Proceedings of the 14th International Space Syntax Symposium

079

Mean Aggregate Isovist Cascade Analysis

A temporal approach to spatial analysis.

Sam McElhinney, Canterbury School of Architecture, University for the Creative Arts

smcelhinney@uca.uk

ABSTRACT

Our paper distinguishes between two types of graph structures; 'sparse' and 'dense.' Sparse graphs, such as axial line graphs of cities, can succinctly summate vast spatial networks. In contrast, dense graphs, such as Visibility Graph Analysis (VGA) of architectural spaces, bring nodal redundancy and scalar challenges. As the spatial detail of a VGA grid increases, the computational cost of measures such as integration grows exponentially. We posit that Turner's development of VGA (2001) overlooked qualities of isovists that resolve such issues.

As units of spatial perception, isovists bind inter-visibility relations within their geometry, and encode potential further connectivities through their occlusive edges. We use these insights to generate 'isovist cascades' of visibility relations across architectural plans in real-time; starting from a single isovist in space, seeding new isovists from its occlusive edges, and expanding until all space is covered. We demonstrate how propagation of isovist cascades from stochastically evenly distributed locations, and their subsequent concatenation, produce high definition fields of mean visual depth and integration.

Our approach allows significant gains in computation speed and detail over VGA. Additionally, the temporal resolution of isovist cascade analysis (ICA) facilitates three observations that may offer insights to the cognitive understanding of space. Firstly, the minimum stochastic seed total to establish a mathematically stable field, (our 'sufficient set'), is low, being circa 15 - 50 for all plans. Such findings have implications for cognitive spatial mapping, indicating that surprisingly few orientation points may be necessary for effective navigation.

Secondly, after factoring a logarithmic reduction of field change over time, our case study plans show consistent residual variations, each with distinct amplitude and range, or 'spatial wobble'. We evidence how such values might describe topological qualities of spatial systems in a dimensionless manner; providing a single coefficient metric and a broader prototype relational matrix that classifies spatial systems from the multiaxial to the mono-cursive and from entropic to self-similar.

Finally, we discuss a paradox wherein apparently 'complex' spatial configurations can exhibit 'spatial wobble' factors lower than 'simple' examples. Elaborate spatial systems often contain distinct or recognisable structures that, in turn, afford development of landmark knowledge. Conversely, 'simple' systems often have indistinct and or unexpectedly entropic sub-features, and so afford disorientation. It appears our method reflects such cognitive easements and challenges, suggesting a novel metric for the ease of development of 'sufficient' or functional spatial cognition.

KEYWORDS

isovists, graphs, integration, intelligibility, entropy, complexity, spatial cognition.

1 INTRODUCTION

Much of the contemporary interest in and use of Space Syntax involves the application of computational methodologies to provide empirical insights into spatial configurations (Krenz et al, 2023). Such techniques produce numerous metric measures to evidence and aid understanding of human social behaviours, configurative aspects of space, and infrastructural issues of city and regional conurbations (Hillier et al, 1987; 1993; Peponis et al, 1989). Underlying the calculation of many such metrics is a reliance on graph analysis and links to graph theory.

In our paper we argue that an alternate consideration of partial and temporally changing analytical models might lead to insights and new theory regarding the development of cognitive knowledge of space. We begin by reviewing challenges and certain draw-backs of graph centred approaches when applied to spatial analysis. In so doing we introduce a distinction between types of graph that are more or less appropriate for 'Space Syntax'. We then discuss how a deeper consideration of isovist properties and relations might resolve such issues.

The second part of our paper outlines an alternative temporal approach for spatial analysis of an architectural or urban plan. We introduce a methodology that uses the successive overlap and concatenation of 'isovist cascade' geometries to rapidly produce high definition 'mean visual depth' and 'integration' fields. We subsequently explore the continual-time development of such information within a range of plans, examining how informational stability thresholds relate to originating plan complexity.

In the final part of our paper we outline three conclusions that have both practical bearings for future researchers, and may have philosophical bearings on how occupants construct navigational understanding of novel spaces. Firstly, we identify a threshold of time-development that appears to establish sufficiently coherent and reliable results in any plan. Secondly, we identify how subtle variations in result development (imperceptible to the researcher) can be used to define topological characteristics of space. In so doing, we discuss how elaborate spatial systems often contain distinct or recognisable structures that, in turn, afford development of landmark knowledge. Conversely, 'simple' systems can often include indistinct and or proportionally unexpected entropic sub-features which afford disorientation. It appears our method reflects such cognitive easements and challenges, suggesting a novel relational metric of spatial intelligibility and entropy.

1.1 Integration

'Integration' (HH) is a core measure of Space Syntax. It provides *'a normalised measure of distance from any a space of origin to all others in a system'* (Hillier & Hanson, 1984), or a value for how central a space is to all other spaces. It has been shown at length that integration analysis can be used to interrogate behaviours at all scales of spatial configuration, from individual and team interactions in a building, to pedestrian traffic movements in a city and beyond, and in diverse contexts, even, for instance, including archaeological studies (Hillier et al, 1993; Peponis et al, 1989; Sailer, 2012; Penn et al, 1998; Letesson, 2013; and many more besides)

Superficially, the process of conducting an integration analysis is consistent whether the researcher seeks to consider either a city network or a single room. The standard approach is to translate the subject spatial system into a graph, where 'nodes' represent spaces or junctions, and 'edges' record visual or spatial links between them (Hillier & Hanson, 1984; Hillier et al, 1987; Turner, 2001b). From such a structure the researcher can subsequently calculate the depth (or number of visual steps) between nodes. By doing so from one node to all others, a researcher can find the mean average depth of that node; and by repeating for all nodes, they can calculate the 'all to all' global spatial relations across the entire system. The former is an operation with a computational cost of O(n log n), and so the latter costs O(n^2 log n) in the worst case, or more typically O(n^3), governed by the density of edges between all nodes. The result provides 'mean depth' values per node, which can be normalised to 'integration' values, allowing comparisons between spatial structures (Hiller & Hanson, 1984; Teklenberg et al, 1993; Turner, 2001b).

1.2 Sparse and Dense graphs

We can make our own differentiation of all graph structures into two broad categories; 'sparse' and 'dense' graphs. The former describes a graph in which the number of edges is minimal relative to the number of nodes, with the sparsest having an edge to node count of n-1 (each node having just one edge leading to a new neighbour node). Dense graphs, meanwhile, have high numbers of edges per node, with the densest having an edge to node count approaching n^2, (i.e. each node having edges connecting to all others).

Our distinction between dense and sparse graphs parallels a practical conundrum for Space Syntax researchers. Calculation of global metrics such as integration becomes computationally costly as edge count increases per node. The higher the graph density, the better its fidelity of representation of geometric, experiential space, but the (perhaps prohibitively) greater the analysis time becomes. The sparser a graph, the quicker it will produce results, but the more abstracted the configurative model becomes. In such a way, our distinction aligns to the two

primary types of graph structure within Space Syntax, 'axial line graphs', generally used to consider urban configurations (Hillier & Hanson, 1984); and 'visibility graph analysis' (VGA), generally used for architectural spaces (Turner et al, 2001a; Turner, 2001b) (Fig. 1).

It can be argued that the configurative influence of the geometries of individual spaces are negated in the large-scale axial graph modelling of city networks or regional systems, which consider aggregated movements and behaviours over large distances (Penn, 2003). On the other hand, as a researcher reviews smaller scale spaces, from city plazas to buildings and rooms, geometric qualities become more critical to overall configuration. Here qualitative and configurative aspects of space likely become important to occupant experience and behaviour (Benedikt & Burnham 1985; Peponis et al, 1997).

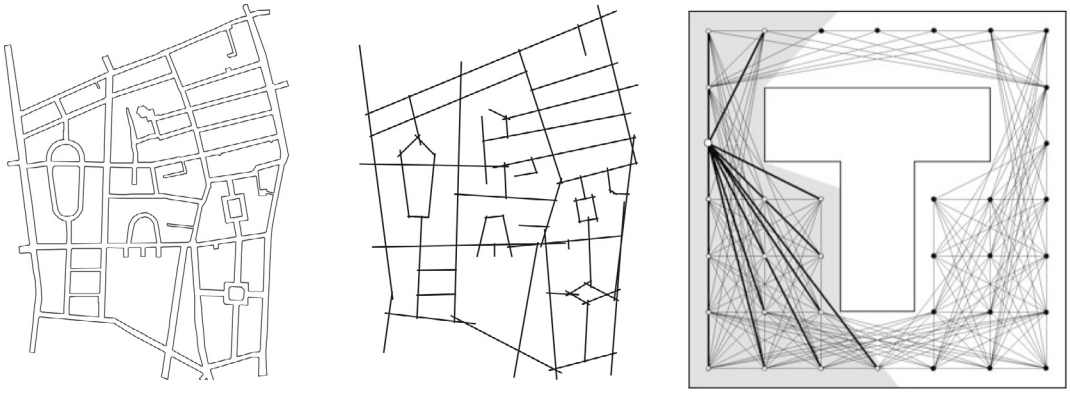

Figure 1: [*Left and centre*] A 'sparse' Axial map as extracted from a street system (after Hillier and Hanson 1984), [*right*] compared to a 'dense' VGA grid in architectural space (from Turner et al, 2001a).

In the case of a city, an 'Axial Line Graph' summates each street via a single, algorithmically derived 'axial line' of the longest vista along said street (Turner et al, 2005). Nodes and edges representing the intersections of all segments of all axial lines are subsequently assembled into an axial map, which becomes an analytic graph (Turner et al, 2005). Such a graph is sparse, with neighbouring nodes usually having distinct edge connections and few duplicated links, reflecting the high degree of change as an individual moves from city junction to city junction. As a result, the computational cost of calculating an axial line based integration analysis is low, and hence such a graph provides an succinct metric summary of a large system such as a city.

Conversely, in the apparently simpler case of an architectural space such as a room, a 'Visibility Graph Analysis' (VGA) is often utilised. In VGA, contiguous architectural space is sub-divided into an orthogonal lattice-grid of 'cells', taken as nodes, and visual connections between cells are

recorded as edges (Turner et al, 2001a). In VGA, neighbouring nodes often have duplicated, near matching sets of edge connections, reflecting the continuity of experience and high consistency of view as an individual moves between locations in architectural space. Such a graph is very dense, with computation costs rising as VGA grid resolution increases. As most VGA inherently requires the entire graph to be assayed to produce a result, a researcher is often forced to wait a long time to obtain integration results, or to compromise and utilise a low resolution of grid cells across their plan. Combined with challenges of normalisation of integration (Dalton et al, 2022a), such issues arguably mean that VGA fails to provide a parsimonious summary structure of space.

In review, at the city scale, the geometric generosities of public space (the width of a boulevard, or the shape of a square) tend to be overlooked, to create a sparse graph that is computationally efficient to analyse. Conversely, at the room scale, the plurality of nodes and edges required to capture architectural space becomes so great that analysis of the graph is inherently flawed. To make an integration analysis viable, the researcher or practitioner must therefore make a pragmatic compromise; either to discard spatial geometry or to impose an arbitrary and distortional, user defined spatial resolution (a discretising grid) upon it.

