Landscape Visualization: Progress and Prospects

Stephen M. Ervin

Abstract:

Landscape Visualization, once the exclusive domain of 'GIS'-style software, has entered mainstream efforts in professional fields such as architecture, landscape architecture, civil engineering and Hollywood movie-effects, and is now enabled by many CAD and animation/rendering systems, as well as GIS and remote-sensing software. The complexities of the essential elements of landscape models – notably landform, water, vegetation and atmospheric effects, and their curved, fuzzy, fractal, irregular surfaces and dynamic and ephemeral systems – present special challenges to computer graphics and spatial information systems. This paper reviews the recent history of landscape visualization technologies and software techniques, critiques the current state-of-the-art in GIS and related software, including ArcGIS 3D Analyst, and identifies a number of interesting current and future challenges and research opportunities.

Figure 1. Digital Landscape rendering produced from DEM data

1. Digital Landscape Visualization

Digital landscape visualization has a relatively short history in the context of other forms of landscape representation – arguably the first efforts were in the 1960's. The development of CAD and computer graphics in general also started at that time, but the majority of those early efforts were focused on the representation and visualization of objects, such as gears, airplanes, teapots, and etc. A specific concern for the landscape was present early on in the development of flight-simulation software (e.g. Evans & Sutherland), and it was during the formative early years of GIS development that visualization of terrain, for example, became a subject of study and development, and grid meshes and TINs, among other useful techniques, were invented. For the next twenty or so years, terrain representation and visualization was predominantly the purview of 'GIS'-style software, with some minor efforts in civil-engineering or computer aided architectural design software. Today, 40 years later, landscape visualization has entered mainstream efforts in professional fields such as architecture, landscape architecture, civil engineering and Hollywood movie-effects, and is now enabled by many CAD and animation/rendering systems, as well as GIS and remote-sensing software. The purposes of these visualization efforts run the gamut of cognitive, functional, and entertainment purposes, from detailed studies of visual impacts of forestry operations to purely impressionistic 'digital landscape paintings'.

In all of these and other efforts, recent developments in computer science and computer graphics have made breathtaking and eye-tricking effects possible; CAD and GIS and image processing and even digital video technologies and techniques have blurred together into a powerful combined system for creating digital landscape visualizations. In spite of these advances, and impressive results, a number of challenging problems remain.

One way to understand the enterprise of landscape visualization is as a combination of '*elements'* and '*abstraction level'*. By *elements*, I mean the six essential landscape elements identified in Ervin and Hasbrouck's *Landscape Modeling: Digital Techniques for Landscape Visualization* (McGrawHill, 2001):

The three quintessential elements of the 'landscape palette'

1. landform; 2. vegetation; 3. water;

and the other three necessary elements that inevitably combine with the first three to create most real landscapes:

4. structures (including 'infrastructure'); 5. animals (including people); 6. atmosphere (including sun, wind, etc.)

These six elements, I have argued, capture most of the constituents of what we usually mean by 'landscapes', in various combinations and proportions – ranging from open ocean to built cityscape, with a wide range of types of landscape in between.

'*Abstraction level*' is an important concept in all representation, as well as many cognitive tasks such as computer programming, planning and design, schema-development, and so on. With respect to visualization, I mean one of the following four levels, which correspond to different inferential purposes:

a. *Diagrammatic* (meant to clarify essential components and relationships, intentionally omitting many details, and with no attempt at 'realism; the most abstract);

b. *Evocative* (meant to bring forth a set of related ideas or feelings in the viewer – often the domain of fine art);

c. *Illustrative* (meant to illustrate typical patterns, configurations, or appearance, or highlight generic elements or combinations, with no effort to exactly portray anything specific);

d. *Realistic* (meant to visually portray some specific, existing or proposed landscape or combination of elements; the least abstract).

These four levels serve to capture a range of purposes for which visualizations are commonly made. This is important, because there is never a single correct answer to any of the many representational and abstraction problems encountered in this enterprise, and so reference to the questions: “What is the purpose?”, and “What is the question?”, is an important touchstone for understanding visualization tasks and evaluating representations.

Diagrammatic representations of landscape, the most abstract, make no pretensions of ‘visual realism’. They are rather, like all diagrams, concerned with ‘data model’ and ‘topology’: what are the classes and instances of interest in this investigation, and how are they connected? A classical diagram of the geologic formations cycle is a good example of this kind of landscape visualization.

