Perceptual and Interpretative Properties of Motion for Information Visualization

Lyn Bartram

School of Computing Science
SImon Fraser University

Technical Report CMPT-TR:1997-15

Visualizing information in user interfaces to complex, large-scale systems is difficult due to an enormous amount of dynamic data distributed across multiple displays. While graphical representation techniques can reduce some of the cognitive overhead associated with comprehension, current interfaces suffer from the over-use of such representation techniques and exceed the human's perceptual capacity to efficiently interpret them. New *display dimensions* are required to support the user in information visualization. Three major issues which are problematic in complex system UI design are identified: representing the nature of change, supporting the cognitive integration of data across disparate displays, and conveying the nature of relationships between data and/or events.

Advances in technology have made animation a viable alternative to static representations. Motion holds promise as a perceptually rich and efficient display dimension but little is known about its attributes for information display. This paper proposes that motion may prove useful in visualizing complex information because of its preattentive and interpretative perceptual properties. A review of animation in current user interface and visualization design and research indicates that, while there is strong intuition about the "usefulness" of motion to communicate, there are few guidelines or empirical knowledge about how to employ it. This paper summarizes types of movement characterization from diverse disciplines and proposes an initial taxonomy of motion properties and application to serve as a framework for further empirical investigation into motion as a useful display dimension. Implementation issues are discussed with respect to real-time display requirements.

Table of Contents

1 Introduction

Complex systems such as those used in supervisory control and data acquisition (henceforth **SCADA**) can be characterized by large volumes of dynamic information which cannot reasonably fit into single displays or even a single computer screen. In such systems the interface must not only represent the data in reasonable ways but should also signal the user effectively when important changes take place and, increasingly important in environments with multiple screens or windows, should provide clear indications when data are associated or related in some way. When appropriately used, graphical representations such as shape, symbols, size, colour and position are very effective in information visualization because they are *mentally economical* [Woo95b] - rapidly and efficiently processed by the preattentive visual system rather than cognitive effort. However, when human perceptual capacity to assimilate all the combinations of codes and dimensions is exceeded interface comprehension increasingly demands cognitive activity and mental economy is lost.

Current SCADA interfaces rely heavily on static graphical dimensions for low-level data display but devote few resources to helping the user integrate information within and across a disparate set of displays and data representations. There is little “perceptual room” left in the standard set of representations to support higher-order gestalt perceptions of system function and state. This can be seen as a *bandwidth* problem: the graphical communication “channels” are cognitively and perceptually overloaded at the user's end. As the amount of data continues to increase and as the visual field in which it is represented expands across more screen space, additional communication and representation dimensions are needed to improve information bandwidth.

One very promising candidate is motion. Advances in graphics technology increasingly support powerful animation capabilities in operator workstations and consoles.However, while there are extensive guidelines on the use of perceptually efficient static graphical techniques in information representation [Ber83] [Tuf90] [Cle93], there is little research on the corresponding uses of motion. Animation is used in visualization and user interfaces in an ad hoc, sporadic manner. Yet evidence from fields as diverse as perceptual science and the performing arts suggest that motion has much richer communication potential. My thesis research is concerned with determining if and how motion may be usefully applied to problems in information visualization. This paper reviews the perceptual and interpretative properties of motion in the context of information display and proposes the basis for a framework of investigation into the usefulness of motion as a display dimension based on a new taxonomy of movement attributes and uses.

1.1 Organization

The paper is organized as follows. Section 2 describes issues in current SCADA information and interface design and identifies certain key problems of information representation which presently used display techniques are ill-equipped to solve. Section 3 discusses the basis of perceptual efficiency in visualization and the approach of ecological design and investigation. The motivations for considering motion as a potentially useful display dimension are considered in Section 4. Section 5 reviews the use of motion as communication in various environments. It begins by describing current and proposed uses of animation in user interfaces and information visualization and presents an overview of how movement is used to represent information in dance, conducting and character animation. Section 6 proposes a taxonomy of motion properties. Section 7 suggests potential areas of application in SCADA interfaces and information visualization. Section 8 discusses implementation issues to be addressed in implementing motion representation effectively and reliably. Section 9 concludes the paper.

