Entropy as a Computational Aesthetic Measure

Richard Cant, Caroline Langensiepen and Andrew Burton
School of Science & Technology
Nottingham Trent University
United Kingdom
richard.cant@ntu.ac.uk

***Abstract* — We derive a measure that describes the complexity and consistency of a 3D virtual scene, so as to be able to compare scenes used for training purposes, and to allow changes in a scene to generate appropriate responses within games. The measure is calculated by constructing an entropy tree that links all the objects within the scene together. The total entropy is thus the sum of all the individual 'object to object' entropies within the tree. The object to object entropy is derived from the nature of the objects (their semantic relevance to each other, termed selection entropy) and to their geometric arrangement. We calculate the selection entropy for a number of examples to show how it quantifies the aesthetic consistency and complexity of the scene.**

Keywords - Entropy, Procedural Content, Scene Modelling, Evaluation

I. Introduction

Within games and simulations, the need to increase the richness of the experience has long been known. Indeed, in military training simulators, one of the key motivations for their use is to provide trainees with scenarios of a richness and complexity that would be expensive, difficult and dangerous to provide in 'real world' exercises. As the quality of graphics and processing in home computers has increased, so this quest for richness has moved into the domain of computer games.

With modern hardware the level of sophistication is now such that the *nature* of the complexity rather than just the level of complexity becomes important. The aesthetic aspects of consistency, relevance and appropriateness contribute to the immersive qualities of an environment, in that wholly consistent scenes will be better accepted by the user as an environment within which they enact their role.

In this paper we will show how scene layout can be interpreted (and indeed calculated) mathematically as entropy. One is usually introduced to the concept of entropy by comparison with 'everyday untidiness' - which is almost immediately discarded in favour of an explanation that can be developed mathematically. This could be based on information theory (in the case of computer science) or statistical mechanics (for physics). However, within a scene, entropy and untidiness can both be understood as being due to geometry (where things lie relative to each other) and content (the selection of objects that are present). The former has been discussed elsewhere [1]. In this paper we consider the contribution to scene entropy that is made by object selection. The ability to express an ontology in mathematical form allows the complexity, aesthetic consistency and appropriateness of a scene to be measured. The entropy thus calculated provides an objective criterion for use in the design of games, or during gameplay.

A brief survey of other measures used in scene analysis and generation is provided in section 2. Section 3 explains the mathematical basis for our measure, and section 4 shows how it applies in some examples of manually generated scenes. Section 5 discusses these results and how the measure may to be used with other criteria to fully automate aesthetically appropriate scene generation.

II. Previous work

Work on measurement of scene complexity has been quite limited. Though 'real-world' cognitive studies often vary scene complexity as part of their experimental setup [2], they do so on a purely heuristic and subjective basis. Measurement of 2-dimensional scene complexity has been done by Feixas et al [3], but based purely on visual aspects such as occlusion within the field of view.

Consideration and measurement of the artefacts within the visual scene is closely linked to the area of automatic scene layout/generation, where more work has been done. One would have thought that this discipline would be relatively mature, since the

DOI 10.5013/IJSSST.a.13.3A.04

ISSN: 1473-804x online, 1473-8031 print

problem of automated layout of circuit boards is an old one [4]. However, the layout limitations of a circuit board and the necessity of laying connective tracks constrains the problem in a very different way to the problem of laying out a human occupied scene.

Semi-automated layout of scenes was done by WordsEye and Seversky and Yin [5] [6], but they took no account of any 'appropriateness' of the elements of the scene. Kjolass [7] considered layout of furniture based on templates of relationships between them, but without adding the possibility of objects placed on them. One of the earliest examples to add this feature was Xu [8]; his aim was to ensure closest spacing of placed objects. This was extended by Akazawa [9], again using contact as constraints, but allowing regular repeated layouts to occur automatically.

Sanchez [10] used a genetic algorithm to achieve a room layout, using explicitly defined rules for constraints on objects and therefore for the fitness function. Vassilas [11] also took the genetic algorithm route, coupled with a declarative method for establishing hierarchical relationships between architectural entities. Although these use some semantic information for layout, this has not been extended into the domain of 'relevance' to the location. In [12] we considered the need to try and make this placement less ordered, and started to develop a placement mechanism controlled by criteria based on regularity and relevance. More recently, [13] used a petri net method of placement based on 'attractor' points within a room around which to base the furniture. Ref. [14] used simulated annealing to place objects provided for a scene in a manner suited to best movement around the scene, but the objects used were pre-selected for the room.

The measures used by the papers on automated placement do not easily relate to aesthetics. Within industrial design, the last 10 years have seen attempts to provide modelling tools that match the perceptions and criteria of the artist, for example by tagging or ontologies to relate the design elements to the intent of the designer [15]. However, numerical analysis is often limited to well-defined areas where some simple criteria can be used and optimised [16]. The other obvious place where consistency is considered is in user interface design, but there the issues relate to consistency of layout, colour etc. or the actions needed to fulfil goals [17], [18].