1.3 Partitioning space; the grid and other stories

DepthMapX's requirement to implement a unit scale for the VGA grid is one attempt to methodologically reconcile the above issues. Turner (2001a) explicitly argues that a grid of 'around one metre' is related to embodied experience, providing meaningful results related to inhabitation and an optimal scale to mitigate edge duplication. Unfortunately, such a scale has a deleterious effect upon finer spatial features (wall thickness, for instance, falls below it, as might window openings), meaning that the uniform application of such a grid can distort architectural geometries, causing variation in integration results. Further parameters such as grid origin point, grid orientation and curve resolution can exacerbate such artefacts (Dalton et al, 2022b). In practice such issues have resulted in adoption of finer grid scales, but doing so leads back to computational challenges and long analysis times.

Given that space is non-discrete, we might reasonably seek a solution that puts aside a cellular grid altogether, in favour of a fundamental spatial basis from which to build integration analysis. Peponis' work on E- and S-partitions does alternately partition space into 'informationally stable' spatial units, derived from geometric change in the continuity of surrounding built surfaces (Peponis et al, 1997). Such subdivisions define regions within which visual information is stable, providing a useful experiential 'unit'. Peponis uses such units as the basis from which to build a graph to examine the integration core of Fallingwater (Peponis and Bellal, 2010), proving the

utility of such an approach. Unfortunately, as the geometry complexity of a space increases, so too does the number of E and S-Partitions generated from it, and thus so too graph density. To a lesser extent, we are returned to the conundrum of meaningful spatial discretisation versus computational effort.

To resolve our dilemma, it helps to consider a hypothetical 'infinitely dense' graph that perfectly describes architectural space. In such a graph, much of the edge density would be duplicated and thus must be redundant. Consider an individual moving between two points in space, A and B, described by nodes on said graph. As the A-B distance reduces, it becomes increasingly likely that the pattern of visibility edges of the two nodes is a match (for an infinitesimal movement, your view won't change). The question therefore becomes not of grid or spatial discretisation, but rather of identification of experiential overlap between A and B. Here lies a path to defining configurative differentiation, through consideration of commonalities of view. We propose that the latter might be better found from true isovist geometries than abstracted graph relations.

1.4 The power of Isovists

As introduced by Benedikt (1979), the 'isovist' provides a single, localised unit of vision, a descriptor for all space that can be directly seen from a given location (Fig.2a). We would suggest that both Turner (2001a) and Peponis (1997) overlooked the combinatory potential of such a unit. Beyond the above simple unitary reading, isovists bind inter-visibility relations within their geometry, and encode potential further connectivities through the phenomena of 'occlusive edges' (Psarra & McElhinney 2014; Gibson, 1979) (Fig.2b). A set of overlapping isovists therefore does not just provide a descriptor of local spatial qualities, but might also inform us about the interlinking of such local visibility networks within a wider global system.

The latter insights might be achieved by studying the commonality and density of isovist overlap. In such a network of overlaps, the need for either a partitioning schema or an arbitrary grid-cell-discretised-graph is negated, and thus the problematic bars to high resolution spatial integration might be circumvented.

Work has already been done that sets a foundation for such an approach. Developing from Benedikt's (1979) original definition of the isovist as 'a bounded system of points visible from a given location', Psarra and McElhinney (2014) identify how they contain 'regions of points that are mutually intervisible' (Fig.2b), suggesting a technique that utilizes commonalities of overlap between single isovists generated at successive stochastically selected locations in space. Such commonalities identify those inter-visible regions within all isovists that recur (or not) in situated

vision, and suggest isovists can be used to both seed and structure a wider spatially relational framework. The stochastic selection, meanwhile, results in a random yet even selection of points through space, circumventing the issues of discretization of space via a scalar grid-lattice, and providing continually observable results, rather than requiring the finite 'finished' moment necessitated by the 'whole graph' analysis of VGA. To date such techniques have only been demonstrated as a method for developing information at local and supra-local levels, describing spatial relations within two visual steps of any given location (Psarra and McElhinney, 2014). The suggestion that isovists become 'dynamic seeds for a global framework of knowledge', whilst hypothesized, is not entirely realized.

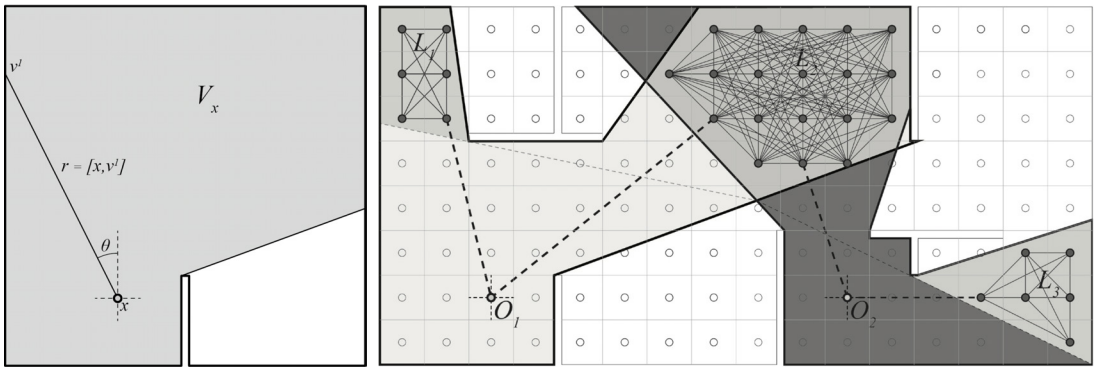

Figure 2: [*Left*], a single radially defined isovist (after Benedikt, 1979), and [*right*], overlapping sequential isovists and their discrete sub-set regions (from Psarra and McElhinney, 2014).

We propose that the above framework can be expanded to global analysis techniques. Below we outline a novel analytical technique which we are calling an 'isovist cascade'. In it, we use occlusive edges and commonality of isovist overlap to calculate and to propagate spatially relational information. Our method repeats the stochastic seeding, continual overlap and informational concatenation, but instead of using single isovists, it constructs chains of linked isovists across a subject plan, facilitating global analyses.

2 MEAN AGGRETAGE ISOVIST ANALYSIS

An 'isovist' cascade is a relational structure within which every point contains information that signifies its own visual depth (or other intervisibility relations), measured across a plan from a given starting location (Fig.3). In calculating an isovist cascade, such information is incrementally propagated from a single origin location out to all points within a subject space. The process of propagation is structured via the occlusive edges of sequentially generated chains of isovists. Such occlusive edges form thresholds of partition and relation between the inter-visible and spatially consistent regions previously discussed and so underpin and enable our approach.

By consolidating shared overlaps among sequential stochastically seeded isovist cascades, we can identify spatial regions characterized by inter-visibility and informational consistency. As each isovist cascade encapsulates globally relational information, such aggregate regions are globally relevant. Regions of overlap that are common to successive isovist cascades will therefore resolve similar global informational values over time. By varying the information that each isovist cascade propagates (such as visual depth, angular depth, or metric depth values), we can thus use isovist cascades to produce a range of globally relational fields, including mean visual depth or integration analysis. Such a process circumvents the problems of definition of spatial partitioning schema or any need to discretise a plan via a scalar lattice-grid to conduct spatial analysis.

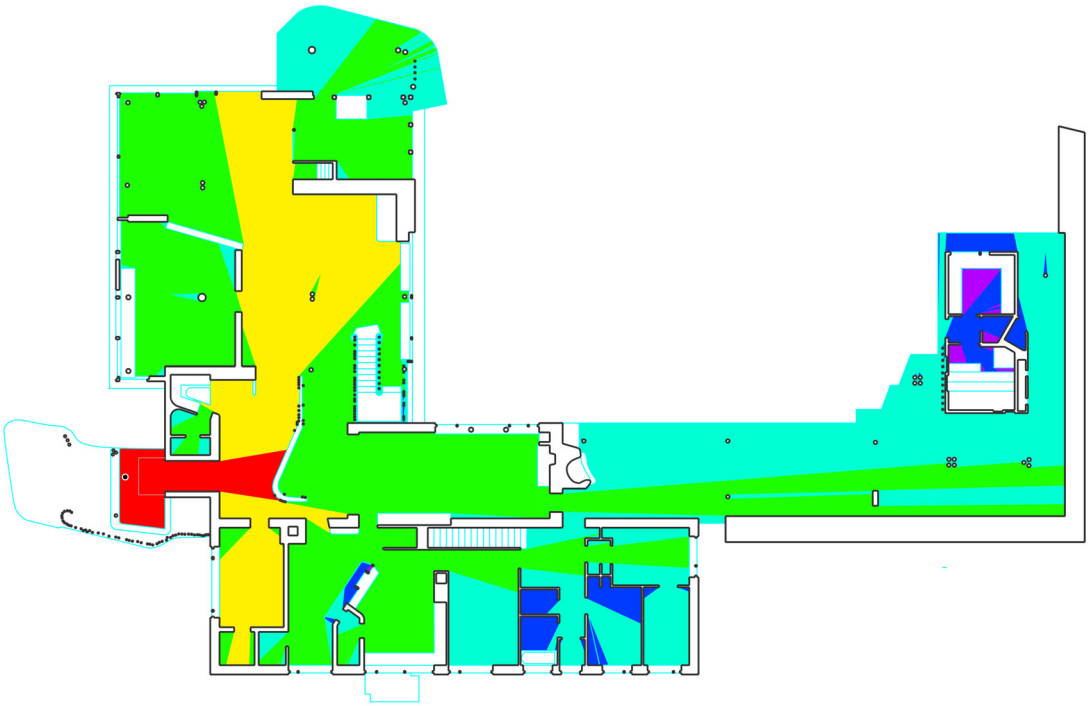

Figure 3: A single Isovist Cascade calculated throughout the Villa Mariea. Cascade is initiated from the left hand side entrance point (red) and progresses through space to the deepest position (purple).

2.1 The isovist cascade process

To initiate an isovist cascade, a seed location is selected from a plan, either by a user, or by machine process. Once a location has been chosen, the isovist geometry from that location is calculated, and the set of all occlusive edges of that isovist are extracted. These occlusive edges provide a 'horizon' of connectivity from the spatial clusters bounded within that original isovist to any neighbouring but 'as yet unseen', spatial regions. A new generation of isovists, seeded along these occlusive edges, yet 'looking' beyond their horizon, can then be calculated. By concatenating these second generation isovists into one compound geometry, a representation

of the 'next visual step' from the seed location can be found. The latter itself potentially includes new occlusive edges, meaning that the seeding process can iterate, from isovist generation, to concatenation, to occlusive isolation, and recursively on to subsequent future isovist generation, until all space has been covered.

When the above process is run in a contiguous architectural space, it produces consecutive wavefronts of isovists. The latter propagate through the subject plan until all points have been 'seen' and added into the isovist cascade. As each new point is included it is assigned a visual step depth from the origin, providing a relational structure (Fig.3). When successive isovist cascades are generated from origin locations that are stochastically evenly distributed throughout an architectural plan, and their point values of relative visual depth are aggregated or averaged over time, a 'mean visual depth' field is incrementally produced. The fields resolved provide mean depth values for all locations throughout a plan that rapidly converge to numeric stability. Such fields are a correlate match for those produced from VGA's intensive grid-cell-lattice, but, as the basis of isovist calculation is true vector geometry, rather than a VGA grid rasterization of spatial geometry, the definition and resolution of such information is essentially unlimited.