1.2 Terminology

I use the following terms in this paper. *Visualization* refers to both the way in which the data is displayed and the "cognitive operation of forming the mental image of the data to facilitate insight" into relationships and constraints [WML94]. A display technique is either digital or analog. *Digital* representations are alphanumeric and portray the exact value of the data. *Analog* representations use graphical coding. Analog *coding dimensions* include colour, size, position and orientation. Dimensionality of *size* obviously depends on the viewing dimensionality (2D or 3D). A coding dimension is also termed a *display dimension*. The *coding granularity* of a display dimension refers to how many distinct meanings, or separate codes, can be efficiently perceived. (See Section 3.1 for a more detailed discussion of perceptual efficiency.) A *display* is a data "container" which may reside on a physical *screen* or in a *window* (virtual screen). A screen can contain one or more displays. A display can be "attached" to one or more windows, or span one or more *pages* (connected screenfuls of information): it may portray groups of devices, sensed or derived data, state and/or events.

2 The User Interface Bandwidth Problem

As the data acquisition capabilities of control systems have increased, the operator's role has evolved from low-level manual control to high-level management and supervision. These complementary trends have resulted in a ballooning of the complexity of the underlying information space and the volume of data used in the operator's tasks[MS97]. The traditional approach has been to add more hardware and software displays to accommodate the explosion of information. The problem is bandwidth: while the display capacity of the system can be increased arbitrarily information transfer is bottlenecked on the limits in the user's perceptual capacity. We define *user interface bandwidth* as the capacity for information communication/transfer between the user and the system at the interface. We are concerned with the communication from system to user, in which information is encoded into digital (alphanumeric) and analog (graphical) forms and the user must interpret (decode) the information in a timely fashion.

The user's cognitive "cost", or effort, of decoding representations is a function of memory access and the mental operations of *search* and *computation*. A *perceptual* operation is carried out by the low-level, preattentive human information processing system (which, for the purposes of this paper, is the human visual system). Perception is a highly efficient, "automatic" parallel process. *Cognitive* operations involve higher-level, "conscious" effort and are serial in nature. The perceptual *coding granularity* of a display dimension refers to how many distinct meanings, or separate codes, can be retained in short-term memory, thus requiring no cognitive recall effort. Digital representations are effective when exact values and quantitative computation are required, but involve the serial, effortful tasks of "reading" and cognitive inference. Graphical representations are useful for qualitative assessment because they can invoke perceptual rather than cognitive inferences [Cas91] and exploit the human capacity to perceive separate discriminable features of objects in parallel (search) and to recognize patterns (computation)[Cle93] [Tuf90] [WC95].

There is, however, a substantial gulf between the knowledge of these basic perceptual building blocks and effective display design in complex information systems. The first problem is one of appropriate display design and over-use of perceptual coding. Information visualization and design literature has addressed issues of graphical perception and appropriate data representation with respect to useful graphic design [Ber83] [Tuf90] [Mac86] and for specific types of information extraction [Cle93]. The guidelines are useful for improving the clarity and usability of visual presentation, but as Casner points out [Cas91] they offer little empirical insight into the fundamental reasons of *why* a particular technique or practice is useful, and they are information- rather than task-centric (concerned with representing information structure as opposed to supporting the user's task-specific needs). Moreover, the proposed techniques exist in a singular context: that is, there is an implicit assumption that all the perceptual and coding resources can be devoted to that representation. However, as will be shown in Section 2.1, users of complex systems are increasingly dividing their resources between a multiplicity of representations and perceptual codes. When the discriminability of a code is exceeded, understanding the representation is reduced to an effortful, serial process of mapping each code to long-term memory and decoding the value.

The second problem relates to the process of distilling information from data. Substantial cognitive effort is required to recognize, retrieve and integrate information from different data in different representations and displays. The fall-out from the over-use of perceptual coding discussed above is that there is no "extra" coding granularity in the currently used display dimensions to off-load these effortful cognitive operations to the perceptual system.

