

# #166

## TOWARDS EMBODIED 3D ISOVISTS

Incorporating cognitively-motivated semantics of `space` and the architectural environment in 3D visibility analysis

---

**JAKUB KRUKAR**

Institute for Geoinformatics, University of Muenster, Germany  
krukar@uni-muenster.de

**CARL SCHULTZ**

Institute for Geoinformatics, University of Muenster, Germany  
schultzc@uni-muenster.de

**MEHUL BHATT**

University of Bremen, Germany  
bhatt@uni-bremen.de

---

### ABSTRACT

Isovist analysis is being increasingly applied to the third dimension. However, this transition seems to be driven by a direct extrapolation of the 2D isovist concept, along with a variety of implicit assumptions that are only valid in the top-down 2D case. This results in definitions that make the 3D isovist equivalent to the geometric volume of `space` visually accessible from the generating point in the 3D building environment. Such a concept neither adequately reflects the strategy of human visuo-locomotive exploration, nor does it account for the resulting mental representation of the explored space. We review two reasons for the dissimilarity between the 2D and 3D conceptualisations of isovists:

- (1) The quantity, and characteristics, of visual information relevant to the human everyday experience is different in the vertical and horizontal plane.
- (2) The left/right symmetry in accessing and interpreting visual information is not comparable to the top/down asymmetry of visual information. As a result, the difference between the experience of being above an object versus being below the object is incomparable, in contrast to experiencing an object from the left side versus the right side.

Thus, a 3D isovist derived by (naively) extrapolating from its 2D counterpart limits its applicability in the studies of human cognition inside 3D environments. We propose a cognitively motivated extension of a general 3D isovist that accounts for these phenomena.

### KEYWORDS

Isovist, visibility, 3D

### 1. INTRODUCTION

Despite being originally considered as a three-dimensional concept (Benedikt, 1979), isovists or viewsheds have typically been implemented as 2D or 2.5D computational models of visibility. Recently, there has been a growing interest in implementing isovists in the third dimension in order to more comprehensively represent visibility in buildings and outdoor landscapes. It is

now important to consider the ways in which the traditional 2D implementation is extended, and what this means to the initially conceived meaning of isovists, as well as related approaches, such as Visibility Graph Analysis (Turner, et al. 2001). Simple extrapolation of the traditional 2D methods yields a 3D isovist equivalent to the volume of geometrical space visually accessible from its vantage point. This is not what isovists were intended to represent. Isovists are assumed to carry information - information potentially accessible (through the visual sense) by, and relevant to, a hypothetical human explorer of that space. Therefore, the overwhelming majority of isovist implementations implicitly or explicitly carry the following key assumptions:

**The Information-Content Assumption.** Isovists can be understood as containers for visual information available from the vantage point of the human perceiver. The shape and size of an isovist is meaningful and interpretable *because* it is assumed to carry information. This can be information as basic as understanding the shape of the surrounding space. Different information can be relevant in the context of specific analyses, for instance: the potential attractiveness of the view (Shach-Pinsly, et al., 2011), co-visibility of settlements (Brughmans, et al., 2015), navigational choice available at a junction (Meilinger, et al., 2012), or emotional associations with the geometric shape of space (Wiener et al., 2007).

**The Human-Standpoint Assumption.** In the majority of applied architectural analyses, the vantage point of the isovist(s) is considered to be corresponding to a possible location of the eye of a human occupant, or to a location potentially visible by an eye of a human occupant. Thus, the properties and dimensions of a human body are relevant in isovist analyses.

It is important to recognise that the existence of such implicit assumptions is not a drawback of the isovist concept. They guide the analysis when the researcher faces an arbitrary choice (such as what resolution of Visibility Graph Analysis to pick; how to fix the height offset in a 2.5D landscape viewshed model; which surfaces should be modelled as penetrable). They also ensure the ecological validity of the isovist analysis, allowing researchers to interpret it with regard to the human experience.

It is the position of this paper that directly extrapolating the traditional 2D analysis into the third dimension (in the form equivalent to the 'visible geometrical space surrounding the vantage point' - from now on referred to as a 'volumetric 3D isovist'), violates the information-content assumption. In consequence, the volumetric 3D isovist computation does not fulfil the same function as its 2D predecessor was envisioned to.

We develop a novel approach of *embodied 3D isovists* that adopts richer notions of 3D visibility by distinguishing semantic, meaningful sub-regions of the traditional volumetric 3D isovist. We derive the semantic sub-regions by explicitly incorporating inherent, qualitatively salient properties of space, primarily orientation, as well as the semantics of the environment itself with respect to the affordances of human visuo-locomotion. We formalise these sub-regions as classes within the Building Information Modelling (BIM) paradigm, by extending the 'range space' *spatial artefact* class (Bhatt, et al., 2009; Bhatt, et al., 2012). We show how novel metrics that analyse changing relationships between these sub-regions capture key distinctions in the subjective visual impression of various environments. We have developed a prototype software tool for computing the embodied 3D isovist, that we use to demonstrate our metrics in exemplary scenarios<sup>1</sup>.