In practice, the above can be achieved through a conceptually simple algorithm, which can be written to access the massively parallel processing capabilities of modern GPUs. As a result, our construction of an isovist cascade can be achieved in near real-time. On a 2023 MacBook Pro, isovist cascades can be calculated across a typical building plan in microseconds, and, consequently, visually coherent analysis fields can be produced in circa 60 seconds.

2.2 Practical considerations and conceptual consequences

It is trivial to translate the mean visual depth values produced in the above to an integration value. Doing so is achieved in line with the equation for integration (Teklenburg et al, 1993):

$$dValue = \frac{2\left\{k.\left[\log 2\left(\frac{k+2}{3}\right)-1\right]+1\right\}}{(k-1)(k-2)}$$

$$Integration(HH)_v = \frac{dValue.(k-2)}{2(MeanVisualDepth_v - 1)}$$

In our adoption of the equation, the lack of a graph means that there is no figure available for *k* (the VGA node count). We overcome the latter by substituting *k* with a unitised area value for the subject plan, set by the researcher to whatever scale appropriate. Doing so does not address

problems in VGA normalisation discussed by Dalton et al (2022a), but it is a pragmatic step that provides parity of result with established VGA. One consequence of note is that such a fix explicitly couples integration values to an area unit. Such a relation is perhaps already an underacknowledged implicit consequence of using node count in VGA, the latter being a product of area of plan divided by unit-scale of grid cell.

A practical and conceptual consequence of our method is that calculation of integration shifts from an absolute process to a temporal one. Established VGA integration involves the complete analysis of the whole of a graph that represents a spatial complex. Doing so produces an exact, 'repeatable' result, but requires researchers to wait until the graph is completely assayed. In contrast, our approach develops incrementally over time, with results continually changing and reviewable throughout. Practically speaking, such a process means that a researcher can obtain rough estimate results far more rapidly than waiting for a complete and finite outcome, with consequences for both workflow expediency and potential uptake in real-world design practice (Anklesaria, 2022). Conceptually, our approach opens the potential to consider and measure how a global descriptor such as mean visual depth or integration emerges from layers of local, partial knowledge (Penn, 2003). Work has been done to examine the development of integration from a gradually expanding graph (Dalton et al, 2022a), and link artefacts of such development to qualities of spatial configuration and mechanisms of cognition, particularly via the notion of an 'aha moment' (Charalambous et al, 2021, McElhinney et al, 2022). We expand such work below.

3 MEASURING THE METHOD

Here our aim is to provide evidence and insights into the performance and characteristics of our proposed method. Unlike typical graph analysis, mean aggregate isovist cascade analysis does not produce an absolute outcome. Instead, the values that it calculates converge to stable results over time, and as so it is prudent to be cautious regarding result repeatability. To establish reasonable certainty and methodological guidance regarding the latter, we have considered two primary questions in the following section of our paper.

Firstly, we examine how many cycles of our method are needed to produce consistent results for multiple plans. Such a benchmark is crucial to establish when our analysis converges to a relatively stable set of numeric values that can be relied upon. Secondly, we consider whether characteristics of the plans analysed might themselves affect the mechanics of calculation, and thus cause variation in any moment of convergence on a stable, repeatable result. Here we seek to test whether certain types of plans may require more or fewer iterations of our analysis and examine how such differences may be related to aspects of plan complexity.

3.1 A set of experimental plans and grids

To address the above questions, a range of plans were assembled, covering 36 hypothetical grid spaces, 16 real building spaces, and 7 macro cityscapes. Figures 4a-c provide the integration fields produced for each via our isovist cascade method.

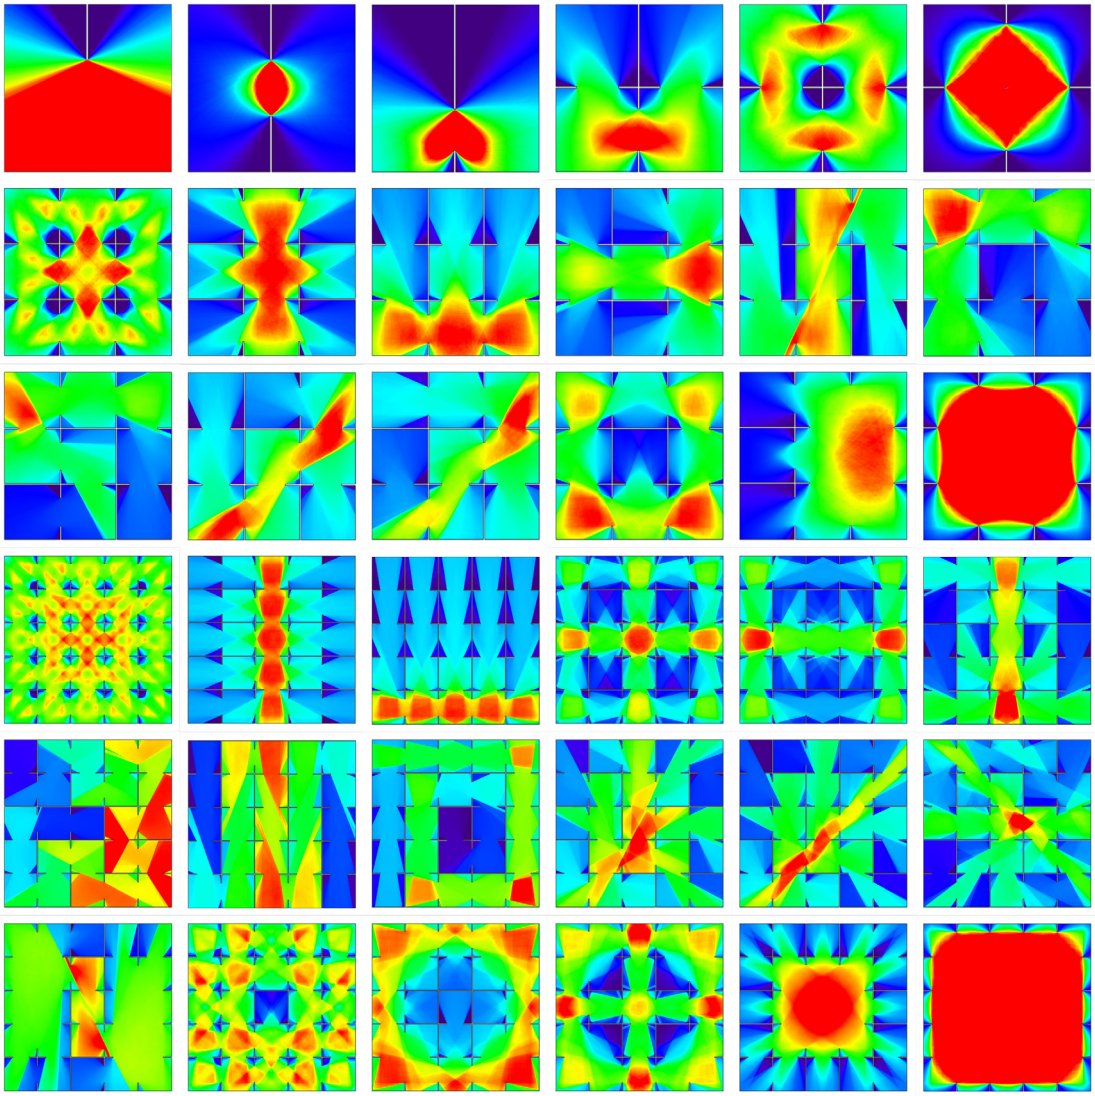

Figure 4a: Integration results fields produced in 36 sample hypothetical grids. [*Top row*], 2-4 cell grids (layouts X1 – X6 left to right); [*second and third rows*], 9 cell grids (layouts X7-X18 left to right); [*final three rows*], 25 cell grids (layouts X19-X36 left to right).

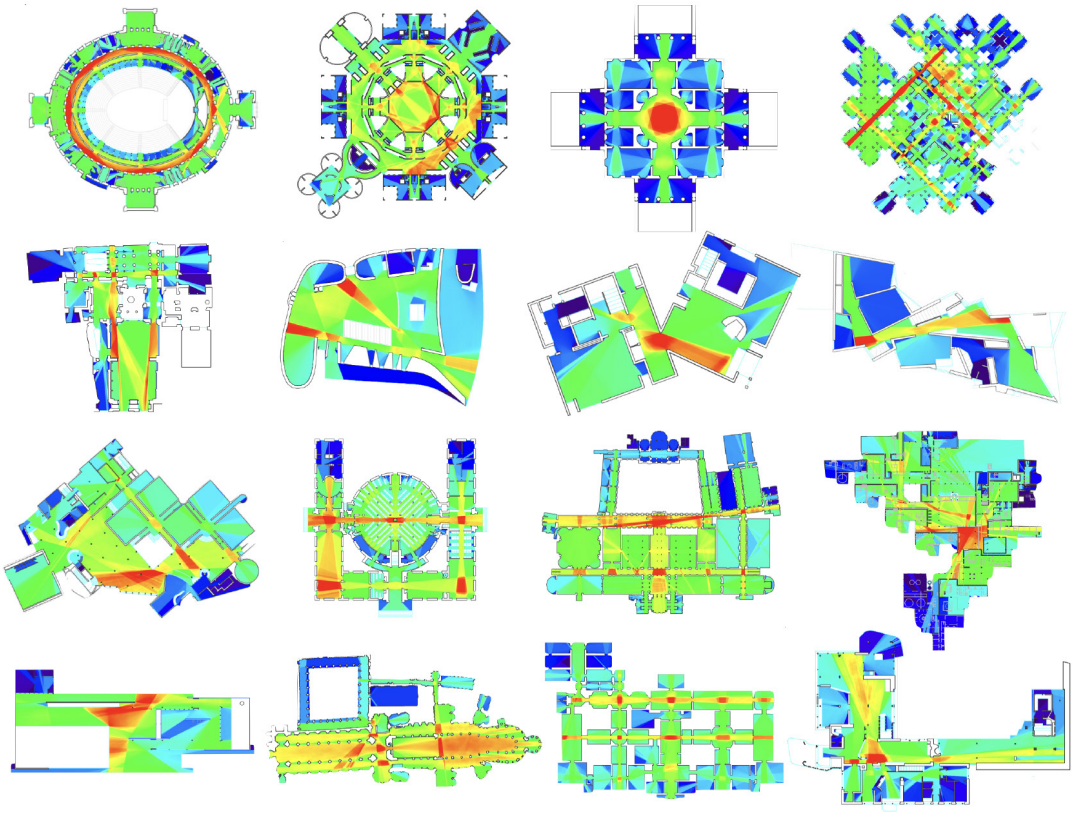

Figure 4b: Integration results fields produced in 16 sample building plans. [*Top row*], left to right; Royal Opera House, Dakha Parliament, Villa Rotunda, Beheer Centraal Offices; [*second row*], left to right; Soane Museum, Ronchamp, Fisher House, Garau Augusti; [*third row*], left to right; Munchengladbach Museum, Stockholm Library, Victoria and Albert Museum, Amsterdam Orphanage; [*final row*], left to right; Barcelona Pavilion, Canterbury Cathedral, Generic Gallery, Villa Mariea.