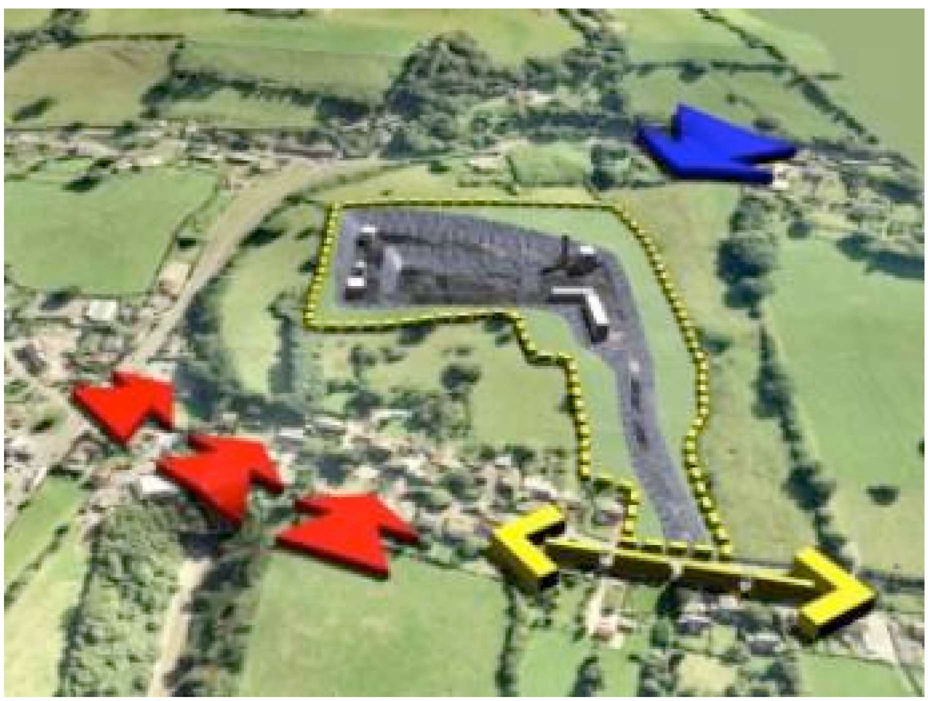

Figure 2. A Diagrammatic Visualization

Evocative representations may appear less abstract, but they make no appeal to any real or concrete phenomena; rather they are intended to evoke moods, feelings, images, and thoughts. Evocative landscape representations may be made to convey a sense of place, as artistic expressions, or as a fictional virtual landscape. Computer game backgrounds, and many kinds of digital landscape representations found on the web, made by such software as Bryce, Vue d'Esprit, et al, are of this ilk.

Figure 3. An Evocative Visualization

Illustrative representations are less abstract yet, and typically have crossed the 'realism' threshold. While not meant to be specific, they are meant to be realistic enough to base inferential judgments on, ranging from 'will it work?" to "Do I like it?" Artists' renderings of proposed architectural, landscape, and urban developments are all of this kind. They don't claim to show the very trees that will be planted, but that there will be some trees, of approximately this size, shape and color. The term '*geo-typical'* has been used to describe these renderings in GIS.

Figure 4. An Illustrative Visualization

Realistic representations are perhaps the most commonly undertaken –and often the least successful. This is simply because a representation is a re-presentation, and not the thing itself (the map is not the terrain, etc.), and closing the gap is an asymptotic challenge! So-called ‘photo-realism’ has been a rather sought-after goal in computer graphics for several decades, but as any photographer can attest, photos are not always, or even ever, ‘realistic’. There are still considerations of lighting, framing, focus, resolution, etc, that greatly influence the perception and value of the photograph. The great landscape photographer Ansel Adams didn’t ‘simply capture’ landscapes, he transformed them in the process.

