# Chapter 5

# Creative Workflows for the Media Artist

## 5.1 Overview

The ways in which artists create art has changed dramatically over the past hundred years. In the 1950s, following the invention of the computer, a group of scientists and artists began exploring the production of images by digital means. From these beginnings, computers have dramatically come down in cost and software tools for making art have become widespread. This has opened up a range of expressive possibilities, and the variety of tools available cover a range of new techniques available to the media artist. These possibilities include low-level programming, information aesthetics, design, interactive art, artificial life, and rendering. This thesis considers the major choices in workflow currently available to media artists and proposes an integrated tool to explore these various dimensions simultaneously.

The goal of this work is to develop new creative workflows for media artists. In 2005, the National Science Foundation sponsored a workshop to consider Creative Support Tools across several different domains. Discussions covered creativity in tools for chil-

dren learning to program, engineers creating novel designs, users navigating the web, and new media artists exploring visual and interactive art [Shneiderman et al., 2005]. They found that, more than other groups, media artists explored "risk-taking and subversive attitudes" to express cultural acts, often resolved by artists developing their own tools, where each of these tools embodies a particular technique for creating art.

The term *workflow* derives from a sequence of steps in manufacturing, a means of production. While this description is appropriate in some ways, as the media artist creates aesthetic objects using technology, the definition of workflow for the purpose of this thesis is broadened to cover all the ways in which media artists explore or express an idea through technology. A *creative workflow* is a way of generating, exploring, or resolving a particular idea through technology. While a programming language might be used to express an abstract sequence of steps, a *workflow* is defined more broadly to resolve all the possible constraints available to the artist through creative choices.

One of the most common choices for the media artist is the programming language being used. According to the definition above, however, the choice of language is just one of the many dimensions along which a creative workflow may be resolved. Some media artists choose to work with a particular language while others develop their own [Reas and Fry, 2006]. However, there are many other dimensions along which a tool can be expressive. A particular tool may support three-dimensional forms better than another, or may be more suitable to interactive art. A primary contribution of this

work is a consideration of specific dimensions of creative choice which are relevant to media artists. The dimensions considered here include:

- 1) Programming and Language
- 2) Modality and Media
- 3) Performance and Computation
- 4) Motion, Complexity and Autonomy
- 5) Structure and Surface
- 6) Image and Idea

A particular *workflow* resolves each of these dimensions through a set of choices. As discussed in Chapter 2, some of these choices may be inherent constraints determined by the tool, while others may be creative constraints resolved by the artist. Workflow is defined as the combination of these constraints and choices across the dimensions above.

The contribution of this work is to conceptualize tools for media art that function along these dimensions simultaneously. The first generation of media artists had no choice but to develop their own tools, since the first computers did not know how to make images [Dietrich, 1986]. This drive to make ones' own tools is now considered an essential part of how media artists work, but should not be an exclusive requirement for making art. Ideally, tools for the media artist should allow programming when desired while also supporting the other available dimensions for expression. One motivation is that this eliminates the repetitive work required in engineering ones own tools, allowing the artist to work at a higher level. Another motivation is that tools that consider the

variety of workflows used by media artists cover a potentially broader expressive range than any one particular tool or language. Thus a central question is how best to support the creative work of media artists.

LUNA, a Language of Natural Animation, is developed to show that these creative dimensions can be resolved together in an integrated tool, and its contribution is to support a wide range of different techniques employed by media artists. LUNA is not necessarily an ideal tool for media artists since every tool introduces its own inherent constraints [Candy, 2007]. For example, LUNA does not necessarily go into great *depth* either in particular structures (e.g. its support for highly detailed characters is limited) or in particular modalities (e.g. its support for audio is limited). I assume these details may be easily added later. Other potential limitations of LUNA will be considered later. Instead, this work focuses on providing a context in which the dimensions above are simultaneously resolved and made available.

### 5.1.1 Motivation

Motivations for this work come from several different directions. My own background in art and computer science involved an exploration of a variety of forms. Early on I studied generative art based on simple rules, imitating many of the early algorithmic systems such as particle systems and Conway's Game of Life. Later on I explored physical sculpture and became interested in digital representations of form, volume, and surface. This led to an interest in rendering and lighting, but also to automatically

generating three-dimensional forms using dynamic algorithms. Unfortunately, while sophisticated tools exist for three-dimensional forms (Maya, 3D Studio), and others exist for algorithmic art (Processing, Max/MSP), I found few tools capable of combining these in real-time systems. In addition, I also began to explore images, video, and other media which I viewed as potential sources for generating unique sculptural forms. The development of LUNA thus captures several dimensions and ideas which I found initially difficult to separate as an artist.

The history of media art, modern art and computer graphics helps to resolve these aspects. Early artist-scientist pioneers, such as Michaell Noll, Ken Knowlton, George Ness and Freider Nake, developed the first computer images using simple shapes [Dietrich, 1986]. Shapes such as lines and curve, Figure 5.1a, are the starting point for several new directions. From this point, the vocabulary of shapes can be extended in three-dimesions to curves, surfaces and volumes, as was done by Bresenham, Pierre Bezier, Ed Catmull (Figure 5.1b) [Foley et al., 1997]. Once surfaces are established, other groups focused on the rendering and lighting of these surfaces, including Bui Phong, James Blinn, Don Greenberg, and Turner Whitted (Figure 5.1c). However, one can also take these basic shapes and study their *motion*, resulting in dynamic objects as explored by William Reeves (Figure 5.1d). With the presence of motion, unique behaviours are explored by Craig Reynolds, James Whitney, and others (Figure 5.1e) [Levy, 1992]. *Interaction* with basic shapes leads to real-time systems as investigated by Ivan Sutherland, Doug Engelbart, and Alan Kay (Figure 5.1f). Treating the digital image itself as a structure,

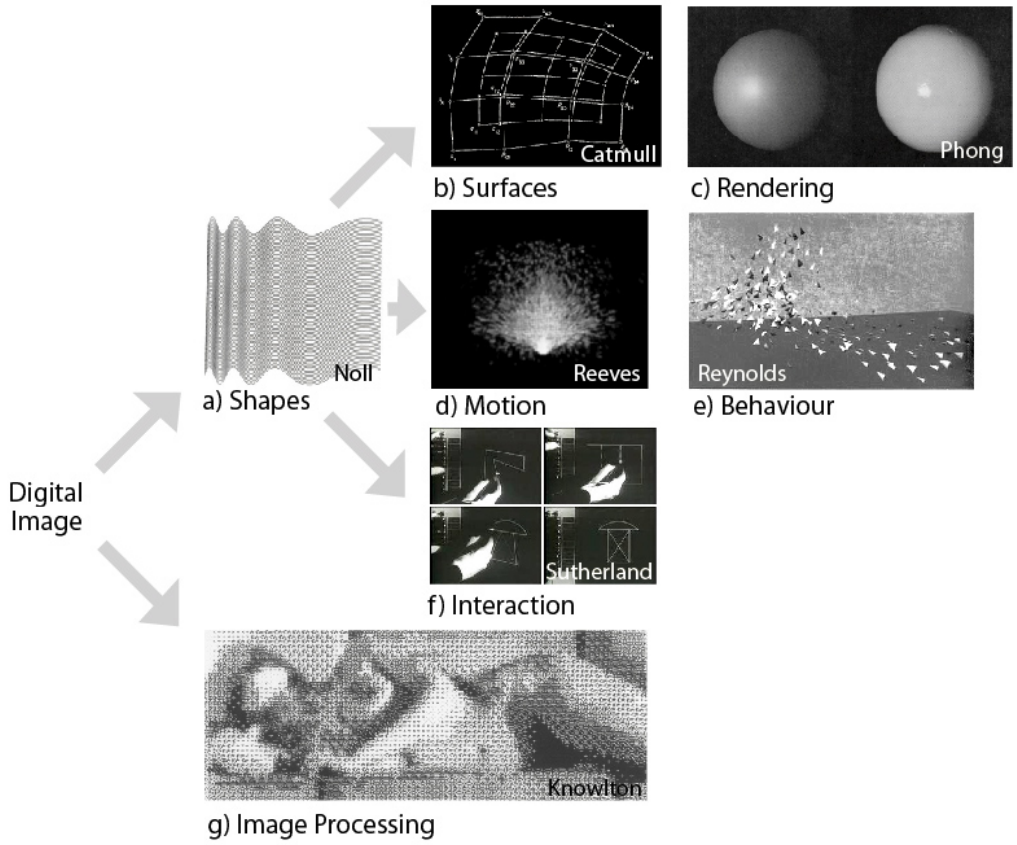


Figure 5.1: Early history of the digital image, including a) shapes (Michael Noll), b) surfaces (Ed Catmull), c) rendering (Bui Phong), d) motion (William Reeves), e) behaviour (Craig Reynolds), f) interaction (Ivan Sutherland), and g) image processing (Knowlton).

in which shapes are contained, results in the study of *image processing*, Figure 5.1g. [Rosenfeld, 1983].