Our hypothetical grid-space plans were designed to examine variations of spatial separation and connection across a system, and provide insight into how the development of 'complexity' in a plan might affect our method. Based upon underlying grids that move from 2-cell, to 4-cell, to 9-cell and finally 25-cell frameworks, the grid configurations explored families of analogous spatial morphologies, repeated and expanded across the increasing density of frameworks. The 'real world' plans, meanwhile, were collated to provide a broad contrast to such hypothetical systems, and ranged from simple architectural layouts to extended city figure-grounds.

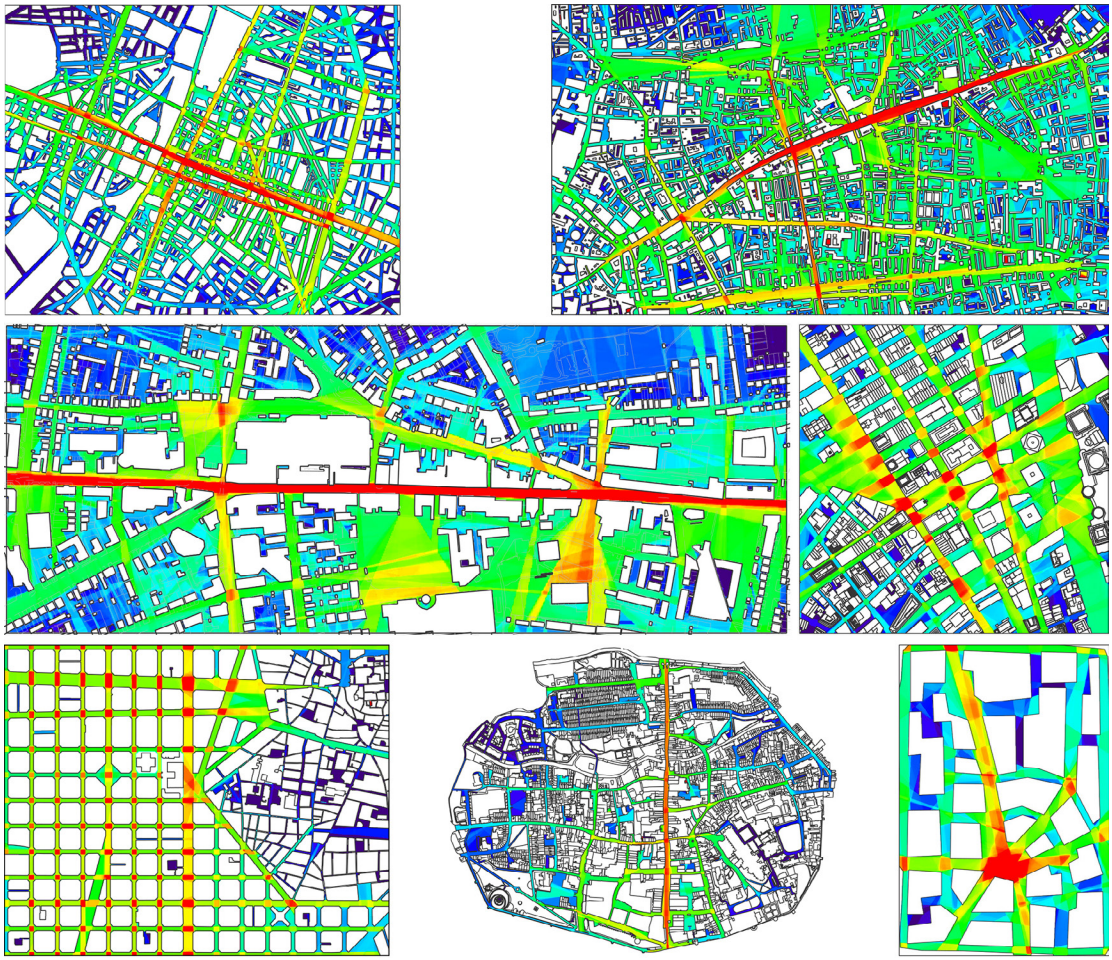

Figure 4c: Integration results fields produced in 7 sample city fragments. [*Top row*] left, Paris in 1890, right, Whitechapel area of London; [*middle row*] left, Houslow, right, Manhattan; [*final row*] left, Barcelona, centre, Canterbury, right, Hillier's 'Intelligible City'.

3.2 Field stability

Here we examine the stabilization of results within mean visual depth or integration fields generated by the isovist cascade method. We do so by recording value variations across the field during repeated analysis cycles.

To quantify field variation over time, we measure absolute change in mean depth value at each point in the plan as each new isovist cascade is aggregated into the overall result. We then calculate a single average 'total field variation' value (in units of step depth) per cycle. Repeating for multiple cycles yields a sequence of average absolute variations. Finally, we normalize our values against the actual average mean depth of each plan. The latter articulates how much the variation in any given cycle meaningfully affects the value of results measured at that point.

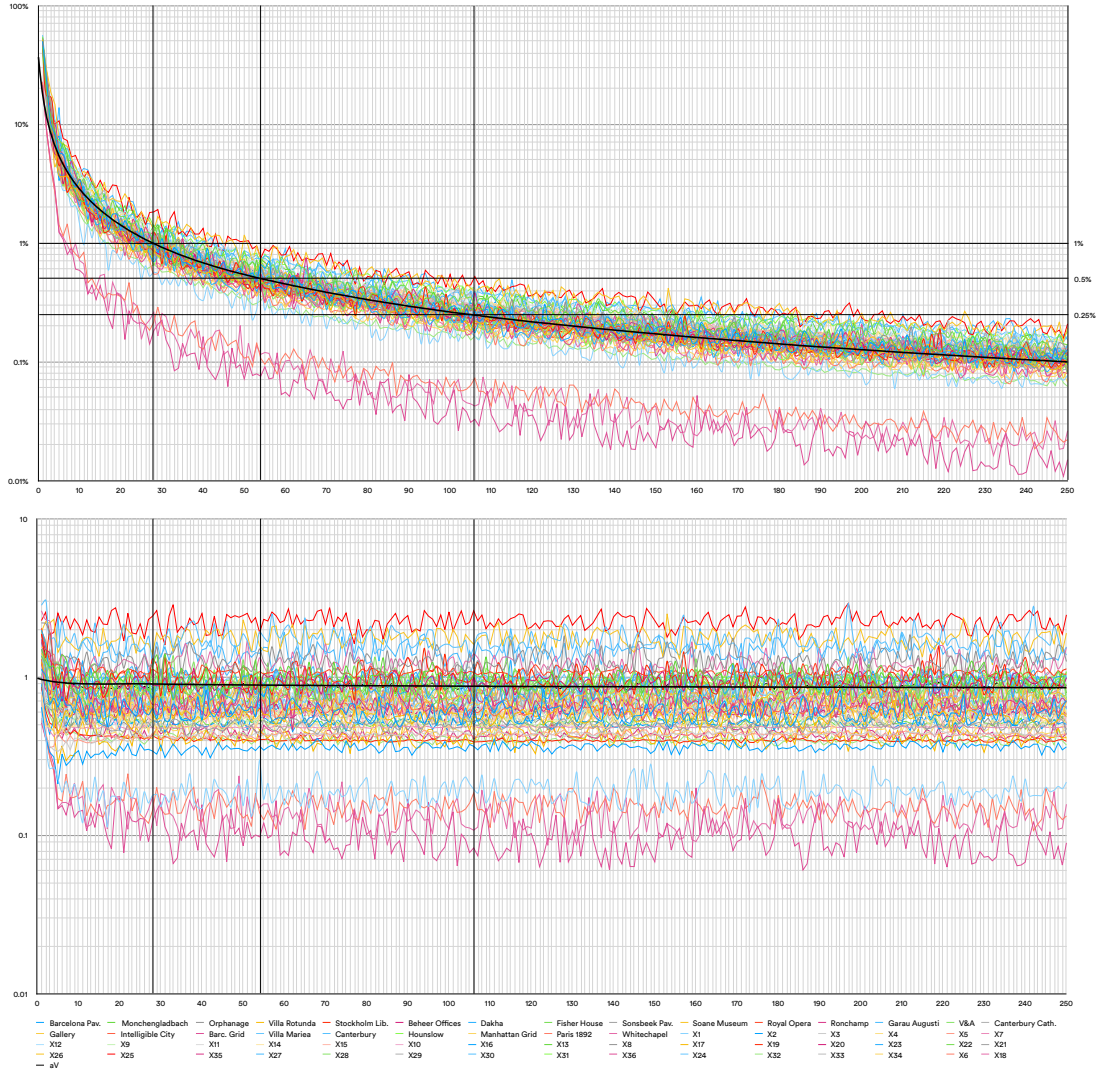

Figure 5: [*5a, above*] Plots of the curves for variation in field values in each plan; [*5b, below*] Plots of the normalised latent variation in field values in each plan. Plots use the same key; average curve is shown in each in black. Outliers are the grid plans X-1, X-6, X-18 and X-36 (refer to Figure 4a).

We repeated our data recording exercise ten times, across 250 isovist cascade cycles, for all 59 of our subject plans, each containing circa 500,000 data points. Figure 5a provides the resulting variation curves for each plan and a summary average curve for all plans.

3.3 'Spatial wobble'

We soon realised that our variation measurements conformed to a power curve, rapidly decreasing absolute variation over time. Such a curve is logical, as it emerges from the mathematical function of the averaging that occurs as the depth values of each new isovist cascade are added to the overall result. To better examine how plan differences might affect our

analysis, we sought to remove the averaging function by multiplying the mean variation value at each cycle by the total number of cycles completed at that stage. Doing so produces values for each plan that describe what we are calling 'latent variation'. Plots of the latent variation curves for each plan are provided in Figure 5b.

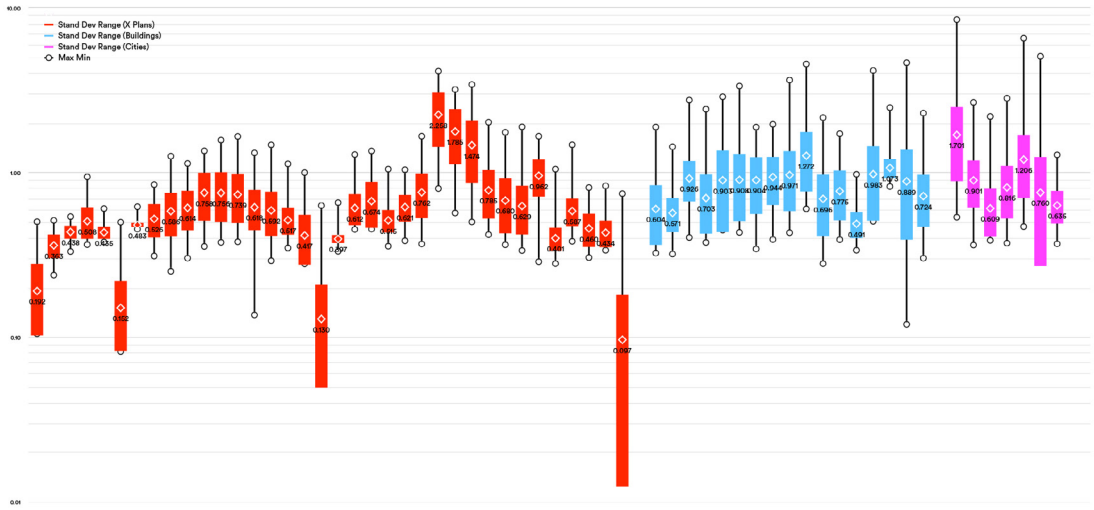

Figure 6: Whisker plots that summarise the information in Figure 5b. Red cluster represents hypothetical grid layouts (in numeric order from left to right); blue, building plans, and pink, cities.

