Six St Jeromes: notes on the technology and uses of computer lighting simulations

SIMON NIEDENTHAL

I would like to begin with a quotation from St Bonaventure: 'Light is the substantial form of bodies; by their greater or lesser participation in light, [bodies] acquire the truth and dignity of their being'.¹ By 'bodies', Bonaventure refers not to our physical being but, rather, to matter itself; his argument takes place in the context of the great medieval metaphysics of light. In a similar vein, Robert Grosseteste begins *De Luce*, his thirteenth-century treatise on light, with the assertion that:

Light of its very nature diffuses itself in every direction in such a way that a point of light will produce instantaneously a sphere of light, unless some opaque object stands in the way. [. . .] Corporeity [that is, the capacity of matter to take form and occupy space], therefore, is either light itself or the agent which performs the aforementioned operation and introduces dimensions into matter in virtue of its participation in light.²

Light is portrayed by both theologians as the mechanism by which the cosmic x, y and z axes are defined, the first of a number of intriguing echoes of medieval thought on light and optics that one encounters in digital image synthesis.

In a less metaphysical sense, an understanding of the power of light to define things underlies foundational lighting curricula everywhere (including my own). A quick survey of some notable photographs demonstrates the way in which

¹ Umberto Eco, *The Aesthetics of Thomas Aquinas* (Cambridge, MA: Harvard University Press, 1988), p. 110.

² Ibid.

light and value are central to the perception of pictorial form, surface qualities, space, and volume. Qualities of light, moreover, have the capacity to wield emotional influence through association with primal visual needs. And as with all visual experience, what we see and the way we think about light is informed by cultural traditions. As sensory phenomenon and interpretative touchstone, light yields to both scientific and cultural analysis, and serves as a useful starting point for a consideration of new applications of digital technologies to the subjects of art and architectural history.

In his seminal work *The Reconfigured Eye*, William J. Mitchell traces the development of digital image synthesis from simplest to most complex techniques, with reference to corresponding changes that occurred in painting during the Italian Renaissance and afterwards. His opening argument establishes a dialogue between photography and painting, two media with differing manners of engaging value and representing light effects. 'Synthetic-shading procedures' Mitchell writes, 'are used to develop perspective views into closer approximations to—even simulations of—photographs. (Thus they are closely analogous to the procedures Renaissance artists employed to convert line cartoons into tonal and coloured paintings)'.³ The process of creating a digital image based on a 3D scene begins with the fashioning of a wire-frame model to express the physical boundaries of surfaces and objects. In the early days of digital image generation, the boundaries were denoted as a list of vertices in 3D Cartesian space, though now we have become accustomed to manipulating wire-frame models through interfaces as graphical representations. Once the objects and spaces of a scene have been modelled, surface characteristics are associated with each formal element, light qualities and positions are specified, and the scene is rendered. The rendering program takes into account the geometry, surface qualities, and lights of a scene and renders a two dimensional representation of that scene from a given viewpoint.

Mitchell goes on to explore digital image rendering from simple flat shading to the most complex radiosity calculations, referring along the way to artists as diverse as Reynolds, Vermeer and Paul Gauguin. The easiest way to render a surface in the computer is to simply specify a single colour for each point on the surface, but over time rendering algorithms have been refined to allow the calculation of shading and light effects of progressively greater complexity and fidelity to visual experience. One conceptual breakthrough was the development of ray-tracing. The rationale for ray-tracing is that while there are billions of light photons bouncing around in any visual environment, the ones that really matter are those that enter the eye.⁴ In the ray-tracing process a virtual picture

³ William J. Mitchell, *The Reconfigured Eye* (Cambridge, MA: MIT Press, 1992), p. 137.

