

# Rendering and Affect

D.J. Duke<sup>1</sup> P.J. Barnard<sup>2</sup> N. Halper<sup>3</sup> M.Mellin<sup>3</sup>

<sup>1</sup>Department of Computer Science, University of Bath, Bath, UK.

<sup>2</sup>MRC Cognition and Brain Sciences Unit, Cambridge, UK.

<sup>3</sup>Department of Simulation and Graphics, Otto-von-Guericke University, Magdeburg, Germany.

---

## Abstract

*Previous studies at the intersection between rendering and psychology have concentrated on issues such as realism and acuity. Although such results have been useful in informing development of realistic rendering techniques, studies have shown that the interpretation of images is influenced by factors that have little to do with realism. In this paper, we summarize a series of experiments, the most recent of which are reported in a separate paper, that investigate affective (emotive) qualities of images. These demonstrate significant effects that can be utilized within interactive graphics, particularly via non-photorealistic rendering (NPR). We explain how the interpretation of these results requires a high-level model of cognitive information processing, and use such a model to account for recent empirical results on rendering and judgement.*

Categories and Subject Descriptors (according to ACM CCS): I.3.m [Computer Graphics]: Miscellaneous

---

## 1. Introduction

Psychophysical experiments into colour perception and visual acuity are ongoing in computer graphics. They provide insight into human perception that informs development of new methods for realistic rendering, and an empirical methodology for assessing progress in areas such as virtual environments. Perceptual studies have also underpinned some developments in non-realistic rendering (NPR). This includes the use of NPR to convey shape and detail in technical and artistic illustration<sup>10, 21</sup>, and work in visualization, for example techniques to differentiate regions<sup>14</sup>. NPR research also draws on wider sources of inspiration. Analyses of the perceptual merits of artistic techniques have underpinned new rendering styles for multi-dimensional data<sup>16</sup>, and are currently an area of active research<sup>9</sup>.

Communication, however, is about more than just perception. The use of NPR raises issues about interactions within higher-levels of cognitive process. For example, sketch rendered images of architectural designs have been shown to result in qualitatively different dialogue between architects and clients, compared with that elicited through more realistic images<sup>19</sup>. This effect is not just a matter of perception, but involves also levels of cognition connected with judgement and articulation. In psychological terms, the image has an *affective* content, separate from its structural detail. Sec-

tion 2 of this paper summarizes a series of empirical results that suggest this is not an isolated effect.

Individual results do not in themselves provide insight into why such effects occur, or how they might be utilized within graphics applications. Our first contribution, in Section 3, is to provide an explanation of these results, using the concept of *invariants* to establish a link between properties of representations and the cognitive resources involved in their interpretation. The significance of this link for computer graphics is highlighted in Section 4, in which we summarize key results from a series of recent experiments assessing the effect of NPR style on judgements. These results, details of which have been reported in a separate paper<sup>12</sup>, provide empirical evidence that affective content can be used to guide choice and decision-making in interactive graphics. Using the model developed in Section 3, we are able in Section 5 to explain the nature and variability of the invariants accessible through NPR. Through this, we gain insight into how to scale these results up to graphics applications.

Running through the paper is an argument that to understand how rendered images are used requires models that can accommodate not only perception, but also the dynamics of how meaning is processed in a task environment. The conclusion sets out a case for how such knowledge can contribute to open challenges in computer graphics.


## 2. Connotation in Graphics

Vision is one of five sensory modalities that we have evolved to (firstly) survive in and (secondly) make sense of our environment. Computer graphics provides the capability of synthesizing virtual environments, and in describing user interaction with and within such environments it is usual to think in rational, reflective terms: how well can features of the environment be discerned, how realistic do they appear, and how effectively does the environment support specific tasks.

Our understanding of the world is not based just on surface phenomena perceived directly. It also involves deeper levels of representation that capture and reflect relationships and regularities at higher levels of abstraction. One accessible example of these is in art forms, the use of onomatopoeia in poetry, or the ability of cartoonists to convey meaning and emotion through figures that bear little resemblance to their inspiration. An understanding of such interactions involves the concept of *invariant*, developed by the work of J.J. Gibson (see<sup>5</sup> for an overview). An invariant is a property common to (shared by) a range of entities or behaviours. Three aspects of invariants are significant here:

1. they can reflect stable properties of entities, and properties of behaviours over time;
2. they can apply across a range of abstraction levels;
3. they can be restricted in some way by the context (culture, task, or environment);

Even within a single modality such as vision, human response is affected by qualitatively different kinds of invariants extracted from the incoming data streams. Only some of these invariants correlate with what might be thought of as reasoned responses. For example, Figure 1 contains a sketch of a radar display, originally due to Provins et al<sup>18</sup>. When asked to glance at the display and immediately identify which group (circles or triangles) of aircraft are the 'hostile' forces, nearly all subjects identify the triangular symbols<sup>18, 12</sup>. This is not a reasoned response, there is no overt information in the scene that 'triangles' are hostile. It is, rather, an example of an affective (emotional) response.




