

Hanna S. (2006) Representing Style by Feature Space Archetypes: Description and Emulation of Spatial Styles in an Architectural Context, In Gero, JS (ed) *Design Computing and Cognition '06*. Springer. pp. 3-22. ISBN: 978-1-4020-5130-2

REPRESENTING STYLE BY FEATURE SPACE ARCHETYPES

Description and Emulation of Spatial Styles in an Architectural Context

SEAN HANNA
University College London

Abstract. Style is a broad term that could potentially refer to any features of a work, as well as a fluid concept that is subject to change and disagreement. A similarly flexible method of representing style is proposed based on the idea of an archetype, to which real designs can be compared, and tested with examples of architectural plans. Unlike a fixed, symbolic representation, both the measurements of features that define a style and the selection of those features themselves can be performed by the machine, making it able to generalise a definition automatically from a set of examples.

1. Introduction

At its core, style is what distinguishes one group of works from another. This paper proposes that we can define a style using an *archetype*, an ideal model comprised of the features that exemplify the style. This concept differs from the description of a type, or category into which particular examples can fall, and from that of a prototype, precedent or case, which are actual instances on which later examples can be modelled. An archetype is something between the two, a generalisation that can not exist materially, yet matches and is compared to many actual instances. This is almost certainly not a real example, but an abstraction made up of only those features necessary to differentiate it from other archetypes.

Many approaches to style are based on explicit symbolic representations (where fixed concepts are mapped to named variables) or rule systems. These can tell us useful things about designs and how they can be made, but are inflexible. They reveal some of the ways we learn about styles pedagogically, but typically fixed, historical ones. By contrast this work proposes a method to automatically derive representations from real examples of design.

It is based on the mapping of design examples in a high dimensional feature space, and uses methods of dimensionality reduction of this space to yield an archetype that describes the style. This can be used to classify, and

as a measure to generate new designs. The use of a feature space agrees with our own intuitive ability to evaluate designs as being stylistically nearer or farther from one another, and is commonly applied in machine learning, in which a space is constructed in which each dimension is a measurement of a particular feature, and so each example can be represented as a single point. The nearest neighbour algorithm (e.g. Duda et al. 2001), for instance, classifies an unknown example of data by simply measuring its distance to previously known and labelled examples, or prototypes.

Two innovations are proposed over such existing methods. The first is that the archetype is a generalisation that combines both the concept of the ideal example and the particular space in which it is measured. In the nearest neighbour algorithm, a prototype is a real example of data, and all examples are measured within the same space. The archetypes presented here are measured in a lower dimensional space consisting only of the features relevant to that style, and each archetype may be measured in a different feature space. The archetype, then, is a point in a feature space consisting of dimensions in which examples of a particular style are closely clustered, and examples of other styles are distant. It is comprised of both point and space.

This provides a method for analysis of existing designs, but not synthesis of new ones. Rule-based definitions can be useful because they can be followed to produce new designs, whereas a classification algorithm by itself clearly cannot. The second innovation incorporates the notion of affordances (Gibson 1979), to consider the design process as a continual evaluation and choice between possible alternatives as the development of the design progresses. These choices can be made by repeated measurement against the ideal that the archetype represents.

The approach was tested within the design domain of architecture, using the plan in particular. This paper implements the two major aspects of the approach in sections 3 and 4. The first deals with analysis, and begins by providing methods by which plans of real buildings can be embedded in a feature space such that those that are similar fall near to one another. This yields a way in which all styles might be understood and represented by a computer, which is not based on any predefined symbolic representation. The second part refines this spatial embedding, and combines a very simple generative process to synthesise designs. An archetype is defined from a set of examples and used to guide new designs in the same style.

2. Related techniques

Several techniques from related fields of machine vision and space syntax are relevant to this work. They are outlined here along with a discussion of various existing approaches to style.

2.1 APPROACHES TO ARCHITECTURAL STYLE

The need to differentiate is fundamental to communication. Gombrich (1960) suggests that art provides categories by which to sort our impressions: ‘without some starting point, some initial schema, we could never get hold of the flux of experience.’ Like a game of ‘Twenty Questions’, where the actual identity of a concept is gradually communicated through trial and error, the symbols of art do not represent reality in all its detail, but only what is necessary to convey the message. In his view the set of symbols used constitute the style, and one general approach to representing style is rooted in identifying the equivalent of these symbols, either as features or generative rules of the work.

Architectural style has been represented a number of ways, both for analysis of existing designs, and for synthesis of new ones. Analytical methods have been proposed that model similarity or group designs based on a count of pre-defined features. Chan (1994) uses a list of material and compositional features to evaluate buildings as to their style, and finds a correlation with the judgements of test subjects. Experiential qualities of architecture have also been structured or mapped to rule sets to guide architects (Alexander et al. 1977; Koile 2004, 1997), and this approach has been implemented by Koile (1997) in an expert system that is also able to recognise style as a count of defined design characteristics.

Rule systems have also been developed to generate new designs in a particular style, such as Hersey and Freedman’s (1992) computer implementation to create possible Palladian villas. Shape grammars (Stiny 1976, 1980) are possibly the most widely studied, and have been successful in providing generative descriptions of many styles. A descendant of Noam Chomsky’s linguistic generative grammar, they provide an explicit rule-based method for producing final designs, and have yielded examples in the apparent styles of Palladian villas (Stiny and Mitchell 1978) and Frank Lloyd Wright’s prairie houses (Koning and Eizenberg 1981). Recent approaches have expanded the method to allow decompositions on finished or existing designs to generate new rules for design exploration (Prats et al. 2006).