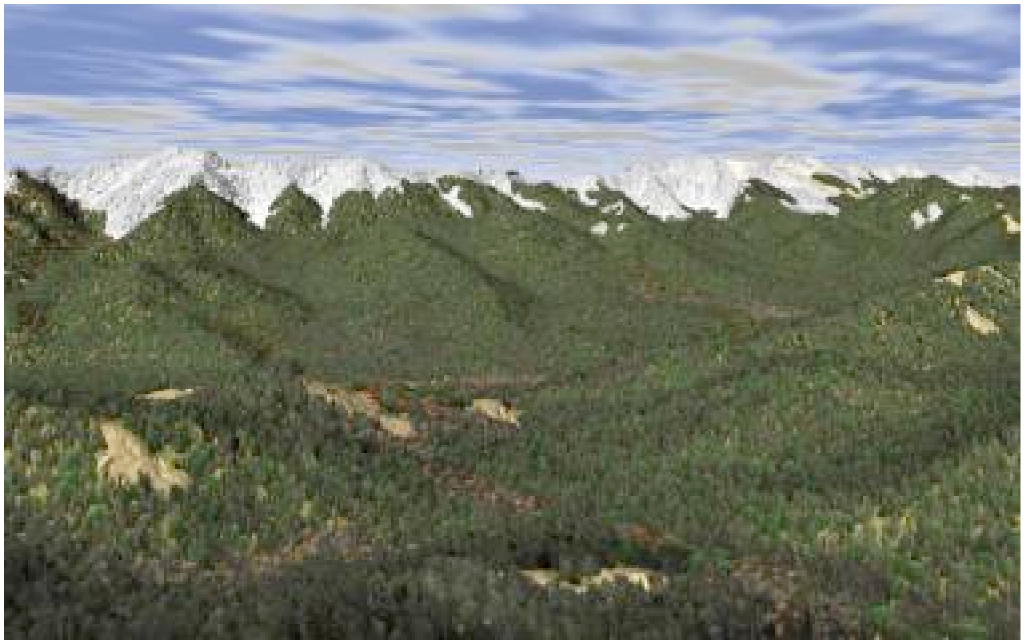

Figure 5. A Realistic Visualization

Nonetheless, the temptation to try to create 'accurate' visualizations is a powerful magnet, and the tools for doing so are increasingly effective. And for some inferential purposes, such as evaluations of visual impact or visibility analyses, a degree of realism, precision and accuracy is important to the task. For these kinds of visualizations using GIS tools, the term '*geo-specific'* has been used [ref].

In the approximately four decades of experimentation and development of landscape visualization, starting with the infancy of GIS (for example, Carl Steinitz' work at the Harvard Lab for Computer Graphics and Spatial Analysis), the trajectory has been from primitive illustrative (geo-typical) towards realistic; with a substantial branch of the evocative in the last two decades. (Diagrammatic representation have never been high on the list or very well developed – partly because they are in some ways the most difficult – even while being the graphically simplest!)

In the next section, I outline some of the issues for illustrative and realistic representation of the essential landscape elements, as they may be visualized using GIS software.

2. Landscape Elements

2.1 Landform

Landform typically requires large digital data sets, and so terrain and site modeling have historically been associated with GIS, rather than CAD, and assumed to require larger computers, more memory and faster processors than other modeling tasks. Whereas CAD was initially concerned with points, lines, planes, cubes and cylinders (an oversimplification, to be sure), terrain modeling has depended upon a variety of representations and mathematical abstractions, from spot elevations and contour lines, to 2D grids and 3-D meshes, ruled surfaces, triangulated irregular networks (TINs), regular and irregular solids, Boolean operations and NURB surfaces. The grid representation, the basis of all Raster GIS, is the 'mother of all landform representations', and enables a wide range of analytic calculations, including slope, aspect, visibility, drainage, and others. As a visual representation, grids suffer from two basic flaws: a requirement for constant spacing which is inefficient when surface variation details are few, and inadequate when details are many; and the fact that the four points of a grid cell may not be planar, but more often form a complex curved surface (hyperbolic paraboloid), which poses problems for simple computer graphics rendering algorithms that depend upon flat planes for rapid calculation of, for example, surface normals and shading. TINs overcome both of these problems, and so have mostly overtaken grids as the representation of choice for visualizations.

In addition to the shape of the surface, as determined mathematically by mesh, TIN or otherwise, for visualization purposes a visible surface texture and coloring are also required. Asphalt, concrete, brick and other regular tilings are the easiest, but even these present difficulties in rendering. Often the 'edge-match' of a tiled surface is visible, or the repetitive pattern is visible enough to be disturbing. Most rendering programs don’t have good texture-scaling capabilities, so that textures that work well in the mid- and background, look distorted, out of scale and out of focus in the foreground. Getting a simple 'grass' surface is not something that is yet commonplace in digital landscape modeling. One solution for these texturing problems is presented by 'procedural' texturing approaches, that generate surface features (or pixels) 'on-the-fly', rather than depending on a simple 2D image to be used. Typically, these require greater processing power and longer rendering times, but promise greater control over scale-dependent-detail and (possibly random) variation.