**Figure 1:** Radar Display (after Provins et. al., 1957)

It has been argued<sup>2</sup> that such responses utilize latent knowledge, in this case an invariant of sharp shapes invokes a connotation of threat. In evolutionary terms, rapid reaction to sharp shapes (e.g. in the form of fangs) may have conveyed a survival advantage, and become encoded into the

neural architecture. The link between cognition and invariants forms a substantial part of this paper.

Whereas Gibson viewed invariants as relationships 'in the world', our notion is more general, encompassing relationships and regularities within and across a range of mental representations. The pattern of properties captured by such invariants can be multimodal, drawing together perceptual, cognitive and affective attributes. An example of this is demonstrated by the following experiment, due originally to Davis<sup>7</sup>. When asked to associate the shapes in Figure 2 with the sound of two words, 'Ulloomo' and 'Takete', most people associate 'Takete' with the shape on the left of the figure.



**Figure 2:** Two connotative shapes (after Davis, 1961)

Davis' result can be explained by positing that certain visual characteristics (i.e. invariants) of the shape match invariants of the sound. When asked to explain the correspondence, subjects find it difficult to articulate *why* these qualities match; this is a characteristic of latent knowledge, it is automatic, and not necessarily subject to conscious introspection. Further, response to latent invariants is extremely rapid, as shown by Guthrie and Weiner. In their experiment<sup>11</sup> subjects were asked to make character judgements based on an outline image of a seated figure. The subjects were shown an extremely brief, subliminal image of a complete outline, followed, for a longer duration, by an outline with missing segments. The study showed that the figure was judged as more hostile and aggressive when the subliminal image contained angular features, as opposed to curved features.

How invariants come into play depends on the resources available to the human information processor. An example can be seen in studies of behavioural attribution. If a subject unfamiliar with brownian motion is shown an instance of this phenomenon, they will often interpret it as, for example, a 'lot of bugs'; the random motion of the particle ensemble is an invariant that maps to a familiar behaviour, in this case a cloud of insects. However, for someone with experience of brownian motion, the invariant properties of the motion can also be mapped to that internal knowledge. Similarly, the Heider-Simmel experiment<sup>13</sup>, in which viewers provide an account of interaction between abstract shapes, elicits behavioural invariants (e.g. 'one thing chasing/controlling another') that are mediated by cultural expectations.

The experiments described above show invariant properties in the display being mapped to judgements, but involve only small-scale displays and low-level tasks. More recent research has shown that invariants linked to rendering style

also result in qualitatively different responses in larger-scale tasks, involving higher levels of judgement. A study on the role of sketch rendering in architectural CAD<sup>19</sup> was mentioned in the introduction. It has also been noted that sketch characteristics can be used to draw attention to/from particular parts of a scene<sup>20</sup>. Taken together, these results suggest that there are qualities of images, independent of overt structure, that influence interpretation. These qualities involve invariants that are not bound to a particular modality, but rather come into play at a high level of cognitive process, at which affect and ideation interact. In the next section we describe how such invariants impact cognition; later in the paper we visit results that show how such invariants arise from, and can be controlled through, NPR styles.

### 3. Rendering Style and Affect

Early psychological experiments, described in Section 2, show that graphic displays carry more than just structural information (i.e. 'that object is there'). Images also evoke a latent, affective response. Recent experiments, summarized in Section 4, have shown that such effects (a) are not limited to simple images and tasks, and (b) can in principle be utilized to guide high-level judgements. By themselves, these results are isolated data. What we need is a basis for 'interpolating' from these experiments, understanding what is causing these effects, how they might be deployed in graphics applications, and what freedom is available to designers in manipulating their use. In other words, experiments are only useful in so far as they lead to a generalizable understanding of the underlying phenomena.

Use of latent knowledge, including affect, has previously been explored in the context of interactive systems<sup>2</sup>. This section returns to the theoretical framework used in that paper, and demonstrates how it accounts for the newer results summarized in the previous section. The concept of invariant, introduced in Section 2 is useful in this context, in that serves as a 'common currency', bridging between computing and psychological models. Invariants can be understood both as properties of graphics, and as phenomena utilized in cognitive processing. By understanding the operation and principles of this processing we can better appreciate how such invariants can be utilized in graphics applications.

