

# **Design agents and the need for high-dimensional perception**

**Sean Hanna**

*Bartlett School of Graduate Studies, UCL, UK*

Designed artefacts may be quantified by any number of measures. This paper aims to show that in doing so, the particular measures used may matter very little, but as many as possible should be taken. A set of building plans is used to demonstrate that arbitrary measures of their shape serve to classify them into neighbourhood types, and the accuracy of classification increases as more are used, even if the dimensionality of the space in which classification occurs is held constant. It is further shown that two autonomous agents may independently choose sets of attributes by which to represent the buildings, but arrive at similar judgements as more are used. This has several implications for studying or simulating design. It suggests that quantitative studies of collections of artefacts may be made without requiring extensive knowledge of the best possible measures—often impossible in real, ill-defined, design situations. It suggests a means by which the generation of novelty can be explained in a group of agents with different ways of seeing a given event. It also suggests that communication can occur without the need for predetermined codes or protocols, introducing the possibility of alternative human-computer interfaces that may be useful in design.

## **1.0 Introduction**

Examination of the act of design by an individual agent, whether human or artificial, frequently involves an attempt to define the way in which that agent perceives the world. This paper suggests that the specific attributes an agent may perceive are relatively unimportant, but rather it is a high dimensionality of perception or input that is necessary.

In studying design or implementing an artificial agent, therefore, the attributes of a design artefact to be measured need not—indeed, should not—be determined a priori. While the suggestion that any set of attributes will do may seem counterintuitive, this paper will attempt to show that

there is a more effective alternative strategy. This is to consider a large number of possible attribute dimensions, even if arbitrary, and allow the agent to select the relevant subset or subspace from these. This effectively allows for interpretation and reinterpretation on the part of the agent.

The strategy will be demonstrated with respect to a real set of design artefacts: building plans taken from various neighbourhoods. By taking a number of quantifiable measures of the shape of each, it is possible to classify the buildings such that each is identifiable as belonging to its particular neighbourhood. In brief, it will be shown that while some measures may be more or less useful in this, the correct identification of buildings improves as more measures are taken.

This has implications with respect to design creativity both at the level of the individual agent and of the group. For the individual, these concern the level at which symbolic representation occurs. Approaches to representation in Artificial Intelligence can be broadly positioned with respect to two extremes: a classical approach considering intelligence to be the manipulation of “physical symbol systems” directly representing the world [1], and a radically embodied one in which the world need not be represented at all [2], [3]. While the latter has strong merits, there are many aspects of design, from words to drawing conventions to standardised CAD representations, that appear strongly symbolic at least as far as communication is concerned. These symbolic elements are characterised by an interface that is clearly defined and comparatively low-bandwidth [4]—it is a reduction of the full dimensionality of possible measurements of the world. The classical assumption (famously made by Simon [5] in his description of an ant on the beach) is that this interface is identical to (or possibly external to) the boundary of the creative agent. What is suggested here, however, is that to the extent a symbolic interpretation or reduction of dimensionality exists, it must be internal to the creative agent. Perception is high-dimensional, then interpreted internally.

For a collective of many agents, this implies there may be at one time a variety of different interpretations of any observed event, a phenomenon that is arguably necessary for the generation of novel ideas. Many theories of creativity take the essential moment of insight as “seeing [something] as” something else [6] or changing “frames of reference” [7]; even within the extreme symbolic stance, Newell and Simon [1] mention the potential advantage of “moving from one representation to another”. Clarke [8] explicitly notes from extensive archaeological data that novelty arises from small changes naturally inherent in the population; and reflective, hermeneutical [9] and systems [10] approaches to creativity or design likewise suggest that this novelty arises naturally, without being artificially imposed.

Hillier and Hanson [11] introduce the concept of *morphic languages*, in which the linguistic expressions may be the designed artefacts themselves, but the lack of a single, shared symbol system extrinsic to the agent raises a potential problem for communication. This paper aims to show that changes can happen as a result of different interpretations, as above, but communication is still possible. It will outline how agents can still make similar decisions due to patterns inherent in the observations, and demonstrate that this is possible for at least one set of data relevant to the design of architecture.

Clark and Thornton [12] make a distinction between two types of machine learning problems: type-1, in which the relevant patterns in data are immediately apparent; and the more difficult and complex type-2, in which any number of arbitrary patterns may be seen, and the data must be re-coded before the relevant regularities are visible. The latter type are apparently far more prevalent in real world data (and interesting design situations), but by structuring our thought via language, social custom and other observation external to the data, humans demonstrate an ability to turn a type-2 problem into a tractable type-1. This paper will go a step further, to suggest that the data itself, in instances relevant to design, may gradually approach type-1 as more dimensions or attributes are observed. In this way, different agents may differ slightly in their independent judgements, yet overlap enough that communication via the morphic language of the artefacts themselves becomes possible.

## 2.0 Relevance

If it can be demonstrated that for many instances of design the particular choice of attributes/dimensions is of less relevance than the number used, this will impact at least three broad areas.