We emphasise that our refined 3D visibility model is situated within Building Information Modelling (BIM), in the form of *spatial artefacts* (Bhatt, et al., 2009; Bhatt, et al., 2012) – surfaces in the built environment are not simply abstract geometric shapes, but are *objects* within a BIM: walls, slabs, ramps, doors, openings, furnishing, artefactual spaces, and so on. Thus, our visibility model is not only used to numerically quantify the geometric volumes and surface areas of visibility spaces, but can identify the classes (and unique identifiers) of objects that interact with our semantically meaningful visibility sub-spaces.

1 We have implemented our prototype tool as a C++ plug-in to the BIM-based *InSpace3D* architectural design analysis system (Schultz & Bhatt, 2013). We implemented the 3D volumetric isovist generator using ray tracing. 3D visualisations were created using *glc\_player*: <http://www.glc-player.net/>

## 2. LIMITATIONS OF VOLUMETRIC 3D ISOVISTS

The construction of our body and the workings of our visual system have consequences on how we move through the world, accumulate the visual information, and perceive the world. Our visual field is wider than higher (consider how we would see space if our eyes were aligned vertically on our forehead), and we move along the horizontal axes more than along the vertical ones. The result of this arrangement is that:

- (a) **Horizontal visual information is easier to access than vertical information.** It is also more often relevant to our everyday operation. We *see more* horizontally than vertically and for the majority of the time we care more about what is located horizontally than vertically.
- (b) **The right/left symmetry of our body and of our visual field is countered by the top/down asymmetry of our body and of the expected visual input.** Moving (or perceiving visual information) from left to right is relatively equivalent to moving (or perceiving) from right to left. Moving (or perceiving) from top to down is not equivalent to moving (or perceiving) from bottom up.

The above two features of human visual perception are not reflected in the volumetric 3D isovist analysis for two key reasons:

- (a) Volumetric 3D isovists treat 1 horizontal unit of spatial information as equivalent to 1 vertical unit of spatial information.
- (b) Volumetric 3D isovists do not consider any form of vertical asymmetry.

Both issues likely result from its direct extrapolation from 2D isovist implementations, as those are linked to the symmetric (horizontal) way of looking at the world and only afford quantification of the horizontal geometry of the visibility polygon. As a consequence, the volumetric 3D isovist is unable to model visual information in a manner relevant to the human three-dimensional experience of space. The way humans perceive the world horizontally is not equivalent to the way they perceive the world vertically. In practical terms, a person faced with a real-world task inside a building, be it dynamic (e.g. 'find room 103'), or static (e.g. 'express how much you like the entrance hall'), will not explore, visually sample, or remember horizontal information in a manner equivalent to the vertical information.

A three-dimensional isovist representation aiming to describe the 'experience-in-space' (Benedikt, 1979, p. 63) must account for the embodied limitations of human spatial perception, and the resulting cognitive strategies of spontaneous information acquisition. The volumetric 3D isovist approach fails to satisfy this requirement.

## 3. EMBODIED 3D ISOVIST

We define the *embodied 3D isovist* that enables the analysis of three-dimensional space grounded in the way humans perceive and explore information present in the visible shape of that space. In this section we present a sample operationalisation of the concept. We then demonstrate how the embodied 3D isovist differentiates between features of architectural space unaccounted for within the 'classic 2D' and the 'volumetric 3D isovist' approaches.

Our operationalised definition is based on three steps.

**Step (1)** Generate the volumetric 3D isovist (any standard method from computational geometry will suffice).

**Step (2)** Classify surfaces on the boundary of the isovist that meet surfaces in the environment according to absolute directions (top, down, left, right, front, back). We define the orientation of each surface based on the magnitude of the vertical component of the surface normal, with respect to the environment's spatial frame of reference<sup>2</sup>. Vertical surfaces are symmetrical - a

<sup>2</sup> Intuitively, the surface normal is an arrow pointing perpendicular, away from the surface, e.g. the top of a flat floor slab has a surface normal of  $(x,y,z) = (0,0,1)$ , where  $z$  is the vertical component. The length of surface normals are 1 i.e. they are normalized vectors.

left-hand wall affords the occupant the same set of actions, and perceptual experiences, as a right-hand wall. However, horizontal surfaces are not - floors are associated with different affordances than ceilings both in terms of actions they allow and information they are correlated with. For this reason, we consider vertical surfaces to be qualitatively equivalent; horizontal surfaces are classified as downwards facing surfaces (usually ceilings) and upwards facing surfaces (usually floors). In our prototypical implementation we defined: *top* surfaces (e.g. ceilings) having normals with  $z < -0.8$ , *bottom* surfaces (e.g. floors) having normals with  $z > 0.8$ , and *vertical* surfaces having normals with  $|z| < 0.1$ .

**Step (3)** Calculate the volumes of three-dimensional figures defined by the isovists' vantage point and the three types of surfaces (all walls, all floors, and all ceilings). In a cubic space, these figures will have identical 'pyramid' shapes. Their apex will always be defined by the vantage point. This is demonstrated in Figure 1.