#### 3.1. Interacting Cognitive Subsystems

Cognitive science attempts to understand the capabilities and operation of the mind in terms of its ability to process information. One form that this understanding takes are models, at various levels of detail, that posit mechanisms by which observed phenomena can be explained. These observations include narrowly focused experiments, clinical results, and more recently detailed neurophysiological studies of the mind operating *in vivo*. Because of the large number of observations available, differences in methodology, and different approaches to modelling, there is no single 'model' of

cognition. Rather, there are a collection of models, accounting for a different range of phenomena. Within this diversity, however, a number of trends have emerged. One is that cognition is a distributed process; damage to (say) one region of the brain does not completely destroy one kind of faculty, and may leave other faculties untouched. Similarly, different areas of the brain process particular kinds of information.

Interacting Cognitive Subsystems<sup>2,3</sup> (ICS) is one of a number of models that provide a comprehensive account of human information processing (EPIC<sup>15</sup> is another). It is one of the more developed models that can account for and explain a broad range of phenomena, including task performance and affective properties<sup>4,1</sup>. This breadth of application, and in particular its treatment of affect, motivates its use here. In ICS, cognition is viewed as an activity distributed between nine subsystems, shown in Figure 3.




Figure 3: ICS Architecture

Briefly, these nine systems are:

Sensory subsystems

- VIS** visual: hue, contour etc. from the eyes
- AC** acoustic: pitch, rhythm etc. from the ears
- BS** body-state: proprioceptive feedback

Structural subsystems

- OBJ** object: mental imagery, shapes, etc.
- MPL** morphonlexical: words, lexical forms

Meaning subsystems  
**PROP** propositional: semantic relationships  
**IMPLIC** implicational: holistic meaning

Effector subsystems  
**ART** articulatory: subvocal rehearsal and speech  
**LIM** limb: motion of limbs, eyes, etc

A feature of the model is its distinction between two levels of meaning, propositional and implicational, a feature demonstrated recently through experiments using affective properties derived from auditory perception<sup>22</sup>. The propositional level is concerned with referential knowledge, that is ‘facts about the world’, including properties of and relationships between entities. In the case of a scene from a computer game, this is facts such as ‘another player is standing next to a door’, or that ‘the door is opening’. In contrast, the implicational level is concerned with schematic models, capturing among other things ideational and affective content. In the game scenario, representations at this level would include plans that associate passing through the door with achieving some goal, and also awareness of any threat posed by the second player in the scene. Implicational level therefore involves a higher set of invariances. This is the level of knowledge representation that we need to be concerned with, as it is here that invariants related to affect and the analysis of consequences come together.

All subsystems have the structure shown in Figure 4. Incoming data streams arrive at an input array, on which a number of processes operate in parallel. A single ‘copy’ process in each subsystem takes the incoming data and lays down a record of this in the image record, an episodic memory. Other processes are working on the input array, trying to transform the available representation into other codes. Here and in the remainder of the paper, the notation ‘A-B’ is used to denote the process in subsystem ‘A’ that transforms its input into representations for subsystem ‘B’; thus ‘obj-prop’ is the process in the object subsystem that produces propositional-level representations. For transformations such as obj-prop that produce a higher-level representation, this involves the abstraction of particular invariants from the input. In the case of ‘obj-prop’, the processes is attempting to output a semantic interpretation of spatial (structural) invariants over its input. Transformations producing a more ‘refined’ code, for example obj-lim, involve implementing abstract invariants in terms of lower-level capabilities, for example instantiating the muscular control needed to position a limb at some point in space.

In general a subsystem is receiving data streams from a variety of sources. Only some combinations of source are coherent, that is, the invariants within the incoming representations are compatible. Transformations can only operate on coherent products, and if none are available the output may become unstable and effectively unusable by other processes<sup>8</sup>. If the incoming data is insufficient to form a stable response, a transformation can disengage from the in-




Figure 4: Subsystem

put array and instead take input from the image record. In this mode of processing, the transformation can operate either on the proximal end of the record, the data most recently copied from the input array, or on data revived from deeper within the experiential record. Either way, entering this ‘buffered’ mode has a significant effect on the transformation. Buffering brings in additional attributes that offer greater affordances for subsequent processing. When a process is buffered, an experimental subject will be focally aware of the extended attribute space associated with that subsystem, and the effect of any invariants in the input streams may be influenced by patterns of invariants formed over long time scales or involving different combinations of attribute. In contrast, in non-buffered (direct) mode, the effect of invariants is exerted in the background, but will dominate processing. Significantly, a constraint on the architecture is that, within a configuration, only one process at a time can be operating in buffered mode.