We derive a complexity measure based on entropy which includes semantic elements to assess scenes. Previous work on the measurement and use of entropy and information theory for graphical scenes includes [3], [19] and [20], but (as stated earlier) these relate only to visual complexity of the field of view, with no semantic content. However, a study by [21] assessing similarity of human artefacts and relating this to their levels of entropy found a “strong and linear” relationship between entropy and diversity, suggesting that it is possible to use quantifiers of similarity to calculate entropy for a scene. We would therefore argue that any aesthetic measure of a scene should take into account the semantic relationships between the objects.

III. Formulation of the Entropy Measure

A. *Basic Entropy Definitions*

In this section we review the mathematical formulation of our entropy measure and show how it can be used as a quantitative aesthetic measure. An initial discussion of this material can be found in [22], however we clarify it here.

In physics, entropy is a well-defined thermodynamic quantity. If a system is in a macrostate that has N possible microstates (macroscopically indistinguishable), then the entropy is defined as:

$$E = k_b \ln(N)$$

Shannon [23] defined the concept of entropy as a way of measuring the density of information in a signal. It is usually expressed in terms of the probabilistic rule that generates the data as:

$$E = -\sum_i p_i \ln(p_i)$$

where p_i is the probability of the *i*th state.

Shannon's version of entropy is the same as the physical one if all states have equal probability. When evaluating the entropy of a scene we face a problem because we do not have access to the rules that generated the scene and we can only see a single microstate. Thus we must infer the rules from the scene.

For a given scene it may be possible to deduce many different rules that could have generated it. The standard method to resolve this type of situation is the principle of maximising entropy (see for example [24]). However a simplistic application of this method lead to an impasse, because given only a single sample of the system, the inevitable result would be to choose the rule that allows all possibilities equally and hence has maximum entropy.

We resolve this problem by recalling that we are considering scenarios that have been generated by human agency as opposed to natural forces. In the latter case it is reasonable to assume randomness as the baseline and attempt to establish deviations from randomness. Instead, we compare the observed arrangement against an ideal, perfectly ordered, arrangement. The principle of maximising entropy is then implemented by identifying the entropy of the observed state with that of a rule that includes all states that lie between the observed state and the ideal state, with equal probability. Note that this implies choosing the lowest entropy representation of the observed state, since that is the one that lies closest to the ideal. Thus we initially choose the most ordered interpretation of the observed state but then maximise entropy in our choice of probability distribution.

B. General Entropy Formulation

In this section we will discuss the general principles of entropy based on sets of rules. In the following section we will show how these general principles can be applied to the selection of objects.

In general any system will have a large number of potential states. Most of these will be unlikely if the system is controlled by human beings because humans have a tendency to impose structure on their surroundings. For example if you move through a typical dwelling house you will find that objects are segregated by function.

Let T be the set of states of the system. We define a rule R as a boolean function on T and so:

$$W(R)=\{t \in T \,|\, R(t)\}$$

is the set of states that are observed to obey the rule. One example of such a rule might be "all pots and pans are in the kitchen", so states of the house in which there was a pan in the dining room would disobey the rule and fall outside W.

Sometimes the differences between states are either not observable or not important. We can handle this situation using the concept of a symmetry, Z, which we define as an equivalence relation on T.

$$[a]=\{t \in T | t\, Z\, a\}$$

$|a|$ is the set of states that are deemed to be indistinguishable from a.

In the kitchen example, $|a|$ might be the set of states corresponding to re-ordering of identical pans.

For Z to be consistent with R then for two states a and b

$$a\, Z\, b \rightarrow R(a)=R(b)$$

Formally we can define the entropy S for the rule R, given the symmetry Z to be the logarithm of the cardinality of the quotient set of W(R) by Z:

$$S(R,Z)=\ln(|W(R)/Z|)$$

In future we will write the cardinality:

$$|W(R)/Z|=\Omega(R,Z)$$

or more simply as $\Omega(R)$

Next we must move from defining the entropy of the rule to defining that of the states. Most of the states $t \in T$ will satisfy a number of different rules, and so we need a criterion for choosing between them. We implement the reasoning in subsection A, above, by identifying the entropy of the state with that of the lowest entropy rule that it satisfies.

In a practical system we would not expect to have just a single rule, but rather a set of individual rules that combine together to produce a composite overall rule. We therefore need to consider the ways in which rules can interact with each other. Rules may be characterised as incompatible or intersecting. Intersecting rules can be ordered or orthogonal.

C. Incompatible Rules

Given two rules R_1 and R_2, if $W(R_1) \cap W(R_2)=\emptyset$ then we say that R_1 and R_2 are incompatible.