As an approach to style, a style is often (e.g. the examples above) encoded with a specific grammar, unlike linguistic grammars that generate a particular language with any number of possible styles. A creative human then works with the shape grammar to make a specific design within the style. As a tool for analysis, the grammar or rule set is constructed by a human designer, a fully automatic process seen as undesirable or impossible (Knight 1998). In its generative capacity it is then followed by a user choosing which rule to apply at each stage in the process to create designs comparable to originals.

Another approach to style proposes that it is not defined by clear and predetermined features or rules, but can be quantified by various measurements taken from examples of the works. More general analytical techniques using information theoretic measures have been used to measure distance between individual plans (Jupp and Gero 2003), and to derive values for measurements such as entropy (Gero and Kazakov 2001), morphology and topology (Jupp and Gero 2004) that can be used to evaluate examples of a style. These have the advantage of quantifying designs of any style as real values on the same scales, so that variations within or between styles can be measured uniformly.

It is this second approach that is preferred in the present work, for several reasons. While the setting of explicit rules or feature definitions can tell us interesting things about a style, they are often a simplification that will either produce some designs that would be considered outside the style, or fail to produce all possibilities within (Hersey and Freedman 1992, ch. 2). The boundaries of this group are impossible to set exactly, and styles are subject to gradual evolution. Even the seemingly clearly defined classical orders, before being codified by Vitruvius (2001), were the result of centuries of experiment and still subject to debate on particular proportions throughout the Renaissance. Most importantly, styles are based on examples. They are only discernable after several similar designs exist, and it is through the emulation of examples that a style is communicated. In the continuum of stylistic differences styles can always be further subdivided, but only down to the level of the individual work.

But there may be no particular class of measures that we can specify in advance to contain the description of all styles. While style is often considered (as in Gombrich 1960) the method of expression as opposed to the content, Goodman (1975) argues for a broader definition of style to include aspects of both *what* is expressed in addition to *how*. In his definition, the style ‘consists of those features of the symbolic functioning of a work that are characteristic of author, period, place or school’, and these can only be determined in relation to the works, not beforehand.

This present work differs from previous approaches in that design examples will be evaluated by a general analysis, then the relevant features determined automatically in the definition of the archetype. By so doing, both the processes of defining a style and building examples of it can be performed by a machine.

2.2 FEATURE DESCRIPTION BY DIMENSIONALITY REDUCTION IN OTHER FIELDS

To automatically generalise a description of a style, either as a set of relevant features or a generative rule, from a set of given examples is more difficult than setting it in advance, but this approach is beginning to be explored in

other stylistic domains, such as musical style (Tillmann et al. 2004). It is based on more firmly established techniques of machine classification and learning in other fields, particularly machine vision, as used in applications such as face recognition and robot navigation.

Dimensionality reduction is often used in applications from face recognition to linguistic analysis to infer distinguishing features from a given set of high-dimensional data. Principal component analysis (PCA) provides a new set of axes aligned to the characteristic vectors, or eigenvectors of the covariance matrix of the data set. The dimensions in which the data varies least are discarded to yield a lower dimensional subspace of the data that can then be used to make classifications. The principal components of face images, for example, referred to as 'eigenfaces' (Turk and Pentland 1991), are used by face recognition software to effectively quantify a new face by how it measures against each, and its best match found from existing data.

More closely related to our experience of architecture is the problem of a robot visually navigating through a real environment. Rather than explicit labelling, it has been found preferable to allow the computer to come up with its own concepts: an unsupervised learning of the visual features of the environment. In work on robot navigation of unstructured environments, Durrant-White (2004) has used dimensionality reduction on the image data recorded by the moving camera of a robot's visual system. When thousands of single-frame images of the environment are plotted on a graph defined by the set's characteristic vectors, they fall into distinct clusters. Each contains a set of images of a particular object that can thereby be considered meaningful to the robot, given the makeup of its visual and 'nervous' system. Without supervision, these also generally correspond to objects such as 'tree' or 'rock': the same kinds of distinctions and classifications that we would make.

2.3 REPRESENTATION OF SPATIAL FEATURES AS GRAPHS

To represent stylistic features of space, the computer requires an appropriate substitute for the sense data provided by the images or sound digitisations in the applications above – a method through which to perceive experiential qualities of the building. Two related space syntax techniques both provide an approximation of how people actually move through or use a space, using only an analysis of the plan. *Visibility graph* analysis quantifies the connectivity of a set grid of points within a space by the unobstructed sightlines that exist between them. From these, various measures such as integration, connectivity or mean depth of points can be used to derive a statistical analysis of the space based on the plan. (Turner et al. 2001) *Axial graph* analysis quantifies the connectivity of sightlines themselves, derived from plan vertices (Hillier et al. 1983; Hillier and Hanson 1984).

Properties of visibility and axial graphs have been shown to be strongly related to both spatial perception and resulting behaviour of people within spaces. Strong correlations have been found with measures of visibility graphs and observed way-finding, movement and use in buildings (Turner et al., 2001), and urban pedestrian movement (Desyllas and Duxbury 2001). Axial graphs have likewise been shown to be closely related to directly observed movement (Peponis et al. 1989; Hillier et al. 1993), building use and social interaction (Spiliopoulou and Penn 1999), and indirect behaviour such as land values and crime (Desyllas 2000; Hillier and Shu 2001).

Of the two methods, axial graphs are used in this work as the sightline endpoints are not predetermined and therefore invariant to plan scaling or rotation. The details of the algorithm are beyond the scope of this paper (see Turner 2005), but measurements such as connectivity and integration taken from the graph in axial analysis quantify the kinds of experience of that space as inhabited by a large number of bodies. Rather than these predetermined features, the raw graph itself can be easily taken as the machine's generic experiential input.