Diagrammatic landform visualizations are usually found in geology or hydrology texts; and usually they are handcrafted, since they don’t actually come from an existing data set.

Evocative landform is the domain of the great sculptural landscape architects, such as George Hargreaves, Michael van Valkenburgh and others, as well as ‘earth artists’ including Robert Smithson and Michael Heizer. While none of the above has visibly leveraged GIS tools for their efforts, a great population of amateurs is involved in the creation of imaginary, often evocative, landforms and landscapes. The early mathematical experiments with ‘fractal terrain’ were bolstered by the evocative qualities of some of those early renderings. And off-the-shelf software such as Bryce 3D, Vue d’Esprit, Terragen, and others – often designed for the purposes of video-games environments and special-effects backgrounds – have put simple tools within reach of millions for creating such visualizations. A simple web search for ‘imaginary terrain’ will turn up a number of examples of galleries of such images.

2.2 Vegetation

Figure 6. Early Forest Service rendering of vegetation on landform

Representing vegetation – whether trees, forests, or grassy groundcover – is a daunting challenge, in any medium, but especially when counting polygons! Whereas an ordinary building might be well represented with several thousand polygons and simple geometric primitive solids, no part of a plant is flat, square or even really cylindrical. Millions of polygons – or even greater orders of magnitude – are required to begin to capture an ordinary tree, with its fractally complex branching and articulation – and a forest represents a greater numerical and visualization challenge yet!

While some software exists for generating 3D solid models of plant forms, (such as Onyx TreePro, AMAP, and others) these are typically not integrated in to GIS, and are typically used only by a subset of specially interested or motivated illustrators. More common by far is the 'bitmapped billboard' technique, in which a photographic image of a tree (or ground texture, etc.) is projected onto a transparent flat plane in the scene. This method is available in every CAD or rendering package today, and provides a very good general facility for bringing photographic detail into a scene without too great a polygon count. The method is general – the same method can be used for people, cars, streetlights, and etc. –

and extendable: a forest can be rather simply modeled by a sufficient number of these texture-mapped planes. This is the most widely used technique, by such software as 3Dnature World Construction Set, Terragen, and others, for generating vegetated and forested landscape scenes.

A major limitation of this and most other techniques is that they cannot capture or reflect dynamics in the landscape—trees growing, dying, blowing in the wind, or co-evolving to maximize exposure to sunlight. Growth of trees or shrubs – a simple fact of the landscape, and a major contributor to visual quality over time - is hardly addressed at all in any modern software except in the most cartoonish way. Thus, it may be possible to substitute different tree symbols (circles of different sizes in plan, or texture mapped photographs of different growth stages in elevation) as function of time, but no tree in any CAD or GIS system ever died from crowding or natural causes; and none has yet grown lopsided from environmental factors.

Many early attempts at representation of vegetation were in effect purely diagrammatic, due to their limited graphical techniques (such as the Forest Service 'Perspview' and other such images). Evocative vegetation visualizations generally come from the same camp as evocative landform scenes – amateurs and artists using the power of computer graphics to create trees, forests, shrubs and often fantastic landscapes.

Most 'realistic' visualizations are truly only 'illustrative' with respect to vegetation; very few have photographs of the actual plant or plants in question, relying instead on 'typical' representative plant photographs. Only a virtual arboretum, for example, might actually bother to have virtual representations of specific plants. Most software for creating visualizations for the forestry industry, such as SMARTFOREST and others, rely upon libraries of typical plants, although varying by species, height and age, for example.

2.3 Water

Water seems deceptively simple compared to the patent complexity of vegetation. On a still morning, a lake may be modeled as a flat reflective plane, just a mirror in the landscape. Look a little closer, though, and you see refractivity and transparency mixed in, and with just a little wind, ripples and waves. Modern rendering systems have techniques for reflectivity and refraction (technically requiring ray-tracing, a time and memory-consuming set of calculations) and ripples and waves are mathematically simple perturbation s to model.