These fields define the basic ways in which the digital image may be manipulated while, over time, their separation has resulted in distinct communities. Modern motion pictures and video games take advantage primarily of three-dimensional modelling, lighting and rendering, relying on technical artists to manually create characters and

textures to support a particular narrative. Algorithmic artists such as Harold Cohen, Jean Pierre Hebert, Roman Verostko and Vera Molnar, have explored programmatic, non-representational aspects of image making by looking at motion and behaviour [Verostko, 2006]. Information artists work with the database as a source for the layout of shapes and forms. Similarly, interfaces in film and gaming change very gradually due to their means of distribution, while interactive artists such as Ken Feingold, Golan Levin, Michael Naimark, and Simon Penny tend to incorporate or even invent new interfaces regularly in installation works.

A primary motivation for this thesis is to develop workflows that allow these different practices to be combined in a single tool, to demonstrate that techniques among these communities may be shared. In LUNA, for example, one goal is to incorporate modelling and rendering typically found in animation software and to combine these with processes for algorithmic, interactive, and performance art. Although the dimensions covered are not designed to address media arts completely, for example aspects of database and web art are not considered, the goal is to bring together several areas of digital art which have evolved into distinct tools by integrating the dimensions along which these choices are made.

This list of dimensions appears to share some similarities to studies of formalism in art history. For example, according to Hatt, in the *Principles of Art History* Wolflinn presents formal properties "meant to provide general descriptive terms, which would capture the development of artistic vision across countries and ages."

[Hatt and Klonk, 1992] [1]. Roger Fry develops a familiar set of formal elements in observing line, mass, space, light, color, and plane [Fry, 1926]. However, a critical difference from these and the above list is that the formalist elements in art analysis are based on style in order to understand art created using similar techniques across civilizations and periods of time, not only to distinguish the works of art but to offer explanations for why and how they were created by cultures in time [Preziosi, 1998]. While I also seek to understand the historic aspects of media arts the dimensions proposed here are based essentially on technique itself, and my primary contribution in not a reflective analysis of works in media art but a workflow for integrating its techniques. Although many of the design choices in LUNA are inspired by historic works of art, these dimensions relate to choices in technique made by the artists during the creation of the work rather than styles derived from looking at art after the fact.

When a sufficient number of examples are considered over time, it may be possible to analyse style in a specific area, such as organic art for example. Steven Levy, in his book Artificial Life, explores some of the historic aspects of this type of art and studies the culture in which they developed [Levy, 1992]. Christian Paul, in Digital Art, surveys several of the active areas of media art [Paul, 2003]. Although it is difficult to distinguish style from technique since both are constantly changing in media arts, such a study could be an interest area for future research.

[1] Wölfflin's five dualities include 1) Linear versus Painterly, 2) Plane versus Recession, 3) Closed versus Open, 4) Multiplicity versus Unity and 5) Absolute versus Relative Clarity

Each of the current areas of media arts offers a different perspective on the construction of the digital image. Language, computation, modality, autonomy (behaviour), structure, image and idea are the primary dimensions which are considered here from the perspective of technique. The cultural significance of the digital image is not addressed here, nor the social effect of these abstract and scientific ideas. In addition the impact of these techniques on society, which is the study of media theory, is generally not addressed; although certain aspects of structuralism are relevant in examining technique. Finally, it is useful to mention that this is not a goal-oriented process since the creation of a tool for media artists can never have an ideal or final form. Rather, the goal of this work is to propose an integration of several of the dimensions of digital image making which have diverged into separate disciplines over time, and to show that it is possible to invent new tools which combine these. The motivation for this work is to bring together communities engaged in creative image making by reducing the natural boundaries that form between tools evolving in particular domains.

### 5.1.2 Evaluation

The evaluation of tools to support creative tasks can be difficult. The most immediate form of evaluation of any tools is that it achieves the functionality it claims to. With regard to the LUNA, the six dimensions above should be realized through specific examples found in the system. Since each of these are treated as a dimension along which choices are made, I consider each dimension as a collection of at least two

examples showing a range of behaviour. However, my argument is not only that tools can be developed to generate specific examples, but that their integration will provide better overall support for creative workflows among media artists. What I mean by better support is a general increase in expressiveness, creativity, and ability to explore interesting areas. How would this be evaluated?

Among the outcomes of the Creative Support Tools workshop was the discovery that creative tools cannot be easily evaluated for the quality of the outcome. It is impossible to say which tools support creativity to a greater or lesser extent:

> "One important issue with the design of creativity support tools is how they can be evaluated. How do you know if a tool is being helpful? Human-computer interaction professionals are used to measuring the effectiveness and efficiency of tools [for specific goals], but how do you measure if it supports creativity? As discussed above, tools that are not effective and efficient will probably hinder creativity, but it isnt clear that the reverse will hold. To try to measure creativity, the Silk system designers evaluated many different properties, including the number of different designs produced, the variability of the components used, the variety of questions about the designs from collaborators, etc. [Landay 1996], but these still do not really get at the *quality* of the solution. It is still an open question how to measure the extent to which a tool fosters creative thinking." [Resnick et al., 2005]

While the authors explored the issue of evaluating creative tools in other domains, this finding is especially relevant to media artists since the outcome of the tool is itself a subjective work. Thus there is the added complexity that the quality of the tool cannot be distinguished according to the objects it produces on a creative level. Fortunately, the workshop participants found three criteria which were agreed to be a good relative measure of creative tools. These are: 1) low threshold, 2) high ceiling, and 3) wide walls. Low threshold means "that the interface should not be intimidating, and should give

users immediate confidence that they can succeed", high ceiling means that "the tools are powerful and can create sophisticated, complete solutions", and wide walls means that "creativity support tools should support and suggest a wide range of explorations" [Resnick et al., 2005].

How do these concepts map to creative tools for media artists? While we can agree the visual interface should be simple to use, low threshold may also be taken to mean that the underlying language expressed by the tool is also simple since media artists may wish to work either in a programming language or the visual interface along the dimension of *language*. The idea of 'high ceiling' can have several interpretations. Does this mean powerful in its ability to support different modalities (images, sound, etc.), powerful in the structural detail it can achieve, or powerful in the processes it can perform? I will consider each of these as they arise through examples. Finally, there are several levels on which artists can explore a range of expressions. This may be through the parameters to a particular system, parts of a particular object, or the combination of these systems. Despite distinctions which must be resolved during their evaluation, these criteria provide useful guidelines for evaluating creative tools relative to one another.

LUNA is proposed as an integrated tool for media artists, which suggests that its greatest improvements will occur over time with the active participation of a community of users and developers.[2]. In the current design, a number of features are presented as

[2]To enable this, the core of LUNA will be released under the open source LGPL license for both Windows and Linux, with the interactive editor released as a free executable for personal use (with dynamically loaded user modules).

potential directions. For example, along the dimension of modality LUNA has icons for audio, video and data. In these areas integration has been considered and designed into the system so that practical details may be completed more easily in the future. During the evaluation of LUNA I thus distinguish several areas of development: 1) inherent limitations of the language, 2) potential ideas which were not explored in any way, 3) partial features which were considered and designed into the basic language of LUNA for future growth, but not fully implemented, and 4) features which were completed.