For all plan types the latent variation has two properties, a mean trend that is linear and constant, and a degree of oscillation around that mean. Figure 6 provides a whisker plot summary of these properties, including a mean variation, a standard deviation of variation, and a maximum to minimum range of variation for each plan. Collectively such properties describe a degree of 'spatial wobble' or underlying volatility in the isovist cascade analysis for each plan.

Calculation of 'coefficient of variation' provides us with a metric by which to further interrogate our spatial wobble phenomena. Coefficient of Variation is the ratio of the standard deviation to the mean, ($CV = \frac{\sigma}{\mu}$) and as such is a dimensionless number, (freed from issues of plan scale or depth complexities), allowing us to compare results between spatial systems. Figure 7a shows said value of coefficient of variation calculated and plotted horizontally for each plan, whilst Figure 7b compares the same to the space syntax metric, 'intelligibility' (Hillier et al, 1987), calculated for each plan.

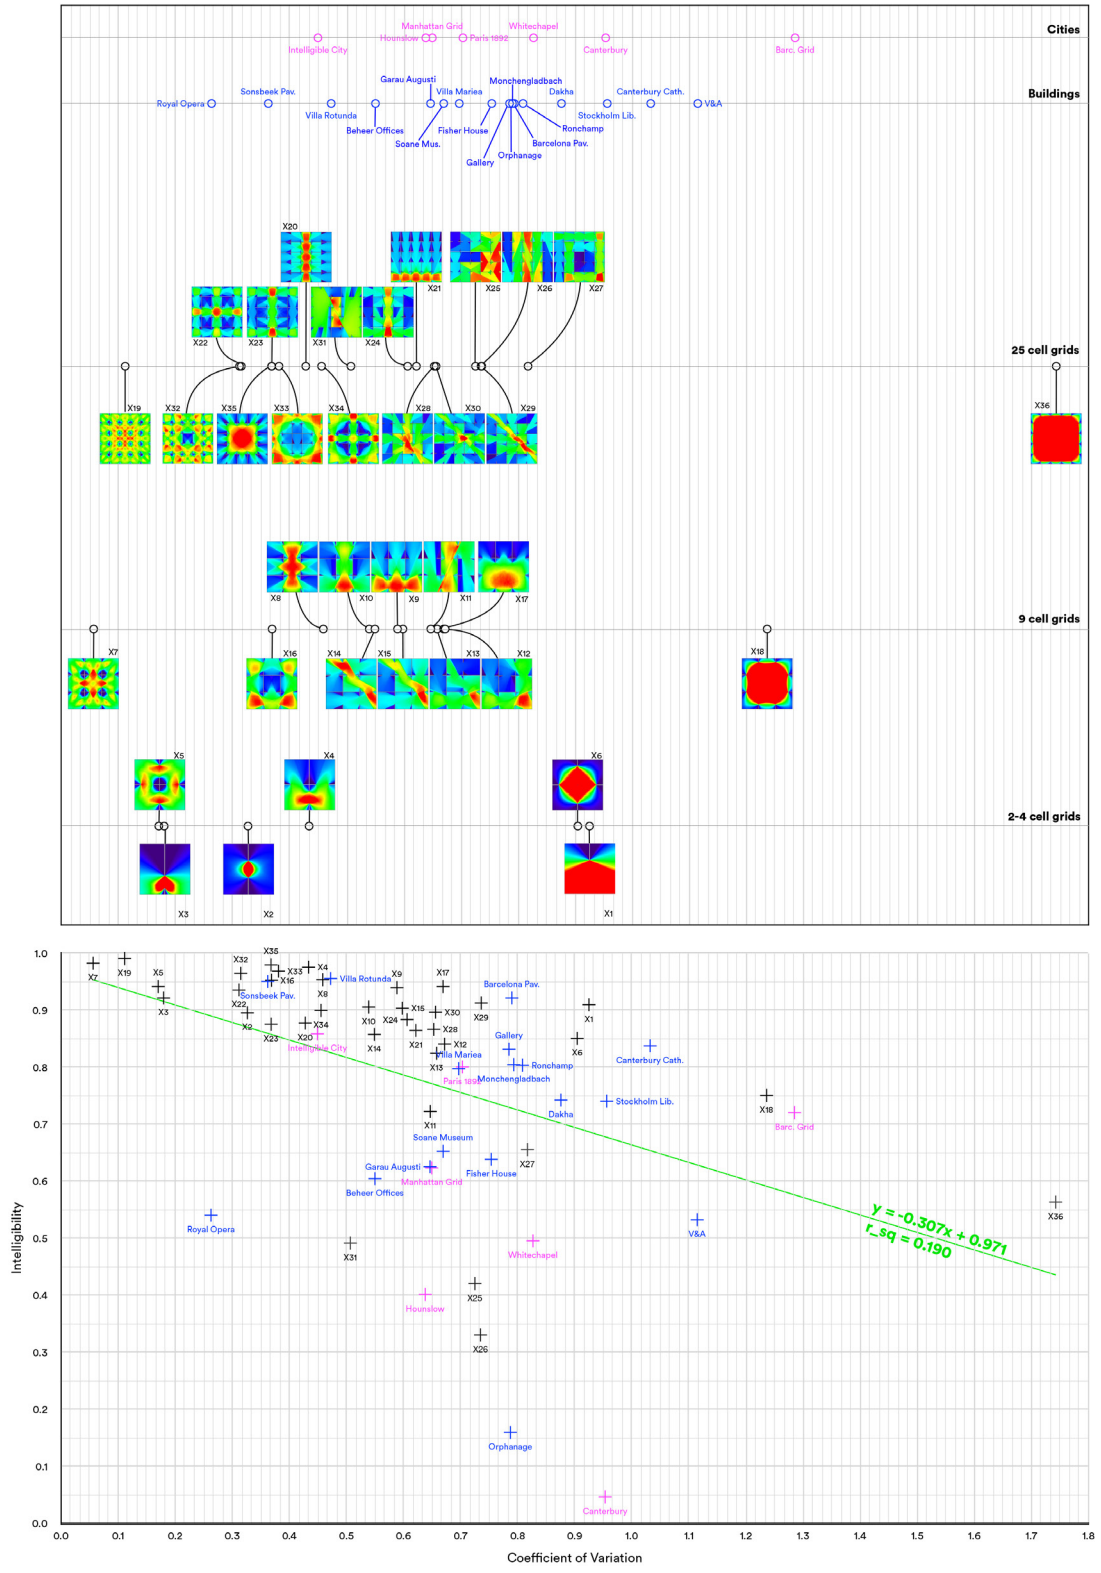

Figure 7: [*a, above*] a linear plot of Coefficient of Variation for each plan; separated into related sub-sets; [*b, below*] a plot of the same against intelligibility for each plan (r^2 = 0.19). X axis is the same in each plot.

We can also directly plot our standard deviation of latent variation against the mean of the same, to explore the relations of these measures that underly our coefficient of variation metric (Figure 8a). Finally, to remove any effect of 'size' (or more accurately, available depth) of spatial system, we repeated the latter exercise but with normalised values of mean and standard deviation. The former we normalised by dividing it by the average mean visual depth measured in the plan, expressing the mean variation as a fraction of available depth. The latter is divided by the range of variation, expressing the standard deviation of variation relative to the total potential for variation. Figure 8b shows the outcome for all plans.

3.4 Observations

In Figure 5a we see that all plan types exhibit a rapid fall of relative variation over time, conforming to a power curve. There appears to be little difference between plan types, with the only outliers being extreme case hypothetical grid spaces (largely empty 'room' spaces). The latter are anyway instances of rapid drop of relative variation. Working from the average 'curve of all curves', the amount of change or uncertainty in result measured at any given moment drops to 1% of the result value by about 30 cycles. It falls further to 0.5% by 60 cycles, at which point the result field typically appears visually stable to the human eye. Once past 250 cycles of analysis, the average curve is down to 0.1% of result, and all plan types fall under 0.2%.

Once we remove the power-function of averaging from our data, we establish a series of 'latent variation' values that oscillate around level constants for each plan (Fig. 5b). Such constants are best tangibly expressed in terms of visual step values, and omitting the previously mentioned outliers, appear to range from around ±2.5 visual steps, to ±0.3 visual steps of latent variation. By plotting such mean latent variations, the standard deviation of variations, and their min and max ranges in whisker plot form (Fig. 6), we might expect to see differentiations between plan types. Surprisingly such distinctions do not appear (beyond the extreme outliers), with the whisker plots produced by the finer grain hypothetical grids not obviously distinguishable from those of real buildings, or city figure grounds.

We can however make some more intriguing observations when we consider our coefficient of variation metric plots (Fig. 7a). Here we can identify morphological sets, particularly across our hypothetical grid-plans. Lower values of coefficient of variation are observed for more interlinked grid structures, and high value outliers reflect larger 'room-like' spaces. There also appear to be sub-set sequences of spatial change that translate to increases in the coefficient result. At the same time, our coefficient values for each plan show no correlation with the established Space Syntax metric of intelligibility, with an r^2 result of only 0.19 (Fig. 7b). Discussion is provided below.

