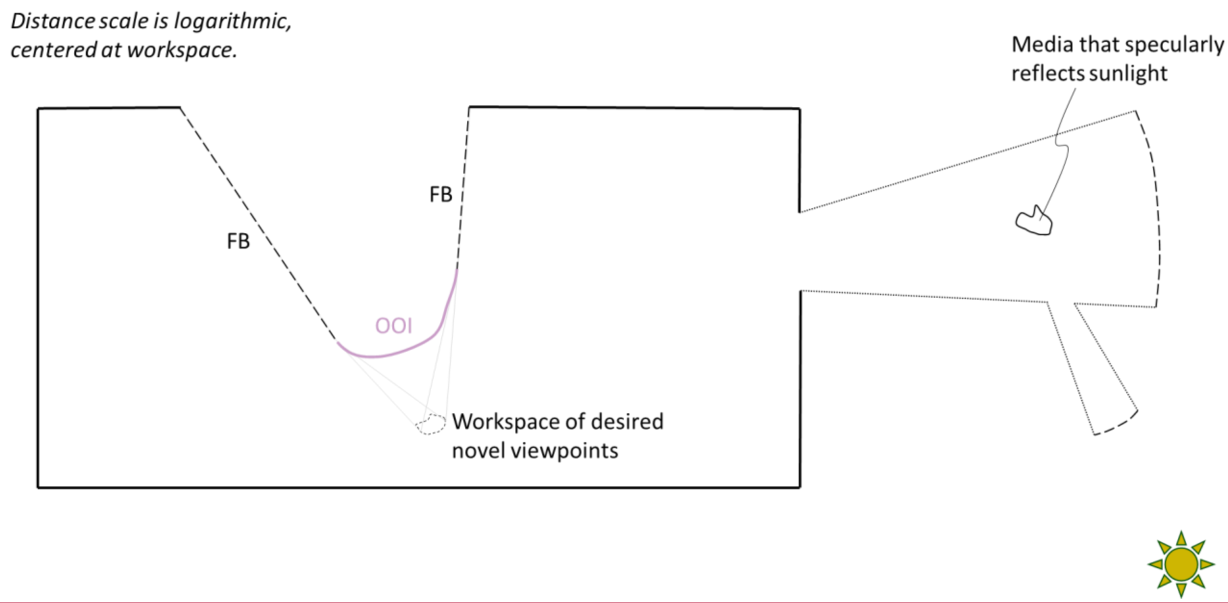

Figure 6.4-2: After extraction – Parsimonious subscene extracted for high-fidelity rendering from novel workspace

6.5 Reconciliation

Further prominent attributes of subscene reconciliation:

- Reconciliation honors the target fidelity specified for each region in the subscenes being reconciled.
- QR is able to reconcile subscenes that overlap to any degree in space, time, and fidelity (especially resolution).
- Reconciliation between subscenes will minimize exitant radiance loss (residuals) between some subset of the elsies in each subscene.
- One natural way to reconcile a subscene is to let lower-uncertainty elsies in the subscene and its parent superscene further refine each other's elsies.

6.6 Maintenance of hosted scene models

A hosted scene model (typically LSM or XSM) undergoes periodic maintenance that includes refinement and other updates to scene regions, as well as reconciliation between scene regions. To load-balance these maintenance tasks, QR extracts subscenes and dispatches them for parallel HWA processing. Updated subscenes will then typically undergo some degree of reconciliation against the rest of the hosted scene model.

7 Optimization

7.1 Overview

Many QR processing operations, notably reconstruction, are based on mathematical optimization (aka just “optimization”) that minimizes the deviation between values of quantities predicted by a working scene model[77] and reference values of those quantities:

$$\underset{\substack{\text{subset } \mathbf{p} \text{ of} \\ \text{scene parameters}}}{argmin} \sum_{\substack{\text{predictable} \\ \text{quantities } \mathbf{q} \\ \text{of interest}}} w_{\mathrm{i}}\left(q_{\mathrm{i,pred}}(\mathbf{p}) - q_{\mathrm{i,ref}}\right)$$

where:

- **p** (vector of scalars) is a chosen subset of the parameters that characterize a (region of a) scene.[78]
- **q** (vector of scalars) is a chosen set of quantities (each of whose values) QR can predict from the values of parameters **p**.
- $q_{\mathrm{i,pred}}(\mathbf{p})$ is the predicted value of the *i*-th quantity predicted by (calculated from) parameters **p**.
- $q_{\mathrm{i,ref}}$ is the reference value of the *i*-th quantity. This reference value is often, but not always, a radiance value measured by an observing camera.
- w_{i} is the weight of the deviation (aka prediction error, aka residual) between $q_{\mathrm{i,pred}}(\mathbf{p})$ and $q_{\mathrm{i,ref}}$. These weights are often based on a measure of the (inverse) uncertainty of $q_{\mathrm{i,ref}}$.