⁴ This insight is taken from the Lightscape Visualization System Getting Started, pp. 1–7.

plane composed of a grid of pixels is defined, and a ray is traced from an eye point through each pixel into the scene until it reaches a surface. The light environment informs the surface characteristics of each point, and the resulting hue is used to colour the pixel. This process continues until all pixels in the grid have been coloured. Ray-tracing is, as Mitchell points out, analogous to the way in which a perspective is constructed, and he quotes Leonardo:

'Obtain a piece of glass as large as a half sheet of royal folio paper and fasten this securely in front of your eyes, that is, between your eye and the thing you want to portray. Next, position yourself with your eye at a distance of two-thirds of a braccio from the glass and fix your head with a device so that you cannot move at all. Then close or cover one eye, and with the brush or a piece of finely ground red chalk mark on the glass what you can see beyond it'. Thus each ray [Mitchell continues] in the viewing pyramid projects a colour from a point in the scene to a point on the picture plane. . . .⁵

Besides the similarities to perspective construction, the process of projecting rays from the eye to a surface beyond is reminiscent of Euclid's theory of vision, which was considered viable at least until the optical work of Al-Hazan, though popular long into the medieval period.⁶

Although ray-traced renderings allow the simulation of more sophisticated visual phenomena, such as accurate reflections and refraction, the products of this sort of rendering still lack a sense of the real presence of light. Specifically, the interreflection and colour bleed from one surface to another are not calculated. In a room like the one in which I write this sentence, we could say that there are areas of direct illumination, where lights are shining directly on surfaces, and other areas of indirect illumination that are lit by the light bouncing from surfaces. The indirect illumination is not calculated in a ray-traced image. Practically speaking, simulating accurate light presence in a ray-traced image calls for a procedure that combines a photographic process—involving setting lights within the software—with one in which areas of light and dark value are applied directly to surfaces in a manner more analogous to painting. In the words of the old cinematographer's saw, 'if you can't paint with light, light with paint'.

So, following Mitchell's lead, I have rendered a simple scene using techniques of progressively greater complexity (fig. 1, plate 6). This is a simulation of a room with a skylight and several cubes in it, rendered under default lighting. Default lighting in *Alias* is the computer graphics equivalent of Grosseteste's formless universe. Anyone who has looked at light for a while can play around for a few minutes and produce something that demonstrates a more convincing impression of light in the environment (fig. 2, plate 7). The Phong shading

⁵ Mitchell (as in note 3), p. 139.

⁶ David Park, *The Fire Within the Eye: A Historical Essay on the Nature and Meaning of Light* (Princeton: Princeton University Press, 1997), p. 78.

algorithm appears complex, yet it renders a fraction of the effects of light in a real space. In fact, it is so limited a model that I had to paint many of the subtle shadows and gradations in this figure directly on the surfaces, rather than rely upon the algorithm's inadequate simulation of surface-to-surface interreflection.

The advent of radiosity rendering—which models light behaviour much more accurately than other rendering modes—signals a major advance in the capabilities of computer simulation (fig. 3, plate 8). Radiosity rendering derives originally from methods developed by thermal engineers in the 1960s to simulate heat transfer between surfaces, as a means of designing jet engines, among other things. In the mid-1980s, computer scientists at Cornell University began to apply these techniques to the modelling of light with the computer. Radiosity calculations are based upon the principle of conservation of energy. When light is introduced into an environment, it illuminates some surfaces directly, and, depending on surface characteristics, some of the light is absorbed and the rest reflected back into the environment. Most current renderers do not calculate the contribution of surface-to-surface interreflection to the illumination in a scene. Radiosity software models interreflection by creating a mesh which stores the illumination values associated with each surface. After distributing energy from the light sources, the algorithm checks each surface and determines which has the most energy to reflect back into the environment. The progressive distribution of light energy continues until all of the energy in the environment has been absorbed, or until the process is stopped.

Architectural historians eager to simulate unbuilt or destroyed buildings have seized upon rendering technologies of greater sophistication, and examples abound of reconstructions of everything from Frank Lloyd Wright buildings to the castle of Mad King Ludwig of Bavaria.⁷ Notable works that have pioneered the use of radiosity rendering technology in an architectural context include Kent Larson's 1993 simulation of Louis Kahn's unbuilt Hurva synagogue.⁸ Larson's most recent work continues the exploration of light and surface in Kahn's unbuilt projects for Scripps at La Jolla.⁹ Also beautifully realised is MIT colleague Takahiko Nagakura's visualisation of Terragni's unbuilt Danteum.¹⁰ These projects are not the rule, however, for although computer-generated images have become commonplace in recent years, just as commonplace is the absence of rich and compelling illumination qualities and surface definition.