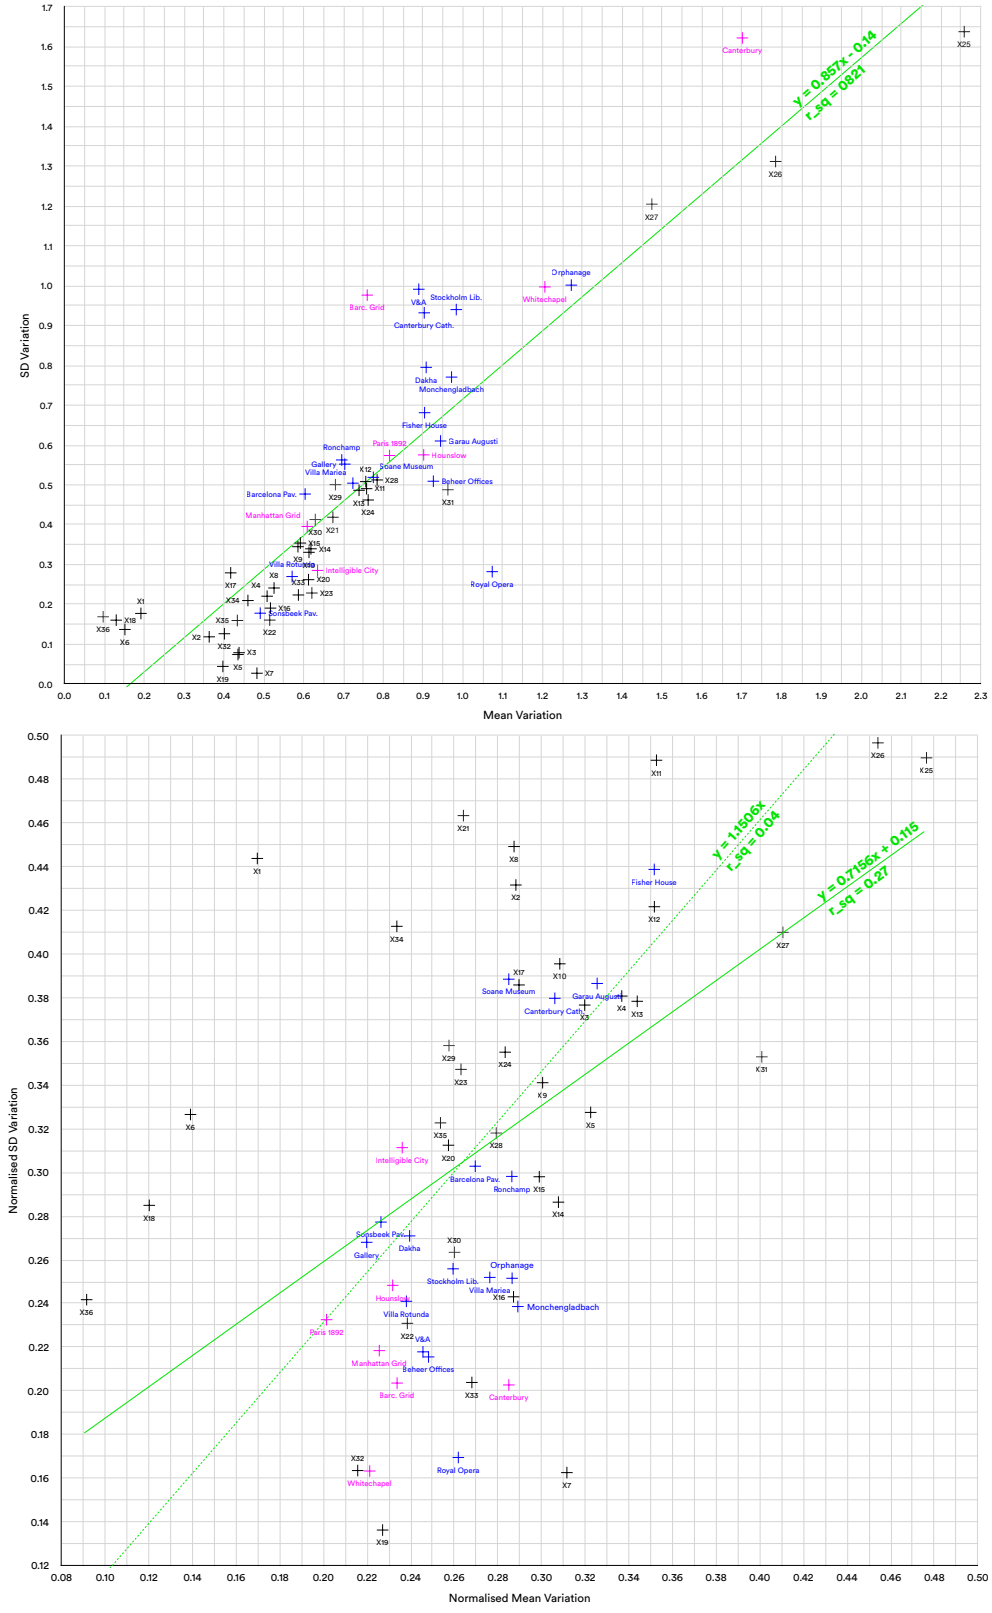

Figure 8: [*a, above*] Mean variation plotted against standard deviation of variation for all plans ($r^2 = 0.82$); [*b, below*] Normalised mean variation plotted against normalised standard deviation of variation for all plans ($r^2 = 0.27$, and drops to $r^2 = 0.04$ for a trend line constrained to the origin).

We can also identify patterns in our plots of plans' mean variation against standard deviation of variation. Unnormalized, the chart of comparison of these values (Fig. 8a) shows a relatively tight linear spread of subject plans. Here the line of best fit provides a high correlation for the hypothetical grid spaces (r^2 = 0.910) and a slightly less so for all spaces including buildings and cities (r^2= 0.821). The dispersal of subject plans along the line of best fit appears related to their inherent 'size-depth-complexity', with 'simpler' hypothetical grid spaces closer to the origin, and more spatially convoluted grid, building and city forms progressively further from it.

Comparing normalised mean latent variation and normalised standard deviation of variation breaks the linear correlations observed above, suggesting a successful controlling for said scalar size-depth-complexity effect (Fig. 8b). Doing so appears useful, producing a much broader spread of points in the latent space of the chart, and hence greater potential to identify morphological sets. Cities (in pink) seem generally proximal to one another; there appear to be two or more groups of building plan (blue) and several distinct clusters of the hypothetical grid plans across the chart. We interpretatively illustrate and discuss the potential of such patterns below.

4 INTERPRETATION AND CONCLUSIONS

Here we provide answers to our two primary questions; *how many cycles of isovist cascade achieve reasonable result stability?* and; *do plan types meaningfully affect the threshold of stability attainment?* We then expand a conceptual discussion on two emergent findings; the apparent 'complexity' gradation of plans via coefficient of variation; and the potential for morphological classifications via comparison of normalised (dimensionless) latent mean and standard deviations of variation.

4.1 Result stability and influence of plan types

Our results demonstrate that a low number of isovist cascade cycles establishes coherent mean visual depth or integration results. We have found no significant difference in result development between plans of differing type or complexities (none, that is, has a discernible impact upon the results themselves), meaning that the threshold for result stability appears to be consistent across all plans. Combined, these findings have a practical and a conceptual significance.

Practically speaking, for numeric results, we suggest that 250 isovist cascade cycles provides an acceptable upper 'reliability-threshold' for all plans. That many cycles have typically taken us between 60 – 180 seconds to complete on a 2023 MacBook Pro. Such time requirements are a significant improvement over VGA, for a much higher resolution of field outcome. Some

researchers may be content with a higher degree of variance in their result, or may wish for a faster outcome, and so can choose their own threshold levels accordingly. Visual patterns in the results field are certainly apparent much earlier in the analysis process, typically at circa 30 cycles. For many researchers such visual information alone may be sufficient for early exploratory work or quick assessment of working hypotheses. It can also be illuminating (and enjoyable) to simply watch a field becoming coherent. Applying our analysis method in such a manner could expand Space Syntax's attractiveness and reach to industry and practice settings, where architects and designers often favour rapid evaluation and 'fast twitch' feedback loops that align with intuitive design approaches (Anklesaria, 2022).

Philosophically, our isovist cascade and concatenation method offers a novel window into how humans and animals might build cognitive maps of unfamiliar spaces. By efficiently aggregating local information into incrementally improving global representations, it suggests a layering process potentially allegorical to how the brain may build navigational knowledge of new environments (Penn, 2003). Notably, even with a limited number of sampling cycles, our method achieves a reasonable degree of spatial coherence. Such a finding implies that, to establish 'sufficient' global knowledge of a space that allows effective navigation, far fewer orienting landmark points may be necessary than typically assumed, even in complex environments. The latter challenges the prevalent approach of analysing the whole of a spatial graph, advocating instead for partial or 'sufficient' analyses to capture and emphasise salient landmark information.

4.2 A classifier for 'spatial wobble'

Beyond the above conclusions, our coefficient of variation metric sorts plans into a form of morphological order. Our results here have no correlation to intelligibility, suggesting that another spatial property is being identified. What the classification represents can be interpreted by consideration of the factors in the coefficient; our 'latent standard deviation' and 'latent mean' of variation. Each is driven by qualities of the overall visual depth available in the spatial system, measured in each successive isovist cascade.

Our latent mean variation describes the average change in depth relations to be expected from each new isovist cascade. It estimates how much each new isovist cascade might be expected to affect the overall understanding of a space. In experiential terms, mean variation therefore describes how intensely one's cognitive understanding of a spatial system's structure might vary as one takes differing positions within it. In a very simple, low depth system, this would be low (moving affords little change in extant spatial relations); but in system with high depth, it would be high (moving might cause increased change in spatial relation and thus knowledge). Latent

standard deviation, meanwhile, describes how widely the above change in depth relations themselves vary. In uniform spatial systems, the standard deviation would be low (moving results in repetition of view structures, and therefore produces cascades with similar depth relations); but in a highly convoluted system, it would be high (moving provides multiple, unexpected, 'alternative perspectives' and thus cascades with wide-ranging depth relations).

Our coefficient of variation compares the above factors. It is high when the variation between isovist cascades is greater than our mean predicts, and low when it is less than our mean predicts. As such, the effects of actual depth in the system are normalised out, giving a dimensionless metric for how a system's spatial sub-structures facilitate or frustrate overall knowledge of space.

When standard deviation in the system is lower than mean variation predicts, movement to new locations is likely to afford patterns of spatial relations similar to those previously encountered. The latter can be seen in our how regular grid plans, (Fig. 07, X5, X7, X19) produce a low coefficient. In the latter layouts, multiple different positions have similar structures of relative visual depth, caused by the overlapping and interlocking pattern of axial views within the grid. As grid cells are closed off, causing spatial differentials such as dead-ends, and severing such visual and spatial linkages, the coefficient accordingly rises (Fig. 07, X5 -> X4, or X19->X32). In cognitive terms, we might therefore say that a low coefficient of variation identifies spatial morphologies in which local sub-structures reinforce global spatial comprehension. In short, regular opportunities for 'aha moments' afford useful potential reorientation points during navigational exploration.

When standard deviation is lower than mean variation in the system predicts, movement to new locations is likely to afford patterns of spatial relations different to those previously encountered. As our hypothetical grids increase their branching qualities (Fig. 07, X8 -> X10 -> X9, or X20->X21), and ultimately become mono-cursive paths, or non-linear spatial sequences, (Fig. 07, X25 -> X26 -> X27), the coefficient rises. In doing so, it reflects how opportunities for visual loops and spatial linkage are removed, and so movement to new locations results in a high degree of change in relative depths. In cognitive terms, we might therefore say that a high coefficient of variation identifies spatial morphologies in which sub-structures frustrate global spatial comprehension. In short, the removal of revelatory 'aha moments' increases the potential for dis-orientation during navigational exploration.

Finally, convex 'room-type' scenarios provide an interesting condition, providing a low standard deviation, but a comparatively even lower mean. Doing so produces a high coefficient of variation (Fig. 07, X6, X18, X36). It is tempting to discard such cases as outliers, but, by extension, room like

features, included within a more involved hypothetical grid, can invert the expected coefficient by 'compounding' such effects within the overall. The latter seems at variation with our model. How can inclusion of a simple room space cause 'more dis-orientation'?