The methodology used in this study is to examine each dimension, to explore its significance for media artists with historic examples, to see how well LUNA integrates that dimension into the overall workflow of the tool, to evaluate the criteria of low threshold, high ceiling, and wide walls, and to evaluate the system according to the actual and potential choices it makes available to the artist.

## 5.2 Programming and Language

### 5.2.1 Procedural Languages

The first two exhibits of computer generated images showed works by Bela Julesz and Michael Noll at the Howard Wise Gallery in New York in 1965 (Computer Generated Pictures), and by George Nees and Frieder Nake at the Galerie Niedlich in Stuttgart, Germany the same year [Dietrich, 1986]. All of the early pioneers were also scientists working in institutions such as IBM and Bell Labs in order to provide access to the large, costly computers need to make these images. For the first generation of computer

artists, direct programming of the computer was a necessity. Michell Noll attempted to simulate constructivist and abstract works of art using very basic mathematical shapes, in images such as *Bridget Riley's Painting Currents*, 1996 where he reproduces the op art of Bridget Riley. Using the density of type to reproduced tone, Ken Knowlton and Leon Harmon created *Studies in Perception I*, 1996, the first digitized and reproduced nude figure. Initial collaborations between artists and engineers took place in Cybernetic Serendipity, an exhibit at the Institute of Contemporary at in London curated by Jasia Reichard [MacGregor, 2002]. Computing at this time had a very high threshold, and artist-scientists were required to program in low-level languages using basic mathematics.

To facilitate making graphical images scientists began by developing extensions to generic programming languages. George Ness and Leslie Mesei added graphics commands to ALGOL 60 and to Fortran in 1969. Eventually, work by Ken Knowlton and others led to more complete graphic languages. Although they simplified the work for scientists, they were still found to have a high threshold for artists.

> "As an animation language it provided instructions for several motion effects as well as for camera control. Knowlton had initially hoped that artists would learn the language to program their own movies, but he came to realize that they usually wanted to create something the language could not facilitate, and they also shied away from programming... None of the graphics languages mentioned received widespread use, partly because their implementation was machine dependent and also because each language was restricted in scope." [Dietrich, 1986]

Interactive sketching introduced by Ivan Sutherland in 1963, which would allow for direct drawing of shapes using a pen, would not be widely available for another

decade [Sutherland, 1963]. For the algorists in the 1970s, these constraints were not a major barrier as the artist's interests contained a strongly constructivist element, so the use of algorithm and mathematics coincided with their creative goals. For artists such as Harold Cohen, Roman Verotsko, and Manfred Moher, programming became the abstract language "through which they created a new reality" [Verostko, 2002]

In a relatively short time, computing reduced in cost and better graphics hardware became available. Through the 1980s graphical languages proliferated. GINO, Graphical Input/Output System, abstracted the concept of logical output and input devices. Languages such as GPGS and PHIGS supported hierarchical object arrangements, allowing for descriptive scenes [van Dam, 1998]. In the 1990s, OpenGL and DirectX were the first languages to be widely used on new graphics hardware, built first as extensions to the generic C/C++ language. Graphical languages finally became widespread in the mid 1990s.

Programming directly in text-based languages continues to be an important way to create visual images as these languages can allow the artist greater control over the specific rules used for indicating new behaviours. To reduce the high threshold of learning full programming languages, Casey Reas and Benjamin Fry introduced the Processing language to promote image making as a way to teach programming in 2001 [Reas and Fry, 2006]. A sample program in Processing is shown in Figure 5.2. In LUNA text-based programming in C/C++ allows authors to create new interactive nodes, discussed in further detail with the introduction of visual languages below.

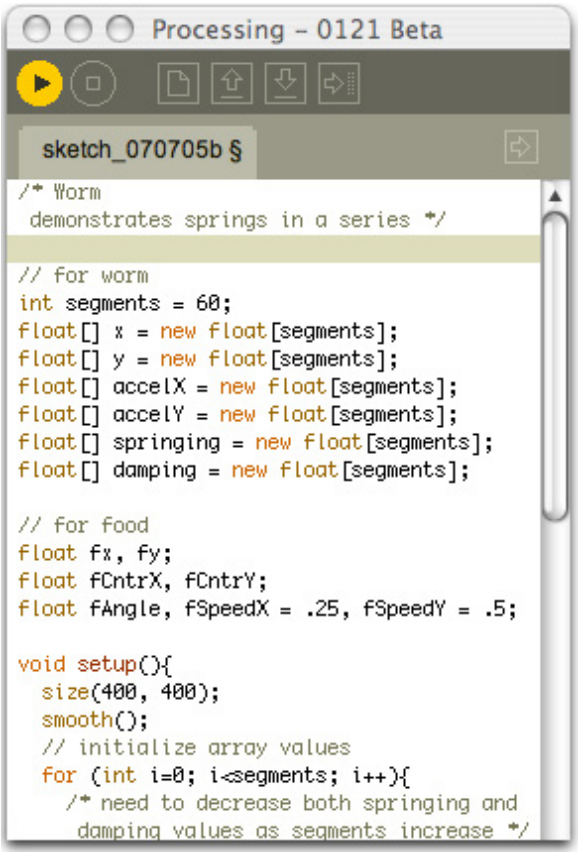


Figure 5.2: Text-based authoring environment in Processing.

Graphics libraries are now available in many other generic text-based languages such as Python, C/C++, Flash and Java (on which Processing is based). Although the use of text-based languages continues to grow as more media artists learn to program, the threshold for achieving complex forms using these languages is still high. While it may be that media artists should learn programming as a part of an education in media arts, it does not necessarily follow that all media artists wish to create exclusively by algebraic or procedural programming.

### 5.2.2 Visual Languages

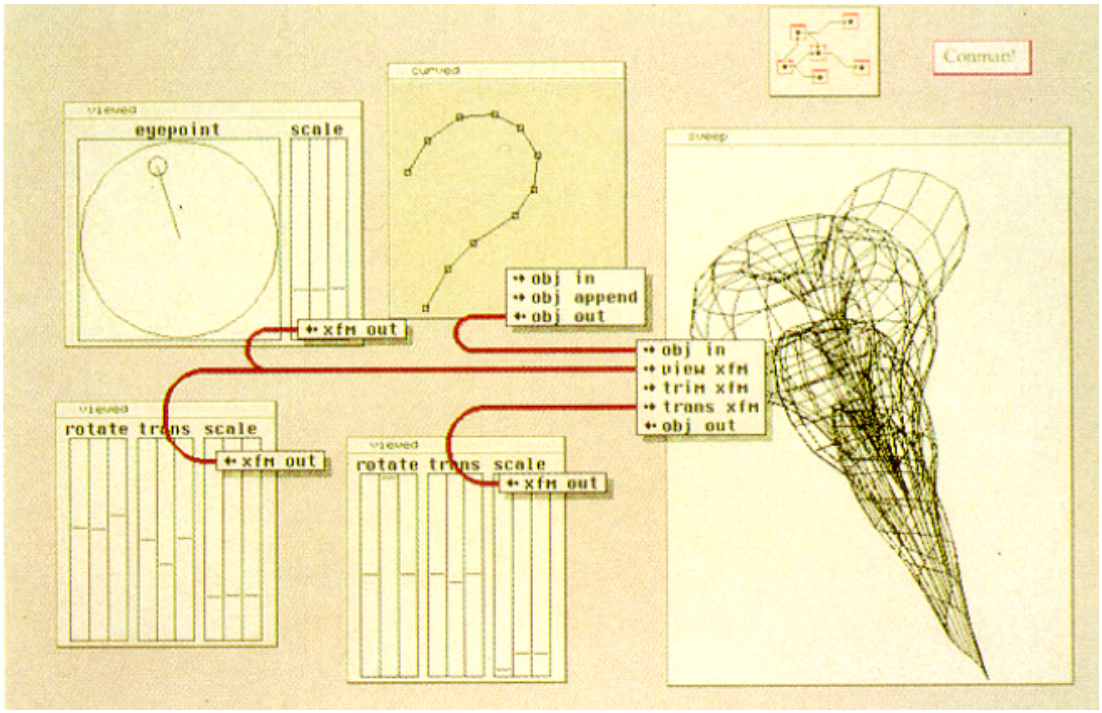


Figure 5.3: Con Man, an early visual language for graphics by Paul Haeberli in 1988.