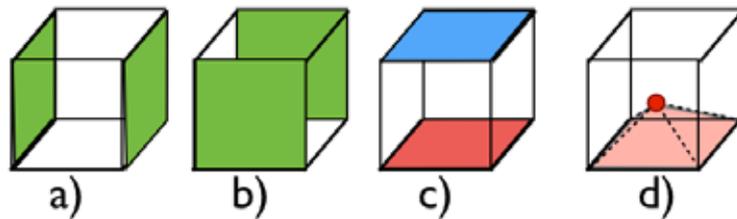


Figure 1 - Wall surfaces (a) and (b) are considered as equivalent, floors and ceilings (c) are considered to be distinct surfaces. Together with the vantage point, oriented surfaces are used to define the volume of a figure constituting each directional part of the 3D isovist (d).

### 3.1. APPLYING THE EMBODIED 3D ISOVIST

In the previous section we divided the isovist volume into semantically distinct regions based on the orientation of visible surfaces. We will now define relations *between* those regions.

Consider a cube with an isovist vantage point located in its centre. This defines 6 identical pyramids which share one apex. Each wall, floor, and ceiling surface defines the base of a single pyramid. The *height* of each pyramid is the distance from the base to the apex, in the direction of the normal of the base. The volume of a pyramid with *base area* and *height* is:

$$volume = \frac{1}{3} \cdot base \cdot height$$

We refer to pyramids with a floor or ceiling base as *vertical pyramids*, and we refer to pyramids with a wall base as *horizontal pyramids*. The relation of the average height of two vertical pyramids to the average height of all four horizontal pyramids equals  $0.5 / 0.5 = 1$  (further referred to as 'vertical-to-horizontal', or *v-h ratio*). The relation of the top-surface pyramid volume to the bottom-surface pyramid volume is  $0.5 / 0.5 = 1.0$ . This ratio will be referred to as 'top-down ratio' (*t-d ratio*).

$$v-h \text{ ratio} = \frac{\text{average height of vertical pyramids}}{\text{average height of horizontal pyramids}}$$

$$t-d \text{ ratio} = \frac{\text{sum of top-pyramid volumes}}{\text{sum of bottom-pyramid volumes}}$$

A person 170 cm tall, standing in an average room of 2.5 m height, will have their 'vantage point' located above the room's center yielding the *t-d ratio* lower than 1<sup>3</sup>. In the example case of a (very small) cube-shaped room with side length 2.5m, the *t-d ratio* is calculated as: top-volume / bottom-volume =  $((1/3) \times 2.5 \times 2.5 \times (2.5-1.7)) / ((1/3) \times 2.5 \times 2.5 \times 1.7) = \sim 0.47$ .

In most common everyday situations, human height above the floor is fixed and rooms are rarely lower than 2.5 m. Therefore, *t-d ratio* much lower than 1 is unusual (although such a proportion can be associated with the observer standing on an indoor balcony near the ceiling - this situation will be reviewed later). Moving into spaces with high ceilings, on the contrary, is associated with a growing *t-d ratio*. Consider a larger cube of dimensions 50 x 50 x 50 metres with a vantage point fixed at 1.6 m above the centre of its floor. The *v-h ratio* in such a space also equals 25/25 = 1.0, reflecting the fact that the perceived shape of the environment, and the relation between the accessible vertical and horizontal information, is the same as in the smaller cube. The *t-d ratio*, however, is  $((1/3) \times 50 \times 50 \times (50-1.6)) / ((1/3) \times 50 \times 50 \times 1.6) = (50 - 1.6) / 1.6 = 30.25$ , reflecting the perceived verticality of the larger space. Neither of these numbers change when the observer moves into the corner of the cube. Similarly to the measurands proposed by Benedikt for the 2D quantification, additional values can be obtained to reflect the position of the vantage point with respect to the boundaries of such a figure (e.g. based on the length and variance of distances to the isovist's boundaries). These are not explored in detail in the current article.

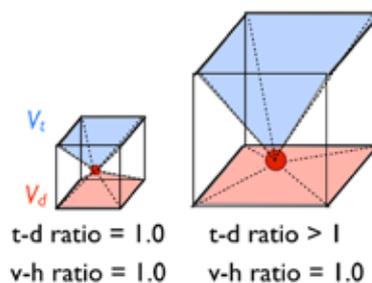


Figure 2 - Two cubes of different dimensions with a vantage point located either in the centre (left) or close to the floor (right). The 'horizontal pyramids' extending to wall surfaces are not visualised, but in both cases the average height of the two vertical pyramids divided by the average height of the horizontal pyramids yields the *v-h ratio* of 1.0.

Jointly, the *v-h ratio* and the *t-d ratio* can therefore describe the shape of perceived space, while differentiating between horizontal and vertical information, as well as between the information visible upwards and downwards (Figure 3). A space low and wide will have the *v-h ratio* much lower than 1, and the *t-d ratio* will depart from 1.0 as the observer's vertical position changes: the *t-d ratio* will grow if the observer has more information above, and will decrease when the observer is located closer to the ceiling. Conversely, a space which is narrow and high will result in the *v-h ratio* much greater than 1.