2.1 More Data, More Displays...

Control system user interfaces typically still reflect an elemental design approach [Woo95b] ("one sensor, one display"). Users often work with dozens of displays and thousands of data points. In control rooms where space is not limited (such as those used in telecommunications and power distribution [Byb92] [DBB91]) it is not uncommon to see interfaces physically distributed across banks of CRT stations (on which users "page through" windows or displays), large wall-mounted shared displays, static maps/mechanical status boards, shared message boards and video screens from remote camera feeds. In one typical power distribution system (TransAlta Utilities[DBB91]) the CRT-based system comprises 2507 possible pages and users routinely manage a few hundred. Even in space-constrained environments like the cockpit, the interface is spread across six to eight screens, some dedicated to particular displays and some which can be configured by the pilots [CCMG97]. Thus users are constantly "flipping through" displays in space and time to build up a mental image of the system state and behaviour [Byb92] [Bai91].

Displays tend to be densely populated with data to optimize the use of screen real-estate and reduce the number of sequential accesses a user makes. Digital (alphanumeric) displays, which force operators to read information in a serial, cognitively effortful fashion, are used when exact values need to be known, but are increasingly augmented by graphical displays. Graphical representations include charts, graphs, diagrammatic displays (such as maps and schematics) and the depictive *mimic displays* [GGB89] [Bai91] which reflect the physical appearance and organization of system components. Most graphical displays are still two-dimensional, although there is increasing interest in exploring 3D as a more effective use of visual space [WAS86] [Alv93] [WML94] [WMT96] [MOW97]. Symbols and icons are heavily used. The most common display dimensions for coding value and state are colour, position and size (where position refers to spatial proximity and alignment, and size is the height or width of an element.). The most common indication of fault, or *alarm*, conditions is blinking or flashing the relevant display element [GGB89] [DBB91] [Woo95a].

The efficacy of these representations is constrained by screen space and perceptual resources. There is some debate on the number of symbols which can be perceptually decoded (Gilmore suggests a limit of 15 to a "symbol alphabet" [GGB89] while Bainbridge gives evidence to show competence up to 33 [Bai97]). However, process and network displays typically use significantly larger symbol sets, consisting of the symbols associated with the underlying physical system (e.g. [oA86]), the symbols which refer to the control system itself, and finally the icons which are related to the user interface (such as iconified windows, display identifiers and cursors).

Factors which limit perceptual effectiveness of a display dimension to which data values are assigned include its *coding granularity* and visual acuity. In particular, colour is over-used in most systems [WAS86] [Bai91] with fully saturated hue as the dominant code. Yet only 7-10 hues can be usefully distinguished [Hea96] in the narrow foveal range of vision. Moreover, certain colours (such as red and orange) appear much brighter than others of equal luminance and thus attract attention, distracting from the decoding operation. Hue changes in the periphery are not well perceived. Relative position and size are decodable only within a common frame of reference (i.e., plotted against similar axes and scales), and are difficult to assess when the representations are not close to each other [Cle93]. Thus, for example, it is difficult to determine whether two tank levels displayed on separate screens (or non-aligned displays) are equivalent without calculating each value separately and comparing them (a cognitive operation) as opposed to more efficient perceptual inferences that can be made by comparing two adjacent heights. Shape, position and size are difficult to decode in conditions of *visual noise*: that is, when the displays are too cluttered and the resolution of the representation poor [Bai91]. Gilmore suggests a general guideline of 50% density for SCADA displays [GGB89]; it is not clear, however, how density is measured. What is apparent is that most displays are too densely populated and the subscribed display dimensions over-used, complicating rather than facilitating user comprehension and causing a sense of data overload [WAS86] [Byb92] [Woo95b].

Flashing or blinking is a particular example of data overload. An abrupt change of luminance or onset of motion automatically attracts visual attention to the area [STK81]. This effect is negated and in fact interferes with search when multiple elements on the displays begin flashing. As one NASA Mission Control operator stated, after the Apollo 12 spaceship was hit by lightning. "...all the lights came on. So instead of being able to tell you what went wrong, the lights were absolutely no help at all" [Woo95a].

2.2 ...Insufficient Information

In practice, SCADA interfaces often suffer from "too much of a good thing": too many colours, shapes and visual cues are combined injudiciously in an effort to represent increasing volumes of data, resulting in interfaces which are more rather than less effortful to use. We believe the bandwidth problem is exacerbated by too much direct **data** and not enough **information.** Woods defines information as a link between the data (*referent*), the symbol (*representation*) and the *user* (knowledge, context and expectations) [Woo95b].