D. Intersecting Rules

When two rules are not incompatible then they can be described as intersecting. In that case there will be states that satisfy both. Two situations arise. It is possible that satisfying one of the rules guarantees that the other will also be satisfied. In that case $W(R_1) \subset W(R_2)$ we say that the rules are ordered. When that is not true then a new, lower entropy, composite rule can be created by combining the two:

$$R_3=R_1 \wedge R_2,\; W(R_3)=W(R_1) \cap W(R_2)$$

E. Ordered Rules

Ordered rules may be defined as follows. If two rules R and P exist such that a state s which satisfies R must also satisfy P we write $R \le P$.

If $R \le P \wedge P \le R$, then $P = R$ and the rules are equivalent. If there are states that satisfy P but do not satisfy R then we write: $R < P$ or $P > R$.

For example, any object that satisfies "is a tablespoon" will also automatically satisfy "is an item of cutlery".

Frequently, ordered rules will occur in sets. Let

$$R=\{R_i \ldots\}\, i \in \mathbb{N},\; 1 \le i \le N$$

DOI 10.5013/IJSSST.a.13.3A.04

ISSN: 1473-804x online, 1473-8031 print

be a totally ordered set such that $R_i < R_j$ *iff* $i < j$. And so:

$$W(R_i) \subset W(R_j) \text{ iff } i<j \quad .$$

We define the exclusive rule-set:

$$\begin{aligned}
&\hat{R}_i \text{ where } \hat{R}_i(t)=R_i(t) \wedge \overline{R}_{i-1}(t) \\
&W(\hat{R}_i)=W(R_i) \cap W(\overline{R}_{i-1})=W(R_i) \setminus W(R_{i-1}) \\
&R_0(t)=false \\
&W(R_0)=\emptyset
\end{aligned}$$

If we observe the system to be in a state *t*, satisfying the exclusive rule $\hat{R}_i$ then the "strictest" rule that it satisfies from the original set is R_i. We define the entropy of *t* based on R_i rather than $\hat{R}_i$ based on the principle of maximising entropy. The entropy for *t* is thus:

$$S(t)=\ln(|W(R_n)|) \quad t \in W(R_n)$$

$$n \leq i \ \forall\ i \text{ such that } t \in W(R_i)$$

This can be re-expressed in terms of the exclusive rule-set as:

$$S(t)=\ln\left(\sum_{i=1}^{n} |W(\hat{R}_i)|\right)$$

We can define a maximum compatible symmetry for the exclusive rule set using the equivalence relation Z_m, as follows:

$$a Z_m b \text{ iff } \hat{R}_i(a)=\hat{R}_i(b) \ \forall\ i$$

It is obvious that:

$$|W(\hat{R}_i)/Z_m|=1$$

and so

$$S(t, Z_m)=\ln(n)$$

In many cases each member of the exclusive rule-set is satisfied by the same number of elements i.e.

$$|W(\hat{R}_i)|=m \ \forall\ i$$

and so

$$S(t)=\ln(mn)=\ln(m)+\ln(n)$$

F. *Orthogonal Rules*

Consider two sets of rules $R_1 \ldots R_N$ and $R'_1 \ldots R'_M$ and their exclusive rule-sets $\hat{R}_1 \ldots \hat{R}_N$ and $\hat{R}'_1 \ldots \hat{R}'_M$. Define the composite rules:

$$\hat{R}^c_{ij}(t)=\hat{R}_i(t) \wedge \hat{R}'_j(t)$$

If $W(\hat{R}_{ij}) \neq \emptyset \ \forall\ i, j$, then the equivalence relation Z_m^C defined by:

$$a Z^C_m b \text{ iff } \hat{R}_{ij}(a)=\hat{R}_{ij}(b) \ \forall\ i, j$$

can be used to create the quotient set: $T^C_m=T/Z^C_m$
So if:

$$\begin{aligned}
&t \in T^C_m=(p, q) \quad p \in T_m \, q \in T_m{}' \\
&\text{then} \\
&S(t)=S(p)+S(q)
\end{aligned}$$

i.e. the entropy can be calculated separately in each domain and then added:

$$\begin{aligned}
S(t)&=\ln\left(\Omega(R_1)\Omega(R_2)\right) \\
&=\ln\Omega(R_1)+\ln\Omega(R_2) \\
t &\in W(R_1) \cap W(R_2)
\end{aligned}$$

$$t \in W(R_1) \cap W(R_2)$$

We will use this result frequently. In particular, given a set of objects in a location (eg dining room, study), it is possible to compute the entropy for each of the objects individually and then derive the entropy of the overall scene by simply adding the results together. It also allows the entropy associated with the geometric arrangement of objects to be separated from that derived from the choice of objects. As this paper is primarily concerned with the latter "selection entropy", we will now consider it in detail.