In the first case, it determines the possibility of quantitatively studying design via its artefacts without having to be sure about the validity of the particular measuring system used. If two significantly high-dimensional systems will converge on the same results, either one may be used effectively. This is particularly relevant as most real design situations deal with what Rittel and Weber [13] term “wicked” problems—a set of problems which can never be clearly defined and have unforeseeable implications and effects. Thus in studying design to make recommendations for real design practice, one cannot rely on knowing enough about the problem in advance to inform the particular choice about the most relevant attributes to measure.

In the second case, it appears necessary for a proper understanding of creativity, with respect to reinterpretation [6], [7], [14] and social interaction [9], [10], that the mechanisms for variance within a population be investigated. If creative leaps are ultimately rooted in small changes, a model that imposes these stochastically via straightforward random number generation (as occurs in “creative” models from genetic algorithms to populations of agents) may miss a crucial feature. Investigating how differences in interpretation occur may outline and quantify how much larger changes occur in a social system.

Finally, there is the very practical issue of how a designer can interact with the computer that is increasingly necessary in practice. Almost all current interfaces are constructed on the assumption that communication is based on predetermined protocols, often via agreed symbol systems, but this need not necessarily be the case. If two distinct agents can make similar judgements about an observed event via independently and arbitrarily chosen means of measuring it, then that event stands as effective communication. In the case of design, where an important element of communication is via sketching and similar methods that are both difficult to codify and easily reinterpreted, this may allow systems of interface with future design tools that are much more akin to the way designers interact with one another.

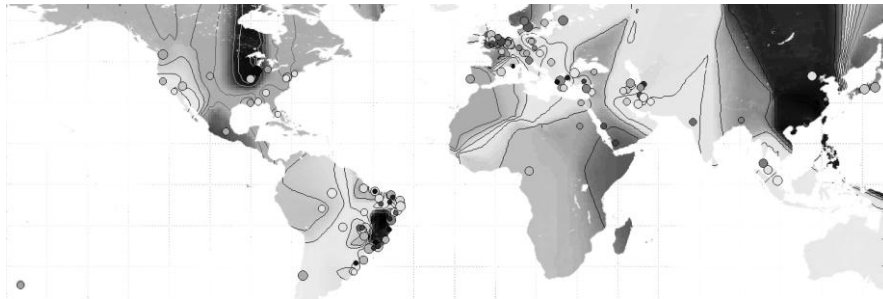
### **3.0 Example systems: observing types in architecture**

The task of recognising and identifying distinct *types* of designed artefacts is taken as a primary subject of investigation relevant to design. Several approaches to type exist, sometimes distinguishing it from style in referring to objective matters of utility rather than subjective judgement [15]. As the main aim of this paper is to demonstrate that predefinition of relevant attributes is unnecessary, type will here refer broadly to all potential characteristics. In addition, the notion of a type is sometimes treated as clearly definable [16] and permanent [17], or sometimes recreated in every generation [18]. The latter view is taken again, for the same reason.

Clarke [8] sets out a clear working definition of type and demonstrates its effectiveness. An artefact type consists of a set of measurable attributes that is not monothetic, in that every member of the group displays all of the set of attributes, but polythetic, in that there is a looser overlap between attribute subsets. This correlation between individual members varies by context and scale, as Clark also uses the polythetic set to describe assemblages and cultural groups at higher levels. The effectiveness of this

in an archaeological context is particularly relevant in that (as Clarke frequently notes) the attributes available to the archaeologist are necessarily limited and arbitrarily selected by the gap in time. This definition also lends itself to multivariate and computational methods, cluster analysis, and unsupervised learning.

The use of high-dimensional input has been shown effective in revealing types in artefacts at many levels of scale. For architectural and urban examples, spatial configuration is frequently represented topologically by a graph—the edit distance between these has been used, for example, to identify differences between Turkish and Greek house types in Cyprus [19]. At a larger scale, in a data set of 150 cities distributed around the world, the spectra of the entire street network graph was used to identify each as a vector in a 100 dimensional space, from which a subspace was extracted to represent the set [20]. The identification of a given city's geographical location was then found to be largely predictable purely by its form (fig 1).