In one important class of QR optimization problems, the **q** are radiances of exitant radiels faring in directions of interest, often but not always directly toward observing cameras.

[77] QR’s postulated model of a scene at the current stage of processing
[78] The elements of **p** can be parameters of elsies and/or other entities in the scene model.

7.2 Key attributes

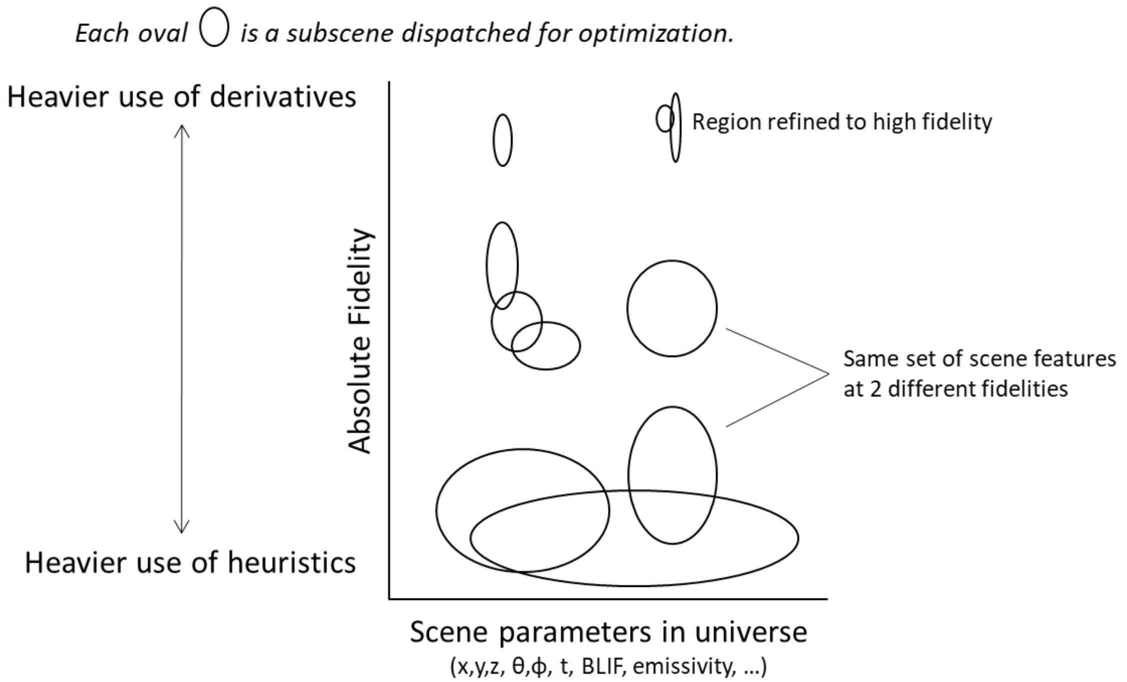

Figure 7.2-1: Typical optimization problems QR formulates and solves throughout a QRM universe

Referring to the preceding figure[79], QR optimizes (optionally initializes, and then refines) one or more parameterized subscenes in a QRM universe. Key attributes of such optimization:

- QR can dispatch a wide variety of subscenes for HWA optimization.
- Each dispatched subscene can arise in one of several ways, including:
  - First-time virtualization of a real scene "from scratch"
  - Extraction from an existing (super)scene
  - Novel composition of existing scene components
  - Other scene creation by an automated or manual process
- In a single dispatched subscene, the set of scene parameters can be highly heterogeneous, spanning any combination of position, orientation, time, BLIF, emissivity, analytic entity parameters, and more.
- A dispatched subscene can span any extent in (each dimension of) parameter space.
- A dispatched subscene can span any extent in resolution, uncertainty, and power (the dimensions of fidelity).
- A dispatched subscene can overlap, to any degree, any number of other dispatched subscenes in [parameter, fidelity] space.
- QR can optimize a dispatched subscene using any combination of loss minimization methods, optionally making use of derivatives (e.g., calculated analytically, numerically, and/or using automatic differentiation).