Current SCADA interfaces are information-deficient in three areas.

2.2.1 System Behaviour and Change

Increases in magnitude and complexity emphasize the need for *higher-order* system behaviour information.Key examples are summary views (integration of system functions over many variables) [DBB91], meta-information and representations of change (nature, rate and history).

Performance and predictive summaries are important because the lack of explicit indication of high-order system function and state forces users to expend significant cognitive effort and time in assembling a mental image of how well or poorly the system is doing (and has done) from a multitude of lower-level data views. What few summary views exist tend to be separate displays which are heavily used but which impose the mental burden of spatial and temporal information integration with the more detailed data displays [BODH94]. In particular, case studies of network managers in telecommunications and power distribution operators emphasized that users want more effective alarm grouping, prioritizing, sorting and summary mechanisms without paying the extra overhead of more displays to manage and to relate to their detailed views [DBB91] [Byb92].

Meta-information is directly concerned with situation awareness, especially where automated operation is involved [ATP95] [SW95]. New levels of automation in complex systems have resulted in a proliferation of *modes* or operational contexts. This places new cognitive demands on the human controller, who must track the automation processes and carry out different actions depending on system mode [SW95]. Modal information is often presented to the human as simply a small flag or textual field in one part of the display which is often missed or ignored [CCMG97]. When several operators are working on a shared part of the system they need to explicitly warn each other about potential overlapping actions. Often this is done by voice [Byb92] but in high-tempo workloads this adds to situation confusion and overload.

Perhaps the most crucial requirement to understanding a dynamic system is effective *representation of how the system changes.* Users rely on temporal reasoning to understand dynamic system behaviour. Decortis et. al. describe control system operators reasoning about time as an explicit variable (duration, time of occurrence) and implicitly (how the state of the system at time *t* relates to the state of the system at time *t+1*) [DdKCV91]. Current representations of change are either direct and immediate (i.e., the event simply causes a change in the value displayed) or indirect and persistent, in which time is explicitly represented (e.g., trends). Trend graphs, which plot time as an explicit variable along an axis, are extremely useful as process histories but require substantial screen space and thus are typically used as separate displays with limited groups of devices or values [DBB91].

Not only must the event of change itself be obvious but the magnitude and the rate of change must be immediately apparent. Consider the example of the Apollo 13 accident which was caused by an explosion in one of the oxygen tanks. Mission controllers monitored a screen of digital numbers in which the catastrophic change in tank pressure showed up only as a sequence of three values for one tank over four seconds (996 psi., 1,008 psi, and then 19 psi). As a result the event was missed and it took 54 minutes of investigation into subsequent system failures and explored hypotheses before the problem was detected [Woo95b]. An explicit indication of the rate and magnitude of this change would have immediately alerted the controller to the explosion in the tank.

2.2.2 Integration of data across displays

Our particular concern is with representation congruence and coherence across the whole user interface. Display proliferation has led to the keyhole or lost in space phenomenon [Woo84] and imposes significant burdens in mental integration and the assembling of information for problem-solving across space (disparate displays and surfaces) and/or time (sequential views) [Bai91] [DBB91] [Byb92]. Our previous work with a multi-screen interface [BHD95] identified the need for integrative cues to perceptually connect data in disparate displays. It is common for users to inspect various views of data in different contexts as part of monitoring and problem-solving. For example, a power distribution operator investigating a line overload problem may need a schematic of the relevant switches, a trend of the line loads and an alarm history. The relevant data typically reside in various separate displays and the user has to visually “collect” the appropriate items[1]. Given the dynamic nature of the system, this visual collection involves not only obtaining

---

[1] One BC Tel network manager referred to this task as “inviting all the right pieces of information to the party”.

a particular set of static values but also continual attention to those areas on the display: in effect, the user is maintaining a set of “visual pointers” [PBF+93]. Woods defines the information space as a *virtual perceptual field* over which we have limited viewpoints and emphasis the needs for mentally economical (perceptually efficient) orienting cues to alert the user to the fact that something interesting is going on another part of the perceptual field [Woo84].