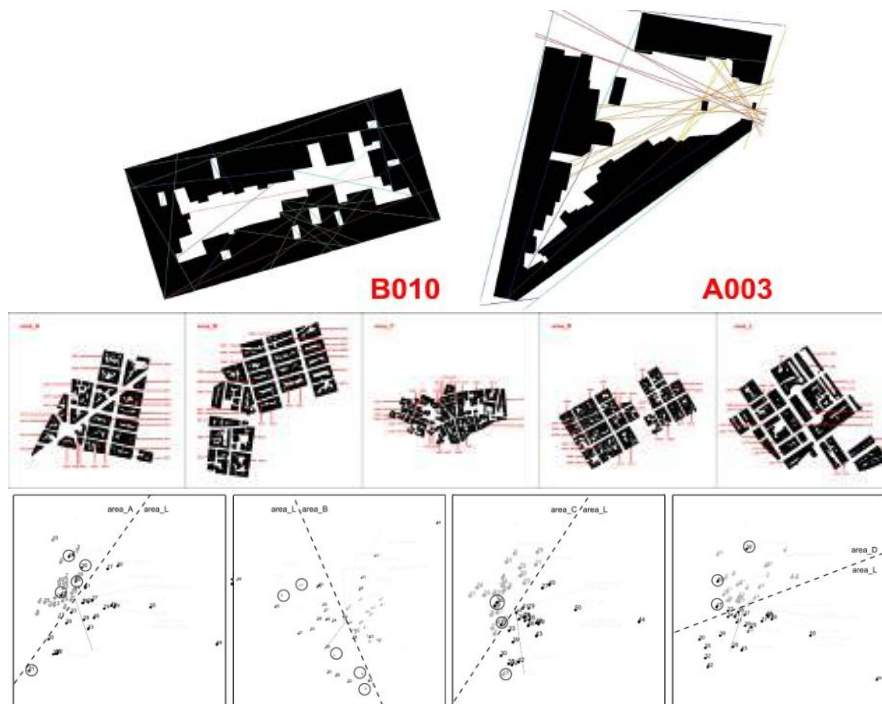


**Fig 1.** Cities represented by their graph spectra can be placed geographically based on the form of their street network [20].

The use of such numerical type definitions has also been used to effectively guide a search in design generation or optimisation, by defining an objective function for a genetic algorithm to produce desk arrangements for the layout of workplace interiors [21]. Here, the objective is not set explicitly, but derived independently by a supervised learning algorithm based on a set of precedent examples. The algorithm derives the relevant features from the input set (e.g. convex groups of desks, clusters of a certain size) without any prior definition of these features, and generates plans to match these. In each of these cases, input to the algorithm is high dimensional, then reduced as required.

A similar set of types is used as the example for the work described here. The data is taken from a study of the properties of the building footprint and block configuration of four distinct neighbourhoods in Athens,

and one in London [22]. In this, it was demonstrated that a set of measurements of arbitrarily selected attributes of the plans of each block were sufficient to classify them by neighbourhood, even though the particular features relevant to this were not known. A set of thirteen individual measures were used, including topological features such as number of courtyard voids and geometrical features such as fractal dimension. Principal component analysis (PCA) of this thirteen dimensional space then revealed a distinct clustering of blocks by neighbourhood (fig 2).



**Fig 2.** A set of measurements taken of the shape of urban blocks (top) allow them to be clustered into distinct neighbourhoods. Image: Laskari et al. [22].

This example set of buildings has been chosen partly because its design scale is familiar. More importantly, while the dimensions of properties such as graph spectra are quite abstract, the particular measures used to describe the samples are each distinctly comprehensible, clear and distinct, even though their selection was arbitrary. The following section will investigate the reasons why such an arbitrary selection of attributes results in a correct classification into distinct neighbourhood clusters, and in particular the main hypothesis of this paper—that this is a result of there being a sufficiently high number of such attributes in the data set.

## 4.0 Method and results


Under almost any circumstances, increasing the number of dimensions will permit an increased number of possible allowable classifications—this need not be tested. What will be tested here is classification on the particular space given by PCA. This is an unsupervised method which will yield a fixed subspace determined by the variance of the set as a whole, regardless of any class labels that might be assigned. The degree to which such an unsupervised (PCA) analysis of real samples allows classification into separate groups will therefore indicate the degree to which independent agents may observe the same phenomena without prior labelling.

Each sample is quantified by measurement of thirteen attributes, taken from [22] (the first four relate to changes in direct sight lines from different locations on the perimeter of the internal voids, see [23]):

1. mean connectivity value for the perimeter of the internal voids
2. mean distance between subsequent points of mean connectivity
3. vertical standard deviation of perimeter connectivity
4. horizontal standard deviation of perimeter connectivity
5. fractal dimension
6. perimeter of all voids internal to the block
7. number of voids internal to the block
8. total block area
9. building footprint area
10. number of buildings in the block
11. number of disjoint building clusters in the block
12. ratio of internal voids open to the street
13. number of vertices on the building contour

In each of the following sections, the method will select subsets of attributes, of sizes varying between  $a=1$  and  $a=13$ . These will constitute the maximum possible input to a theoretical agent. Although far greater dimensionality would be possible, this range will suffice to show a clear trend of improved classification as more dimensions are used. PCA will then be performed on these to yield a feature space  $\Phi$  of reduced dimensionality  $d$  (typically three-dimensional). It is within this space  $\Phi$  that classification will be performed, and the effectiveness of this determined by the minimum linear classification error within this reduced PCA space.

Classification errors will be calculated as the ratio of incorrectly classified samples within the total set of 125 samples (25 building groups in each of 5 classes). Errors will therefore be shown to 3 significant digits.



**Fig 3.** The classification error decreases as the space in which samples are classified increases in dimensions (bold indicates mean error, grey the range from min to max error).