[79] Note that the diagram does not show processing stages ("algorithm time", as opposed to real clock time represented in the scene model and denoted "t" in the horizontal axis label).

- Refinement performed at higher fidelity (smaller / narrower ovals higher up in the above diagram) can subsequently influence parameters at lower fidelity, including via subscene / superscene reconciliation.

7.3 Example cases

7.3.1 Normal vector refinement

In this example, QR refines the (postulated) direction of a surfel's (surface elsie's) normal vector from initial direction $\mathbf{n}_{init}$ to final direction $\mathbf{n}_{fin}$:

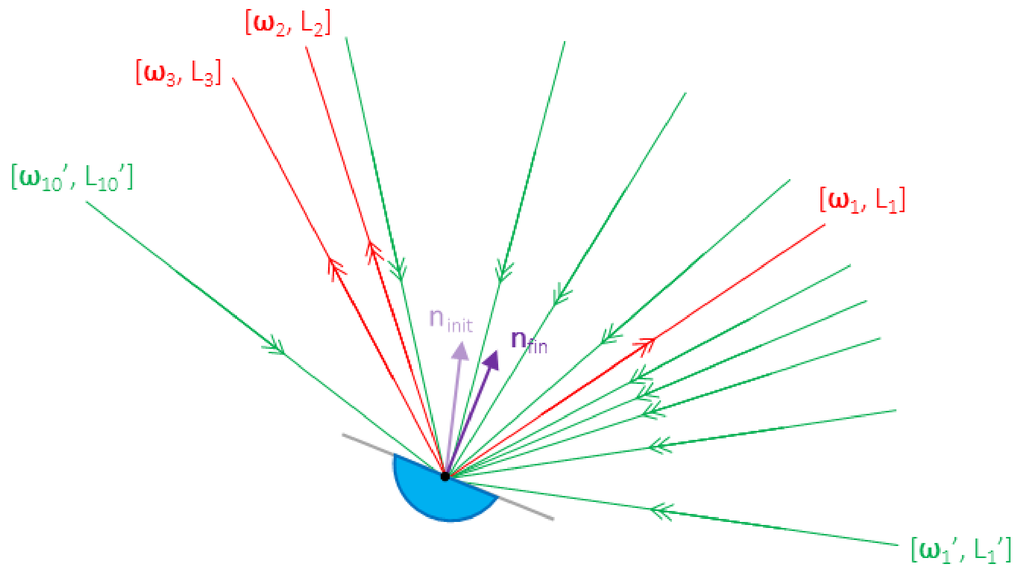

Figure 7.3.1-1: Refinement of a single elsie's normal vector direction by minimizing exitant radiance loss – 10 importance-sampled incident radiels flow through the surfel's BLIF to yield 3 predicted exitant radiels in directions of interest at which loss residuals can be calculated and used to adjust (perturb) the surfel's normal vector.

Key points about this example:

- QR performs the normal vector refinement on a single surfel, in contrast to the familiar case of refining a (usually contiguous) large set of surfels / elsies.
- Exitant radiance can be predicted (and loss residuals can be formed) at positions other than exact elsie origins by interpolating the macrogeometry and BLIF parameter values of nearby surfels.
- Though they often will, exitant radiels $\boldsymbol{\omega}$ do not necessarily terminate at directly observing cameras. In the vein of §7.1, it suffices to calculate each $L_{i,ref}$ by a process independent of the forward BLIF interaction that yields $L_{i,pred}$. Example processes that can yield $L_{i,ref}$:
  - Upstream[80] propagation from pixel values of a directly observing camera
  - Upstream propagation from a fenestral light field

[80] Describes algorithmic propagation of information in the direction opposing the temporal flow of light / photons in the real scene.

    - Upstream propagation from an observing camera via reflection / transmission at sufficiently specular intermediating matter (i.e., "indirect observation" by a camera)
- QR calculates the (2-DOF) angular uncertainty u($\mathbf{n}_{fin}$) of the refined normal vector using uncertainty propagation logic that accounts for the uncertainty of each used parameter of the BLIF, incident radiels, and exitant radiels, as well as the sensitivity[81] of $\mathbf{n}_{fin}$ to (changes in) each of those parameters. Calculation of such propagated uncertainties is paramount in QR's ability to provide interpretable measures of the accuracy of reconstructed scene characteristics.
- In calculating u($\mathbf{n}_{fin}$), QR accounts for the degree to which the exitant radiels represent complete angular coverage over exitant directions (2π steradians in the case of an opaque surface).