⁷ B. J. Novitsky, 'Reconstructing Lost Architecture', *Computer Graphics World* (December 1998), 25. See also the April 1999 issue of *CGW*.

⁸ Vincent Scully, 'A Virtual Landmark', *Progressive Architecture* (September 1993), 80–87.

⁹ Kent Larson, *Louis I. Kahn: Unbuilt Masterworks* (New York: Monacelli Press, 2000).

¹⁰ These images are taken from the 1998 Siggraph reel.

fects of light in a
many of the subtle
s, rather than rely
the interreflection.
behaviour much more
advance in the
rendering derives
1960s to simulate
lines, among other
iversity began to
computer. Radiosity
energy. When light
scatters directly, and,
absorbed and the rest
do not calculate the
diffusion in a scene.
in which stores the
storing energy from
determines which has
progressive distribu-
environment has been

rayed buildings have
been, and examples
of light buildings to
have pioneered
include Kent
Larson's
Kahn's unbuilt
MIT colleague
These
images have
the absence of

Mitchell's contention that the development of techniques of digital imaging is analogous to painting procedures is both encouraging and problematic: encouraging because it suggests that processes we are currently engaged in can be informed by art history; problematic because it fails to address significant differences between digital and analogue. Making precise distinctions is central to thinking critically about media, yet we continually witness the bleeding of art terminology into the realm of the digital. Early on I was struck by the way in which filmic and photographic concepts were embedded in software. Consider, for example, the camera icon from Maya software (fig. 4, plate 9). But the analogies in software go far beyond the photographic. In computer graphics we refer to renderings, we employ software called Painter and Piranesi, and, while waiting for an action to be performed, my laptop executes a little version of the Mona Lisa instead of showing the usual clock face. Another amusing example of the intersection of painting vocabulary and interface design is the creation icon in 3D Studio Max, a hand with a spark at the fingertip that is a cross between Adam's creation in the Sistine Chapel and ET. But anyone who has performed a task as elementary as 'drawing' a curve in a 3D package understands just how limited are the analogies between digital and other media. I always think of defining a c.v. curve as being more like an abstract type of wood bending, of applying forces to a ductile material and editing the points until the desired profile has been achieved. Any notion of the gestural component of drawing pretty much goes out the window in 3D software.

The construction of painting space in a digital medium is one means of exploring Mitchell's painting/rendering analogy, and, in fact, the last few years have seen a number of examples. Though the simulation of a painting in 3D space is often seen as striking a blow against the ascendancy of photorealism in computer rendering, some projects seek a wider scope, including 'Rouen Revisited' by Paul Debevec at the University of California. The Rouen project uses Monet's 30 or so canvases of the cathedral as well as historical and contemporary photos, to allow a viewer to reconstruct the front façade from any number of possible viewpoints at any time of day or season. Debevec describes his project as follows:

Fascinated by the play of light and atmosphere over the Gothic church, Monet systematically painted the cathedral at different times of day, from slightly different angles, and in varied weather conditions. Each painting, quickly executed, offers a glimpse into a narrow slice of time and mood. We are interested in widening these slices. . . .¹¹

¹¹ From Paul Debevec's Home Page, alberti.cs.berkeley.edu/~debevec/. Another project involving the 3D reconstruction of a painting space is the stomach-churning fly-through of Van Gogh's Yellow Room on the artmuseum.net web site: www.artmuseum.net/vangogh/3dloginpage.asp

While this project undoubtedly advances the technology of image generation, it raises questions about the uses of digital construction. Architecture at least aspires to three dimensions in its built form. There are many reasons that the Kahn and Terragni projects, for example, remained unbuilt—often financial ones—but there is still something to be learned about the modulation of light that can be usefully approximated by computer simulation. Construction of painted architectural space is a bit more problematic. After all, a painting is quite content to remain in two dimensions.