#### 4.1 Effect of overall attribute dimensionality on classification

Superficially, the number of dimensions used in any supervised classification task will have an obvious effect on the accuracy of the results—more dimensions yield a greater variety of hyperplanes for drawing distinctions between classes, and if all samples are appropriately labelled, the system has more opportunities to select the appropriate ones. This can be clearly seen in the data (fig 3) with an expected decrease in error as dimensions increase from 1 to 13. (The mean error decreases monotonically while the minimum error reaches its optimum point at 8 dimensions, before the overall variance decreases due to a decreasing number of possible permutations.) What is not obvious, however, is whether an increase in the overall number of *available* attributes is of any further benefit beyond this. The above fact gives no reason to expect any improvement in classification whatsoever when:

- a) the examples are not labelled (unsupervised learning), as would be required for autonomous reinterpretation, etc., or




- b) a subspace of constant, reduced dimensionality is available, as is always necessary in practice—an infinite set of attributes theoretically exists but is impossible to observe.

This section tests the hypothesis that the overall quantity of attributes is only relevant inasmuch as it provides a greater dimensionality in which arbitrary classification can take place (and therefore more possible classification hyperplanes), and finds it to be false. Rather, a pattern that may be considered intrinsic to the data set itself becomes progressively more evident as more attributes are used.

The effect of varying numbers of attribute dimensions within the data set was tested by taking the mean errors of classifications performed within a subspace of constant dimensionality, derived from principal components. A low dimensional subspace  $\Phi$  of dimensionality  $d=1$  to  $d=5$  is taken from the principal components of the entire data set as specified by a given set of attributes  $a \geq d$ . Classification errors are then compared from linear discriminant analyses performed in  $\Phi$  of equal dimensionality  $d$ . By varying the number of attributes  $a$  used in each subspace  $\Phi$ , the effect of available attribute dimensionality can thus be compared independently of the classification dimensionality  $d$ . All possible permutations of  $a$  attributes are used from the total set of 13, classification performed on the resulting PCA subspaces  $\Phi$ , then the overall mean error is recorded (fig 4) for increasing sets of attributes  $a=d$  to  $a=13$ .

Bold lines show the mean classification errors in  $\Phi$  of dimensionality  $d=1$  (top) to  $d=5$  (bottom). Where  $a=d$  (the minimum possible, with no reduction in dimensionality), classification is identical to that in the space of the original attributes and cannot be further improved. To the extent that the attribute dimensionality  $a$  is relevant only by virtue of increasing the dimensionality of the classification space  $\phi$ , each of these mean errors should show no further improvement, and tests with randomly labelled data sets (dashed and thin lines, fig 3) show this to be the case. However, all tests show improved classification. For a single component (top) this improvement is negligible, but for 3, 4 and 5 dimensions  $d$  there is significant improvement, approaching that of the optimal classification in the full dimensionality  $a$  of the original attributes.

This demonstrates that an increased number of attributes is clearly of value in describing the structure of the data, even when only a fixed number of components are used. As more attributes are used, the dimensions in which the data is naturally most varied overall more closely approximate the dimensions most useful for distinguishing the separate subclasses. This result is far from inevitable, as the attributes in question were chosen arbitrarily and so may have turned out to be redundant or conflicting.



**Fig 4.** Classification error decreases as more initial attributes are used, even if the space of classification remains of a constant dimensionality. Bold lines indicate spaces  $\Phi$  of 1 to 5 dimensions. Thin and dashed lines show no improvement for random data sets.

What this suggests is that the subgroups or clusters revealed by PCA are inherent to the data itself, rather than arbitrary designations imposed by the particular labelling scheme—they are gradually revealed as Clark and Thornton’s [12] type-1 as more attribute dimensions are used. By contrast, the narrow lines in figure 4 indicate the effect of arbitrarily chosen classes, with dashed lines showing the mean errors for the same data in which only the labels were sorted at random, and solid lines showing the same for data in which each attribute value was resorted independently. In all cases the error rates are not only poor, but fail to increase as more attributes are used to define them.

This contrast indicates that the labelled clusters within the data set are intrinsically meaningful, in that they are discovered by unsupervised and unlabelled PCA, and are simply revealed by the measurement of larger sets of attributes. The following sections will unpack this observation by investigating whether a relationship is discernable for particular subsets of attributes.

#### 4.2 Effect of particular sets of attribute dimensions

The possible subsets of attributes of any fixed dimensionality  $a$  are not equal in terms of accurate classification—there is some variance in the error rates from which the means above were taken. If constrained to a limited dimensionality of measurements to be taken, one would hope to be able to choose the optimal set of attributes to yield the best classification. This section examines whether the improvement in classification with larger sets above is a result of particular combinations of attributes, and whether these optimal combinations can be determined beforehand. It tests:

- a) whether any *particular individual attributes* can be found to contribute to the overall reduction in misclassification error
- b) whether there are similarities between *particular sets of attributes* that reduce the error
- c) whether these *attribute sets can be determined prior to performing the classification* itself.


*Particular individual attributes* were found not to contribute significantly to the overall reduction in misclassification error. Contribution to errors for each attribute was calculated by taking the error rate (mean misclassified examples) in a constant subspace of dimensionality  $d=3$  for all possible combinations of  $a=6$  attributes that included the attribute in question. A significant variation in these would indicate specific attributes responsible for error or correct classification. However, while errors using subsets of  $a=6$  attributes had a considerable range from 0.296 to 0.616 overall, the range in mean errors due to particular attributes was minor, as shown in table 1.