3 One might be tempted to consider the relationship between the volumes of vertical and horizontal pyramids. However, this does not usefully reflect the shape of an environment. For example, in a unit cube the ratio of the volume of vertical and horizontal pyramids is:  $2 / 4 = 0.5$ . If we vertically *stretch* the cube into a tall rectangular cuboid then we might expect this ratio to increase reflecting greater vertical information, however *the proportion of horizontal pyramid volume to vertical pyramid volume stays the same*. This is because, while the height of the vertical pyramids increases, so does the surface area of the walls (i.e. the *base* of the horizontal pyramids).

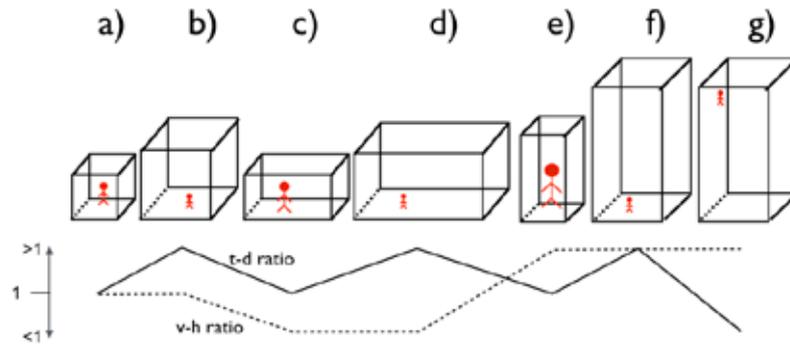


Figure 3 - A schematic (not to scale) representation of the relation between the *v-h ratio* and the *t-d ratio* in convex environments. Cases (a), (c), and (e) represent a vantage point located exactly in the middle of the space's height.

An example reflecting the need for the upwards and downwards distinction are patios with vista openings into multiple floors or balconies. Consider a simplified environment presented in Figure 4. The proportion of horizontal information available to the viewer is limited (i.e. quantified as the relative volume of the horizontal pyramids), therefore the *v-h ratio* reaches a value comparable with very tall but narrow spaces without such balconies. The environment presented in Figure 4 has the dimensions of  $40 \times 20 \times 20$ , but its *v-h ratio* (derived for a 'person standing' on the floor) equals  $\sim 2.6$ . A space without such balconies would need to have the dimensions of  $40 \times 20 \times 78$  in order to achieve the same *v-h ratio*.

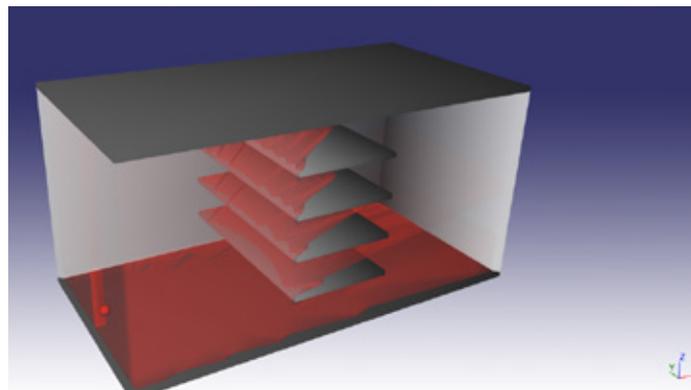


Figure 4 - A screenshot of the volumetric 3D isovist on a simplified environment of dimensions  $40 \times 20 \times 20$ , with 'balconies'. The vantage point is located 1.6 units above the floor, near the left hand-side wall. Note how the balconies occlude the visibility of the opposite wall. This results in decreased volumes of 'horizontal pyramids' (i.e. the portion of the 3D isovist that hits wall surfaces).

Distinguishing between very high and narrow, but convex spaces and those which are low but have many vertical occlusions is possible with 'vertical jaggedness' - a measure analogous to the inverse of *circularity* defined for the 2D case (Benedikt, 1979). This can be defined as the cubic root of summed vertical isovist volume, divided by the square root of the summed surface area of all upwards and downwards facing surfaces (Figure 5).

$$\text{vertical jaggedness} = \frac{\sqrt[3]{\text{sum of vertical pyramid volumes}}}{\sqrt{\text{sum of vertical pyramid bases}}}$$

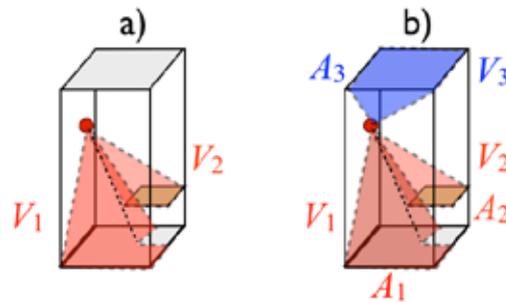


Figure 5 - (a) The total 'downwards volume' is the sum of  $V_1$  and  $V_2$ ; (b) Vertical jaggedness jointly considers downwards and upwards volumes ( $V_1+V_2+V_3$ ), together with all upwards facing and downwards facing visible surface areas ( $A_1+A_2+A_3$ ).