In a fully convex 'room type' space all space is in view. Spatial relations in visual depth terms therefore remain identical as one moves within the space. In a global sense, the latter means that movement within the space does nothing to aid wider cognitive knowledge, likely affording few salient configurative features or landmarks that might trigger recognition of a global location (beyond the room itself!). Instead, positional knowledge must be taken from local visual perceptions such as perimeter or drift (Dalton & Dalton, 2001). It is said lack of potential for broader cognitive orientation that we believe produces a high coefficient result here.

4.3 A tentative morphological matrix

Finally, our plotting of normalised standard deviation of variation against normalised mean of variation merits further discussion. The results form a broad spread of points with little overall correlation (Figure 8b). Nevertheless, within the latter spread, clusters of points emerge, and morphological structures in the subject plans seem to drive their chart location. To allow a simpler consideration of the said phenomena, we rotate the plotted points relative to two baselines (Figure 9). The new x-axis represents the original forced line of best fit, and the y-axis aligns with the original mean x value (a circa 45° rotation). Figure 9a shows the hypothetical grid plans plotted, whilst Figure 9b includes all results.

Perhaps most notably, in Figure 9a, 'room-type' layouts (Figure 9a, X-6, X-18, X-36) fall in the top-left quadrant, mirrored by the sequence of 'grid-type' layouts (Figure 9a, X-5, X-7, X-19) in the bottom-left quadrant. Both sets trend away from the mean (y axis) baseline and towards the x axis baseline, perhaps toward a point of convergence upon it. The latter makes logical sense, as an infinitely fine, interlinked version of grid space would be identical to a square room. Elsewhere in the matrix we see a central cluster of simple systems (Figure 9a, including X-9, X-3, X-4 and X-10), that trend towards the top-right quadrant as morphologies become more sequential (i.e. movement is corralled through successive spaces). In the right-hand quadrant above the x-axis, such sets seem to be influenced by characteristics of branching leading to alternate destinations, whilst below the x axis, plan types exhibit spaces linked on a single path (Figure 9a, X25-27).

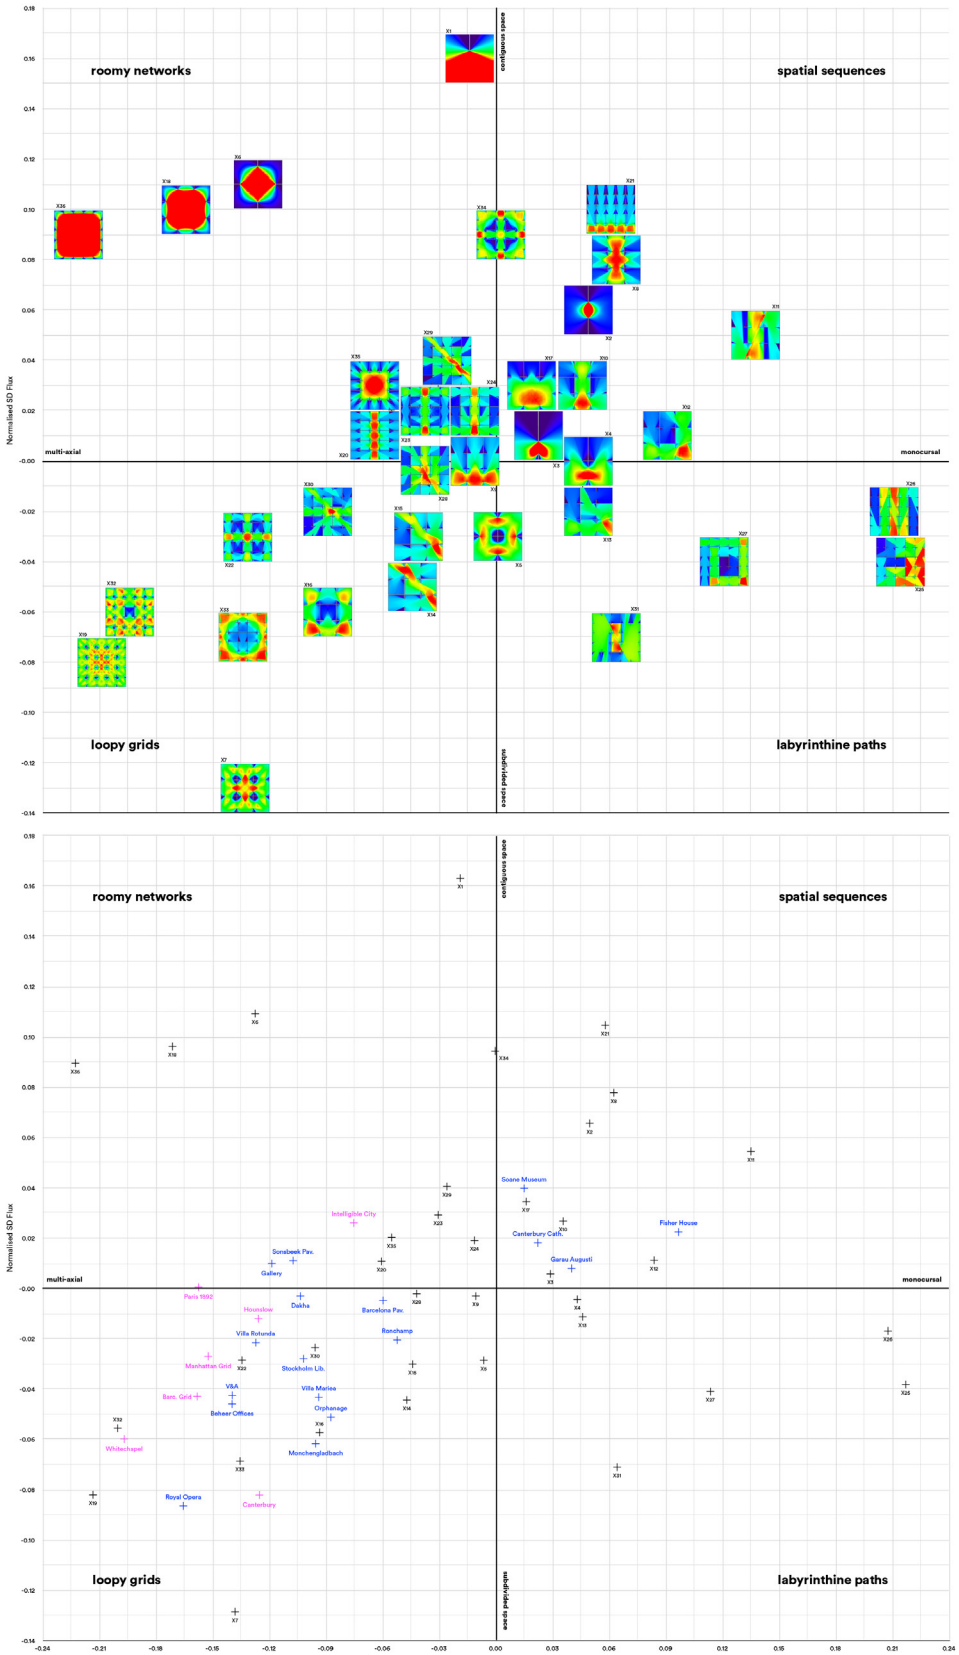

Figure 9: Translation of Figure 8b data by rotation along the line of best fit, and to mean x value. With hypothetical grids [*a, above*]; with buildings and cities [*b, below*]. Axes should be read relative to Fig. 7.

What then is the mechanism that might be underlying such proto-classifications? In our normalised plot, the factors of standard deviation and mean are themselves now dimensionless, allowing comparison between varying plan sizes and depth-complexities in more detail than with coefficient of variation. We can now view where the mean variation of a plan falls above or below that which would be expected from the mean depth contained within the plan, and secondly, where the standard deviation of variation of that plan falls above or below that expected from the maximum range of depth available within the plan.

Based on the above, a first overall distinction can be made between the left- and right-hand sides of our matrix. On the left, systems have a lower mean than expected and, on the right, higher. It seems logical that mean variation is lowered by numerous interlinking views (reinforcing locational knowledge) or loops in a spatial structure (hence our grids and rooms occur on the left-hand side) and increases if diverging branches of space or singular sequential paths prevent such means of spatial re-orientation (hence these structures occur on the right-hand side). The x-axis in our matrix thus charts topological qualities of linkage, from multi-axial to mono-cursive systems. A second distinction can be drawn between the upper and lower halves of our matrix. In the upper half, systems have a broader standard deviation than expected and in the lower, a more constrained one. The former appears to be driven by higher spatial variation in the system, reflected in a high degree of structural difference between successive isovist cascades. In the latter, spatial variation is lower, with relations within isovist cascades tending to be similar. In line with such an interpretation, spatial systems where there is a recognisable uniformity of spatial feature, or commonality of structure and experience throughout (grids and sequential mono-cursive paths), fall in the lower half of the matrix. Spatial systems with variegation of local structures and experiences, or occurrences that run counter to the general spatial condition of the system (branching systems and those with room spaces within them) tend to be found above the horizontal base line. The y-axis in our matrix thus charts entropic qualities of spatial arrangement, from consistent spatial systems to variegated ones.

From the above two distinctions we can advance a tentative classification of our matrix's four quadrants into 'loopy grids', 'roomy networks', 'branching sequences', and 'labyrinthine paths'. Without making an in-depth discussion of the differing real-world plans charted in Figure 9b, they do seem to corroborate such a taxonomy. A small cluster occurs in our branching sequence quadrant, including Soane's Museum, Kahn's Fisher House, and Canterbury Cathedral; all highly choreographed spatial sequences. Most cities and plans fall in the looping grids quadrant; particularly notable here is the Royal Opera House's proximity to X-33, with both systems being an underlying grid with a strong orbital loop. Intriguingly, both the space syntax 'invented-real'

plans, the Gallery Space, and Hillier/Dalton's intelligible city (Dalton, 2001), are positioned just within the roomy networks' quadrant, perhaps reflecting an intrinsic (perhaps deliberate) balance between an interlocked multi-axial set of spaces that retain room-like (convex) qualities.

4.4 Limitations, and a final note.

The above discussion of course comes with important qualifications. We should emphasise that our observations and classifications are tentative, requiring further investigation and confirmation. The base lines of our final matrix have been identified experimentally as estimates. Whilst they seem a good fit (see, for instance, Figure 9a, X-5, X-34 and X4 as 'balanced' systems appear correctly located on a base line), further additions to the matrix might improve these. So too would a more fundamental means of definition.

It should also be remembered that most real-world plans are combinatory, involving features of multiple morphological types within their whole. Given the latter, it may be more of value and import to the researcher to read plan positions on our matrix relative to one another, rather than in an absolutist manner; for example, by saying that Beheer Offices exhibit a more interlinked and multiaxial grid structure than the Amsterdam Orphanage.