2.2.3 Data relationships

Improving interface bandwidth implies that information visualization needs to be considered as an integrated whole - displaying not only data but also the *relations* between data in perceptually efficient ways. Traditional display design concentrates on data organization to convey certain qualitative relationships between the data (e.g. “same as”, “higher than”, “earlier than”, “part of”). There are no well-established techniques of displaying the semantically richer dynamic relations between elements both within and across displays of *association*, *dependencies, sequence*/*order* and *causality.* Association can include user- or system-defined grouping (all alarms of a certain type, all devices under maintenance). Dependency illustrates how processes and data rely on each other in system functions (especially useful in “what if” scenarios and contingency analysis). Order may refer to temporal or hierarchical ranking. Perhaps the most pressing need is for effective representation of causal information. Causal data is increasingly available from diagnostic sub-systems but is delivered to the user in textual form, requiring a mapping to other state displays as opposed to an intuitive comprehension in place.

3 Issues in Display Design

3.1 Perceptual Principles for Visualization

3.1.1 The Proximity Compatibility Principle

Wickens and Carswell [WC95] suggest that displays relevant to a common task should be perceptually “close”. Their *proximity compatibility principle* (PCP) depends on two dimensions of similarity: perceptual proximity and processing proximity. It proposes that close task proximity is best supported by close perceptual proximity; conversely, independent processing requires distant perceptual proximity. *Perceptual (display) proximity* defines how close together two display channels are in the user’s perceptual space (i.e., how similar they are). For example, two sources will be perceived as more similar (in closer proximity) if they share colour, physical dimensions, code (analog or digital) or are spatially near. Pr*ocessing (mental) proximity* defines the extent to which sources are used as part of the same task. Integrative tasks, in which two or more data must be computed or compared to arrive at the needed information, have high mental proximity. Nonintegrative processing of similar tasks has lower proximity. Similarity depends on the sharing of certain features: e.g, metric (information portrayed in same units); functional (same operational or device group); processing (same computational routine on different sources); or temporal (dissimilar sources processed concurrently involving frequent visual transitions and contributing to the same goal). Finally, nonintegrative processing of dissimilar tasks, in which the user is switching attention between disparate tasks with no common goal, has the lowest mental proximity.

The authors suggest several display manipulations to decrease information access by increasing perceptual proximity: put the objects close together; group them visually (perhaps by enclosure); use the same display source (colour, texture, orientation) to associate information source (e.g., all tank gauges are bar charts); use the same display property to indicate value (height of the bar); and object integration. Object integration arranges information sources so they appear to be part of a single object and exploits the perceptual processing mechanism that decodes the “separable” features of objects in parallel [KT92]. Examples are connecting a series of dots with a contour, or the common dimensional integrality of a point in an *(x,y)* graph rather than parallel measures of extent.

3.1.2 Emergent Features

Wickens and Carswell [WC95] emphasize that Information integration (as opposed to access) is well supported by *emergent features*: properties inherent in the relations between raw data encoding which serve as a direct cue for a an integration task which would otherwise require computation or comparison of the individual data values. An example is alignment of bar charts of similar value, which is not a property of any of the individual bars in isolation, or volume of a rectangle whose height and width are mapped to separate sources. They caution that emergent features, while effective for integrative tasks, can interfere with focused attention on decoding individual values, especially in noisy environments where adjacent or overlapping images will lead to decreased discriminability.

3.1.3 Directed Attention

Focusing attention in a visually noisy field with many data channels whose values are constantly changing (for example, in fault management situations) requires the user to maintain control of where she is attending at the same time as being aware of potentially interesting areas as conditions change. Woods has defined the need for a set of cognitive tools to support control of attention in fault management situations where the user must sort through an overload of raw data [Woo95a] which should exploit preattentive reference. He identifies several criteria for such attention-directing signals:

- accessibility (i.e., the user should be capable of picking them up without losing track of current activities;
- partial information: the signal should carry enough partial information for the user to pick up whether to shift attention to the signalled area; and
- mental economy: the representation should be processed without cognitive effort.