I would argue that 3D simulation of painting space is useful when it goes beyond merely expanding the subject to another dimension and instead adds an additional facet to our understanding of the work at hand. Applying photo-realistic simulations of light behaviour to a painting can serve, for example, as a means of throwing the painter's value choices into relief. Thus the tasks of the 'Six St Jeromes' project: first, to reconstruct some of the intellectual context of a painting, and, second, to rethink the dangers, uses, and legitimate pleasures of this sort of simulation. The goal of this project as it was initially conceived was to put Mitchell's thesis to the test, to collide a fifteenth-century painting with the most recent rendering technology, and to see what emerges.

Antonello da Messina's 'St Jerome in his Study' (National Gallery, London) dates from somewhere between 1450 and 1475 (fig. 5, plate 10); there seems to be some uncertainty as to whether the painting represents Antonello's early work, or was produced during his later stay in Venice.¹² The French novelist Georges Perec, who kept a print of the painting on his wall, nicely grasped the snug bookishness of the study. 'The whole space' Perec writes,

is organized around the piece of furniture (and the whole of the piece of furniture is organized around the book). The glacial architecture of the church (the bareness of the tiling, the hostility of the piers) has been cancelled out. Its perspective and its vertical lines have ceased to delimit the site simply of an ineffable faith; they are there solely to lend scale to the piece of furniture, to enable it to be inscribed. Surrounded by the uninhabitable, the study defines a domesticated space inhabited with serenity by cats, books and men.¹³

The St Jerome painting offers itself for several reasons. First, because da Messina was among the first Italian artists to work with oil paints and glazes, and the new materials allowed him a greater range of tonal and light effects. He was also active in a period in which the changes documented by Mitchell were becoming fully felt. Moreover, the painting responds well to conceptions of light in the art theory of da Messina's contemporaries like da Vinci.

¹² Cf. N. Little, 'A Note on the Date of London "St Jerome in his Study" by Antonello da Messina', *Arte Veneta* (1976), 154–55.

¹³ Georges Perec, *Species of Spaces and Other Essays* (New York: Penguin, 1997), p. 86.

The project began with the construction of a wireframe model of St Jerome's study, corresponding to the perspective of the painting (figs. 6 and 7, plates 11 and 12). A preliminary shaded version of the scene with default lighting and basic texturing (fig. 8, plate 13) was brought into Lightscape software for radiosity processing. Once the model was set up in an environment that roughly corresponds to the illumination qualities in the painting, it became possible to begin a quick series of 'what if?' renderings: what would it look like if sunlight was coming from the other side? (fig. 9, plate 14). What if the light coming in the doorway was very cool and flat? (fig. 10, plate 15). What if soft light was coming in exclusively from the clerestory windows above? (fig. 11, plate 16). What if Jerome decided to burn the midnight oil? (fig. 12, plate 17).

At this point it struck me that I had achieved only a sort of inverse appreciation of the light of the painting. That is, none of my variations seemed as satisfying as the light in the original. So in my fifth and sixth versions of St Jerome I attempted to reconstruct the lighting environment suggested by the painting (fig. 13, plate 18), and used whatever techniques necessary to replicate the values on the surfaces (fig. 14, plate 19). In doing so, I learned that the painting manifests a close relationship to theories on painting and light of near contemporaries like da Vinci. The painting exhibits, of course, a wonderful chiaroscuro, as defined by Leonardo:

The first intention of the painter is to make a flat surface display a body as if separated from this [picture] plane, and he who most surpasses others in this skill deserves most praise. This accomplishment, with which the science of painting is crowned, arises from light and shade, or as we may say *chiaroscuro*.¹⁴

Further, the treatment of the saint's face echoes other writings of Leonardo on the rendering of facial form:

The utmost grace in the shadows and the lights is added to the faces of those who sit in the darkened doorways of their dwellings. Then the eye of the beholder observes the shaded part of the face thrown into deeper shade by the shadows from the aforesaid dwellings, and sees brightness added to the illuminated part of the face by the radiance of the atmosphere. Because of such increases in the shadows and light the face acquires great relief. . . .¹⁵

Besides his 'St Jerome', Antonello's portraits also frequently demonstrate this strategy. In his treatise on light in Renaissance theories of art, Moshe Barasch distinguishes two schools of thought on the concept of illumination. In the first, a 'functional' conception, the major task of light is to illuminate bodies and render them visible, while the function of shadows is to produce relief and the semblance of corporeality. The second approach emphasises the relation of light

¹⁴ Mitchell (as in note 3), p. 139.

¹⁵ Ibid, p. 143.

to contents and themes.¹⁶ 'St Jerome in his Study' responds to both of these conceptions. Close scrutiny reveals careful delineation of the form of the desk, apparent in the touches on the panel that do not reference the suggested light environment—gradations, for example, where one surface meets another. And as we have seen, the treatment of Jerome's face echoes da Vinci's advice on how to make a body emerge from the picture plane. But the presence of light in the painting also references iconographic themes. In her study of the painting, Penny Howell Jolly notes the presence of Marian symbolism and the similarities of the Jerome painting to an Annunciation. 'Like Mary at the Annunciation', Jolly writes, 'Jerome is bathed in a celestial light and visited by a divine presence'.¹⁷

Ultimately this project suggests possible uses for radiosity rendering as a tool of art scholarship. One could use computer simulation and radiosity processing to check a 2D representation against a 3D structure upon which it was based, to help visualise the behaviour of light and simulate the space of unbuilt or lost projects, to aid in the process of conservation or historic preservation, or to study the illumination environment of works of art. This is especially useful in cases where the conditions of illumination in a space have altered over time.

As a field of study in itself, however, computer-generated light has barely been explored. One good critical touchstone to apply to the generation of new digital projects is the question: could this project have been done if the computer did not exist?¹⁸ Could it even have been thought of? One thing that sets computer simulations apart from light in real space is the light algorithm itself, and gaining an understanding of computer-generated light calls us to contemplate the differences between the algorithm and real light behaviour. For example, one unique quality of light produced by algorithm is that the algorithm can be reversed: if the intensity of a source in most 3D packages is set to a negative number, the source extracts light from the scene. 'The light that shineth in darkness' becomes 'the light that shineth darkness', inverting the process with which I opened this paper: 'Light of its very nature diffuses itself in every direction in such a way that a point of light will produce instantaneously a sphere of light'. . . . The capacity to simulate the inverse of light behaviour—to create spheres of darkness—is clearly one path to explore with algorithmic light, and offers the possibility of new modes of spatial visualisation that philosophers like Grosseteste could never have conceived.

¹⁶ Moshe Barasch, *Light and Colour in the Italian Renaissance Theory of Art* (New York: New York University Press, 1978), pp. xi–xiii.

¹⁷ Penny Howell Jolly, 'Antonello da Messina's St Jerome in his Study', *Art Bulletin*, 65 (June 1983), 252. Jolly's argument is based upon a reading of Jerome's Epistle 22 to Eustochium, in which Jerome counsels imitation of Mary as the basis for conduct.

¹⁸ I am indebted to an editorial by John Lansdown in *Leonardo Magazine* (1998) for this test.

Plate 6: Cornell box, default lighting.

Plate 7: Cornell box, raytraced rendering.

Plate 8: Cornell box, radiosity rendering.

Plate 9: Maya camera icon.

Plate 10: 'St Jerome in his study', Antonello da Messina, c. 1450.

Plate 11: Wireframe.

Plate 12: Untextured three-dimensional form.

Plate 13: Default light and textures, no shading.

Plate 14: Jerome.

Plate 15: Jerome.

Plate 16: Jerome.

Plate 17: Jerome.

Plate 18: Jerome.

Plate 19: Jerome.

Plate 20: Haus Altenkamp, räumliche Darstellung des Entrées überlagert mit Schadenskartierung, Ansicht von Nordwesten.