Often water, whether flat or rippled, with waves or vortices, can be effectively modeled with a TIN representation, just as landform – only with a different texture (and different physical and optical properties). Lakes, ponds, and oceans are easier for most GIS and other software, than streams and rivers, since these latter have non-flat, highly variable and dynamic surfaces. When water is not a separate surface, but rather part of n ecosystem, as in a wetland, or has 3D complexities, as in the case of a waterfall, then GIS data can not directly represent the water, and any visualization will have to use secondary sources of information for texture, shape, etc. Because of its ephemeral and fluid nature, water in the landscape remains perhaps the most difficult element to render.

Diagrammatic visualizations of water are rarely seen outside of hydrological analyses, (or perhaps illustrations of drainage or irrigation systems.) On the other hand, every flat simple blue expanse may be considered essentially diagrammatic, as the details of water texture are usually too complex to display. Evocative water is easier to do, since water, reflections, fog, and associated atmospherics are among the most romantic visual cues for many people. ‘Realistic’ water is surely the hardest form to generate, and is rarely seen in landscape visualizations, outside of high-tech Hollywood effects (as in the ‘Finding Nemo’ film, and others).

2.4 Structures (including ‘infrastructure’)

Modeling 3D structures is inherently the domain of CAD, rather than GIS software. GIS has long been constrained by the essentially 2D or 2.5D data sets at its core; few buildings are so constrained. Many however, are simple extrusions of polygon plans; some even have flat tops, and many have simple roof structures that can be approximated with TIN like representations. So, with the addition of some texture mapping, the 'extruded plan' approach found in some software like ArcGis 3D Analyst, works well for many cases. For 'real' buildings, CAD models are necessary, but for illustrative geo-typical representations, of neighborhoods, cities, or regions, the extrusion method works quite well. Software such as CommunityViz and others rely upon these simple representations. Mostly, in GIS-based visualizations, the 'illustrative' is the goal; but with rendering software coupled to CAD modeling systems, the geo-specific is clearly possible, and some very detailed representations of campuses and cities have been made.

\

Figure 7. Visuallization   of City of Boston by extruded polygons

2.5 Animals (including people)

Animals including people are essential elements of most natural or built landscapes, even in indirectly or invisibly. They are almost never actually included in any database, except indirectly (as implied residents of a structure, or creatures in a habitat), and yet their presence is an important part of real, ‘fleshed-out’ visualizations. Bitmapped billboards are the most common human inhabitants of most landscape visualizations, though some ‘TiIN-men’ (3D models composed of TIN-like 3D meshes) may be found, too! Modeling the appearance of animals is just as hard as other landscape elements, for many of the same reasons: fuzzy, curved, complex, dynamic features. And their behavior is all the more difficult to represent, much less comprehend! For almost all animals, only the ‘illustrative’ is ever attempted in landscape visualizations.

Figure 8. Cow and People in the Landscape

2.6 Atmosphere (including sun, wind, etc.)

As much as every landscape model is made up out of the previous five elements, especially for visualization purposes they are all dependent upon the 'atmosphere' within which they are situated. In many rendering programs there are rendering parameters that substitute for the (ethereal but real) atmospheric qualities of lighting, fog, haze, and others. In this category also are included the motion of the air (wind), and other 'objects' which are present but so distant they must be substituted for in most landscapes (the sun, the moon, clouds, e.g.). Atmospheric haze – the effect of lightening, blue-ing and blurring of objects in the distance – is perhaps the most useful of these effects for landscape-scale GIS-based visualizations, and many renderers incorporate some form of this effect, though often without much real control over its parameters.

A visible sky, with clouds, is another feature of most landscapes, and that is usually achieved by a simple background (2D) bitmapped image; or possibly an image texture mapped onto an enclosing 'sky-dome'. Neither accommodates the need for perspective shifts in rendered images from different viewpoints. Some systems now incorporate 'procedural' methods for rendering clouds, ranging from dense gray stratus to fluffy white cumulus, which add an element of realism to landscape images.

Lighting in general is an important part of the landscape and perceptions of it. Bright direct sunlight, overcast diffuse daylight, and moonlight through clouds all cast quite different light – not to mention shadows – on the landscape. The vast majority of landscape visualizations (and computer renderings in general) are created with simple, diffuse light, and so rarely capture highlights, shadows, sparkles or other ephemeral and essential visual elements of landscapes. Shadows alone can consume over half of the rendering time for a complex image, and so the urge to suppress them in most visualizations is understandable, but regrettable.