3.2 An Ecological Approach

Woods [Woo95b], Wickens [WC95] and Casner [Cas91], among others, emphasize that data representation which is not relevant to the greater task environment is inevitably ineffective. Woods calls this the “decoding” problem [Woo95b]: domain data can be cleverly mapped into visual attributes, but unless the user can decode the representation in the actual task (under conditions of attention switching, risk and time pressure) the representation fails. This it is not sufficient to design a display with emergent features which are not mapped to variables of importance to the current task [WC95]; we need representations which are *ecologically valid,* a concept which draws from Gibson’s principles of *ecological perception* [Gib76].

Gibson theorized that we perceive our environment directly as ecological entities and movement rather than as abstractions of light and shadow which must then be internally computed into meaningful components. The composition and layout of objects in the environment constitute what they can *afford* the observer. *Affordances* can be thought of as the possibilities, opportunities and indeed meaning of objects in the environment: since affordances can be directly perceived, it follows that meaning and value can also be directly perceived (rather than computed).

Gaver [Gav93] has employed this approach in determining salient properties of sound to convey information about events and meaning in the environment, which he characterizes as the study of *ecological listening* as opposed to that of audio perception. Similarly, the *ecological* *interface design (EID)* approach of Vicente and Rasmussen [VCP95] emphasizes the representation of higher-order function, state and behaviour information as task-relevant variables integrated over lower-level system data. Ecological design can exploit the perception of emergent features. The power of this approach is that the act of integrating the values into knowledge of the system's functional performance becomes a perceptual rather than a computational operation.

3.3 Design Challenge

There are two complementary directions which must be addressed in ameliorating information overload in the user interfaces to complex systems:

1. explore new perceptually effective ecological representations which may increase information dimensionality and thus interface bandwidth; and
2. determine whether these new coding dimensions can extend the integrative effect across displays and representations separated by space (and possibly by time). An important property must be explicit support for *visual momentum*, the user's ability to effectively extract information across displays [Woo84].

4 Motivation

We believe motion to have great potential as an display dimension for three reasons:

4.1 Low-level perceptual efficiency.

Motion perception is a preattentive process: motion can elicit "pop-out" effects in which moving objects can be searched in parallel by the visual system [WS91]. Ware reports studies that show moving objects can be searched in parallel for targets with different direction and different rates of rotary motion[JH70] [WL94]. Psychologists believe that motion, like colour and form, is handled by a dedicated visual processing mechanism [Cut86] [PBF$^+$93], indicating that it is a "separable" feature of an object [KT92]. (The reader is directed to [WS91] for a more detailed description of motion perception.) In his extensive review of temporal factors affecting information transfer from visual displays, Sekuler [STK81] reports motion detection times as low as 50 msec. Like all visual functions, the periphery is less sensitive than the centre of the visual field to motion perception. However, motion response degrades less than spatial acuity or colour perception in the periphery. Movement is reported as improving the visibility of targets embedded in "random or cluttered" fields, especially away from the centre, where detection time is considerably reduced. Thus, unlike hue or shape discrimination which require the visual acuity of the foveal range, movement is suited to extracting information from "noisy" environments across the entire visual field. This is important in interfaces in which operators are not guaranteed to be looking straight ahead at a display all the time.

The human visual system is very good not only at perceiving but also at tracking and predicting movement. Recent research has shown we can track multiple motions in parallel [PBF$^+$93] without effortful context-switching. Eilan et. al. state there is "compelling evidence for the internalization and semi-automatic use of quite specific physical principles that generally yield very accurate [mental] representations of object trajectories" [EMB93]. They suggest that humans employ a low-level "intuitive physics" which correlates geometrical properties of distance and size with the physical properties of velocity, mass and acceleration. Subjects are able to interpolate and judge trajectory points when shown "interrupted" motions, manifesting an inherent ability to predict the current and continuing motion of objects in space. Studies reported by Cooper and Munger suggest this is done by internalized *kinematic* rather than *dynamic* principles. Kinematic principles link position, velocity and acceleration without regard to mass; dynamic principles, on the other hand, employ concepts of forces and mass to explain changes in rest and movement states [CM93].