G. *Selection Entropy*

Using the result above, the selection entropy of the whole scene can be derived by creating a tree structure

DOI 10.5013/IJSSST.a.13.3A.04 ISSN: 1473-804x online, 1473-8031 print

that relates each object to another object, with the "root" object being linked to the overall scene. Each link in the tree structure will have an entropy and the overall entropy will be the sum of all of these.

Consider first the overall scene. Let the set of all available objects be *U*. The set of all possible collections (multisets) of *n* objects is *T* where

$$t \in T = (a_1, a_2, \dots, a_n)\, a_i \in U$$

In the absence of any rules relating the objects to each other or to the scene location, the overall entropy is given by the multinomial coefficient:

$$S(t \in T) = \ln(|T|) = \ln\left(\frac{(|U|+n-1)!}{n!(|U|-1)!}\right)$$

Now we assume that the total set of available objects is extremely large compared to the size of the collection so we can use Stirling's approximation to compute the factorials giving:

$$S(t \in T) \approx n\ln(|U|) - \ln(n!) \text{ assuming } (|U| \gg n)$$

The second term relates to the possible orderings of the members. The process of choosing the lowest entropy representation of the observed state effectively picks a unique ordering, so this second term can be disregarded. The overall selection entropy then becomes simply the sum of the selection entropies for individual objects. It is therefore sufficient to consider the entropy of linkage between individual objects and the location itself or between pairs of objects,

$$S(t \in T) = \ln(|T|) = \sum_{i=1}^{n} S(a_i) \text{ where } S(a_i) = \ln(|U|)\, \forall\, i$$

If we know something about the location then there will be rules that relate the objects (individually) to that location ie an ontology or a taxonomy. For example, in ontological terms, a kitchen "could contain" a cooker, a refrigerator etc., while in a taxonomy, they could be considered members of the class "kitchen furniture".We can form a hierarchical set of rules if U is a totally ordered set, with the "most likely" objects at the start.

Thus if:

$$U = \{u_1, u_2 \dots u_N\}$$

then the associated rule-set is

$$R = \{R_i \dots\}\, i \in \mathbb{N},\, 1 \le i \le N \text{ where } R_i(u_j) \Leftrightarrow i \ge j$$

and $S(t \in T) = \sum_{i=1}^{n} S(a_i)$

now $\exists\, j$ such that $a_i = u_j$ which obeys $R_k\, \forall\, k \ge j$

and hence

$$S(a_i) = \ln(|W(R_j)|) = \ln(j)$$

Further simplification can be achieved by replacing the total order $\le$ with a strict weak order $<$. For each location *L* we identify C_L, a strict weak ordered subset of *U* consisting of those objects that are identified in the ontology as being associated with *L*. Similarly for each object *O* we define C_O, a strict weak ordered subset of *U* consisting of the other objects that are identified in the ontology as being associated with *O* in some way (eg a desk and chair are associated).

Given < we define the equivalence relation ~ as:

$$\neg(a < b \vee b < a) \rightarrow a \sim b \quad (a, b \in C_L)$$

A total ordering $\le$ can now be defined on the quotient set $C_L/\sim$ based on <:

$$C_L/\sim = \{c_1, c_2 \dots c_N\} \;:\; c_i \leqslant c_j \Leftrightarrow i \leqslant j$$

And a rule-set *R* based on $\le$ can now be defined :

$$R_i(a)(a \in c_j) \Leftrightarrow i \geqslant j$$

These can be used to compute individual object/location entropies based on the results in the previous section:

$$\begin{gathered}\forall\, a \in c_j \in C_L/\sim : S(a) = S(a, L) = \ln\left(\sum_{i=1}^{j} |W(R_i)|\right) \\ \forall\, a \notin C_L : S(a) = \ln(|U|)\end{gathered} \qquad (1)$$

A similar process can be followed to compute object/object entropies:

$$\forall\, a \in c_j \in C_O/\sim : S(a) = S(a, O) = \ln\left(\sum_{i=1}^{j} |W(R_i)|\right) \qquad (2)$$

For many objects there will now be more than one way to compute the entropy depending on which related object is used. To resolve the situation we use the principles outlined above and use the object that gives the lowest entropy result.

$$S(a_i) = Min(S(a_i, a_j)(\forall\, j \ne i), S(a_i, L), \ln(|U|))$$

However, because associations between objects are symmetrical it is likely that:

DOI 10.5013/IJSSST.a.13.3A.04

ISSN: 1473-804x online, 1473-8031 print

$$S(a_i)=S(a_i,a_j)$$
$$S(a_j)=S(a_j,a_i)$$

In other words, the two objects each justify their presence in the scene in terms of each other. To overcome this, we compute the entropy of the objects in a directed acyclic graph. Thus the entropy of each new object is calculated only with reference to those objects that have already been calculated.