2.3.1 Measuring between graphs

Several approaches to similarity measurement have been based on small graphs of adjacency or connectivity of spaces in plan. Dalton and Kirsan (2005) use the number of transformations necessary to derive one such graph from another to measure the similarity between buildings, and Jupp and Gero (2003) suggest an analysis based on similarity and complexity measures of semantic graphs. With larger and more complex graphs as generated by axial lines, calculation of similarity becomes more difficult, but this can be overcome with graph spectral analysis. In image recognition applications, this has been used to effectively query a large database of logotype images for matches (Robles-Kelly and Hancock 2003). Spectral analysis of a graph uses the eigenvalues and eigenvectors of its connectivity matrix, and is relevant to the kinds of analysis considered here. The leading eigenvector, for instance, gives the result of a steady state random walk on the graph, equivalent to the connectivity of sightlines as described above. The spectrum of a graph, or ordered set of eigenvalues, is useful in that it can be used to represent the graph as a single feature vector.

3. Analysis: representation of a style in feature space

The archetype feature space must be derived initially from a general input capable of containing features relevant to all styles. This section tests both the acceptability of the axial graph as sense input, and the general use of a feature space for real building plans.

It is proposed that the description of a style is inherent in the examples of that style, and so examples of one style are objectively more similar to one

another than to examples of other styles. The stylistic description should therefore not be needed a priori to make this distinction, but the algorithm should be able to make the same kinds of classifications that we would, without explicit training. In this section a feature space is found for a group of buildings by unsupervised dimensionality reduction as in the examples in section 2.2. The examples are found to naturally fall into recognisable clusters within the space without the need for explicit labelling, and these correspond to the buildings' inherent similarities. This will provide the basis from which to derive an archetype from the examples.

3.1 DESCRIPTION OF PLAN STYLES IN A FEATURE SPACE

An initial test to describe styles in a feature space used a set of 24 plans, taken from various twentieth century iconic buildings (Weston 2004). This involved a simple dimensionality reduction of a feature space to confirm the hypothesis that proximity does indeed indicate stylistic similarity.

Axial graphs were constructed for each of the 24 samples, and this data – in effect a set of binary adjacency matrices – was taken as the computer's raw sense data, or experience of each of the spaces. Analysis was performed using Depthmap software, which constructed a minimal axial graph based entirely on plan input and an objective algorithm. (See Turner et al. 2001) Figure 1 displays the lines of sight in Frank Lloyd Wright's Fallingwater shaded to indicate the degree of spatial integration. Darker lines clearly reveal the higher traffic zones that link the rooms of the house. The spectrum of this graph was taken to form the raw feature vector so as to deal directly with the connectivity in the axial line graph. In doing so there were no subjective choices between further measurements provided by the software (of integration, moment of inertia, etc.) or of the selection of units in which to




Figure 1. A complete axial graph of Frank Lloyd Wright's Fallingwater (left) and the reduced minimal graph (right).

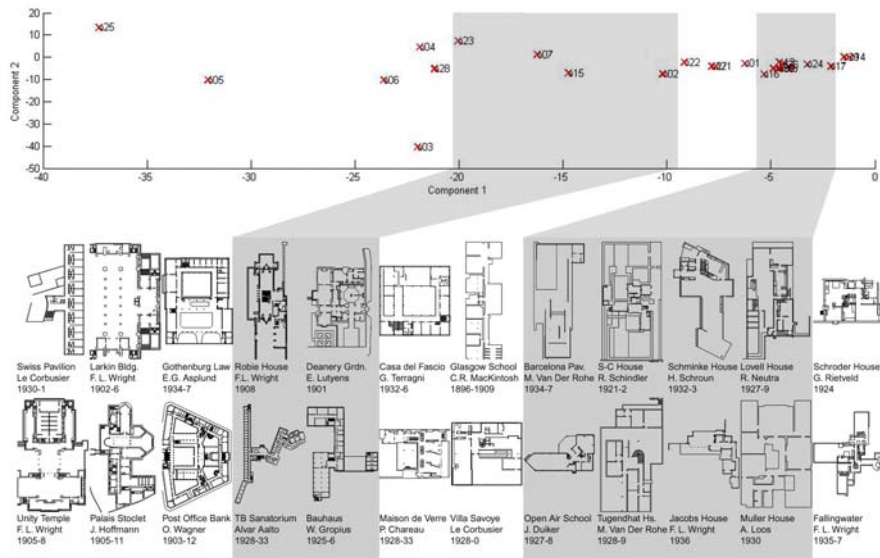


Figure 2. The example buildings plotted in feature space.

measure. The spectra were created by ordering the eigenvalues of these adjacency matrices in order of decreasing magnitude, thus yielding a feature vector for each plan.

With the spatial analysis of each plan quantified as a feature vector, the example buildings can then be plotted in a high dimensional feature space, with each value in the spectrum on a different dimensional axis. PCA determines, for a given set of data, the combination of dimensions in which it is most likely to vary, and these are used as new axes in a reduced version of the space that captures the essential statistical features of the data set. A reduction of the plans' feature space based on the first two principal components of the set is shown in Figure 2. The dimensions of this new feature space are strictly computational, and are meaningful only in a statistical sense, rather than in the sense that they could be easily described. The first component, on the horizontal axis, represents a combination of the features in which the plans are judged by the *algorithm* to differ the most.