In addition, we use motion to derive structure and animacy from very sparse cues. In his seminal work on biological motion, Johannson [Joh73] found that subjects could identify characteristics and structure of human figures from only 10 moving "dots" (lights attached to the bodies). Subjects could not identify any meaningful structure from static presentations of the dot groupings. However, even sparse movement gave instantaneous rise to the recognition of a moving body. Subjects not only identified the full body from as few as 10 lights (and the legs from only 5); they also identified the gaits and the quality of the gait characteristics from the motion (walking, walking with an injury, running, etc.) When the motion was rigid it was identified as mechanistic: non-rigid motion gave the sense of *animacy.*

Subjects in experiments conducted by Bassili [Bas78] identified facial structure and emotion from similarly sparsely placed lights when motion was present but were unable to extract any information when the stimulus did not move.

Finally, motion has the effect of grouping. Things that move together are seen as grouped or associated. Cutting showed that rotating disconnected objects around a common axis led to the perception that they were connected into a rigid structure by invisible "rods" [Cut86]. Gibson looked at moving "patches" or closely bunched textures of dots: differences in their speed resulted in a perception of "twoness" [Gib76]. Realistic simulations of herd and flock behaviour have been produced by ensuring some communality between the movements of the individual actors (see [Rey87] for an early example]).

4.2 Interpretative scope.

Motion is cognitively and ecologically ***rich.*** Gibson defined motions as *ecological events* to do with the changes in the layout and formation of objects and surfaces around us. In his approach all perception is motion perception: the flow of such information through the optic array is what gives us the information about the 3D world [Gib76] [Joh75]. It is obvious motion can convey information that static representations cannot: it is difficult to imagine an intuitive static display of causality, for example. Decortis found that spatio-kinematic representations helped users reason more effectively about temporal data in continuous processes [DdKCV91]. In Gibson's terms, motion *affords* behaviour and change. In our real virtual field things are constantly moving. Evidence from perceptual psychology indicates that the onset or change in movement in our visual field grabs our attention involuntarily, suggesting that multiple, constant and irregular movement should be highly disorienting. Yet obviously as functioning actors within our environment we are able to somehow manage and selectively attend to the visual information without constant conscious effort.

At the other end of the cognitive spectrum, we derive very rich meaning from movement. Relatively simple combinations of motions can be interpreted as highly sophisticated behaviour. The arts of drama, dance and music map very complex emotions and motivations on to gestures and movement. Moreover, we tend to anthropomorphize movement sequences in which there are several "actors", no matter how abstract the representation. Jetha's thesis project, which investigated how people carry out complex design tasks, had subjects generating dance sequences from an initial "abstract" motion sequence involving a cube, pyramid and a sphere in which all but one subject mapped the object simple movements onto articulated human figures with many more degrees of freedom [Jet93].

Character animation relies on the exaggeration of movement to deepen our understanding of behaviour and motivation [TJ81]. Moreover, basic spatiotemporal properties of movement elicit impressions of intention and actor [Kas81]. Very simple actions which are computationally inexpensive to animate can produce complex psychological impressions [LW90]. (This is discussed more in Section 5.2.3.) People construct complex emotional interpretations of behaviour and intentionality from different patterns of motion. This contribution of motion to *social perception* has been investigated by Heider and Simmel [HS44], Kassin [Kas81] and Berry and Springer [BS93a].

Heider and Simmel [HS44] investigated the perception of *attribution* using an animated film technique to show people a cartoon in which a large triangle, a small triangle and a circle moved around a rectangle with a "door". Subjects anthropomorphized their impressions and attributed complex behavioural states, motives and personalities to the objects, such as timidity, aggression, protection and affection. Berry and Springer [BS93a] conducted a study with preschool-age children in which they replayed three versions of the Heider film: one in which the structural object information was intact but the movement was disrupted; one with intact movement but structural distortions; and one in which both the structural and movement properties were disrupted. Their findings confirm that attribution and causality perception was based on the patterns of motion rather than structural information. Kassin further showed that this informing of social perception holds across diverse populations [Kas81].

Michotte reported extensively on the contribution of motion to the direct perception of causality [Mic63]. He found solid objects were unnecessary for creating a causal impression: rather, it arises from specific combinations of motion which our perception unifies into a single, compound "causal" movement. "Pure" causality was perceived when the causal object A was totally responsible for the subsequent movement of the passive B (*launching* and *entraining*). "Weaker" effects (i.e., in which the effecting object generated some latent behaviour in the effected object B, such that B's behaviour was autonomous without being spontaneous) were identified as *triggering*, *attraction* (such as iron filings to a magnet) and *transporting*.