3.2. Affective and Structural Graphics

Although all ICS processes are attempting to operate on incoming data, the products of only some of these processes are relevant to the current task, and/or are linked to focal awareness. The collection of processes that are deployed on a given task is called a *configuration*. Figure 5 shows a configuration corresponding to a task involving interpretation of a graphical display, using a thick black line to show the flow of representations between the processes involved:

In this configuration there are processes extracting objects from the retinal image (vis-obj), identifying structural relationships (obj-prop), and utilizing those relationships in interpreting the situation relative to goals and plans (prop-implic). Essentially all of these processes are using invariants derived from the structure of the graphic. If the subject’s attention is focused on the meaning of the task, buffering will be located at the propositional level.

In parallel with this structure activity, the vis-implic process is also operating, building a high-level schematic model directly from the raw visual input. This gives rise to a second component of the configuration, shown in Figure 6 by a grey line, connecting vis to implic.




Figure 5: Visual Processing Configuration




Figure 6: Affect in the Configuration

In early organisms with a simpler cognitive mechanism (essentially just sensory and implicational systems), such a model might indicate a basic ‘fright’ response on the detection of hostile features, or an aggressive response in the presence of food! In our evolved form, this stream is only one voice in a fugue. What it does do, however, is to provide an affect-related invariant to the implicational system. In the case of the radar display (see Figure 1), rather than responding just on the basis of the structural content (‘circles’ and ‘triangles’), processes in the implicational system, operating in direct mode, are accessing an input array in which the invariants of a ‘threat’ schema (from vis) are correlated with spatial and task-relevant properties of the display (from prop). When judgement is elicited, the blending of the triangle/threat invariants leads to propositional output, (implic-prop, carried on the second grey flow in Figure 6), that identifies the triangles with the threat. Many processing tasks require reciprocal exchange of information between the propositional and implicational levels, and the existence of such data flows, particularly when coupled with buffering moving between the prop and implic levels, is referred to as the ‘central engine’ of cognition.

The results of Davis’ connotative shape experiment, also described in Section 2, are derived by a similar mechanism, this time though involving blending between products derived from both visual and acoustic sources. The flows carrying the affect-related invariants are shown by black lines in Figure 7. The possibility of affective invariants derived from body-state information (bs-implic) is also shown. For example, if a subject was asked to touch an object and select a word (or visual shape) that corresponds

to the object, it would be expected that a sharp object such as a thorn would map to triangles/‘takete’, while a soft object, e.g. a tennis ball, would map to rounded shapes/‘ulloomo’. Although we cannot be exact, we note in this example that abrupt changes in form or constituents are among the features shared by one class of representation, but are lacking in the second. In interpreting the meaning of images, the central engine also has available structural information carried on the data streams vis-obj and obj-prop, shown by grey lines in Figure 7. These too generate invariants at the implicational level, via prop-implic. Thus invariants formed at implic involve correlation over general patterns within the sensory and proprioceptive environment, coupled with those derived from referential meaning. Consequently the implicational level is critical for utilizing and matching perceptual invariants, and that multimodal patterns of invariants form an essential part of the process of interpreting rendered images.

#### 4. Evidence of the Nature and Variability of Invariants in NPR

The results presented in Section 2 form the baseline for a recent series of experiments designed to explore how rendering style affects users’ decisions and judgements within interactive graphics. The experiments also demonstrate that the effects of the type described above are not accidental, and, in particular, that they can be accessed and controlled through NPR techniques, for example stylized silhouettes.

Empirical results from nine experiments are reported and interpreted in a separate paper<sup>12</sup>. In each experiment, the subject has, under time pressure, to make a judgement based




Figure 7: *Blending of affect*

on a displayed image. Subjects' response and response times are recorded, along with demographic data (including gender and first language, as the subject population covered multiple nationalities). These confirm that rendering style can have a significant effect on interpretation, and that this effect can be controlled. Some of the theory needed to explain these effects has been set out in the previous section, but there are a number of results that point to more complicated interactions between levels of cognitive process, with important consequences for the use of rendering style. The experiments are summarized below; unless otherwise noted, all results are statistically significant.

#### 4.1. Assessment of Danger and Safety

The example of aircraft representation<sup>18</sup> was the starting point for three hypotheses exploring the jaggedness of contour as an invariant to denote danger.

- **Experiment (a):** A radar mock-up, similar to that shown in Figure 1, was displayed, and subjects were asked to determine on which side the hostile forces were located. (This experiment was intended to re-establish the results of<sup>18</sup>, and provide a test for our methodology.)  
**Result:** 84% of subjects selected the side marked by triangles as being hostile.
- **Experiment (b):** Generalizing (a), it was hypothesized

that sharp contours carried an invariant connoting a general sense of threat. Subjects were presented with three representations of a door, one in a simple outline, one outlined with jagged edges, and one with a wavy outline. They were told that there was danger behind one of the doors, and were asked to identify which door this was.