Yet it can be seen that these include meaningful features such as typology (most houses toward the right and larger, public buildings to the left) as well as general stylistic trends. The shaded groups indicate the proximity of most of the axially symmetrical, pre-Modernist buildings to one another, as well as rough zones of early, and of high modernist buildings, typically with spaces that flow into one another articulated by shifting orthogonal planes. There is an approximate chronological order from left to right, seen clearly in the buildings by Wright, the Villa Savoye and contemporary Maison de Verre are next to one another, and van der Rohe's two works are virtually overlapping.

The divisions in the diagram are drawn simply for clarity, and are not meant to suggest a division into distinct groups. The points as plotted at this point represent a continuum in a uniform space. It is true that van der Rohe, Le Corbusier and Wright can be considered to have quite different styles, as can different buildings by Wright alone, but proximity in this feature space is meant to suggest those buildings that are more similar to one another by degrees. The only significant outliers in this regard seem to be those caused by typology: the private houses vs. the public buildings, but at this point no attempt has been made to draw the distinction. (The machine learning algorithms to be described in section 4 will allow this.) The fact that buildings of similar styles do fall near to one another in the reduced feature space confirms that the features indicated by the principal component are at least statistically related to style. Archetypes based on such a space may be used as style descriptors.

3.2 CLASSIFICATION OF PLANS INTO DISTINCT SETS

The plans above form a sort of continuum as they are taken arbitrarily from various styles over a forty year period, but the same process of graph analysis and dimensionality reduction can be used to plot the features of a particular group, and thereby define it. The above method was used again on a more focussed set of only two contrasting types in which there could be no ambiguity as to the label of each.

A set of 40 sample plans (Figure 3) was used containing examples of two building types: modern offices, and neoclassical museums. In overall shape the instances of these two types are quite similar to one another. Many, for example, consist of a ring of usable floor plate around one or two central voids: in the case of the offices this void represents a central core of lifts and services, in the case of the much larger museums based on neoclassical models it is a courtyard emitting natural light to interior galleries. This similarity in shape between the two groups ensures that the classification is based on more subtle details.

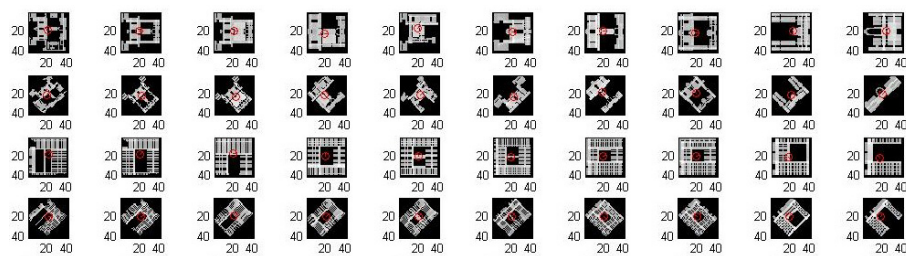


Figure 3. 20 plans representing museums (upper), and 20 offices (lower).




Figure 4. The plans are classified by the machine by PCA (left) and multi-dimensional scaling (right). The horizontal axis is the principal axis of the data set; the vertical axis is the second.

3.2.1 Results of plan classification by axial line analysis

In the plot above (Figure 4, left), PCA is performed on the spectra derived from the axial graphs, and the results plotted against the first two components (again, derived statistically from the data). It can be seen that the two groups of plans can be almost entirely classified even on only the first principal component (horizontal axis). There is a clear separation into two clusters, with a large margin between. The museums, marked by ‘o’s, are distributed over a region to the left of the origin, and the offices, marked by blue ‘x’s, are toward the right. An outlier from the office group can be accurately classified as well using another dimension of the second principal component. A multi dimensional scaling method (Figure 4, right) was also used as an alternative to PCA, resulting in an even clearer distinction.

No supervised training was required at this point. It simply indicates that in all the most obvious ways (obvious to the unsupervised computer) that the plans differ from one another overall, the plans within each type are more similar to one another than to those of another type, and that this distinction can be made by the algorithm.

3.3 DEFINING ARCHETYPES BY THE CLUSTERS

The resulting plots show points in space that are each representative of the spatial and perceptual features of a building plan. It is therefore possible to quantify such features as spatial relationships in a meaningful way. There is variation in their placement, but they fall into two clear groups corresponding to the two classes of buildings used.

The setting of an archetype from this requires only the selection of the point that best describes the region of reduced feature space in which the examples of that style lie. Nearest-neighbour algorithms for analysis define prototypes as the generating points of the Voronoi tessellation that separates the classes (Duda et al. 2001), but for the synthesis to follow in section 4 this

would bias the result when these lie close to the edge of a cell. Also, for only two classes there are infinitely many such points. The point that minimises the distance to all examples in the cluster is simply the *mean*, which can be applied as easily for two style clusters or two thousand. This mean and the mapping to the reduced feature space together constitute the archetype.

The lower dimensional feature space that results allows a description of style that is convenient for analysis and measurement – in that any plan example can be evaluated, and compact – in that only a few dimensions need be used. Because most of the dimensions have been removed the space itself comprises only those features that are relevant to differentiate examples of one style from another, and the mean point of each of the clusters above can be said to be the archetypal ideal of each group. Like any archetype it is a generalisation, and examples that fit it exactly may not ever be found in the real world, but it can serve as a model to guide the creation of new designs.

4. Synthesis: Production of new designs

The feature space and point that together define an archetype can be clearly used to measure example designs. In this section more sophisticated classification algorithms are used in place of PCA to derive the features, and methods for improving stylistic fidelity will be investigated. The use of *supervised* learning with *labeled* examples will imply a reduced feature space that is not just statistical, but meaningful. Analysis will also be combined with a generative method to synthesize new designs, and these will be used to evaluate the success of the archetype representation.