[81] Quantified, for example, using partial or directional derivatives

8 Dynamism

8.1 Key attributes

Key attributes of dynamism represented in scene models:

- A scene model can represent dynamism in its matter field and/or light field.
- Supported types of matter field dynamism (all Newtonian in nature) include:
  - Rigid body motion (translation and rotation)
  - Scaling (uniform or non-uniform in X, Y, Z)
  - Elastic deformation
  - Plastic deformation
- Dynamism can exist at multiple, overlapping time scales, ranging from gradual change over years to continuous change of significant magnitude many times per second. Examples:
  - Flicker from manmade light sources is one common kind of periodic light field dynamism where a significant change in radiance happens many times per second.
  - The surface of a concrete building weathers over 5 years, leading to a significant change in the BLIF of its matter field.
  - Over the course of a day, the outdoor light field changes gradually as the Sun makes its way across the sky. A QR user virtualizes different parts of a building exterior at different times throughout this day. During one capture session around noon, a small cloud covers the Sun for 10 seconds. In reconstruction processing, QR represents the light field dynamism at both time scales.
- Represented dynamism can have discontinuities, always occurring at elsie origins in spacetime.[82] Examples:
  - Someone turns on a second bank of ceiling lights during a QV capture session. The sudden (discontinuous) brightening of the room's light field is represented by elsies with their origin's time coordinate set to the instant of brightening.
- The dynamic behavior of a set of related elsies can be represented (often more compactly) using one or more dynamic analytic entities associated with those elsies. Examples:
  - The matter field motion of a rolling bowling ball of surfels is represented by a sphere center translating with some velocity, the sphere's radius, and the ball's orientation (rotation state) at some reference time. QR then calculates the position and orientation of any of the ball's surfels for any subsequent time point during the ball's motion. BLIF calculations could be done to predict exitant radiance at ball surfels, for reconstruction or other processing goals.
- Scene boundaries can change as dynamic scene components change.
- When a postulated scene model represents dynamism in certain scene characteristics, QR can calculate loss residuals for observations of those characteristics that have been recorded at different (real) times. To form residuals, the dynamic model can be used to predict values of the (parameters that quantify those) characteristics at the capture times.

[82] Origins of elsies representing dynamism (aka "dynamic elsies") are defined in spacetime [**x**, t].

9 Display

9.1 Key attributes

Key attributes of display of scene models:

- QR can display any entity in a scene realistically, aka in a realistic style, aka photorealistically.
- QR can display any entity in a scene "analytically", using one or more non-realistic representations that symbolize characteristics of the entity.
- In a single displayed image of a scene, a single entity can be displayed realistically and/or analytically.[83]
- When a single entity occurs both realistically and analytically in a single displayed image of a scene, the analytic representations of the entity typically overlie its realistic view.
- In a scene, only its plenoptic field can be displayed realistically. Scene model entities outside the plenoptic field[84] can only be displayed analytically.
- QR offers rich configurability and control over display modes.
- QR defaults to a display mode that presents key information helpful to human the current processing goal.
- In both realistic and analytic display, a scene's plenoptic field and frontier blend seamlessly.
- QR uses analytic display to present various measures of scene model fidelity, especially uncertainty.

[83] Exaggerated pixel radiance deviations used in Intuitive Capture display fall into the analytic category.
[84] See figure 4.6-1 "Quidient Reality Model".

Appendix A – Scene Reconstruction Accuracy

Measurement uncertainty is a challenging domain. Complex mathematics, software, sensing, procedures, and concepts are often used to address it properly. See the “Guide to the Expression of Uncertainty in Measurement (GUM)” for an introduction to internationally approved metrology concepts. In this exhibit, working definitions of accuracy applied to the emerging domain of scene reconstruction are proposed. These definitions along with AI/ML based “perceptual” definitions will be refined / added as appropriate.

Introduction

Kinkaku-ji (the Golden Temple) and the sunlit grounds that surround it is an example of a real scene. A digital model of such a scene can be created using sensors and computers through a process called “scene reconstruction.” To assess the accuracy of reconstructed models of a scene like Kinkaku-ji, scene reconstruction accuracy metrics are proposed.