**Table 1.** Contribution to errors: the mean errors for all sets of 6 attributes that include the attribute in question.

A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12	A13
.434	.464	.469	.459	.460	.465	.431	.457	.475	.447	.464	.451	.464

This very slight contribution of each attribute to error rates became more significant, however, when *particular sets of attributes* were considered. While there was a negligible correlation ( $R=0.16$ ) found between the similarity between any two sets of attributes and their error rates, the very best sets resulting in the lowest errors (0.296–0.313) were found to contain four of the same attributes in common: [1 7 12 13], so the effect of limiting the sets to particular attributes was tested next. Attributes were ordered based on their contribution to error in table 1, and both the ‘best’ [7 1 10

12] and ‘worst’ [9 3 6 13] were progressively removed from the available attributes. The mean and range of error rates for the remaining combinations of attributes is shown in figure 5. A noticeable change in the classification errors is evident here: rising monotonically when the ‘better’ attributes are unavailable and vice-versa.



**Fig 5.** Error rates of attribute subsets with specific attributes withheld. Error rates decrease when the ‘worst’ attributes are not used, and rise when the ‘best’ are withheld.

In both tests above the optimal attributes for measurement were determined only by knowing the correct classification results—an impossibility if one is attempting a classification on unlabeled examples. The third test of attribute sets is whether these optimal subsets can be determined by any *measurable diversity within the data as a whole*, and therefore can be found before performing the actual classification itself. The increase in classification accuracy as dimensionality  $a$  increases (§4.1) suggests the hypothesis that the optimal attribute subsets are those that are most diverse, in terms of each attribute independently providing more information about each example. If this is true, it would both explain this improvement

for large sets and suggest a method by which the appropriate attributes for any given data set can be found.

The simplest measurement of the independence of two attributes with respect to the data is the correlation between their respective measures of all the examples in the data set. A high product of these coefficients for a given set of attributes should indicate greater independence or diversity and therefore a lower classification error. This was found not to be the case, with almost no correlation between attribute dimension diversity and error rate ( $r^2 < 0.0025$ ). A number of other measurements of attribute set diversity were also taken (product of minimum angles, sum or product of dot products, convex hull volume), using the projections of principle axes for each into the 3-dimensional classification space  $\Phi$ , with similar results. Within each of these tests, neither the overall correlations nor the sets of attributes known to be error minimising revealed any discernable pattern—these latter, ‘best’ sets were sometimes of high diversity, but just as often medium or low.

These results suggest that while some subsets of attributes of a fixed dimensionality  $a$  are clearly superior in terms of providing an accurate classification, these do not have any relationship to the distribution of the data that can be found prior to classification itself. One therefore cannot identify a priori which ones to use for the classification.

### 4.3 Different sets of attributes for different classes

Classification errors for all five classes have been used together to examine the attribute subsets above, but the resulting suggestion that particular attributes (e.g. 1, 7, 10, 12) are optimal is somewhat misleading, as different classes may be better distinguished by different attributes. Table 2 shows the attributes that best classify each individual class in isolation (within a space  $\phi$  of  $d=3$ , again found by PCA) for attribute sets of  $a=3$ ,  $a=4$ ,  $a=5$  and  $a=6$ . The top row in each section shows the errors in each class for the optimal attribute set for Class 1, the second row for Class 2, etc. Those that are optimal for one class are not for another, as evidenced by the greater classification errors for other classes (grey cells), and the fact that many of these errors even increase as more dimensions are used. There is some overlap in the dimensions used—attributes 7 (no. of voids) and 8 (block area) occur repeatedly—but there is also a great deal of difference. When limited to three attributes, [7 8 10] provide an excellent classification of Class 1 samples, but the best subset for classification of Class 2 samples, [3 9 12], contains none of the same attributes. Nor are these particular sets of much use in classifying the other classes within the data set: while [3 9 12] yields a successful result (error=0.125) for Class 2,

it misclassifies nearly half of the Class 1 examples (error=0.42) and does little better for the others. The same effect can be seen for all classes and subset dimensionalities in the table.

**Table 2** The subset of attributes that best classify one class differ from those that best classify another. Bold indicates attributes that do not appear in the previous table.

Best 3 attributes	Class 1	Class 2	Class 3	Class 4	Class 5
[7 8 10]	0.13	0.225	0.28	0.22	0.245
[3 9 12]	0.42	0.125	0.31	0.35	0.27
[7 8 13]	0.18	0.25	0.17	0.335	0.24
[3 8 10]	0.245	0.18	0.27	0.16	0.36
[1 7 12]	0.285	0.17	0.23	0.435	0.135

Best 4 attributes	Class 1	Class 2	Class 3	Class 4	Class 5
[2 7 8 10]	0.125	0.215	0.28	0.24	0.2
[4 5 8 12]	0.365	0.12	0.31	0.215	0.34
[1 7 8 13]	0.235	0.27	0.165	0.405	0.19
[3 5 11 13]	0.38	0.33	0.29	0.16	0.31
[1 4 7 12]	0.28	0.145	0.26	0.285	0.115