4.1 DESIGN AS SELECTION OF AFFORDANCES

Style can be considered a choice between given alternatives (Gombrich 1960), but rather than seeing this as a choice between generative rule systems, it can also be a choice of design moves within a system. While it may not be conscious (Goodman 1975) the act of creation implies the repeated choice of one design decision over another. One might consider this the ongoing selection of what Gibson (1979) terms *affordances*: opportunities for action provided by the environment, or in this case the design, as the design evolves over time. At any stage of the design of a building there are only certain possibilities open to the architect, and the act of adopting one style over another can be seen in this sense as a selection from the afforded alternatives. Tang and Gero (2001) suggest the act of sketching, with constant drawing and re-evaluation is such a process, choice of rules is the inherent design activity of shape grammars, and the explicit representation of choices as design spaces has also been proposed for CAD environments (Brockman and Director 1991). In this section a generic

system will be used to generate designs, but the style will be expressed by the choices made within it.

4.2 A BASIC OPEN GENERATIVE SYSTEM: BUILDING AGGREGATION

No generative algorithm is a direct counterpart to the axial graphs used in section 3, but there is one that precedes them. Axial analysis was developed initially to analyse the variation in settlement layouts, particularly in French and English stylistic variants of what was termed the 'beady ring' form (Hillier and Hanson 1984). This form itself was found to be the natural result of a very simple but open generative system of aggregation, the basis of which will be used to produce new design examples in this section.

The study of existing 'beady ring' settlements in southern France is an exposition of how many small decisions generate a recognisable order. (Hillier and Hanson 1984) As growing hamlets reach a certain size, a regularity appears in the configuration of the public spaces between the buildings. It takes the shape of a rough ring of joined spaces around a central clump of buildings facing outward, and several inward-facing groups of buildings around the perimeter. While this pattern appears in each of the hamlets studied, and continues to grow with the town, it need not be the result of central planning, as the simple aggregation model shows.

A minimal building unit is made up of two face-wise adjacent squares of a grid, with a closed building cell facing on to an open void space in front. The model allows these pairs to aggregate such that each new pair must join its open cell to at least one other open cell already placed, and the closed cell does not join another closed cell only at the vertex. The choice of available positions and orientations of each new unit is completely random, but each time the model is run, the emergent global structure forms that of the beady ring settlements studied, with a chain of open spaces onto which inner and outer groups of buildings face.

More important for the question of style as differentiation are the specific differences between towns. In Hillier and Hanson's own study they note the differences between the beady rings of France, and their counterpart villages in England that tend toward a more linear arrangement. These cultural differences in global form are also a result of the same uncoordinated local actions over time, yet the decisions of building placement that lead to a circular or a linear arrangement seem somehow to have been instilled into the individual members of the culture, not as contrasting sets of rules but as contrasting choices of application of the same rule set. "It is not simply the existence of certain generators that gives the global configurational properties of each individual [design]. It is the way in which variations in the application of the generators govern the growth of an expanding aggregation." (Hillier and Hanson 1984, pp. 84-85)

The model was initially never run with anything but random choices, but these decisions of application can be made instead by the evaluation of each possibility against an ideal specified by the archetype, and the choice of that possibility which fits best. Learning to build to a particular style is possible because each time a unit is built, the configuration of the new cell pair relative to its surrounding neighbourhood gives an example of an ideal for an individual to follow, another example to build up the archetype.

Although initially applied to town formation, this aggregation model is sufficiently general to represent rooms and corridors in a building or desks and chairs in an open plan office, depending on scale. While it uses a simple and constrained morphology, the grid is still able to represent configurations of very different patterns by the choices made in aggregation, and so it can stand as an analogy to more sophisticated generative methods to demonstrate the principles of this paper. For this reason and for its historical roots in the development of the axial analyses in section 3, the overall structure of this grid aggregation model will be used below to show that stylistic norms can be learned from examples and used to emulate the original pattern.

4.3 TWO STYLES OF AGGREGATION AS EXAMPLES

Two artificial stylistic norms were chosen to be easily distinguishable from one another, and a simple algorithm written to aggregate open/closed pairs of units in the manner of each. The first is a strict arrangement of straight rows rather like highly planned settlements such as Manhattan, and the second is a random arrangement of units joined open cell to open cell (Figure 5).

To learn the two ideals, a classification algorithm is trained on the units as they are built. While the perception of spatial qualities of existing building plan examples in section 3 required the construction of axial graph matrices, this simplified grid model allows samples to be taken directly. Each time a new pair is placed in the plan, its relationship to the 7×7 cell neighbourhood surrounding the open half of the doublet is taken as its input. The 49 cells, each containing either a closed building (indicated by a filled square or 1), a public open space (a dot or -1) or yet unbuilt (empty or 0) are used as the computer's sensory experience of that particular built example.






Figure 5. Two styles: strict rows and random aggregation.

Figure 6. Examples from a 49-dimensional feature space.

As in the case of the plan graphs, these neighbourhoods are points in an initial feature space. Each unique example can be represented by a point in a 49-dimensional space, a 2-dimensional projection (by principal components) of which is shown in Figure 6. Neighbourhoods of the straight rows are indicated by ‘×’, and the random style by ‘○’ markers in the centre. Clear clusters are less visible than in section 3, but this can be overcome through the use of supervised learning algorithms in the following sections to perform the mapping to lower-dimensional feature space.

In this space, the mean point of the cluster will represent an archetype of that style to be used in a straightforward building algorithm. At every step a given number of positions and orientations are available to be built, and the decision is simply the act of choosing which one of these affordances best fits the ideal by measuring proximity to the lower-dimensional archetype.