Approaches

There are two approaches that can be used to assess the accuracy of reconstructed scene models:

Comparison Approach

A universally accepted way to assess the accuracy of a model is to compare it to a more accurate model. According to this approach, the model being tested is called the Model Under Test (MUT) and the model that is serving as the “reference” is called the reference model (REF). If REF is very accurate relative to the MUT, it may be thought of as “true” or “real” and is sometimes called a “ground truth” model. Notwithstanding the accuracy of REF, it is helpful to keep in mind that it is a model, too, also subject to error.

Self-Consistency Approach

An alternative approach to accuracy assessment is the use of a self-consistency metric. For example, in photogrammetry, the repeatability of multiple observations of the same features can be used to compute a self-consistency metric. In machine learning, cross-validation can be used to compute another type of self-consistency metric. Suffice it to say (i) self-consistency metrics are basically measures of repeatability rather than accuracy, (ii) they can be unexpectedly poor estimates of accuracy if used improperly, and (iii) comparison approaches should be used to validate self-consistency approaches before the self-consistency approaches are used “in production”. In the remainder of this document, we will discuss only metrics based on the comparison approach.

Adjustment

Before accuracy metrics can be quantified, corresponding features in the MUT and the reference model (REF) must be located. Then using corresponding feature locations in the MUT and REF model, an adjustment, for example by the method of least squares, is performed to appropriately align the two scene models. Different sets of parameters can be fixed or freed to perform the adjustment depending on the application. Examples of parameters include rigid body parameters (e.g., X, Y, Z, roll, pitch, yaw, scale), other affine parameters, or even warping parameters (e.g., pure images for human consumption that are a little warped can work just fine). Datums can be used as needed to initially align the two models in a common coordinate system. Weighted deviations between certain features are used to perform the adjustment. To the extent that the prescribed adjustment is degenerate, appropriate reductions in degrees of freedom (DOF) are taken in order to perform the adjustment. To the extent that the models are plenoptically ordered in the same way, plenoptic elements are constrained to be in order.

Accuracy Metrics

The weighted root mean square deviation (RMSD) is the performance metric used in this document. Once the adjustment is complete, deviations can be used to compute performance metrics as follows:

$$RMSD = \sqrt{\sum_{i=1\ to\ N} w_i d_i^{\,2}}$$

Where:

RMSD is the <u>absolute</u> accuracy metric

N = the total number of deviations

w_i = the i^{th} weight normalized by the sum of all the weights[85]

d_i = the i^{th} deviation

The <u>relative</u> RMSD is normalized by (divided by) the characteristic size of the space in which the accuracy metric is quantified (the "measurement volume"). It is expressed as a ratio in parts per decade (e.g., 1 part per 1000 or 1 "PPT").

Scene Reconstruction Accuracy

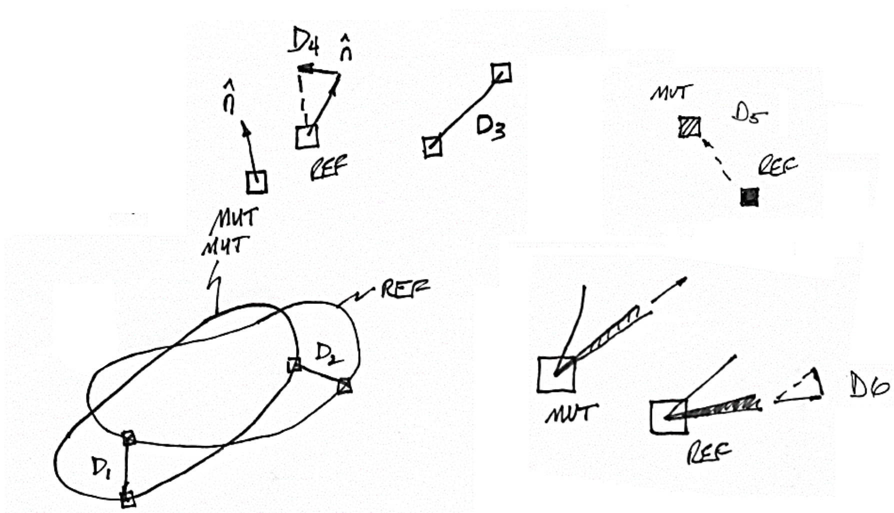

[85] When no weights have been explicitly specified, all weights w_i will default to being equal when calculating RMSD figures.