The development of Ivan Sutherland's Sketchpad in 1963 led to the first generation of direct interaction technologies. Drawing programs allowed designers and animators to manipulate images manually. In 1968, Ken Pulfer and Grant Bechtold of the National Research Council of Canada created the first hand drawn computer movie entitled *Hunger* by drawing each frame using a wooden mouse. "Markup" and "Superpaint" were the first drawing programs by William Newman and Dick Shoup from research based at Xerox PARC. Myers surveys a number of other parallel developments in early computer interaction [Myers, 1998]. While drawing and sketching are the first obvious uses of direct interaction with machines, visual interfaces for computational tasks

began with visual programming languages (VPL). Early visual languages had similar constructs to text-based languages, and contained iconic representations of mathematical operators, loops and conditionals. A key concept which distinguishes VPLs from text-based programming is the introduction of *nodes* which allow logical operations to be encapsulated and modularized, so that data is viewed as *flowing* through a visual representation of a set of tasks. An early system which experimented with visual languages to perform image processing is ConMan [Haeberli, 1988]. A survey of generic visual languages for goal-oriented tasks, such as ARK, VIPR, Prograph and IBM Data Explorer, can be found in Boshernitsan and Bownes [Marat and Downes, 2004].

The idea of the media artist as programmer is an on-going trend, so the need for visual languages for artists is relatively new. Interactive artworks, such as *A-Volve* by Sommerer & Mignonneau (1994), use programming to achieve a particular, dynamic relationship between the art work and the viewer. This usually requires artist/engineering collaborations since the programs are specific to the installation. A generalized approach to media arts would be to treat each *type* of interface device as a building block with which an exhibit is constructed. Edmonds et al. explore the idea of what tools might be best for media artists:

> "A fundamental question that we have been considering is, what kind of environments best support the development of digital art? There is one answer to this question which, although it may sound a little strange, is, nevertheless, appropriate. In art and technology environments, we need environments for building environments." [Edmonds et al., 2004]

An interactive system mentioned by Edmonds is Max/MSP, created by Cycling '74 and originally written by Miller Puckette (author of PureData), Figure 5.4. Designed for audio synthesis, Max/MSP is a visual language that can generate interactive art using an extension called Jitter. One key contribution of Max/MSP/Jitter is that it allows for many different devices to control audio-visual events by treating all data in the system as a *signal*. David Wessel and Matthew Wright add support for gestural interfaces to Max/MSP by creating Open Sound Control, a communications protocol that allows signals and events to move from device to synthesis, even between remote computers [Wessel and Wright, 2002]. This enables the media artist to work with many different *input devices*, discussed in more detail in the next section.

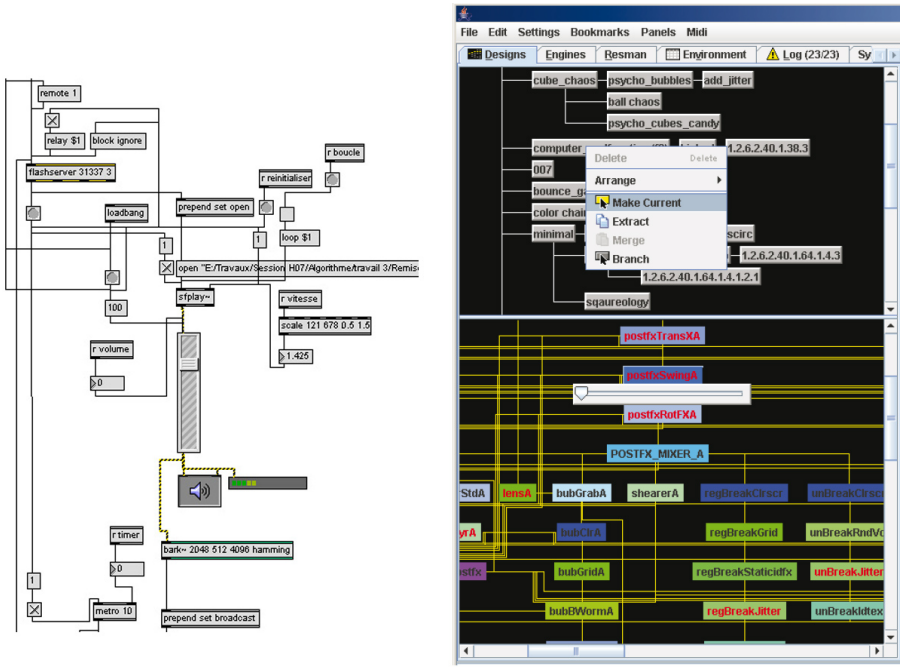


a) Max/MSP, Cycling '74 b) Soundium, Pascal Müller et al.

Figure 5.4: Two visual languages for media artists, Max/MSP and Soundium.

For artists interested in three-dimensional visual forms, the generalization of data as signal may present some problems. First, three-dimensional forms are represented by computers in a particular way, by using verticies, edges and faces (one possible way), so they are not easily encoded as *signals*. For example, although it is possible to represent a three-dimensional tree structure in Max/MSP, the designer must author a special object to 'encode' each type of geometry. In addition, this encoding requires that the artist-user must continually remember the type of the data as it flows through the system. Finally, the signal processing metaphor results in a visual language which is low-level, consisting of operators and expressions as found in generic programming languages. Although this makes it expressive, it also requires mathematical thinking to understand a Max/MSP patch.

Several top-level decisions drove the design of LUNA, with a special emphasis placed on designing an intuitive interface for artists unfamiliar with mathematics or logic. This is embodied in several workflow principles during the project: 1) creative expressiveness in the interface should not require mathematical expressions or logic, 2) the interface should express high-level concepts first, and particular details only on demand, 3) the language should be capable of complex structures *and* multimedia.

These constraints, especially the first, led to a visual language inspired by the board game Scrabble, which achieves a huge combinatorial variety in word letterings purely through the relative placement of tiles. This suggested a node-based workflow incorporating large iconic tiles with a minimum of extraneous information (i.e. no numbers

or values on the graph). The goal of creating a high-level interface helps the design by supporting the idea that each node should represent a particular aesthetic task. A key design of the language which led to its implementation was the discovery that nodes can represent both their *structure* and *behaviour* in a way that leads to natural extension of the language. The structure of the data is carried from node-to-node in the graph, while the behaviour determines how each input is processed. In fact, these two ideas are the only textual information presented in the graphical interface. The icon itself expresses a pictographic idea of the *behaviour* that will occur, also written in large type above the tile, while the structure of the output is shown in smaller type, depicted in the overall color of the tile.

The graph design of LUNA is not the only way to express structure. In computer graphics the relationships between objects, their physical proximity and orientation in two or three-dimensions, is commonly represented using a *scene graph*. The scene graph became widely used after IRIS Inventor in 1993 to describe scenes consisting of different objects [Strauss, 1993]. Separately from the audio-signal processing community, the graphics community found that objects in scene graphs can also be expressed in visual languages to facilitate on-screen interaction. Objects in these visual languages, such as Maya's Hypergraph, represent geometric relationships and transformations [Bar-Zeev, 2007]. Procedural systems such as Houdini express both behaviour and structure, but the language does not make it clear what these structures are, and may require a series of intermediate steps to change geometry type (see Chapter 4).

A recent approach to this duality is found in Soundium/Decklight, developed by Corebounce [3].