**Result:** 67% of subjects selected the jagged-contour door.

- **Experiment (c):** The hypothesis here was that the sharp contour invariant could be used to connote danger, even for objects for which there might normally be a latent cultural invariant mapping to safety. Subjects were shown one of three images. Each image consists of a house on the left, and a clump of trees on the right. They were asked to identify the safest location. The rendering of the house and tree were varied: in the first image the house and trees are both rendered with no effects; in the remaining two, one of the house and trees are drawn with jagged edges.  
**Result:** Given a simple house and tree, 87% of subjects selected the house. This corresponds to a cultural invariant of houses as 'safe'. When the tree was rendered with jagged edges, 88% selected the house for safety; however, when the house was rendered with jagged edges, 81% selected the trees for safety. This demonstrated that sharpness can be used to connote threat, and this connotation can override a cultural safety invariant.

#### 4.2. Assessment of Strength and Weakness

Tests were conducted to investigate whether silhouette thickness is an invariant that can be mapped to an attribution of 'strength' or 'weakness' of graphical objects. Although these invariants are not in themselves forms of affect, the interpretation required for a positive outcome requires interaction with affect-related invariants at the implicational level.

- **Experiment (a):** It was hypothesized that a dotted silhouette would carry a connotation of weakness compared to a thick, solid silhouette. Subjects were shown a radar display with two groups of friendly and hostile forces, with one group of each rendered with a solid silhouette, and one group with a dotted silhouette. They were asked to determine which of the friendly groups was the weaker.  
**Results:** 88% of subjects selected the dotted circles as representing the weakest friendly force.
- **Experiment (b):** Subjects were shown three images of men, identical apart from silhouette strength, and were asked to judge which was the strongest. The silhouettes used were dotted, normal, and thick. The hypothesis was that the image drawn with the thick silhouette would be identified as the strongest.  
**Results:** Subjects selected the normal and strong silhouette figures as stronger, at levels above what could be attributed to chance. Reasons for the split between the thick and normal contours are suggested in<sup>12</sup>. This experiment shows how judgements can be constrained by factors other than display invariants; further cognitively-grounded issues are raised in the next section.

- **Experiment (c):** The images from test (b) were re-used, with subjects now asked to identify the *weakest* man. The hypothesis was that the man rendered with a dotted silhouette would elicit the most responses.

**Results:** 67% of subjects selected the dotted silhouette as being the weakest figure.

### 4.3. Goal-directed Interaction

The third set of tests investigated goal-related judgements concerning selection between alternative paths and objects. In this context, ‘approaching’ an entity represents a positive primitive of gain, avoiding an entity or path reflects a sense of indifference. As inquisitiveness is a key skill for young animals, there is a rationale for thinking that invariants of detail or content will operate latently.

- **Experiment (a):** Subjects were shown an image of a scene with two divergent paths. One path was rendered in a high level of detail (LOD), the second at a low level. Subjects were asked to choose which path to explore first; it was hypothesized that the high LOD path would be chosen. To control for culturally-based preferences<sup>6</sup>, subjects were divided into two groups, with one group shown the mirror image of the scene.

**Results:** The population of first-language speakers showed a strong inclination (74%) to choose the high LOD path. Second-language speakers show a much higher variation and their results do not exhibit this bias. This result is important, in that differences between the cognitive processes deployed by these groups can provide guidance on how and where NPR effects can best be utilized. Further details of the empirical test results can be found in the interpretation section of<sup>12</sup>.

- **Experiment (b):** Subjects were shown a scene with two paths in the foreground, and a path in the distance leading to a house. They were asked to judge which of the two paths, would lead to the house. Again, the paths differed in the level of detail at which they were rendered, and two groups of subjects were used, with the image being mirrored between groups. The hypothesis was that the high LOD path would be preferred.

**Results:** As in the previous case, first language speakers showed a demonstrable preference for the high LOD path.

- **Experiment (c):** The third test was intended to explore the relationship between rendering style and attention. Subjects were shown a scene of a shelf with twenty objects. Ten objects were rendered in an NPR style, ten in a realistic style. Subjects were asked to choose three objects, with the hypothesis that subjects would tend to select the NPR-styled objects. As a control, for one group the background scene was rendered in an NPR style, for the second the background was drawn realistically.

**Results:** A significant majority of subjects selected two or more NPR objects.

The rationale for hypothesizing that subjects would prefer the high LOD paths requires further elaboration. At one level, it might appear obvious that more detail is ‘more interesting’, but this just pushes back the question to one of *why* detail should be interesting, that is, what invariants are present within the detailed pathways. Our explanation is concerned with rate-of-change. Data streams derived from ‘regions’ of high variation (detail) will exhibit a high rate of change, as attention passes through the structure. High rates of change can be linked to excitement or arousal; dually, low rates of change with indifference. Inquisitiveness, mentioned earlier, is triggered by change.