Note how a change in vertical jaggedness affects the meaning of the *t-d ratio* (Figure 6). In a highly convex space, a low *t-d ratio* has a limited functional merit but clearly offers a different perceptual experience from a high *t-d ratio*. On the contrary, in a vertically 'spiky' environment, its functional value is much more evident. Being able to easily see a large proportion of walkable surfaces has a different value in an environment that occludes them more often. It can also be speculated that seeing 'a lot of air' above one's head has different aesthetic connotations in those spaces which are: (1) narrow, compared to those which are wide; (2) in those which are convex compared to those which are jagged; as well as (3) in those situations when we have a lot of information underneath us and those when we stand on the lowest bottom floor. Most importantly, interpretations of these measurands are not convertible: changes in *t-d ratio* values larger than 1 might have a different impact on human experience compared to changes in *t-d ratio* values lower than 1, even if the change is of the same numerical magnitude (Figures 6 and 7).

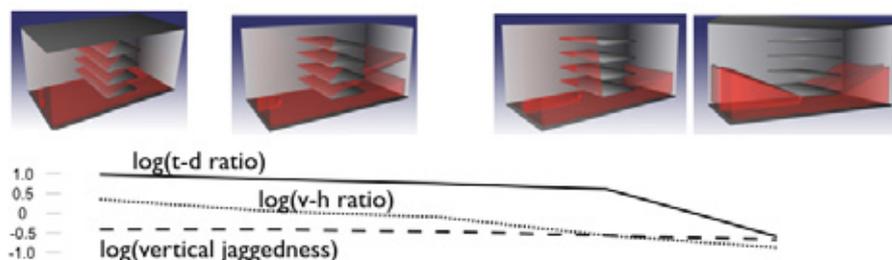


Figure 6 - The relation between three metrics depending on the vantage point's position in a mock up environment with balconies.

he implications of the distinction presented above will depend on the context of the analysis. Benedikt (1979) considers the Guggenheim Museum in order to demonstrate how a person moving along the ramp experiences almost no change in isovist (unless they step closer to the ramp's edge). While the volumetric amount of 'range' visually accessible to the user indeed does not change throughout such a walk, there is a qualitative difference between being at the top and at the bottom of the ramp. This qualitative difference is relevant to the direct perceptual experience, as well as to the mental representation one can build based on the available visual information. As one moves downwards along the ramp, more and more 'ceilings' become visible at the cost of 'floors'. The navigator stops seeing surfaces from top-down, and is forced to perceive a larger part of the horizontal surface of the ramp from below. Firstly, this can affect the aesthetic experience of individual vistas. Secondly, it affects the amount of navigational choice one can realise from individual viewpoints. Humans can walk on surfaces that are facing up, but not on those facing down. Looking at the underside of ramps, from below, requires an additional cognitive step of inferring navigational actions potentially available on the unseen

top side of the ramp. This can be a non-trivial mental operation in some cases, but it is not equivalent to being completely unaware of the possible existence of these walkable surfaces.

Being able to view artworks hanging on the walls of the gallery can also be accounted for by the embodied 3D isovist analysis. In a gallery of a shape similar to the Guggenheim's, the fact that artworks are hung closer to the floor than to the corresponding ceiling of each wall means that the ratio of visible ceilings-to-floors is positively correlated with the occlusion of the artworks. As the visitor moves downwards along the ramp, they might be seeing the same total proportion of walls, but more of its crucial part (the one containing artworks) becomes occluded by the ramps that are visible from below (Figure 8).

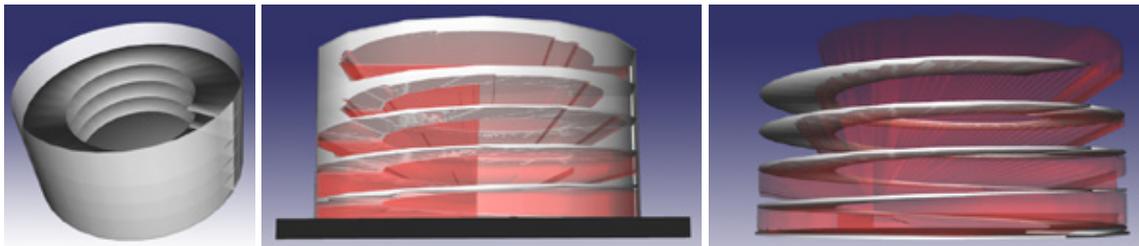


Figure 7 - Simplified 3D model of the Guggenheim museum with 3D isovists.

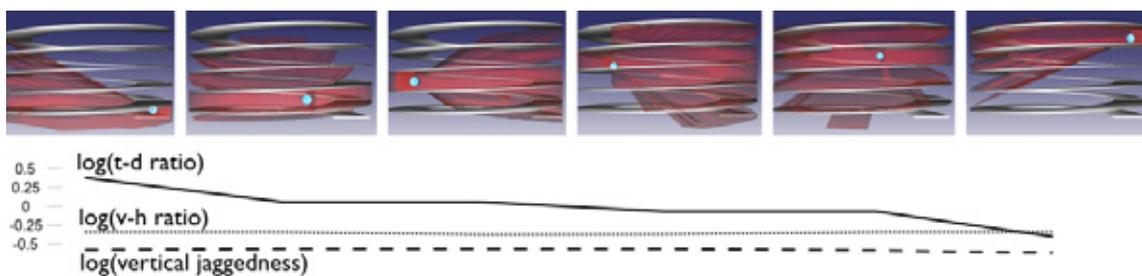


Figure 8 - As the vantage point moves upwards along the ramp, the visible wall surface area remains constant. So does the amount of vertical information (and therefore the ratio between the two). However, the type of the vertical information changes, reflected in the decreasing *t-d ratio*.