Best 5 attributes	Class 1	Class 2	Class 3	Class 4	Class 5
[2 4 7 8 10]	0.125	0.24	0.28	0.245	0.22
[3 4 8 11 12]	0.395	0.115	0.29	0.325	0.395
[1 8 10 11 13]	0.225	0.31	0.15	0.395	0.165
[1 4 5 11 13]	0.315	0.29	0.305	0.155	0.235
[1 2 5 7 10]	0.165	0.25	0.28	0.26	0.105

Best 6 attributes	Class 1	Class 2	Class 3	Class 4	Class 5
[2 4 6 7 8 10]	0.125	0.245	0.29	0.24	0.215
[2 3 4 8 11 12]	0.33	0.11	0.305	0.3	0.385
[1 2 8 10 11 13]	0.225	0.235	0.145	0.33	0.245
[1 2 4 5 11 13]	0.335	0.305	0.29	0.155	0.23
[1 2 5 6 7 10]	0.195	0.315	0.3	0.24	0.08

Worst attributes	Class 1	Class 2	Class 3	Class 4	Class 5
[1 3 6 9 12 13]	0.47	0.19	0.22	0.34	0.335
[1 7 9 10 11 13]	0.155	0.455	0.26	0.33	0.22
[1 2 3 4 5 6]	0.215	0.245	0.39	0.285	0.22
[1 7 8 9 11 13]	0.26	0.27	0.225	0.525	0.215
[3 6 9 11 12 13]	0.41	0.215	0.22	0.385	0.56

There is also less consistency than might be expected as the dimensionality of the attributes  $a$  increases. Attribute numbers in bold show the new attributes added to the optimal subset for each class as the subset increases by one. In many cases, several new attributes replace the previous ones, indicating several of those providing the best classification when only three are allowed are no longer optimal when four are used. The final portion of Table 2 shows the *worst* performing attribute sets for each class, which reveals several attributes (not in bold) that appear also in the corresponding sets of optimal attributes immediately above.

These inconsistencies with respect to optimal attribute subsets for individual classes reinforce the result of section 4.2. While certain attributes or attribute subsets are particularly well suited for classification of specific classes, there appears no way to determine these without knowing the classification beforehand.

#### 4.4 Mutual classification and communication

In considering communication between two artificial agents, or between a human subject and a computer system (or by extension even between two human subjects), the overall error rate resulting from a single attribute subset is less significant than the particular classifications that subset yields. For communication to be effective, it is necessary that the two agents in question make similar judgements on any given piece of data used for communication—that they see the world in a similar way. This section examines the way in which subsets of attribute dimensions matter with respect to the specific examples being misclassified, and the manner of their misclassification.

The mutual classification between pairs of attribute subsets was measured for subsets of  $a=3$  to  $a=12$ . To determine the overall difference in how the pair perform, the total number of examples in the data set that the two classify differently is used; actual classification errors with respect to known labels are ignored. This difference in performance is then compared with the difference between the attributes themselves, where the similarity between any two attributes is calculated by the correlation between their independent classification of the data. The sum of these maximum correlations gives an overall measure of similarity between any two subsets.

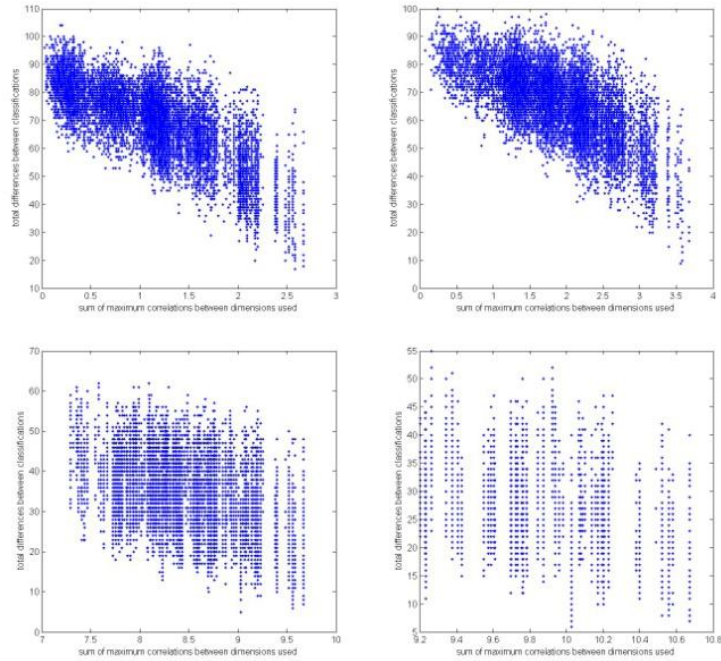
Given the measure of the whole data set by one attribute  $\alpha$  as  $M(\alpha)=\{\mu_1, \dots, \mu_n\}$ , two attribute sets  $[\alpha_1, \dots, \alpha_a]$  and  $[\beta_1, \dots, \beta_a]$ , and corresponding sets of classification results  $[\kappa_{\alpha 1}, \dots, \kappa_{\alpha n}]$  and  $[\kappa_{\beta 1}, \dots, \kappa_{\beta n}]$ , the difference in performance is given by:

$$classDiff = \sum_{x=1}^n \kappa_{\alpha,x} \neq \kappa_{\beta,x}$$

and the similarity between attribute sets by:

$$\sum_{x=1}^{\alpha} \max_{1 \leq y \leq \alpha} \text{corr}^2(M(\alpha_x), M(\beta_y))$$