Given the collection of objects

$$t=(a_1,\ldots,a_n)$$

at any stage of the process, when we have computed the entropy contribution of *m* objects we can divide *t* into two subcollections:

$$(t_m,\bar{t_m}\in P(t)),\ |t_m|=m,|\bar{t_m}|=n-m,$$
$$t_i\subset t_j\Leftrightarrow i<j,t_m\cup\bar{t_m}=t$$
$$t_0=\emptyset,t_n=t$$

(here *P(t)* denotes the power set of *t*).

Now define the entropy of a single object *b* relative to a subcollection t_m as:

$$S(a,t_m)=Min(S(b,a)(a\in t_m),S(b,L),\ln(|U|))$$

We can now define the whole sequence of t's as follows:

$$t_{m+1}\backslash t_m=b_{m+1} \text{ chosen such that}$$
$$S(b_{m+1},t_m)=Min(S(b,t_m),\forall\, b\in\bar{t_m})$$

and the total entropy:

$$S(t)=\sum_{m=1}^{n}S(b_m,t_{m-1})$$

This process effectively constructs an "entropy tree" linking all the objects to each other or to the location. The absence of cycles in the tree guarantees that objects are not bootstrapped off each other; all must link back to the location.

This derivation of the entropy only takes account of the collection of objects and the location. Such a calculation is potentially useful on its own, however it is limited because it disregards the geometric arrangement. If one is trying to measure a pre-existing scene then the entropy is more useful if it takes account of the arrangement of the objects. Similarly, if one is trying to synthesize a scene using entropy then considerations about how the objects will physically fit together are important. As detailed in [22] it is possible to compute a geometric entropy relationship between objects and to build a geometric entropy tree. Once this is done, the selection entropies corresponding to the links in the geometric entropy tree can be computed. This gives a better idea of the true selection entropy of the scene since it now guarantees to link together objects that are in physical proximity.

An even better result can be given by computing a combined entropy tree, using both selection and geometric entropy to guide the process. Since this paper is principally concerned with selection entropy rather than geometric, all the numerical results presented here are pure selection entropy, although we have used geometric entropy to guide tree construction. Thus from the scene shown in Figure 1 we derive a tree structure of which a fragment is shown in Figure 2. Note the way that the links reflect natural expectations about the way that tables, chairs, mats, cutlery and glasses relate together.

H. Practicalities of the Entropy Calculation

We now consider the practicalities of selection entropy calculation. To do this we use two sets of data. The first is an ontology that describes the way in which different types of objects relate together. The second is a library that contains a list of all the objects available to us for scene building together with their detailed characteristics. We showed above that what we need to compute scene entropy is object-object entropy and object-location entropy. Equation (2) above indicates that object-object entropy requires a count of the number of objects that lie "in between" the two that we are trying to link. The ontology provides a framework of links that can be used to navigate from one object to another. The entropy is then computed as the logarithm of the number of objects that are encountered during the navigation.

The ontology includes the following types of linkage for objects:

Goes with
Similar to
Could be contained by
Could contain

The computation of object-location linkage based on equation (1) is similar. The ontology also contains these linkages for object-location:

Could be found in (object-location)
Could have (location-object)

The archetypes in the ontology form a flat structure rather than a hierarchical one. Furthermore, archetypes have to be specific to a location in order to be linked to that location directly. Thus "dining chair" would be linked to dining room and "office chair" would be

DOI 10.5013/IJSSST.a.13.3A.04

ISSN: 1473-804x online, 1473-8031 print

linked to office and study. The archetype “chair”, however, would not be directly linked to any of these locations. It could be indirectly linked because it would be linked to the more specific archetypes via the “similar to” mechanism.

The basic entropy calculation works as follows. In the entropy tree we always link a "new" object to one that is already in the tree. Let object A be the new object and object B the object that is already in the tree. If the two objects have the same archetype (ontology entry) then the entropy is simply the logarithm of the total number of objects available of that archetype. Otherwise we explore one level of linkage from object A and count all the items that exist at that level. Note that in this case we do not count the number of alternatives for object A, only the possible choices for object B. If object B has been found at that level then the entropy is the logarithm of the count. Otherwise we explore a second level of linkage, counting all the objects, and continue with greater and greater levels until the object B is found. The count that is generated in this basic entropy calculation is culled by two further considerations.

The first is what we term the containment level. If we know from the scene that one object contains the other, then we only count linkages that reflect that containment. For a mug contained by a display cupboard, an example of a valid link would be:

Mug "could be contained by" cupboard "similar to" display-cupboard.

In the same situation an invalid one would be:

Mug "could contain" teaspoon "which could be contained by" display cupboard.

In the latter case the *"could contain"* and *"could be contained by"* effectively cancel each other, leaving a peer relationship where a containing relationship was present in the original scene. Similarly if we know that two objects are in a simple peer relationship (next to each other) then links that imply containment are rejected.