These results create a challenge for explanation, one that a straightforward analysis of low-level perceptual mechanisms cannot provide. They represent a significant bridge from the psychological phenomena described in Section 2 to open issues of how and why to utilize non-realistic rendering in computer graphics.

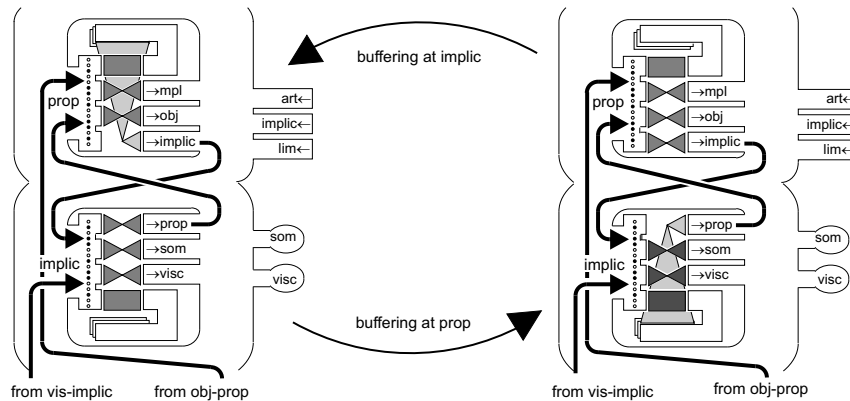
## 5. Invariants, Meaning, and Cognitive Mode

Variations in the experimental results described above, in particular the effect of whether subjects were first or second language speakers, raise the question of how specific invariants are used to meet subjects’ psychological requirements, and how different types of invariant are deployed in performing a given task. Our analysis in Section 3 of early results explained how latent, implicational knowledge is required for successful completion. However, within the cognitive system, content-related data streams, derived from the structural systems (obj and mpl) will be co-present with the implicational streams. The presence of these streams provides an explanation for how and why subjects responses differ from the expected. In particular, we can account for the effect of second-language speakers on the recent path-selection experiments, described in Section 4.2.

### 5.1. Mode Oscillation

Recall from Section 3.1 that a process will try to enter a buffered mode of processing when it is unable to produce usable output directly from the available input streams. We have previously supposed that in the context of the test, buffering will be located at the propositional level, with direct processing of invariants at the implicational level (Figure 8, left). However, in two situations this configuration cannot be sustained:

- If subjects have difficulty interpreting the task requirements, central engine activity will be required to correlate perceptual and task invariants. This requires schematic judgements, involving implicational processing.
- The vis-obj and obj-prop processes are making a semantic representation of the scene available at the propositional level. However, this stream carries no overt information suitable for completion of the task. Consequently, if the



**Figure 8:** *Oscillation in the Central Engine*

prop-implic process engages with this stream, the implicational process will be unable to reconcile this (unstable) stream with the ongoing, high-level task description.

In either case, users will be trying to compare (inadequate) propositional models against implicational judgements. As a result, the implic-prop process will also seek to operate in buffered mode, with the buffer moving to the implicational level (Figure 8, right). However, implic-prop is then operating on invariants in a much richer model than the incoming streams accessed in direct mode, and further, automatic processing, including key latent responses, is suspended. And as implic-prop is then unable itself to generate a task-relevant output, unstable input arrives at the propositional level. Consequently, as shown in Figure 8 the system configuration oscillates over time, with the 'central engine' trying to resolve a task through overt reasoning, but failing to access the relevant latent knowledge. Under pressure of time to complete the task, subjects may then fall back on an arbitrary choice.

Experimental support for this scenario can be found in the empirical results<sup>12</sup>. In the study of radar targets, subjects selecting the triangles as the hostile forces tended to have significantly shorter response time than subjects who selected the circles. In the case of a latent response, subjects will be utilizing the vis-implic stream through direct mode. Just as with Guthrie and Wiener's experiment<sup>11</sup>, this should enable a rapid response. In contrast, subjects consciously attending to the structural stream will override the latent response and enter the oscillatory configuration described above. With no 'rational' reason for selecting one or other of the shapes as hostile, any response will tend to be delayed by the longer cycle of processes involved.

## 5.2. Belief and Focal Awareness

ICS provides insight into an experiment carried out in conjunction with those described in Section 4, but not reported in<sup>12</sup> due to lack of statistically significant results.