> "There is a problem we have not dealt with: As a result from unifying multiple graph-based processing entities in a single graph, we have to deal with different graph processing semantics: For example an audio graph, which is typically flow-based, is processed in a different manner than a scene graph (which is basically an object hierarchy). The question is, how we deal with the coexistence of different semantics in the same graph? Our approach is to segment the global graph into individual subgraphs, which correspond to different semantics. Of course, this segmentation will not be made visible to the application layer." [Arisona, 2007]

The authors introduce a *design tree* to express high-level artistic ideas which are used to generate these processed sub-graphs [Muller et al., 2007], see Figure 5.4. This workflow enables simultaneous audio and live visuals (video), interactive editing during a performance, with examples that focus primarily on transformations such as rotation, translation, and scaling that are common to scene graph languages. This approach to multimedia, while it resolves many challenges in simultaneous audio-visual processing, may make it difficult to address behavioural changes in more complex three-dimensional forms.

The scene graph problem is resolved in LUNA at a high level by allowing each node to contain its own scene graph, while the language conveyed through a combination of nodes expresses dataflow. Unlike Max/MSP the structure of each object is apparent in the tile color as information moves through the graph. Thus each single tile in LUNA represents not just one object, but arbitrarily many, which allows the system to

[3] Pascal Müller, Stefan Müller Arisona, Simon Schubiger-Banz and Matthias Specht, from ETH Zurich, University of Fribourg, and University of Zurich

express complex jointed or articulated structures (a scene graph) while also permitting multimedia processes which operate on these. Due to the storage of these structures (see Chapter 4), modifications to color, position, and orientation can be made at any point in the graph.

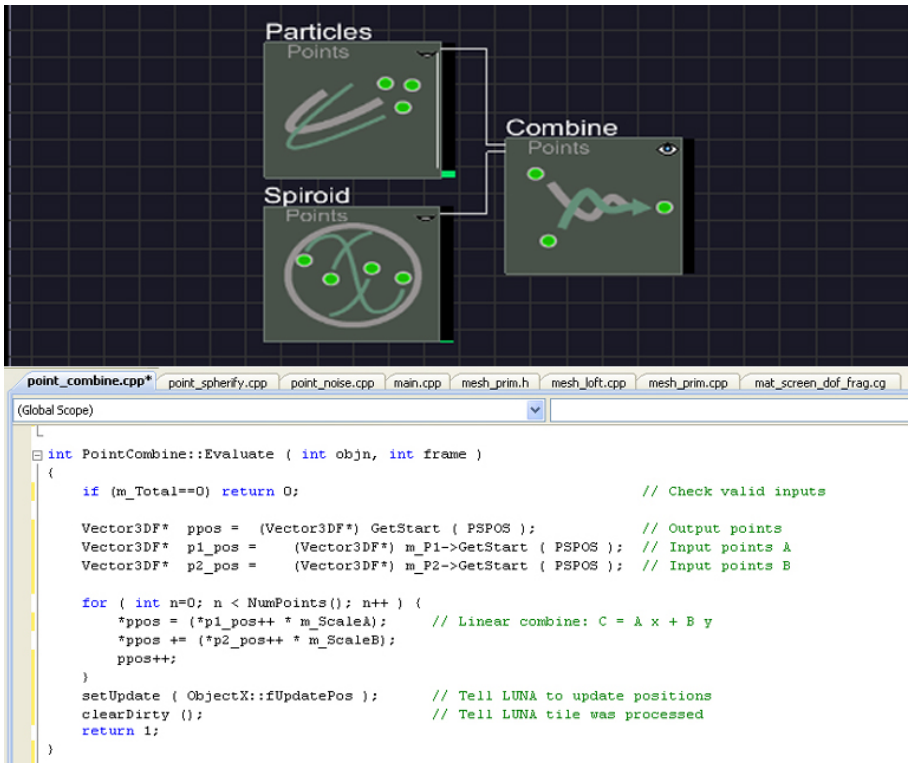


Figure 5.5: Text-based authorship of a point Combine node, and example of its use in the LUNA visual interface.

Users of LUNA may interact with programming at two levels. The first is by authoring new nodes in the text-based language C/C++ to create new fundamental behaviours. Although Processing was designed to have a lower threshold for text-based languages, node authorship in LUNA is simpler than generic C/C++ on which LUNA is based since it provides all the data structures needed to create complex geometries, Figure

5.5. The second means of creative interaction is through the visual language by mixing and combining existing tiles. Using the visual interface may have a lower threshold than any other language currently available to media artists since only a single click-drag motion is required to connect two tiles. [4] The system was even shown to a ten year old, who created complex systems simply by matching up input and output colors of the tiles. From the perspective of creative workflow, this low threshold is achieved through the design constraint of supporting non-technical users.

## 5.3 Modality and Media

Media artists desire to work in a wide variety of ways using different kinds of media, and different interface devices. This dimension of creativity has two particular aspects. *Media* may be defined as the structure of data, while *modality* is defined in human-computer interaction as a physical system or device that generates a particular media [Bolt, 1980]. For example, a video camera is the modality which produces the media of video. While this may seem obvious, devices like video cameras produce both audio and video, while a music keyboard can generate media which is either audio (a signal) or midi (a sequence of notes). In addition, software may process many types of media, or transcode one media into another [Manovich, 2001].

Max/MSP resolves the issue of media by treating every media type as a signal. This simplifies the base design of the software, and focuses attention on the audio signal,

[4] At present, node authorship requires rebuilding of LUNA, although this is expected to change soon as new versions will allow dynamic linking.

but places a burden on patch developers to handle media other than audio. This is partially alleviated by a large development community, but support for fundamental media types such as meshes or materials is difficult as the community must agree to a library standard.

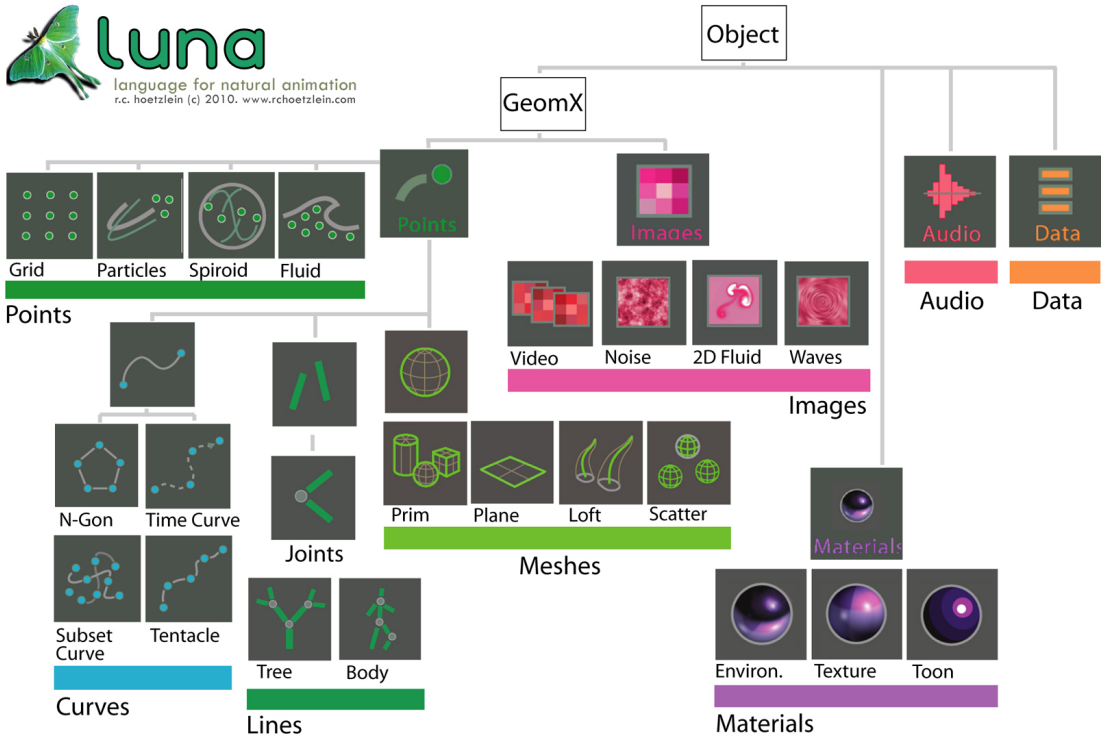


Figure 5.6: Taxonomy of classes and currently implemented nodes in LUNA.