The second is controlled by the tagging of objects in the library. Each object in our object library is tagged with a number of two word (type plus descriptor) tags. For example it could be: *colour red* or *size large* or *age modern*. The use of typed tags is advantageous as it avoids the ambiguities that can occur with single word tags. It also allows us to ignore tags that are not relevant and avoids the need for total consistency in tagging. When calculating the entropy between two objects we only include objects in the count if they match the primary objects as well as they match each other. Thus when linking the modern dining fork to the Victorian place-mat there were ten alternative objects found because there was no match between any of the tags. However, linking the Victorian dining fork to the Victorian place-mat produced only two, the place-mat itself and the Victorian dining table.

IV. Results and Discussion

In order to assess whether our entropy measure has predictable behaviour, and actually captures a sense of the degree of aesthetic consistency of the scene, we conducted a set of experiments. A virtual scene was created which represents a modern dining room as shown in Figure 1. Then successive changes were made to this scene by replacement of assets with others that 'matched' less well, and the entropy tree recalculated at each stage. The important factor here is the tagging of the objects using the mechanism described in the previous section.

Figure 1: Basic dining room layout

The entropy associated with the portion of the tree containing the dining table and chairs, taking all the assets into account, is shown in row 1 of Table 1 (and as illustrated in Figure 2). We then replaced one chair at a time by a Victorian one, until we were left with the scene shown in Figure 3 where all four chairs were Victorian but the table was modern.

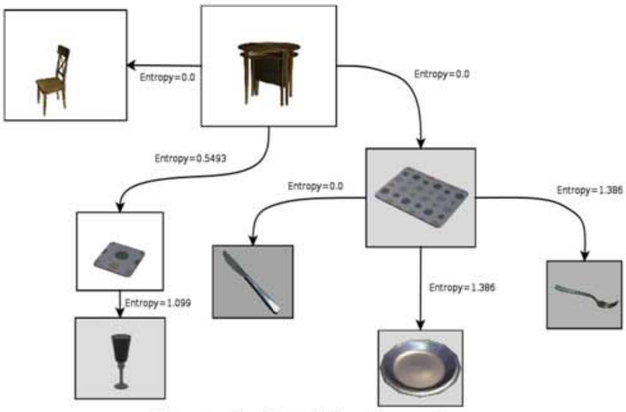

Figure 2: Partial entropy tree

At each of these changes, the entropy was seen to increase, as the chairs were less well suited to the rest of the assets (rows 2,3,4,5 in Table 1). Then the table

DOI 10.5013/IJSSST.a.13.3A.04

ISSN: 1473-804x online, 1473-8031 print

itself was replaced by a Victorian one, followed by the plates and the mats, which were Victorian tartan. These changes are shown in Figure 4, and rows 6, 7 in Table 1.

An attempt to change the modern cutlery for a 'rustic' or Jacobean set made no difference to the overall entropy. This is also understandable, as the rustic set is no more well matched with the Victorian aesthetic than the modern set. This change is shown in row 8 and Figure 5. The next stage was to introduce Victorian cutlery as shown in row 9. This time the entropy falls significantly because there are now no inconsistent objects on the branches of the entropy tree associated with the place-mats.

Figure 3: Chairs changed to Victorian style

Finally, substituting Victorian glasses as shown in Figure 6 (row 10) lowered the entropy even further, since this made every part of the scene consistent. It is noteworthy that the entropy of the pure Victorian scene is lower than that of the modern scene. This occurs because "modern" is a very general category that has many entries in the object library whereas "Victorian" is specialised and so generally does not have as many entries.

Figure 4: Table, placemats, plates changed

Though the change to a Victorian table increased the entropy yet again, it then fell back when the plates and mats were changed to be Victorian. This is understandable in that the whole table is now more consistent aesthetically as it is a Victorian table with Victorian contents.

Figure 5: Rustic cutlery

Figure 6: Entirely Victorian settings

A further test on measurement of consistency was undertaken by experiments on the contents of the wall cupboard, where collections of plates, glasses and mugs are stored. Figure 7 shows the baseline state of this cupboard, although in later images, the cupboard itself is sometimes made invisible for clarity. Table 2 shows the results of the entropy calculations here – although only the entropy contribution from the cupboard contents is shown, as the rest of the scene remained unchanged.. At each of these stages, the previous

DOI 10.5013/IJSSST.a.13.3A.04

ISSN: 1473-804x online, 1473-8031 print

changes were left in place and the new change is additive – for example, the china plates in row 2 were still in place when the willow pattern ones replaced the other modern plates.

As we progress through the sequence of changes, we see the increasing entropy as each change makes the cupboard contents less consistent. In particular, the final change (as shown in Figure 8) where the glasses were exchanged, has a large effect on the entropy. This is not surprising, because we are essentially making two changes at once, one to each group of glasses. It clearly illustrates how consistency and similarity generate much lower entropy values than where a group is disparate.