4.4 LEARNING AND BUILDING TO AN ARCHETYPE

Three experiments test the method of learning an archetype from examples and building to that archetype. The first tests that the style can be learned, with the hypothesis that clearer clustering will lead to a better resulting generation of the style. The second reveals that results can be improved by using a unique feature space reduction for each archetype. The third tests the hypothesis that the results of construction are independent of the choice of learning algorithm and particular representation of the archetype.

4.4.1 Clustering in a feature space and clarity of style: training by Support Vector Machine

A crucial hypothesis to be tested was that the results of learning would allow designs to be produced in a given style. It implies there should be a direct correlation between clear clustering in the feature space and the strength of the style in the resulting design. A support vector machine (SVM) (Vapnik 1995) was used for the initial classification because its easily tuneable parameters allow its mapping tolerance to be adjusted to test this hypothesis.

SVMs operate by finding a maximally separating hyperplane between the two labelled classes in a higher dimensional representation of the input, and that representation is given by a non-linear function with a parameter that can be used to adjust the fit to the data – in this case the width of a Gaussian. Figure 7 shows the results for $\sigma^2 = 5, 15$ and 25 respectively. The SVM output is plotted in the left column with row examples to the left of random examples, such that the vertical axis represents the single dimension of SVM output. The effectiveness of the classification is indicated in the second column images by the shading of each sample, where samples near the mean of the rows style are shaded in light grey and those of the random style are black. Clearly there is a clearer classification for the higher values of σ^2 .

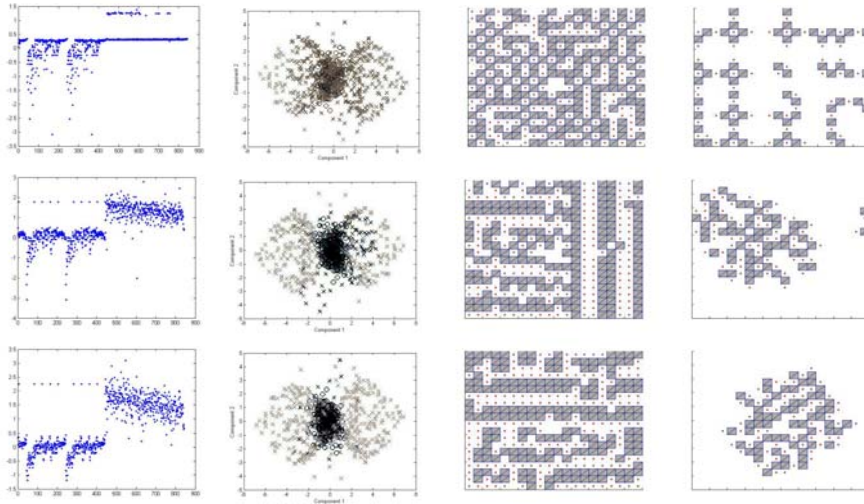


Figure 7. Building results for algorithms trained with a SVM: $\sigma^2 = 5$ (top), 15 (centre) and 25 (bottom). The first image on the left shows the mapping of 900 examples against the vertical axis. The second indicates apparent membership in each cluster by the shading of the points. Resulting building patterns follow emulating rows, then random aggregation.




Figure 8. The same training on a set of 'ideal' examples.

It is evident from the results that as σ^2 increases, the cleaner separation between the two groups by the algorithm results in a clearer construction, as shown in the images to the right. At each construction step, the possible sites and orientations are evaluated by the SVM, and the one closest the mean of either style as learned is selected. The completed constructions over a period of time are shown, one emulating the rows style as learned and the other the random arrangement, and the overall patterns are most easily seen for $\sigma^2 = 25$, particularly for the straight rows.

The initial hypothesis is confirmed, but the separations in Figure 7 are never quite enough, and the classifier can only produce adequate rows with an artificially created set of 'perfect' examples of row neighbourhoods. These are all identical, so that each is exactly perceived as the ideal archetype, and consequently the perfect classification of the two groups results in a stronger expression of the style. (Figure 8.)

4.4.2 Clarifying the archetype feature space: training by neural network

The method thus far performed one analysis for the principal components of all styles. It would yield appropriate archetype definitions if all styles differed in the same features, and thus could be classified in the same space, but this is unlikely. Rather than merely classifying two styles, the benefit of the clear archetype in Figure 8 suggests the choice of a feature space fit to a single style yields stronger results. In this section a unique feature space is found for a single style by training a neural network to find a space in which the points are clustered closely together as differentiated from all others.

A neural network was used to learn the rows style only, with the random examples serving as mere background from which to differentiate the relevant features. A Feedforward Multilayer Perceptron (MLP) (Rosenblatt 1958) was used, with 49 input nodes corresponding to the state of the neighbourhood, 50 nodes in the hidden layer, and a single, linear output that rates each example. Training was conducted by exposing the network to 450 examples from each of the two test styles and backpropagation of errors. Because the goal is to learn the features of the rows style only rather than to classify both, a variation on the typical error function was used. Normally, the error J in classification problems is a function of the difference between an output z and a specified target t . As there was no need for a target for examples outside the style in question however, the target for the rows was set to 0, and the reciprocal of the error used for all other examples, causing the error to fall as examples appear farther away. The advantage of this modified error function was a large separation and an output for most samples very close to 0.

Results of this style-specific feature space were superior to those of the SVM in section 4.4.1. In Figure 9 each of the examples is shown as a single dot in the vertical axis corresponding to the value of the network's single node output. After training, most of the first 450 examples along the horizontal axis (the row units) appear at 0, and most of the others (the random aggregations, to the right) as far away (note the extreme scale of the output axis). The resulting aggregation of open and closed cells produced by the building algorithm very closely resembles that of the original rows from which it was trained.