## 4. RELATED WORK ON 2D AND 3D ISOVISTS

### 4.1 IMPLEMENTATIONS IN 2D, 2.5D, AND 3D

The concept of isovist has been introduced by Tandy (1967) and popularised by Benedikt (1979). Despite considering isovists as three-dimensional on the conceptual level, early computational implementations were, by necessity, two-dimensional. They build on the familiar top-down approach of visualising building plans and operationalise the isovist as a polygon covering the portion of the environment visible from its single vantage point. In his seminal paper, Benedikt lists a number of metrics that can be used to describe such a shape, and suggests their relevance to the spatial experience of a building user.

Perhaps the most influential extension of the original notion has been proposed by Turner et al. (2001) with the development of Visibility Graph Analysis (VGA). In it, multiple isovists are generated from numerous vantage points at a fixed resolution, leading to the creation of a connectivity graph of individual isovists. This enables the study of entire buildings as a continuous space and accounts for visual relations between its subparts. Despite permitting to connect the graph across separate floors, Turner's original VGA is intrinsically two-dimensional, meaning that the method for calculating each individual isovist is based on 'scanning' the space on a single horizontal plane around a vantage point in a top-down model.

As isovists were increasingly being applied in outdoor contexts (such as landscape planning), there was a growing need to account for three-dimensional visibility. Limited by the availability

of three-dimensional data, researchers utilised digital elevation models, in a '2.5-D' visibility analysis. This approach only accounts for a single z-value for each xy-coordinate, limiting its applicability in architecture, where a single wall can consist of many openings, or 'holes'. Bishop (2003) and Llobera (2003) provide an extensive review of these developments.

The availability of new methods for sampling and storing three-dimensional data (e.g. Building Information Modelling and LiDAR) accelerated work on three-dimensional isovist implementations. Derix et al. (2008) presented a set of methods for calculating three-dimensional visibility inside architectural spaces. Based on the idea of the visibility graph, these methods can account for the average person's height, and quantify the changes in how open or constrained the space might seem to a potential navigator.

Loneragan and Hedley (2015) provide a broad review of recent approaches to modeling three-dimensional visibility, together with the contextual differences across these implementations. The authors expand the concept of an isovist's origin and target in order to account for the complexity of three-dimensional geometry in a manner relevant to common visibility analysis applications. For instance, a 'point-to-volume' model can describe how a single human observer might visually access the interior of an apartment located in an opposite building. A 'volume-to-volume' model would specify if any part of the target apartment is visible from any part of the origin apartment. It might be a design challenge to minimise such occurrences, without taking into account which exact parts of one's dwelling are visible from their neighbour's living space. In the similar urban context, Fisher-Gewirtzman (2016) presented a 3D visibility model which accounts for the semantic property of visible parts of the urban environment such as roads and dwellings.

Focusing on the architectural application, Varoudis and Psarra (2014) extended the traditional VGA approach to the third-dimension by accounting for accessibility affordances of floors. Their approach starts with defining 'accessible' and 'inaccessible' spaces in the layout, which are represented as a three-dimensional grid of isovist vantage points. A 'mixed' visibility graph consisting of undirected and directed edges, is then generated. Classical 'undirected' edges are created between nodes representing locations which can serve both as origin and as destination of a potential observer 'looking out' towards the other node; 'directed' edges represent connections between two spaces, of which only one can serve as a potential destination. This can reflect a situation when the observer looks at a void, high in the air above the floor's surface.

The focus on the relation between visibility and movement is traditional to Space Syntax, however the approach that describes human three-dimensional perception as the binary relation of visual and navigational accessibility is limiting. Firstly, it assumes that the human observer always knows whether a distant visible space is traversable or not. This leaves no possibility for modelling uncertainties and inferences involved in the exploration of space (information simply is, or is not available to the viewer). Second, it treats all spaces as equally accessible or inaccessible. This is not consistent with how humans perceive and act in it (e.g. considering all other factors equal, accessing higher ground is usually more effortful). Third, it gives equal weight to vertical and horizontal information, encouraging the assumption that seeing a tall, narrow corridor has the same perceptual effect as a wide, long room with a low ceiling.

Space Syntax as a discipline is believed to have detached the architectural analysis from the single viewpoint, putting the emphasis on the configurational properties of the whole space. But it is important to recognise that the primary units of that configurational analysis have originally been derived from the properties of human body and human cognition: convex spaces represent 'areas potentially co-occupied by multiple people', axial lines are derived from the 'longest lines of sight', and VGA relies on a 'number of moves' required to see the whole space (see [Conroy Dalton, 2005] for a more comprehensive discussion of these implicit human-centred assumptions in Space Syntax). The problem hereby discussed seems to arise from the fact that these basic human-centred assumptions differ in the three-dimensional context and are invalid if simply extrapolated directly from the two-dimensional simplification. It is the goal of this paper to challenge those basic underlying assumptions of the volumetric 3D isovist, and not to turn away from the configurational analysis back into the view from 'the inside' of a

single observer. Further extensions of the presented work, and their inclusion in configurational methods such as VGA, is desired, but not implemented within this paper. The work reviewed above demonstrates how, in the third-dimension, richer notions of visibility are needed that take into account many more inherent properties of 'space' such as orientation, as well as the semantics of the environment itself. To our knowledge, however, the existing implementation of 3D isovists treat space extending in each direction as equivalent, omitting the crucial role that vertical orientation plays in our perception of the environment.