The image in Figure 9 is from an experiment intended to explore whether NPR images convey a high-level invariant connected a sense of disbelief. This was motivated by observations of cartoon art. In cartoons, characters can and are made to perform acts which completely violate normal expectations derived from real-world experience. Yet viewers watch and engage with such features without a sense of disbelief that might be engaged if realistic characters and scenery were present. In the experiment, subjects are divided into two groups, with one group given the upper, realistic, image in Figure 9, and the other the lower, NPR, image. Each subject is then asked whether the character portrayed would be likely to succeed in jumping between the buildings.

In ICS terms, the difficulties with this test stem from attention being drawn to the structural detail of the scene, and (consequently) focusing awareness on matching the semantics of the scene with the task description. For exactly the reasons given above, this increases the likelihood that buffering at implic will bring propositionally-derived invariants into play. The difference with the situation described above is that in the earlier example the task was sufficiently clear to first language speakers that they could perform without implicational buffering, whereas the jump analysis is impossible to resolve satisfactorily for any sample population. A better approach would be to set up an assessment that drew attention to some secondary task, leaving assessment of the physical circumstances to latent implicational knowledge. For example, asking whether the character would be able to retrieve some object (which happened to be located on the other side of the gap) would encourage buffering at the propositional level, allowing direct mode processing of any implicational invariants.

## 5.3. The Application of Affect

As we have discussed, for latent knowledge to be utilized effectively, the user should not be consciously attending to





**Figure 9:** 'Suspension of Belief' Images

affective content of streams. That is, it is desirable that invariants within the external representation that activate latent knowledge should not be modified by invariants extracted by cultural experience or propositional reflection.

There are reasons to believe that these conditions are met in a number of computer graphics applications, specifically games and virtual environments. These combine task-centered activity with a rich visual environment that provides a sense of immersion and presence. In the case where attention is drawn into the visual environment, buffering will move out of the central engine (implic and prop), and into the systems involved in visual interpretation (vis and obj). With the central engine processes 'on automatic', allowing implicationally derived invariants to contribute maximally. And indeed, our experiments<sup>12</sup>, summarized in Section 4, point already to two aspects of games and virtual reality where implicational invariants could be invaluable:

- Games, particularly the new generation of character-based multi-player environments, require users to engage in monitoring their environment for threats and opportunities, and to explore. For novice players in particular, a game is a large, complicated environment in which un-

familiarity may easily lead to disappointment. By using affect-related invariants in rendering a game, it should be possible to provide cues about levels of threat posed by entities within the game, as well as guiding players, when appropriate, towards useful or safe areas.

- Virtual environments are often a facade, presenting the appearance of a detailed environment, but in which the apparent capabilities of realistically rendered objects are in practice matched by few actually capabilities. By using NPR rendering, it should be possible to unobtrusively guide attention away from imposter objects, towards task-relevant or usable features of the environment.

By using rendering style to convey meaning through the implicational level, we support a mode of processing in which the more implic contributes, the more likely it is that a wide range of propositional material will be generated and taken into account within central engine processing. In a configuration where central processes are able to operate directly, the buffer (and focal awareness) can operate at the visual and object levels, supporting the sensation of presence that is a major design objective in these application areas.

## 6. Conclusions

This paper makes contributions to computer graphics at two levels. The first level is the direct contribution. We have set out examples of how rendering style can convey meaning and influence judgement, we have demonstrated (through reference to recent experiments) that these effects are controllable, and have provided a generalized account that explains why these effects occur, how they are subject to variation, and how they might be applied in practice.

At the second level, the paper calls for a re-appraisal of the relationship between graphics and psychology. Traditionally, this has been focused at the level of psychophysics. The most important point of our analysis is that the effects we describe are outside the realm of low-level perception. They involve semantics, affect, and other high-level invariants such as rate-of-change. Worrying solely about the perceptual adequacy or realism of an image is to miss the fact that meaning is ultimately interpreted at higher levels of cognitive process, and involves the integration of more than just structural invariants. Understanding the meaning carried by rendered images requires broader approaches to human information processing. Significantly though, thinking about higher levels of processing does not require thinking in the abstract: effects that exist at this level are measurable, and are accessible through empirical experiment.

There are other links between graphics and psychology that have not been touched on here, most importantly perhaps the phenomenon of 'change blindness'<sup>17</sup>, which may in part be related to the issue of indifference to streams exhibiting low rates of change. There are also links to be explored with the use of sound in virtual environments; as the ul-

loomo/takete example shows, sound also has affective properties, something which cinema uses to good effect in creating and sustaining emotional states. As ICS includes acoustic processing, it provides a natural route for exploring the theoretical basis for multimodal blending.

It is clear that a model is needed to underpin the development and deployment of NPR in a way similar to how psychophysics and models of human vision have supported development of realistic rendering. Such a model may remain a long term goal, but the contribution of this paper represent a first step towards achieving it.