The figure above schematically shows six of many corresponding deviations that might exist in MUT and REF models of a scene, in spatial locations, after the adjustment is completed:

- D1 is an XYZ deviation between the centers of very small corresponding elsies on the boundary of a corresponding shape. This shape is 6-DOF because of its freeform and asymmetric nature. But it might have BLIF differences that would provide 6 DOF even if the shape were symmetric.
- D2 is an XYZ+BLIF deviation at another set of corresponding elsies on the same shape (BLIF detail not shown).
- D3 is an XYZ deviation between the centers of corresponding pointlike features in the two models.
- D4 is the angular deviation between corresponding normal vector features in the two models.
- D5 shows radiometrically different solid angle elements (the MUT having a lighter shade in this particular sael than the REF model). There is no angular deviation in this case.
- D6 shows light field deviations that are both radiometric and directional.

General plenoptic scene deviations can be restricted in any way to compute other metrics. For example:

BLIF Accuracy (BLA)

A BLIF data structure holds (encodes) output radiance to input irradiance ratios for all input to output sael pairs. BLIF accuracy deviations are defined as follows:

- For each input sael in the BLIF …
  - For each output sael associated with the input sael …
    - There are nine radiometric deviations (three spectral, three polarimetric).
    - And a 2-DOF directional deviation (e.g., quaternions).

Geometric Accuracy (GMA)

Matter fields represented by Quidient's scene models consist of media primitives and associated BLIFs embedded in a plenoptic field. If the BLIFs are removed, what's left is geometry. GMA deviations are defined as follows:

- For each elsie (or "query point" interpolated between elsies) to be considered …
  - There are three deviations associated with XYZ.
  - There is a 2-DOF directional deviation if the elsie is a surfel.

Surfel Position Accuracy (SPA)

In the case of surfels, the three position deviations associated with XYZ can be quantified separately from the normal vector deviations (NVA, defined below).

Normal Vector Accuracy (NVA)

In the case of non-split surfels, the 2-DOF directional deviation associated with the surfel's normal vector can be quantified separately from the surfel position deviation (SPA defined above). The 2-DOF deviation may also be collapsed into a 1-DOF deviation by considering the scalar angle between the surfel's normal vector and a reference normal vector. NVA can be defined as follows:

$$\text{NVA RMSD} = \sqrt{\sum_{\text{i}=1}^{\substack{\text{number}\\ \text{of surfels}}} w_{\text{i}} \left(\cos^{-1}\left(\mathbf{n}_{\text{i,mut}} \cdot \mathbf{n}_{\text{i,ref}}\right)\right)^2}$$

where the w_{i} are weights, $\mathbf{n}_{\text{i,mut}}$ is the normal vector of the surfel in the model-under-test (e.g., reconstruction), and $\mathbf{n}_{\text{i,ref}}$ is the normal vector of the corresponding location (e.g., CPA = closest point of approach) on the surface of the reference model.

Matter field Accuracy (MFA)

MFA deviations include BLA deviations and GMA deviations as defined above.

Radiometric Accuracy (RMA)

Radiometric accuracy deviation is defined as follows:

- For each input sael under consideration …
    - There are nine radiometric deviations (three spectral, three polarimetric).

Directional Accuracy (DRA)

Light fields represented by Quidient's scene models consist of radiel primitives located at elsies (in luminel part of an elsie). If the radiometric characteristics are removed, what's left is directional geometry (like a map of Earth, where boundaries matter, but color doesn't matter). DRA deviations are defined as follows:

- For each direction to be considered …
    - There is a 2-DOF directional deviation.

Light field Accuracy (LFA)

There are four types of light field represented in Quidient scenes at the moment: fenestral, emissive, incident, and responsive. LFA is defined for a particular type of light field at one or more points in one or more directions. Light field accuracy deviations are defined as follows:

- For each point location in the light field …
    - There is an XYZ deviation,
    - And a 2-DOF directional deviation.
    - For each sael at the particular point light field that points in the specified range of directions …
        - There are nine radiometric deviations (three spectral, three polarimetric).
        - And a 2-DOF directional deviation (e.g., quaternions).

Size of the Measurement Volume

A "measurement volume" is a region of space in which scene reconstruction is performed to a specified target accuracy and in which a specified set of SRA metrics may be computed. The size of a measurement volume is determined by factors including the reconstruction use case, camera capabilities, expected systematic error in scene modeling (e.g., BLIFs), OOI sizes, chosen camera/lens resolution, camera/system calibration fidelity, and arrangement of optical markers and scale bar(s) that provide a datum ("cage") against which the plenoptic field is resolved.