The design principle for media in LUNA is that each node, individually a) knows what it is, b) knows what it requires, and c) knows what it produces. For example, to make a forest one must know where to plant the trees (point locations), and what the trees should look like (joint structure). To use a computer mouse to change the brightness of an image, one must know the position of the mouse (a point), and the

input image. Figure 5.6 shows a hierarchical map of different media in LUNA. Notice that trees, characters, meshes, and curves all have points as a base class, which implies that any process that operates on points - a bend or twist for example - can operate on all of these objects. There is an interesting similarity between this taxonomy and the early history of the graphics in Figure 5.1, which further supports this architecture as a way of manipulating the digital image in different ways. Although many higher level media types such as audio and events are not yet implemented, as indicated in the figure, this vocabulary of data structures provides a great deal of flexibility.

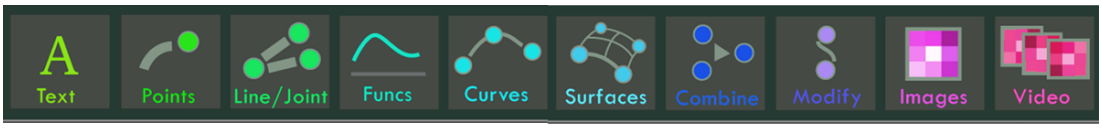


a) Primary toolbar in LUNA specifies media type

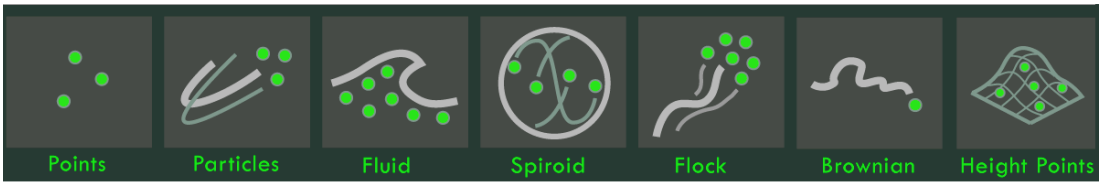


b) Secondary toolbar in LUNA specifies behaviour

Figure 5.7

The primary toolbar in LUNA selects among different media types (structures), while the secondary toolbar provides a choice of processes for generating that particular type (behaviour), as shown in Figure 5.7. Attention was placed on geometric media types, such as points, lines, curves and meshes, to support the author's interest in sculptural form in real-time systems. Choices made in implementing particular processes are discussed in the next chapter. One key goal of the media structures at this level

is the ability to distinguish between disconnected points, articulated structures and surfaces, a feature common to graphics systems for modelling three-dimensional forms such as Houdini, but not currently found in frameworks for real-time multimedia such as Max/MSP or Soundium.

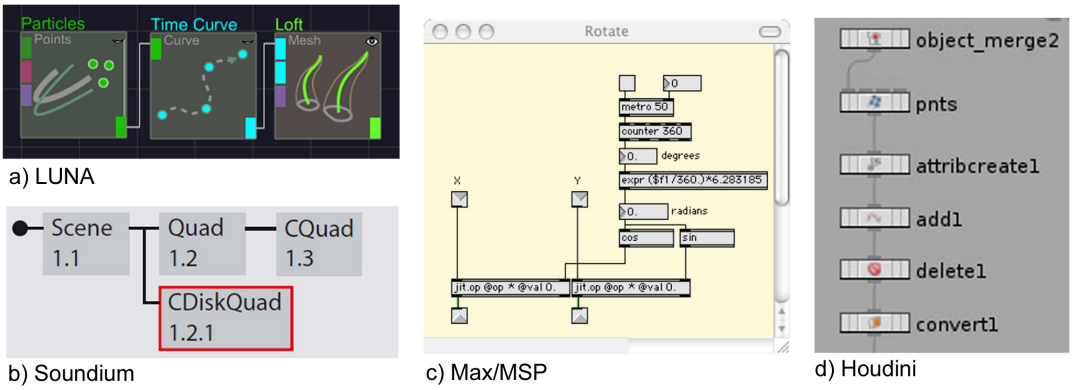


Figure 5.8: Detailed view of connections in various visual design platforms. Colored tabs in LUNA indicate changes in the type of media as it flow through the system.

A unique aspect of LUNA is that inputs are handled by each object. Unlike other media frameworks the type is not fixed by the base system but enforced by individual nodes, thus LUNA may be described as a loosely-typed visual dataflow language [5]. Figure 5.8. shows how inputs are configured in Houdini, Max/MSP, Soundium, and LUNA. What is the type of media flowing through each example in figure 5.8? Color indicates type in LUNA and this is reflected in the design color of each node, making it easy to identify what media is needed for a particular tile to function. This design

[5]The term loosely typed comes from generic text-based languages that do not require the type of a variable to be explicitly stated.

lowers the threshold for use of the system by visually re-enforcing the grammar of the language.

The data types supported in any computer language are a basic part of its specification. Along this dimension, LUNA allows the artist to work with lines, points, curves, surfaces, images and materials (shaders) better than other real-time media systems, already providing a high ceiling in terms of potentially realized structures. Other types such as audio, database input, and volumetric data, are considered by the architecture and may be implemented at the level of text-based authorship by a community of developers in the future. For users of the visual interface, low threshold and wide walls are fostered by the tile design and colored tabs which indicate media type as one works.

## 5.4 Live Performance and Computation

The classical model for the exhibition of the image is the museum, a place where final images are presented as objects of appreciation. More recently, film and gaming have partially displaced the still image with dynamic and interactive forms as a way of communicating with the viewer, while media artists in the traditional of disc and video jockeys explore live performance as a venue for new experiences. For visual art, live performance was made possible by new developments in computer graphics such as the Graphics Processing Unit (GPU), a computer chip dedicated to the function of rapidly generating images. In the traditional model of computing, performance was based on CPU clock rate, which meant that the extent of calculations that could be

performed was directly related to computing power, which limited real-time interaction to all but the simplest tasks. With the advent of the GPU calculations can be performed in parallel so that computation is no longer limited by clock rate.

One of the implications of GPUs, which are now being used for many other tasks besides graphics calculations [NVision], is that computing resources can be targeted toward several goals within the same system simultaneously rather than treating each step as a sequence [Parallel Processing for the Masses]. The implication of this for media artists is that live performance is no longer dependent on how much computing power one has, it is a dimension of choice along which the artist may *focus* various computing resources. This concept of managing computing power is common in the gaming community, where there is a finite budget for graphics, audio, game play, and interaction which must all be realized in 1/30th of a second [Gaming timing pie chart].

Historically, the use of computing power is one of the fundamental distinctions between different types of graphics tools. Systems such as Maya, 3D Studio MAX, and Houdini all interface with offline rendering software capable of producing very detailed, accurately illuminated images using as much time as needed, while live performance tools such as Max/MSP, VVVV, and Soundium use computing resources to perform simpler tasks in a very short time (real-time). In the future, the introduction of GPU computing may allow large computations to also be performed very rapidly, which implies that future tools could reduce the distance between live performance and detailed imaging.