#### 4.2 ISOVISTS AND HUMAN COGNITION

In the two-dimensional context, research investigated the predictive role of isovist properties in relation to human visual perception and spatial cognition. Conroy Dalton (2001) and later Meilinger et al. (2012), linked characteristics of 'partial isovists' to navigational behaviour of participants exploring an immersive virtual environment. Wiener et al. (2007) demonstrated how a small set of isovist measurands correlates with human navigational behaviour (such as finding the most/least visible place in the environment) and experiential ratings e.g. of beauty and spaciousness. In computer-based eye-tracking studies later applied to the outdoor context by Emo (2014), Wiener et al. (2012) demonstrated how the geometry of architectural space visible from the navigators' perspective can explain the preferred choice between two alternative routes. Both works demonstrated how extensively participants explored the horizontal axis of pictures presented to them (focusing on areas offering the longest 'line of sight'), but superficially explored the vertical information further from the horizon. Both indoor and outdoor application of this paradigm demonstrated that the usefulness of vertical information (mostly containing floor and ceiling surfaces) is very limited in a spatial task, and therefore this axis is not extensively explored.

#### 4.3 EMBODIED COGNITION

The concept of embodiment of cognition refers to the fact that human cognitive capacities and strategies depend on the organisation of the body, and therefore actions and processes afforded (but also limited) by it. This approach departs from Cartesian dualism of mind and body, as it assumes that cognitive representations cannot be studied separately from the physical properties of the world. A classic example of some implications are studies which linked positive evaluations of unknown symbols with the action of approaching (flexing arm), and negative evaluations with avoidance (arm extension) (Cacioppo et al., 1993). Robbins and Aydede (2009) provide an excellent primer on this and related concepts, known broadly as situated cognition. For the most comprehensive review of studies demonstrating how a broad range of other spatial cognitive processes is embodied and situated, see (Tversky, 2009). Part of the situated cognition approach is the view of the external environment as acting as 'external memory' to our perception (Myin & O'Regan, 2009). It assumes that in real-life tasks people use the external world as a container for potentially relevant information. Knowing where to (visually) search for something is all we often need (see also (Tatler & Land, 2011)). As demonstrated by Wiener et al. (2012), the search for relevant information is not uniformly or randomly distributed in all space. Not all parts of space are equally expected to carry this required information.

### 5. DISCUSSION AND CONCLUSIONS

We have presented a model of 3D isovists that accounts for the embodied nature of the human perception of space. It also makes it possible to distinguish the situated context of the observer, such as the varying usefulness of top-down visibility in different types of environments. This is a semantic approach to modelling 3D visibility in the sense that we enrich the concept of visible space by defining meaningful sub-spaces (i.e. spatial regions that partition visibility space), based on embodied cognition.

Our model can be used as a basis for extending work known from the 2D implementation in a manner compatible with the available evidence from situated cognition. This is relevant to interpreting well-established isovist measurands in the three-dimensional context. Measurands

such as drift, revelation or jaggedness can have varying behavioral and cognitive correlates when differentiated across the vertical directions. Thus, their interpretation should not be extended directly from the empirical evidence using two-dimensional data and data collected in contexts with negligible need for acquiring vertical information.

A research goal introduced by this distinction is to establish what is the proportional importance of horizontal and vertical information in different contexts. This raises the challenge of representing such asymmetries in a holistic model of larger space. So far, aggregated visibility analyses such as VGA considered and aggregated isovists which all have the same 'rules' of generating them and treated all information uniformly. If any asymmetries were considered, they would be identical across the whole modeled space.

In his review of Space Syntax' contribution to environmental psychology, Montello (2007) points to limitations of the isovist analysis. Applying a 'one-size-fits-all' visibility representation falls short in accounting for known contextual and interpersonal differences in spatial cognition. In this paper we argue that accounting for direction is one of the required building blocks for a concept that Montello calls 'weighted' or 'probabilistic' isovists - a representation that not only conveys what is possible to see, but what is likely to be seen (by differing groups, in differing contexts). This requires reviewing the very basic assumptions of isovist analysis in order to align it with the established evidence on how human users explore, perceive, and cognise architectural space.

#### ACKNOWLEDGMENTS

Jakub Krukar's research at the Spatial Intelligence Lab was supported by the ERC StRG Grant Agreement No 637645.