Figure 6 illustrates that differences between attribute sets do result in approximately corresponding differences between sets of misclassified examples, but only for small  $a$ . As  $a$  increases, this correlation between attribute dimensions and classification results also decreases. For small subsets of dimensions (3,4, etc.) there is a significant degree of correlation (approx  $r^2=50\%$ ), however this decreases to insignificant correlation as more dimensions are added, with  $r^2=12\%$  for  $a=11$  and  $r^2=0.3\%$  for  $a=12$ .



**Fig 6.** Correlations between attribute subsets and classification results indicate the degree to which mutual understanding is dependent on the use of the same attributes by both agents. It is highly dependent for small subsets of 3 or 4 (above), much less so for 11 or 12 (below).



With respect to the situation of two communicating agents, the prior agreement on the particular attribute dimensions used thus appears to matter greatly when the number of attributes is low, as evidenced by the strong correlations for low values of  $a$ . But this matters less the more dimensions we have. At the same time, increasing dimensions decreases the overall difference in performance between agents. If the general strategy for overall error reduction (as indicated in section 4.1) is to increase  $a$ , this also appears to improve the likelihood of different agents making same distinctions (i.e. having similar  $\phi$ ) without prior agreement on the particular sets of attributes to use.

## 5.0 Conclusion

While one should generally expect to find a decrease in classification error as the number of dimensions of the classification space increases, there is no reason to expect this when a classification space of constant dimensionality is derived from arbitrarily varying sets of initial dimensions. As seen in §4.1, this is not the case for a randomly labelled set of data. Nevertheless, classification in spaces derived by PCA (with 1 to 5 components) was seen overall to steadily improve as more initial dimensions were used (§4.1), suggesting that for data sets with similar properties to the one under investigation it is generally beneficial to use as many dimensions of measurement as are available.

There are, naturally, some particular small sets of attributes that yield a better classification than others, or better even than larger sets, but in practice there is no means for determining what these sets are. There appears to be no intrinsic relationship between these attributes (e.g. diversity) with respect to the data as a whole (§4.2), and the best subset of attributes for one class are unsuited to another (§4.3). In practice, in creating a system that is to evaluate any group of artefacts, the remaining viable strategy is not to carefully select the subset of features thought to best describe the relevant classes, but to use the full set of as many features as are available.

The reason larger sets result in better classification appears to be due to the nature of the examples themselves. Clark and Thornton's [12] distinction is useful: type-1 examples would contain regularities that are immediately apparent, while for type-2 the examples may afford a number of equally valid interpretations. The supposed way in which a culturally embedded person deals with the latter is by relying on a structure imposed externally by language or other cues to recode the observations to type-1.

However, there may be intermediate possibilities between the two types. A hypothetical set of examples, measured in infinite dimensions, may be type-1, but appear as type-2 when only a limited subset of these are used. If the classes are determined by a polythetic set of attributes [8] this is almost certain to be the case, because the incomplete overlap between attributes would be less evident as fewer attributes are used. It would appear that the urban block data used here is of this nature, thus as it is viewed in more dimensions, the inherent type-1 pattern becomes gradually more evident.

This may also explain what is occurring when we are able to recognise as obvious distinct types in other sets of artefacts, even when we are unable to describe the specific criteria for the decision. The data set is type-1, but only when considered in many dimensions. If a continuum of intermediate possibilities is considered between type-1 and type-2, the degree to which classifications based on arbitrary subsets of attributes converge as those subsets increase in dimensionality offers a possible means of measuring this for a given set of data.

The observation that the overlap in attributes increases with larger sets implies that communication between distinct agents is possible without having to predetermine which particular dimensions are relevant. The strategy, again, of each agent using as many attributes as possible will ensure that each forms similar interpretations of a given event (§4.4).

However, the use of different subsets of attributes by different agents in a population may help to explain the generation of the novelty that is essential to the creative process in social models of creativity [10]. Many definitions of creativity hinge around coupled notions of novelty and utility [24], but while the latter is clearly justifiable for straightforward functional reasons, the need for novelty is less easily explained. If this novelty actually arises naturally, however, there need be no internal or external motivation required. It seems a fair conjecture that if slightly different sets of attributes can exist within a population that makes similar interpretations of events, then an ostensibly similar population may always contain some variance in its underlying choices of attributes, or ways of seeing. If an occasional event should arise that is interpreted differently among agents, an apparently novel difference of opinion will result. Such a variety of ways of seeing, or “frames of reference” is often stressed as the crucial component of creative insight [7][14], and appears in agent models of creativity [25][26]. In these, agents make design decisions by evaluating available options at any given point against a test criterion given by their own point of view. The result is that this difference in criteria easily drives collective innovation

This also provides a notion of novelty that is not at odds with that of utility. Creativity is often described as a series of paradoxes [27], or explained as seeking a median point between “too similar” and “too different” [25], but the paradox does not exist if more dimensions are assumed. An interpretation of an event by an agent as perfectly normal may simply be seen by the next as highly unusual due to a measurement in a differing dimension. In human terms this conjecture needs further investigation, but it seems natural that we need only do what seems useful and appropriate to us, while others will always interpret things somewhat differently.