Figure 7: Baseline state of cupboard

Figure 8: Cupboard after sequence

As a final test of the variation of entropy, we replaced one of the plates in the cupboard by a model Viking helmet – an artefact with absolutely no relevance to dining rooms, Victorian or modern. (See Figure 9.) This caused an increase in entropy for the cupboard contents to 21.13. This is clearly a far larger increment than any other single change, as compared with row 4 or 5, reflective of the incongruity of the object.

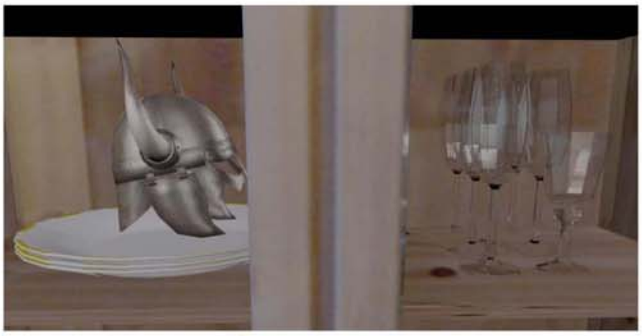

Figure 9: Incongruous viking helmet

Row	Scene	Entropy
1	Baseline Modern	17.787389
2	1 Victorian Chair	19.173683
3	2 Victorian Chair	20.559978
4	3 Victorian Chair	21.946272
5	4 Victorian Chair	23.33256
6	4 Victorian Chair + Table	26.981753
7	4 Victorian Chair + Table + plates, mats	23.008745
8	4 Victorian Chair + Table + plates, mats + rustic cutlery	23.008745
9	4 Victorian Chair + Table + plates, mats + Victorian cutlery	13.798405
10	4 Victorian Chair + Table + plates, mats + Victorian cutlery + Victorian glasses	11.025816

TABLE 1 ENTROPY CHANGE FOR OVERALL ROOM ADJUSTMENTS

V. CONCLUSIONS AND FURTHER WORK

We have developed selection entropy, a mathematical formulation of everyday untidiness, and demonstrated that it can be used to quantify the aesthetic consistency of a scene. In this section we will first discuss an alternative way of implementing the

DOI 10.5013/IJSSST.a.13.3A.04 ISSN: 1473-804x online, 1473-8031 print

scheme and then consider ways in which the work could be used and extended in future. The fact that the entropy is derived from the counting of objects in the library does make it dependent on the library contents. There is also the potential for errors in the ontology or library structure to distort the answers. This could be particularly problematic if the library is small. However our library contains only 81 objects at the time of writing and yet we observe from the results presented above that the numbers are plausible, with consistent changes as the scene is modified. We also note that the relative entropies do not change in any significant way when we add objects to the library. We believe that this shows that our definitions give rise to an adequately robust result, however it is interesting to consider the possibility of using a larger pre-existing library such as Google Warehouse (as was utilised by [25])

Row	Scene Change	Entropy
1	Baseline – modern plates, mugs, glasses, well ordered	10.805274
2	Victorian china large plates	11.498421
3	Willow pattern plates	12.191568
4	Red straight sided mug	12.884715
5	Red mug changed to be tapered (ie different style)	13.801006
6	Wine glass and flute swapped over	15.187300

TABLE 2 ENTROPY CHANGE FOR ADJUSTMENTS IN THE CUPBOARD

Once the library becomes really large the exact details of its contents should be irrelevant, since the entropy is defined logarithmically. Unfortunately Google Warehouse is set up with single word (untyped) tags and so cannot be used directly with our scheme. However it is possible to derive some estimates that show how a large library might affect the results.

As a simple experiment, consider the linkage between our Victorian dining table and the Victorian plate, compared to the linkage between the Victorian dining table and the modern plate. This is a containing relationship (the table contains the plate) so we can assume that the table is in the entropy tree before the plate. This means we consider only the number of alternatives to the table. Our library gives 8 alternatives for the Victorian table linked to the modern plate, giving an entropy of 2.1. Of these 8, 4 are dining tables, 2 are table mats, one is a display cupboard and one is a cupboard Google warehouse gives 3792 hits for dining tables, 1165 for cupboards and 201 for table mats. The total is thus 5158 and the entropy is 8.5. When we restrict to Victorian objects our library gives just 2 and an entropy of 0.7. Google warehouse now gives 37 objects and an entropy of 3.6. The entropy ratios are quite similar, suggesting that the qualitative nature of the results is not very sensitive to library size.