Figure 9. Training of a three layer network on the row samples.

4.4.3 Variations on the representation to learn the same style

Because the style is described by a feature space rather than symbolically, the actual method of feature space mapping in the archetype is quite arbitrary. This section tests that it can be changed and still lead to recognisable output. Interestingly, like Gombrich's game of 'Twenty Questions' in which the initial questions can also be arbitrary, the choice of classification algorithm used to define the style does not appear to matter. In fact one style can be described in many different ways, or feature spaces of different dimensionality.

Figure 10 shows the result of several very different learning algorithms exposed to the same set of examples, each resulting in a very different mapping of features (left) but very similar overall construction of rows. First is a neural network similar to the one in Figure 9, except that only the nearest examples of the random style were used in training. Below this, a different technique is used to train the network: errors from both groups are measured from the mean, but rather than adding the weight updates at each step for the examples from the random style, they are subtracted. The last example is a differently structured network entirely: a Kohonen self-organising feature map (Kohonen 1982). The subtraction training and the Kohonen feature map were found to be the most successful at replicating the overall pattern for this test style.

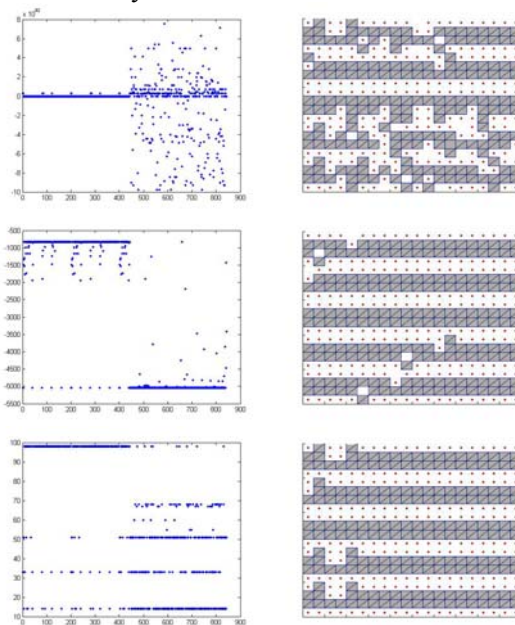


Figure 10. Three completely different algorithms (two double-layer neural networks and one Kohonen network) result in different feature spaces (left), but make similar evaluations and similar constructions.

The similarity of the final constructions indicates a style can be represented many different ways. Even with the constrained grid morphology of design space, there is a drastic difference in the feature spaces (Figure 10, left). These networks have different structures and different training methods, but produce similar results to one another. Feature spaces may differ in detail and even dimensionality, as each of the algorithms is capable of mapping to an arbitrary number of dimensions. The only necessary common process is that each forms an archetype in its unique feature space based on the examples of that group.

5.0 Conclusion

The idea of a style in any discipline is a fluid concept that is always subject to change, and therefore suited to a flexible representation. What is suggested here is that it can nevertheless be accurately represented and emulated. This work has presented an algorithmic method for both deriving a stylistic definition automatically from examples, and using it to generate new designs. Architectural examples were used, and were investigated primarily in terms of their spatial features, but it is intended as a general model in that other forms of input and classification algorithms may be used. Likewise, axial analysis and the aggregation model are not essential to the method, but the principles of feature space reduction and archetype should apply to a variety of analysis and synthesis techniques.

The concept of the archetype proposed is of a defined ideal and of a space in which to measure example designs. It contains only the features most relevant to define that style, but they are not counted as symbolic wholes. Instead one can measure an example's similarity in degrees, on an objective and continuous scale. This results in a definition of style that is flexible, can evolve, and is based on examples.

While fixed, rule-based systems are used as design aids by generating new examples of designs, a flexible, example based method such as this would assist in a very different way. While the archetype may be resistant to symbolic description, so very often are our own mental processes of style recognition, and in many complex problems we can more easily communicate by example than by explicit description. By automatically generalising its representation based on examples presented to it by a designer, such a design aid may propose output based not on rational clarity of process, but on the simple choice of precedents, fashion, taste or a hunch.

The definition of style provided by the archetype is analytical rather than generative, but there is still an obvious role for generative systems to play. The aggregation model in section 4 was chosen for its simplicity and common origin with the analysis in the previous section, but shape grammars and other generative rules could be applied – a likely avenue for

future exploration. Their role in this regard however, is as a framework for exploration of many styles rather than a definition of one.

Creative design is ultimately not a matter of rule following, but of judgement, and the model presented here proposes the flexibility this implies may extend to the definition of styles themselves.

Acknowledgements

I would like to thank professors Alan Penn and Philip Steadman for space syntax related advice, and for introducing some of the background examples presented in this paper. This research has been supported in part by the Engineering and Physical Sciences Research Council, UK and by Foster and Partners.