Motion of the atmosphere – such as the wind, whether gentle zephyr or howling tornado – is yet one more phenomenon hard to model, in part because we lack good physical understanding. At best, it requires super-computers and massively parallel computing. In simplified form, wind effects may be simulated, and some software such as Onyx TreeStorm, Mayand 3DStudioMax have various kinetics and particle modeling systems that are useful for modeling dynamics, motion, and effects such as gravity.

Atmosphere is most important in either evocative or realistic visualizations. In the former, because literal atmosphere is often associate with metaphorical atmosphere; in the latter, because real light and shadow areas, for example, may be important to the appearance or function of the landscape being visualized.

3. Conclusion

Major progress has been made in landscape visualization from the early days of wire frame grids and triangular trees. The special challenges of landscape elements – terrain, water, vegetation, and atmosphere – have attracted the research and development community, and the computer graphics literature is rich with early examples, such as George Lucas' seminal image "The Road to Pt Reyes". But many landscape visualizations produced in the process of planning, design and public review still suffer from a failures ranging from overly coarse landform data to poor integration of built structures and implausibly rendered vegetation. Many other visualizations are never made at all, because of lack of knowledge or access to basic resources.

The challenges of landscape visualization arise in part from the sheer complexity of landscapes – in size, in curviness, in fractal dimension, and so on. Some of the problems come from the need to integrate several different sources of material, or techniques – terrain data from GIS, with models from CAD and textures from image processing software, for example. But to the extent that these challenges can be overcome by ever faster computers, larger disks, cheaper RAM, better software and more clever algorithms – which is surely the roller coaster we are all on, for the indefinite future – a number of important challenges remain, due to the various inferential purposes of landscape visualization. A few of these are summarized below.

Terrain and Vegetation Level of Detail ('LOD')

All representations need to be based upon selections about what to put in and what to leave out. Both terrain and vegetation may have inordinately large amounts of data, which must be simplified and generalized, especially when seen from a distance. This is complicated by the emergence of dynamic models (animations, videos, walkthroughs), in which the distance between a viewer and terrain or vegetation and angle of view may be constantly changing. Sometimes you may benefit from seeing the individual leaves of a tree; some time later that tree may be invisible, or reduced to a dot of color in a valley.

Dynamics

All landscapes are dynamic. Different elements change at different frequencies and time scales, from the instantaneous local scale of water rippling, through diurnal cycles of light and dark, to seasonal vegetation change, annual growth, and geologic spans over which terrain can shift and transform. Procedural techniques and particle systems offer considerable promise for including dynamics in landscape visualizations

Figure 9. Still frame of movie – trees blowing in the breeze

Rule-based geo-typical scenes

The task of constructing illustrative visualizations often involves ‘populating’ areas or locations with objects such as trees, houses, etc. The mechanisms and control languages for doing so are rather crude an d ineffective. Typically, random assignment of a certain total number or average density of objects in a polygon are the only two choices. This may work reasonably well for trees in a forest, but is unacceptable f or houses in a sub-division. Asserting such rules as 'the long side of the house typically faces the street', or 'lots at the corner are typically larger', or trees occur in a double row in the median strip spaced 15 feet apart", are all difficult to achieve without detailed custom programming.

Immersive Visualization

Finally, it's worth noting that digital landscape visualization offers potential for new formats unavailable to landscape painters of old. Digital animations include a temporal element; but new projection technologies make real 3-D immersive environments possible too. VRML is a format for representing environments in 3D so that they can be interactively navigated, and multi-projector technologies with curved screens, like the Imax theater familiar to many, offer the opportunity to experience the landscape 'in the round', rather than through a frame. As these technologies become more available and familiar, landscape visualizations will be become ever more rich at the 'realistic' end of the scale

Figure 10. An immersive multi-screen projection system for landscape visualization

While VRML, or '4D or temporal GIS', or rule-based patterns, or techniques form managing LOD are not new topics, the questions of how best to store, analyze, manipulate and portray complex changing landscapes using GIS remain open ones with rich opportunities for exploration and invention.

Acknowledgments:

Some of these ideas and material first appeared in my paper "*Digital Landscape Modeling Visualization: A Research Agenda*", presented at he Ascona conference on "Our Visual Landscape", and published in the May 2001 special issue of *Landscape and Urban Planning*.

Author Information:

Stephen Ervin
Harvard Design School
48 Quincy st
Cambridge, MA 02138
US
617 495 2682
servin@gsd.harvard.edu