### Acknowledgements

I would like to thank Anna Laskari for the initial analysis of building plans that made this research possible, and for the use of the data she has collected. This research has been supported by a Research Councils UK Academic Fellowship, EPSRC grant no. EP/E500706/1.

### References

1. Newell, A. and Simon, H. (1976). Computer Science as Empirical Enquiry: Symbols and Search. *Communications of the Association for Computing Machinery* 19: 105- 32.
2. Brooks, R.A., (1991) Intelligence without representation, *Artificial Intelligence* 47, 139–159.
3. Dreyfus H (2007) Why Heideggerian AI failed and how fixing it would require making it more Heideggerian, *Artificial Intelligence* 171, pp.1137–1160. Elsevier.
4. Haugeland J (1978) The Nature and Plausibility of Cognitivism, *Behavioral and Brain Sciences*, 2 pp. 215–60.
5. Simon H (1996) *The Sciences of the Artificial*. 3rd ed. MIT Press, Cambridge, MA.
6. Schön D A (1963) *Displacement of Concepts*. Tavistock, London.
7. Akin Ö and Akin C (1996) Frames of reference in architectural design: analysing the hyperacclamation (A-h-a-!), *Design Studies* 17(4), pp. 341–361.
8. Clarke DL (1968) *Analytical Archaeology*. Methuen & Co, London.
9. Snodgrass AB and Coyne RD (1997) Is designing hermeneutical? *Architectural Theory Review*, 2(1), pp.65–97.
10. Csikszentmihalyi, M (1988) Society, culture, and person: a systems view of creativity, in Sternberg, RJ (ed), *The nature of creativity: Contemporary psychological perspectives*, Cambridge University Press, Cambridge, pp. 325-339.

11. Hillier B and Hanson J (1984) *The Social Logic of Space*. Cambridge University Press.
12. Clark A and Thornton C (1997) Trading spaces: Computation, representation and the limits of uninformed learning, *Behavioral and Brain Sciences* **20**: 57–90.
13. Rittel HWJ and Webber MM (1984) Planning problems are wicked problems. In Cross N (ed.) *Developments in Design Methodology*, John Wiley and Sons.
14. Koestler A (1964) *The act of creation*. Hutchinson.
15. Westfall CW (1991) Building Types. In van Pelt RJ and Westfall CW, *Architectural Principles in the Age of Historicism*. Yale University Press, New Haven.
16. Alexander C, Ishikawa S, Silverstein M, Jacobsen M, Fiksdahl-King I and Angel S (1977) *A Pattern Language*, Oxford University Press, New York.
17. Rossi A (1982) *The Architecture of the City*. The MIT Press.
18. Colquhoun A (1967) Typology and Design Method, *Arena*, 83, pp. 11–14. Reprinted in Colquhoun A (1985) *Essays in Architectural Criticism: Modern Architecture and Historical Change*. The MIT Press.
19. Conroy-Dalton R and Kirsan C (2008) Small graph matching and building genotypes. *Environment and Planning B: Planning and Design*, 35 (5): 810–830.
20. Hanna S (2009) Spectral comparison of large urban graphs, In Koch D, Marcus L and Steen J (eds.) *Proceedings of the 7th International Space Syntax Symposium*. Royal Institute of Technology (KTH), Stockholm, Sweden.
21. Hanna S (2007) Defining Implicit Objective Functions for Design Problems, *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO-2007*. ACM Press.
22. Laskari S, Hanna S and Derix C (2008) Urban identity through quantifiable spatial attributes: Coherence and dispersion of local identity through the automated comparative analysis of building block plans, in Gero JS and Goel A (eds) *Design Computing and Cognition '08*, Springer.
23. Psarra S, Grajewski T (2001) Describing shape and shape complexity using local properties. *Third International Space Syntax Symposium*, Atlanta.
24. Boden MA (1990) *The creative mind: myths & mechanisms*. Weidenfeld & Nicolson, London.
25. Saunders R and Gero JS (2001) Artificial creativity: A synthetic approach to the study of creative behaviour, in JS Gero and ML Maher (eds), *Computational and Cognitive Models of Creative Design V*, Key Centre of Design Computing and Cognition, University of Sydney, Sydney, pp. 113-139.
26. Hanna, S. (2005) Where creativity comes from: the social spaces of embodied minds. In: Gero, J.S. and Maher, M.L., (eds.) *Proceedings of HI'05, Sixth International Conference of Computational and Cognitive Models of Creative Design*. University of Sydney, Sydney, Australia.
27. Cropley, AJ: 1999, Definitions of creativity, in MA Runco and SR Pritzer (eds.), *Encyclopedia of Creativity*, Academic Press, San Diego, pp. 511–524.