## REFERENCES

- Benedikt, M. L. (1979). To take hold of space: isovists and isovist fields. *Environment and Planning B: Planning and Design*, 6, 47–65.
- Bhatt, M., Dylla, F., & Hois, J. (2009). *Spatio-terminological Inference for the Design of Ambient Environments*. In Hornsby K.S., Claramunt C., Denis M., & Ligozat G. (Eds.), *Spatial Information Theory. COSIT 2009. Lecture Notes in Computer Science* (Vol. 5756, pp. 371–391). Berlin, Heidelberg: Springer.
- Bhatt, M., Schultz, C., & Minqian Huang. (2012). The shape of empty space: Human-centred cognitive foundations in computing for spatial design. In *2012 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (pp. 33–40). Innsbruck, Austria: IEEE.
- Bishop, I. D. (2003). Assessment of visual qualities, impacts, and behaviours, in the landscape, by using measures of visibility. *Environment and Planning B: Planning and Design*, 30(5), 677–688.
- Bruhmanns, T., Keay, S., & Earl, G. (2015). Understanding Inter-settlement Visibility in Iron Age and Roman Southern Spain with Exponential Random Graph Models for Visibility Networks. *Journal of Archaeological Method and Theory*, 22(1), 58–143.
- Cacioppo, J. T., Priester, J. R., & Berntson, G. G. (1993). Rudimentary determinants of attitudes: II. Arm flexion and extension have differential effects on attitudes. *Journal of Personality and Social Psychology*, 65(1), 5–17.
- Conroy Dalton, R. (2001). *Spatial navigation in immersive virtual environments*. Unpublished PhD Thesis. University of London.
- Conroy Dalton, R. (2005). Space syntax and spatial cognition. *World Architecture*, 185(41–47), 107–111.
- Derix, C., Gamlesæter, Å., & Carranza, P. M. (2008). 3d Isovists and Spatial Sensations: Two methods and a case study. In S. Haq, C. Hölscher, & S. Torggrude (Eds.), *Movement and Orientation in Built Environments: Evaluating Design Rationale and User Cognition* (pp. 67–72). Bremen: Universitat Bremen.
- Emo, B. (2014). Seeing the Axial Line: Evidence from Wayfinding Experiments. *Behavioral Sciences*, 4(3), 167–180.
- Fisher-Gewirtzman, D. (2016). Integrating weighted views to quantitative 3D visibility analysis as a predictive tool for perception of space. *Environment and Planning B: Planning and Design*.
- Llobera, M. (2003). Extending GIS-based visual analysis: the concept of visualsapes. *International Journal of Geographical Information Science*, 17–1, 37–41.
- Lonergan, C., & Hedley, N. (2015). Unpacking isovists: a framework for 3D spatial visibility analysis. *Cartography and Geographic Information Science*, 406, 1–16.
- Meilinger, T., Franz, G., & Bühlhoff, H. H. (2012). From isovists via mental representations to behaviour: first steps toward closing the causal chain. *Environment and Planning B: Planning and Design*, 39(1), 48–62.
- Montello, D. R. (2007). The contribution of space syntax to a comprehensive theory of environmental psychology. In *Proceedings of the 6th International Space Syntax Symposium* (p. iv:01-12). Istanbul.
- Myin, E., & O'Regan, J. K. (2009). Situated Perception and Sensation in Vision and Other Modalities. In M. Aydede & P. Robbins (Eds.), *The Cambridge Handbook of Situated Cognition* (pp. 185–200). New York: Cambridge University Press.
- Robbins, P., & Aydede, M. (2009). A Short Primer on Situated Cognition. In M. Aydede & P. Robbins (Eds.), *The Cambridge Handbook of Situated Cognition* (pp. 3–10). New York: Cambridge University Press.
- Schultz, C., & Bhatt, M. (2013). InSpace3D: A middleware for built environment data access and analytics. *Procedia Computer Science*, 18, 80–89.
- Shach-Pinsly, D., Fisher-Gewirtzman, D., & Burt, M. (2011). Visual exposure and visual openness: an integrated approach and comparative evaluation. *Journal of Urban Design*, 16(2), 233–256.
- Tandy, C. R. V. (1967). The isovist method of landscape survey. In H. C. Murray (Ed.), *Methods of Landscape Analysis* (pp. 9–10). London: Landscape Research Group.
- Tatler, B. W., & Land, M. F. (2011). Vision and the representation of the surroundings in spatial memory. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 366(1564), 596–610.
- Turner, A., Doxa, M., O'Sullivan, D., & Penn, A. (2001). From isovists to visibility graphs: a methodology for the analysis of architectural space. *Environment and Planning B: Planning and Design*, 28(1), 103–121.

- Tversky, B. (2009). Spatial cognition: Embodied and situated. In M. Aydede & P. Robbins (Eds.), *The Cambridge Handbook of Situated Cognition* (pp. 201–216). New York: Cambridge University Press.
- Varoudis, T., & Psarra, S. (2014). Beyond two dimensions: Architecture through three-dimensional visibility graph analysis. *The Journal of Space Syntax*, 5(1), 91–108.
- Wiener, J. M., Franz, G., Rossmanith, N., Reichelt, A., Mallot, H. A., & Bühlhoff, H. H. (2007). Isovist analysis captures properties of space relevant for locomotion and experience. *Perception*, 36(7), 1066 – 1083.
- Wiener, J. M., Hölscher, C., Büchner, S., & Konieczny, L. (2012). Gaze behaviour during space perception and spatial decision making. *Psychological Research*, 76(6), 713–729.