### Acknowledgements

The first author gratefully acknowledges the support of an Advanced Research Fellowship (AF/001275) from the UK Engineering and Physical Sciences Research Council.

### References

1. P.J. Barnard. Asynchrony, implicational meaning and the experience of self in schizophrenia. In A. David and T. Kircher, editors, *The Self in Neuroscience and Psychiatry*. Cambridge University Press, 2003. In press.
2. P.J. Barnard and J. May. Interactions with advanced graphical interfaces and the deployment of latent human knowledge. In *Eurographics Workshop on Design, Specification and Verification of Interactive Systems*, pages 15–49. Springer, 1995.
3. P.J. Barnard and J. May. Representing cognitive activity in complex tasks. *Human Computer Interaction*, 14:93–158, 1999.
4. P.J. Barnard and J.D. Teasdale. Interacting cognitive subsystems: A systematic approach to cognitive-affective interaction and change. *Cognition and Emotion*, 5(1):1–39, 1991.
5. V. Bruce, P.R. Green, and M.A. Georgeson. *Visual Perception: Physiology, Psychology, and Ecology*. Psychology Press, third edition, 1996.
6. S. Chokron and M. De Agostini. Reading habits influence aesthetic preference. *Cognitive Brain Research*, 10(1–2):45–49, 2000.
7. R. Davis. The fitness of names to drawings. *British Journal of Psychology*, 52:259–268, 1961.
8. D.J. Duke, P.J. Barnard, D.A. Duce, and J. May. Syndetic modelling. *Human Computer Interaction*, 13(4):337–393, 1998.
9. F. Durand, M. Agrawala, B. Gooch, V. Interrante, V. Ostromoukhov, and D. Zorin. Perceptual and artistic principles for effective computer depiction. SIGGRAPH 2002 Course. See [http://graphics.lcs.mit.edu/~fredo/SIG02\\_ArtScience/](http://graphics.lcs.mit.edu/~fredo/SIG02_ArtScience/).
10. B. Gooch and A. Gooch. *Non-Photorealistic Rendering*. AK Peters Ltd., 2001.
11. G. Guthrie and M. Weiner. Subliminal perception or perception of partial cue with pictorial stimuli. *Journal of Personality and Social Psychology*, 3(6):619–628, 1966.
12. N. Halper, M. Mellin, D.J. Duke, and T. Strothotte. Implicational rendering: Drawing on latent human knowledge. *Submitted to ACM Transactions on Applied Perception*, 2003. Test details available at [www.opennpar.org/userstudies](http://www.opennpar.org/userstudies).
13. F. Heider and M. Simmel. An experimental study of apparent behaviour. *American Journal of Psychology*, 57:243–259, 1944.
14. V. Interrante. Harnessing natural textures for multivariate visualization. *Computer Graphics and Applications*, 20(6):6–11, 2000.
15. D.E. Kieras, D.E. Meyer, S. Mueller, and T. Seymour. Insights into working memory from the perspective of the EPIC architecture for modelling skilled perceptuo-motor and cognitive human performance. In A. Miyake and P. Shah, editors, *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control*. Cambridge University Press, 1999.
16. D.H. Laidlaw, R.M. Kirby, and H. Marmanis. Visualizing multivalued data from 2d incompressible flows using concepts from painting. In D. Ebert, M. Gross, and B. Hamann, editors, *IEEE Visualization '99*, pages 333–340. IEEE Computer Society Press, 1999.
17. J. May. Perceptual principles and computer graphics. *Computer Graphics Forum*, 19(4):271–279, 2000.
18. K. Provins, H. Stockbridge, D. Forrest, and D. Anderson. The representation of aircraft by pictorial signs. *Occupational Psychology*, 31:21–32, 1957.
19. J. Schumann, T. Strothotte, A. Raab, and S. Laser. Assessing the effect of nonphotorealistic rendered images in CAD. In R. Bilger, S. Guest, and M.J. Tauber, editors, *Proc. Computer Human Interaction (CHI'96)*, pages 35–42. ACM Press, 1996.
20. T. Strothotte, B. Preim, A. Raab, J. Schumann, and D.R. Forsey. How to render frames and influence people. *Computer Graphics Forum*, 13(3):C–455–C–466, 1994. Conference issue: Proceedings of Eurographics'94.
21. T. Strothotte and S. Schlechtweg. *Non-Photorealistic Computer Graphics: Modelling, Rendering, and Animation*. Morgan Kaufmann, 2002.
22. P.G. Walz and R.M. Rapee. Disentangling schematic and conceptual processing: A test of the interacting cognitive subsystems framework. *Cognition and Emotion*, 17(1):65–81, 2003.