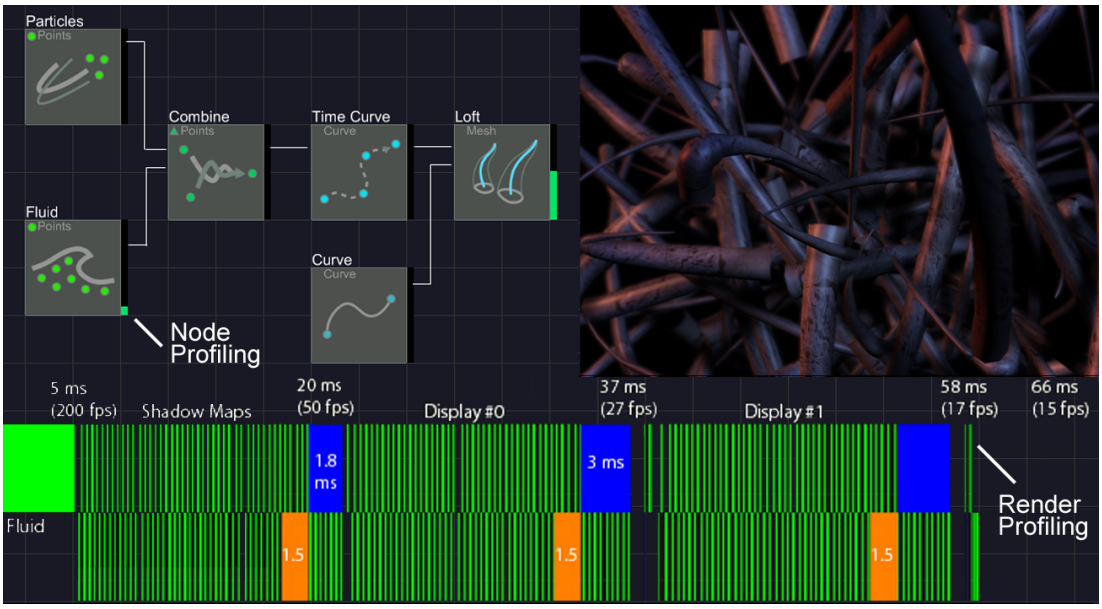


Figure 5.9: Profiling in LUNA. Node profiling shows CPU/GPU resources used for each object in the user graph. Render profiling shows the resources used by the CPU for computation and GPU for rendering.

LUNA introduces several novel features to encourage this convergence. First, rendering in LUNA makes heavy use of the GPU to perform deferred shading, a real-time rendering technique capable of achieving realistic results that were previously only possible with offline methods [Deferred shading]. Secondly, the visual language includes dynamic *profiling*, a technique for measuring computational load. Node profiling shows, using vertical bars, how much computing power is being allocated to each node in the LUNA graph, Figure 5.9. Render profiling shows how time slices are allocated to the CPU and GPU. The vertical green bars represent computation of a each tube shape in the image, orange is a transfer of data from CPU to GPU (and thus lost time), and blue

represents rendering on the GPU (making of the image). These blocks are repeated three times to compute shadow and output for a two-screen image.

To my knowledge, LUNA is the first system for media artists that gives immediate feedback on how computing resources are used. These profiling results already suggest several improvements to the system. For example, the repetition of vertical bars indicates that too much time is being spent on geometric calculation of the Loft node, making this an ideal candidate for authoring this process on entirely on GPU. Some nodes already support this, such as the Fluid system, which can run on either the CPU or the GPU [Hoetzlein].

More importantly, the artist has direct control over how computing power is allocated to different details in the image. Each node supports the concept of a *maximum count*, a set limit on how many objects it will process regardless of how much data it receives. In this example, even though the Fluid system is simulating 4000 particles to the drive the motion of the tubes, the Loft node is generating only 200 tube surfaces, restricting the flow of data to a manageable level while keeping the output interesting. At present, the artist controls the maximum count of each node through the interface.

Although several newer platforms, such as Amira, are offering high performance solutions for interactive visualization of large datasets, see Figure 5.10, one major constraint of artist systems such as Maya, Houdini, Max/MSP or Processing is that it is fairly easy to *stall* the system, making it appear to crash, by simply asking it to compute more than it has power to at an interactive rate. In Maya, for example, creating

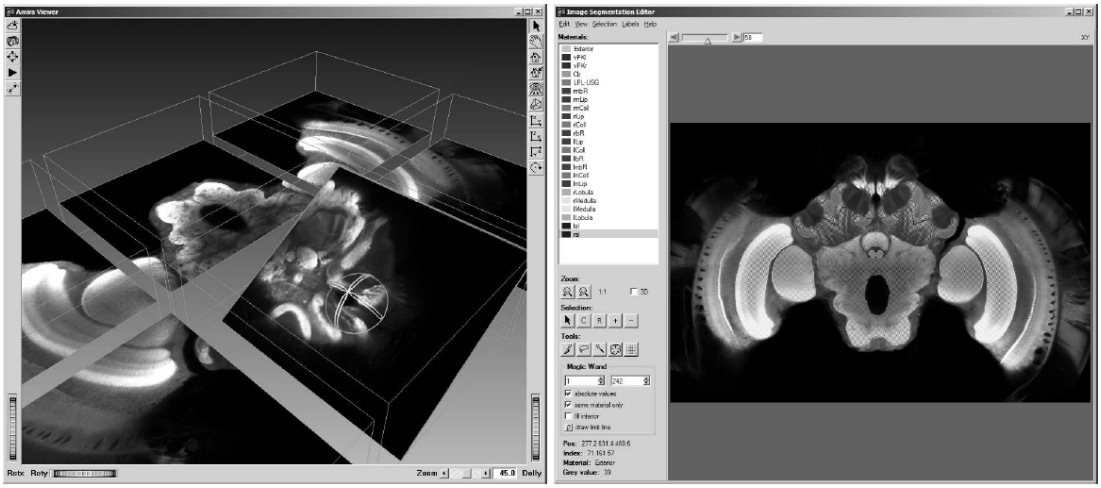

Figure 5.10: High-performance computing in Amira, a system dedicated to interactive manipulation and visualization of scientific datasets.

a particle system of more than 100,000 particles can cause this. In the future, due to LUNA's profiling design, it will be possible to have LUNA itself dynamically adjust maximum counts to automatically determine scene detail so that the system never halts in this way. Since the maximum count can be controlled, and the profiler can calculate the time required for a given node, the rendering system could create a feedback loop where the detail of the slowest object is reduced dynamically.

With the advent of parallel computing using GPUs, it is likely that live performance by media artists will change dramatically in the next decade. In addition to providing tools that give the artist direct control over allocation of resources, LUNA is demonstrably faster than other systems at certain tasks (see Chapter 5) due to its dataflow architecture.

## 5.5 Summary

In retrospect this chapter has focused on three dimensions of media frameworks that *do not* relate to the content of the system. Language syntax, data type (media), and performance establish the basic grammatic rules in which meaning can be *potentially* realized by any computing language. In a traditional view of the media application these rules are found together with specific tools. For example, the syntax of Photoshop is the image, while its operations are all image processing tasks. However, with the development of languages for interactive multimedia the syntax may be so broad that the range of tasks is continually changed by future users, at which point its output cannot be predicted by the inventor since it becomes an open system. At present LUNA generates shapes and tubes with a particular style, but this is because of short-term content decisions, discussed in the next chapter, rather than due to the visual language itself. Interestingly, from the perspective of evaluation criteria, the only criteria directly affected by the 'open ended' aspect of the LUNA language is high ceiling. The ceiling, or expressive power, for a visual language cannot be known in advance since the vocabulary can continue to evolve.

This still leaves a question: How do we evaluate languages for media artists? First, future content and changes in style should be supported by allowing node authorship (making new nouns) in addition to visual authorship within the language (making noun phrases), as is the case with LUNA. Secondly, what is the threshold for using the system? How easy is it to learn the language? The constraint of a non-mathematical interface

in LUNA's visual language led to choices for a minimal design, specific use of color, and simplicity. Finally, how flexible are the range of ideas one can explore? How easy is it to change ideas? Wide walls relates to the flexibility of the language, which is realized in LUNA through the combinatoric connections between tiles.