References

- Alexander, C, Ishikawa, S, Silverstein, M, Jacobsen, M, Fiksdahl-King, I and Angel, S: 1977, *A Pattern Language*, Oxford University Press, New York.
- Brockman JB and Director SW: 1991, The Hercules Task Management System, in *Proceedings of the International Conference on Computer-Aided design*.
- Chan, CS: 1994, Operational definitions of style, *Environment and Planning B: Planning and Design*, vol. 21 pp. 223-246.
- Dalton R C, and Kirsan C: 2005, Small Graph Matching and Building Genotypes. *Environment and Planning B: Planning and Design* (forthcoming).
- Desyllas, J: 2000, *The Relationship between Urban Street Configuration and Office Rent Patterns in Berlin*, PhD thesis, Bartlett School of Graduate Studies, UCL, London.
- Desyllas, J and Duxbury E: 2001, Axial Maps and Visibility Graph Analysis, *Proceedings, 3rd International Space Syntax Symposium*, Georgia Institute of Technology Atlanta.
- Duda, RO, Hart, PE and Stork DG: 2001, *Pattern classification*. John Wiley & Sons, NY.
- Durrant-Whyte H: 2004, Autonomous navigation in unstructured environments, *Proceedings of the 8th International Conference on Control, Automation, Robotics and Vision*.
- Gero, JS and Kazakov, V: 2001, Entropic-based Similarity and Complexity Measures of 2D Architectural Drawings, in Gero, JS, Tversky, B and Purcell, T (eds), *Visual and Spatial Reasoning in Design II*, Key Centre of Design Computing and Cognition, Sydney.
- Gibson JJ: 1979, *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston.
- Gombrich E H: 1960, *Art and Illusion*. Phaidon, London.
- Goodman N: 1975, The Status of Style. *Critical Inquiry*, vol. 1. Reprinted in Goodman N: 1978, *Ways of Worldmaking*. Hackett Publishing Company Inc. Indianapolis.
- Hersey, GL and Freedman, R: 1992, Possible Palladian villas: (plus a few instructively impossible ones). The MIT Press, Cambridge MA.
- Hillier B, Hanson J, Peponis J, Hudson J and Burdett R: 1983, Space Syntax, *Architects Journal*, 178(48), pp 67-75.
- Hillier B and Hanson J: 1984, *The Social Logic of Space*. Cambridge University Press.
- Hillier B, Penn A, Hanson J, Grajewski T and Xu J: 1993, Natural movement, *Environment and Planning B: Planning and Design*, vol. 20 pp. 29-66.
- Hillier B and Shu S: 2001, Crime and urban layout: The need for evidence, in Ballintyne S, Pease, K and McLaren V: *Secure Foundations: Key Issues in Crime Prevention and Community Safety*. IPPR, London.

- Jupp J, and Gero, JS: 2003, Towards computational analysis of style in architectural design. in S Argamon (ed), *IJCAI03 Workshop on Computational Approaches to Style Analysis and Synthesis*, IJCAI, Acapulco, pp 1-10
- Jupp, J and Gero, JS: 2004, A Characterisation of 2D Architectural Style, *Journal of the American Society of Information Science*.
- Knight TW: 1998, Shape Grammars, *Environment and Planning B: Planning and Design*, Anniversary Issue, 1998, pp. 86-91.
- Kohonen T: 1982, Self-organized formation of topologically correct feature maps, *Biological Cybernetics*, 43(1) pp. 59-69.
- Koile, K: 1997, Design conversations with your computer: evaluating experiential qualities of physical form, *CAAD futures 1997*, pp. 203-218.
- Koile, K: 2004, An intelligent assistant for conceptual design, in Gero, JS (ed), *Design Computing and Cognition '04*. Kluwer, Dordrecht. pp. 3-22.
- Koning H and Eizenberg J: 1981, The language of the prairie: Frank Lloyd Wright's prairie houses, *Environment and Planning B: Planning and Design*, vol. 8 pp. 295-323
- Peponis J, Hadjinikolaov E, Livieratos C and Fatouros DA: 1989, The spatial core of urban culture, *Ekistics* 56(334/335), pp. 43-55.
- Prats M, Earl C, Garner S and Jowers I: 2006, Exploring style through generative shape description, *AIEDAM Journal*, Cambridge University Press, 20(3), (forthcoming).
- Robles-Kelly A, and Hancock E R: 2003, Edit Distance From Graph Spectra. *Proceedings of the ninth IEEE International Conference on Computer Vision (ICCV 2003)*
- Rosenblatt F: 1958, The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain, *Psychological Review*, 65(6) pp. 386-408.
- Spiliopoulou G and Penn A: 1999, Organisations as Multi-Layered Networks, *Proceedings, 2nd Intl. Space Syntax Symposium*, pp. 1-24.
- Stiny G: 1976, Two exercises in formal composition. *Environment and Planning B: Planning and Design*, vol. 3 pp. 187-210.
- Stiny G: 1980, Introduction to shape and shape grammars. *Environment and Planning B: Planning and Design*, vol. 7 pp. 343-351.
- Stiny G and Mitchell WJ: 1978, The Palladian grammar. *Environment and Planning B: Planning and Design*, vol. 5 pp. 5-18.
- Tang H-H and Gero J: 2001, Sketches as affordances of meanings in the design process. In J. S. Gero, B. Tversky and T. Purcell (eds), 2001, *Visual and Spatial Reasoning in Design II*, University of Sydney, Australia, pp. 271-282.
- Tillmann, B, Abdi, H and Dowling, WJ: 2004, Musical style perception by a linear auto-associator model and human listeners, *Proceedings of the 8th International Conference on Music Perception & Cognition*, Evanston, IL.
- Turk, M and Pentland, A: 1991, Eigenfaces for Recognition, *Journal of Cognitive Neuroscience*, 3(1) pp. 71-86.
- Turner, A: 2005, An Algorithmic Definition of the Axial Map, *Environment and Planning B: Planning and Design*, 32(3) 425-444.
- Turner A, Doxa M, O'Sullivan D, and 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.
- Vapnik V: 1995, *The Nature of Statistical Learning Theory*. Springer-Verlag, New York.
- Vitruvius: 2001, *Ten Books on Architecture*. Cambridge University Press.
- Weston, R: 2004, Key buildings of the twentieth century: plans, sections, elevations. W W Norton & Co. Inc., New York.