More broadly, we suggest that our work identifies and measures morphological structures that bridge between the local and the global, with meaningful import to considerations of navigation and cognition. Here it is a mistake to overly relate the work to intelligibility (indeed the disassociation here is proven by the lack of correlation in Fig.7b). Intelligibility is the similarity between local configurational properties and the system as a whole; in a sense, the possibility of metonymic orientation (*can I navigate the whole based on the part?).* It thus deals with configurational consistency of a particular type (how local experience is related into global knowledge). Our isovist cascade analysis, and the temporal-stabilisation that it illuminates, instead introduces a novel metric for how a spatial system may directly promote or hinder the development of cognitive knowledge and spatial re-orientation, via the embodied reveal of local-global bridging features such as connective links and loops.

Such features in themselves provide for the successful orientation of a situated explorer though the establishment of landmark knowledge (a.k.a. aha moments), but are more typically found in what one might intuitively call 'complex', or spatially elaborate, systems. Conversely, apparently 'simple' systems often have indistinct or confusing features (as in a winding path), and so may afford opportunity for disorientation. By reflecting such a paradox, our work echoes previous

discussions (both within Space Syntax and without) that highlight the phenomenon of becoming lost in seemingly functionally simple spaces, like single-path labyrinths or curved corridors.

By framing the extent and influence of such embodied cognitive easements and challenges within a spatial structure, our work offers a way to represent the relative coverage of a building or space required to achieve likely sufficient (rather than total) knowledge. In other words, our work may be capturing how quickly someone situated in a space can develop a sense of being able to functionally navigate it, through the knitting together of spatial linkages and landmarks. In turn, the latter may bear some relation to how readily an environment affords the construction of a cognitive map by the individual.

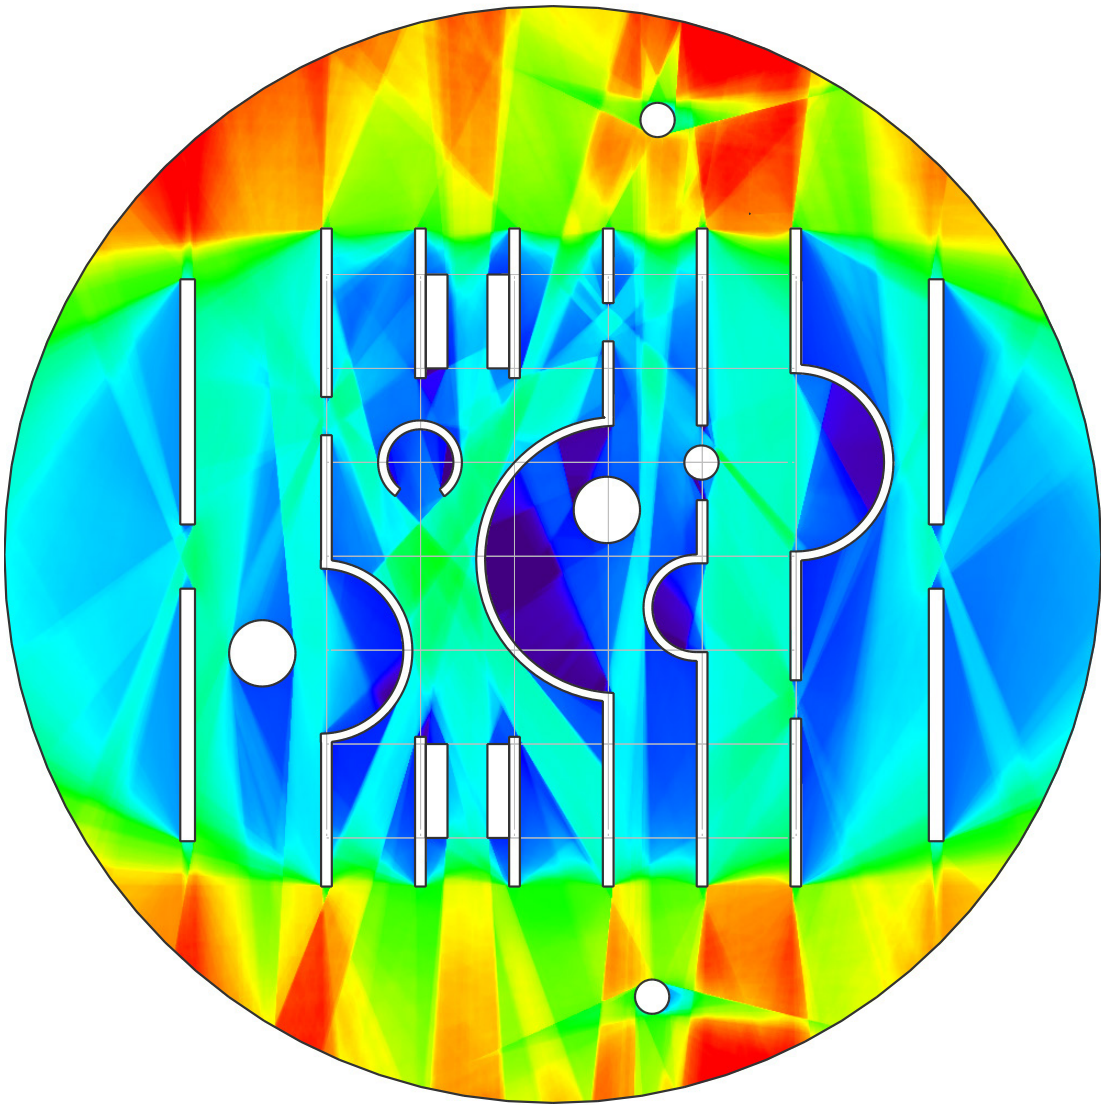

Figure 10: Mean Aggregate Isovist Cascade (MAgIC) Integration analysis for the Sonsbeek Pavilion; classified under our tentative morphological matrix as a multi-axial, semi-roomy network.

REFERENCES

Anklesaria, F, McElhinney, S, (2022). *Syntax as an Iterative Architectural Design Tool: A teaching experiment using spatial syntactic and isovist analysis* Proceedings of the 13th Space Syntax Symposium, Bergen University, Bergen, Norway

Benedikt, M. (1979). *'To Take Hold of Space: Isovists and Isovist Fields'*. Environment and Planning B. v6: pp.47-65.

Benedikt M L, Burnham C A, (1985), *`Perceiving architectural space: from optic rays to isovists',* in Persistence and Change, Psychology Press, London

Charalambous, E., Hanna, S., & Penn, A. (2021). *Aha! I know where I am: the contribution of visuospatial cues to reorientation in urban environments*. Spatial Cognition & Computation, pp.1-38.

Dalton, R. (2001) *'The Secret is to Follow your Nose; Route Path Selection and Angularity',* Proceedings of the 3rd International Space Syntax Symposium, Atlanta, Georgia

Dalton, R and Dalton, N. (2001) *'OmniVista:an application for isovist field and path analysis'* 3rd International Space Syntax Symposium, Atlanta, Georgia

Dalton, N, Dalton, R, McElhinney, S, Mavros, P. (2022 a). *Upper bound projection and restricted random isovists: a solution to the comparison of different environments analysed with visibility graphs*, Proceedings of the 13th Space Syntax Symposium, Bergen University, Bergen, Norway

Dalton, R, Dalton, N, McElhinney, S, Mavros, P. (2022 b). *Isovists in a Grid; Benefits and limitations*, Proceedings of the 13th Space Syntax Symposium, Bergen University, Bergen, Norway

Gibson, JJ (1979), *The Ecological Approach to Visual Perception*, Boston, MT:Houghton, Mifflin & Co.

Hillier, B., Burdett, R., Peponis, J., Penn, A. (1987a), 'Creating Life: Or, Does Architecture Determine Anything? 'Architecture et Comportement/Architecture and Behaviour , 3 (3) 233 – 250. pp.237

Hillier, B., Hanson, J., and Graham, H. (1987b), '*Ideas are in things: an application of the space syntax method to discovering house genotypes'*, Environment and Planning B: Vol 14, 363-385. pp. 364-365;

Hillier, B. and Hanson, J. (1984), 'The Social Logic of Space', Cambridge University Press: Cambridge. pp.108-109

Hillier, B., Penn, A., Hanson, J., Grajewski, T. & Xu, J. (1993) '*Natural movement; or, configuration and attraction in urban space use'*, Environment and Planning B, V.20, pp 29- 66

Krenz, K, Psarra, S, Vinicius, Netto, M, (2023) *'Mapping the conceptual system of an urban theory and its evolution: a text analysis of space syntax conference papers over 20 years'.* Urban Morphology, v27, pp. 161-177

Letesson, Q. (2013). *'Minoan Halls: A Syntactical Genealogy'*. American Journal of Archaeology, vol. 117(3), pp 303-351.

McElhinney, S, Dalton, R, Dalton, N, Mavros, P. (2022). *Detection of intelligibility leaps using isovist-waves: joining the dots to map potential 'aha moment' locations*, Proceedings of the 13th Space Syntax Symposium, Bergen University, Bergen, Norway

Penn, A, Hillier, B, Banister, D, Xu, J. (1998), '*Configurational modelling of urban movement networks'*, Environment and Planning B: Planning and Design, 1998, V25, pp 59-84

Penn, A. (2003). 'Space Syntax And Spatial Cognition: Or Why the Axial Line?' Environment and Behavior. v35: pp.30 – 65

Peponis, J and Bellal, T. (2010) 'Fallingwater: the interplay between space and shape'; Environment and Planning B. v37: pp. 982-1001

Peponis, J., Hadjinikolaou, E., Livieratos, C. & Fatouros, D.A., 1989, The Spatial Core of Urban Culture, Ekistics, 334/335, pp43-55.

Peponis, J; Wineman, J; Rashid, M; Hong Kim, S; Bafna, S. (1997) 'On the description of shape and spatial configuration inside buildings: convex partitions and their local properties'; Environment and Planning B. v24: pp. 761-781

Psarra, S and McElhinney, S. (2014) 'Just around the corner from where you are: Probabilistic isovist fields, inference and embodied projection', The Journal of Space Syntax, v5: pp. 109-132

Sailer, K, and McCulloh, I, (2012) 'Social networks and spatial configuration—How office layouts drive social interaction, Social Networks, v34, pp 47-58

Teklenburg, J; Timmermans, H; Wagenberg, A. (1993) 'Space Syntax: Standardised Integration Measures and Some Simulations'. Environment and Planning B. v20: pp. 347 – 357

Turner, A. and Penn, A. (1999), *Making isovists syntactic: isovist integration analysis*, in the Proceedings of the Second International Space Syntax Symposium, March 1999, Brasilia, Universidad de Brasil, Brazil, Vol 1.08.

Turner, A; Doxa, M; O'Sullivan, D; Penn, A. (2001a). "From isovists to visibility graphs: a methodology for the analysis of architectural space". Environment and Planning B. v28: pp.103–121,

Turner, A, (2001b). 'Depthmap: a program to perform visibility graph analysis', Proceedings of the 3rd International Symposium on Space Syntax, Georgia Institute of Technology, Atlanta, Georgia

Turner, A. (2001) '*Angular Analysis*'. Proceedings of the 3rd International Symposium on Space Syntax, Georgia Institute of Technology, Atlanta, Georgia

Turner, A., Penn, A., & Hillier, B. (2005) 'An algorithmic definition of the axial map', Environment and Planning B v32, pp425-444