Media frameworks may also be evaluated according to their limitations. A simple guide in relation to the goals of media artists concerns modality. Max/MSP and Soundium are limited in their support for three-dimensional forms. LUNA supports such forms but is limited in its support for audio, although this is an area which was designed to be expanded in the future. A deeper question relates to constraints of the language itself: Despite the community authorship of objects, what content might be implicitly unavailable in the system? In media frameworks this is usually not discovered until the vocabulary has been sufficiently explored by others. For example, in the future LUNA could handle complex geometries, real-time displacement, audio, and volumetric data as a source of input through community authorship. Although the following is a guess, the inherent limits of the LUNA language may reside at level of co-dependence between two complex systems, such as a tree structure growing along an abstract geometric surface which is itself dynamically changing. Suffice to say that LUNA achieves the workflow goals of low threshold, high ceiling, and wide walls based on the current interface design without knowing where its ceiling lies. In the future, as with any language, it is hoped that the system continues to evolve beyond the initial content defined by the author.

# Bibliography

[Ades, 1994] Ades, D. (1994). *Andre Masson*. Ediciones Poligrafa, S.A. Rizzoli International Publicaitons.

[Arisona, 2007] Arisona, S. M. (2007). Live performance tools. Digital Art Techniques. ACM SIGGRAPH 2007. Course Notes.

[Bar-Zeev, 2007] Bar-Zeev, A. (2007). Scenegraphs: Past, present, and future. http://www.realityprime.com/scenegraph.php, accessed June 2010. Online article.

[Benjamin, 1936] Benjamin, W. (1936). The work of art in the age of mechanical reproduction. zeitschrift fr sozialforschung.

[Bolt, 1980] Bolt, R. A. (1980). Put-that-there: Voice and gesture at the graphics interface. In *ACM SIGGRAPH*, pages 262–270, New York, NY.

[Candy, 2007] Candy, L. (2007). Constraints and creativity in the digital arts. *Leonardo*, 40(4):366–367.

[Dietrich, 1986] Dietrich, F. (1986). Visual intelligence: The first decade of computer art (1965-1975). *Leonardo*, 19(2):159–169.

[Edmonds et al., 2004] Edmonds, E., Turner, G., and Candy, L. (2004). Approaches to interactive art systems. In *ACM SIGGRAPH*, pages 113–117, New York, NY.

[Enzenberger, 1974] Enzenberger, H. M. (1974). Consitutents of a theory of the media. In *The New Media Reader, ed. Noah Wardrip-Fruin and Nick Montfort (2003).*, Cambridge. The MIT Press.

[Falkenheim, 1980] Falkenheim, J. V. (1980). *Roger Fry and the Beginnings of Formalist Art Criticism.* UMI Research Press, Ann Arbor, Michigan.

[Foley et al., 1997] Foley, J. D., van Dam, A., Feiner, and Hughes (1997). *Computer Graphics: Principles and Practice.* Addison-Wesley Publishing Co.

[Fry, 1926] Fry, R. (1926). *Transformations: Critical and Speculative Essays on Art.* Chatto & Windus, London.

[Greeley, 2006] Greeley, R. (2006). *Surrealism and the Spanish Civil War.* Yale University Press, London, UK.

[Haeberli, 1988] Haeberli, P. E. (1988). Conman: a visual programming language for interactive graphics. *SIGGRAPH Comput. Graph.*, 22(4):103–111.

[Hatt and Klonk, 1992] Hatt, M. and Klonk, C. (1992). *Art History: A critical introduction to its methods.* Manchester University Press, p. 77, New York, NY.

[Huhtamo, 2006] Huhtamo, E. (2006). Twin-touch-test-redux: Media archaeological approach to art, interactivity, and tactility. In *In MediaArtHistories, Ch. 5.*, page 71.

[Levy, 1992] Levy, S. (1992). *Artificial Life*. Pantheon Books, New York, NY.

[MacGregor, 2002] MacGregor, B. (2002). Cybernetic serendipity revisited. In *Creativity & Cognition, Oct 14-16.*, Loughborough, Leic, UK.

[Mandelbrot, 1982] Mandelbrot, B. B. (1982). *The Fractal Geometry of Nature*. W. H. Freeman, 1 edition.

[Manovich, 2001] Manovich, L. (2001). *The Language of New Media (Leonardo Books)*. The MIT Press.

[Manovich, 2007] Manovich, L. (2007). Abstraction and complexity. In *MediaArtHistories.*, Cambridge. The MIT Press.

[Marat and Downes, 2004] Marat, B. and Downes, M. (2004). Visual programming languages: A survey. UC Berkeley. Technical Report: CSD-04-1368.

[Mignonneau and Sommerer, 2005] Mignonneau, L. and Sommerer, C. (2005). Designing emotional, metaphoric, natural and intuitive interfaces for interactive art, edutainment and mobile communications. *Computer Graphics*, 29(6):837–851.

[Muller et al., 2007] Muller, P., Arisona, S. M., Schubiger-Banz, S., and Specht, M. (2007). Interactive editing of live visuals. In Advances in Computer Graphics and

Computer Vision. Communications in Computer and Information Science, 2007, Vol 4, Part 5, p. 169-184.

[Murphy, 2000] Murphy, J. (2000). Ritual resolution. http://www.michaelrees.com/pdf/MURPHY.pdf, accessed June 2010. Lecture notes.

[Myers, 1998] Myers, B. A. (March, 1998). A brief history of human computer interaction technology. *ACM interactions.*, 5(2):45–54.

[Noll, 1966] Noll, A. M. (1966). Human or machine: A subjective comparison of piet mondrian's "composition with lines" (1917) and a computer-generated picture. *The Psychological Record*, 16(1-10).

[Paul, 2003] Paul, C. (2003). *Digital Art.* Thames & Hudson.

[Preziosi, 1998] Preziosi, D. (1998). *The Art of Art History: A critical anthology.* Oxford University Press.

[Puckette, 2002] Puckette, M. (2002). Max at seventeen. *Computer Music Journal*, 26(4):31–43.

[Reas and Fry, 2006] Reas, C. and Fry, B. (2006). Processing: programming for the media arts. *AI & Society*, 20(4):526–538.

[Reed, 1996] Reed, C. (1996). *A Roger Fry Reader.* University of Chicago Press, Chicago.

[Rees, 1999] Rees, M. (1999). Words around objects: The metaphors of new media. http://www.michaelrees.com/pdf/PARIS.pdf, accessed June 2010. Lecture notes.

[Resnick et al., 2005] Resnick, M., Myers, B., Nakakoji, K., Schneiderman, B., Pausch, R., Selker, T., and Eisenberg, M. (2005). Design principles for tools to support creative thinking. *NSF Workshop on Creative Support Tools. June 13-14.*

[Rosenfeld, 1983] Rosenfeld, A. (1983). Picture processing: 1982. *Computer Vision, Graphics, and Image Processing*, 22:339–377.

[Shneiderman et al., 2005] Shneiderman, B., Fischer, G., Czerwinski, M., Myers, B., and Resnick, M. (2005). Creative support tools. *NSF Workshop on Creative Support Tools. June 13-14.*

[Sommerer and Mignonneau, 2001] Sommerer, C. and Mignonneau, L. (2001). Modeling complex systems for interactive art. *Complexity International*, 8.

[Strauss, 1993] Strauss, P. S. (1993). Iris inventor, a 3d graphics toolkit. *SIGPLAN Not.*, 28(10):192–200.

[Sutherland, 1963] Sutherland, I. E. (1963). Sketchpad: A man-machine graphical communication system. In *AFIPS Conference Proceedings 23*, pages 323–328.

[van Dam, 1998] van Dam, A. (1998). Some personal recollections on graphics standards. In *The History of Computer Graphics Standards Development.*

[Verostko, 2002] Verostko, R. (2002). Algorithmic fine art: Composing a visual arts score. In *From Explorations in Art and Technology, by Linda Candy and Ernest Edmonds.*, page 131, London, UK. Springer-Verlag.

[Verostko, 2006] Verostko, R. (2006). The algorists, historical notes. http://www.verostko.com/algorist.html, accessed Sept 2010.

[Wessel and Wright, 2002] Wessel, D. and Wright, M. (2002). Problems and prospects for intimate musical control of computers. *Computer Music Journal*, 26(3):11–22.