In this paper we have only considered selection entropy in any detail. There is considerable scope to use geometric entropy as an aesthetic measure and we intend to pursue this further in the future. There are many ways in which the entropy measurements could be exploited. For example they can be used to assist in the automatic generation of virtual environments with predictable aesthetic characteristics. This would be particularly relevant to a game developer who had built up a substantial library of 3D assets and wished to re-use them in new projects. It would also be interesting to use this measure to drive the behaviour of non-player characters in games. In particular it should enable such a character to react to disturbances that change the aesthetic consistency of the scene.

REFERENCES

[1] RJ Cant and CS Langensiepen, "Methods for Automated Object Placement in Virtual Scenes," in *UKSim 2009: 11th International Conference on Computer Modelling and Simulation*, 2009, 431–436.

[2] M.R. Beck, D.T. Levin, and B. Angelone, "Change blindness blindness: Beliefs about the roles of intention and scene complexity in change detection," *Consciousness and cognition* 16, no. 1 (2007): 31–51.

[3] M. Feixas et al., "An information theory framework for the analysis of scene complexity," in *Computer Graphics Forum*, vol. 18, 1999, 95–106.

[4] T. Hubing et al., "An algorithm for automated printed circuit board layout and routing evaluation," in *Electromagnetic Compatibility, 1993. Symposium Record., 1993 IEEE International Symposium on*, 1993, 318–321.

[5] B. Coyne and R. Sproat, "WordsEye: an automatic text-to-scene conversion system," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001, 487–496.

[6] L.M. Seversky and L. Yin, "Real-time automatic 3d scene generation from natural language voice and text descriptions," in *Proceedings of the 14th annual ACM international conference on Multimedia*, 2006, 61–64.

[7] K.A.H. Kjolaas, "Automatic furniture population of large architectural models" (Citeseer, 2000).

[8] K. Xu, J. Stewart, and E. Fiume, "Constraint-based automatic placement for scene composition," in *Graphics Interface*, 2002, 25–34.

[9] Yoshiaki Akazawa, Yoshihiro Okada, and Koichi Niijima, "Automatic 3D scene generation based on contact constraints," Proc. Conf. on Computer Graphics and Artificial Intelligence. 2005.: 51-62.

[10] S. Sanchez et al., "Constraint-based 3d-object layout using a genetic algorithm," *Proceedings of the 6th International Conference on Computer Graphics and Artificial Intelligence 3IA'03* (2003).

[11] N. Vassilas et al., "MultiCAD-GA: a system for the design of 3D forms based on genetic algorithms and human evaluation," *Methods and Applications of Artificial Intelligence* (2002): 743–744.

[12] R. Cant, C. Langensiepen, and M. Haskard, "Another fine mess: how to generate disorder: to order," in *ACM SIGGRAPH 2006 Sketches*, 2006, 121–es.

[13] J. Taylor and I. Parberry, *Computerized Clutter: How to Make a Virtual Room Look Lived-in* (Technical Report LARC–2010–01, Laboratory for Recreational Computing, Dept. of Computer Science & Engineering, Univ. of North Texas, 2010).

[14] L.F. Yu et al., "Make it home: automatic optimization of furniture arrangement," *ACM Trans. Graph* 30, no. 86 (2011): 1–86.

[15] C.E. Catalano et al., "A framework for the automatic annotation of car aesthetics," *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 21, no. 1 (2007): 73–90.

[16] H.C. Tsai and J.R. Chou, "Automatic design support and image evaluation of two-coloured products using colour association and colour harmony scales and genetic algorithm," *Computer-Aided Design* 39, no. 9 (2007): 818–828.

[17] J.W. Satzinger and L. Olfman, "User interface consistency across end-user applications: the effects on mental models," *Journal of Management Information Systems* 14, no. 4 (1998): 167–193.

[18] J. Mendel, R. Pak, and J.E. Drum, "Designing for Consistency," in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 55, 2011, 2000–2004.

[19] D. Plemenos, M. Sbert, and M. Feixas, "On viewpoint complexity of 3D scenes," in *International Conference GraphiCon*, 2004.

[20] J. Rigau, M. Feixas, and M. Sbert, "An information-theoretic framework for image complexity," in *Proceedings of the Eurographics Workshop on Computational Aesthetics in Graphics, Visualization and Imaging*, 2005.

[21] A.E. Stamps, "Entropy, visual diversity, and preference," *The Journal of general psychology* 129, no. 3 (2002): 300–320.

[22] R.J. Cant and C.S. Langensiepen, "Entropy Measurement within Graphical Scenes," in *12th International Conference on Computer Modelling and Simulation*, 2010, 474–481.

[23] E. T. Jaynes, "Information Theory and Statistical Mechanics," *Physical Review* 106, no. 4 (May 15, 1957): 620-630.

[24] M. Fisher, M. Savva, and P. Hanrahan, "Characterizing Structural Relationships in Scenes Using Graph Kernels," *SIGGRAPH 2011, Vancouver* (n.d.).