4.3 Availability

Motion is under-used and therefore **available** as a "channel" of carrying information, as will be described in the following section. Increases in computing power have made the production of seemingly sophisticated motion relatively inexpensive. For example, Reynolds' 1987 "flocking" simulation [Rey87] has since been rendered in real-time. Basic computer animation techniques like colour table animation and simple forward kinematics [FvDFH90] are accessible to even moderately-configured desktop machines.

We anticipate that a major advantage of motion coding will be its compact use of screen real estate, freeing up spatial display dimensions for other use.

5 Movement As Representation

5.1 Animation

The arrival of animation capabilities at the desktop has provoked interest in the use of known animation techniques for computer-human communication. A common thread to the proposal and inclusion of animation capabilities in user interfaces is a strong intuition that motion and making information objects move should make the interface environment more credible, more "real", less cognitively foreign to users. Baecker and Small [BS90] discussed the potential of user interface animation to reveal process and *structure* (by moving the viewpoint) and introduced the following taxonomy of eight uses of animating *function* to make the interface more engaging and comprehensible.

- *Identification* associates the symbol with its function (“What is this?”);
- *Transition* carries the user smoothly between states (“Where did I come from and where have I gone?”);
- *Choice* shows possible actions (“What can I do now?”);
- *Demonstration* illustrates the capabilities of the tool or service (“What can I do with this?”);
- *Explanation* shows how to employ it (“How can I do this?”);
- *Feedback* provides information on process dynamics and state (“What is happening?”);
- *History* replays previous actions and effects (“What have I done?”); and
- *Guidance* suggests suitable next steps (“What should I do now?”).

Stasko [Sta93] adds four design guidelines drawn from the principles of traditional animation.

- *Appropriateness* dictates that the operation or process should be represented according to the user’s mental model and system entities.
- *Smoothness* is essential since jerky, wildly varying animations are difficult to follow.
- *Duration and control* vary with the type of animation. Demonstrations of unit operations such as selection should be short (not more than a few seconds). Animating continuous processes with a *clocktime correspondence* should be kept faithful to the clocktime. When animation is used as explanation, the user should be allowed to control the rate and replay.
- *Moderation* prescribes judicious application of animation: too much is overdone and too cute.

5.1.1 Animation At the Interface

Motion in its most basic form has long been used in interfaces: much use is made of blinking as a human interrupt to attract and direct visual attention. In many supervisory control systems it is the primary visual cue for alarm conditions. There is some evidence to suggest a limited coding granularity of 4 [GGB89] or 5 [WBKC92] flashing frequencies. Anecdotal evidence indicates that people find blinking excessively annoying and visually ineffective when too many items are flashing (who has not cursed the WWW HTML blink function?) In large-scale systems where alarms tend to propagate rapidly, over-flashing not only reduces effective alarm information but also renders the displays visually disturbing, distracting users from effectively perceiving the needed information from other representations [Woo95a].

Schlueter exploited this *visual dissonance* property of motion to enhance perceptual differences in crowded display environments with overlapping windows: to avoid the effect of window contents seeming to continue, or “bleed”, across borders he “jostled” the windows to enforce perceptual differentiation [Sch89].

Animation is increasingly used as visual momentum to provide smooth transitions between views. The Macintosh [ACem] interface uses “tracers” to draw lines between states when expanding and iconifying windows. The Sun OpenWindows [SM] system uses a similar “telescoping” technique. Later distortion visualization techniques s directly animated the changes in the objects themselves. The Perspective Wall [MRC91] allows the user to smoothly horizontally scroll a linear “sheet” across a magnifying view. In the Continuous Zoom [BODH94] [BHDH95] objects expand and shrink at a minimum rate of 12 frames/sec in response to user controls which magnify certain objects and concomitantly shrink and displace others. The resulting transitions were deemed to greatly reduce the mental integration of continuously shifting display configurations. Cone Trees [RMC91] is a 3D technique for visualizing large sets of hierarchical information in which the user can quickly rotate “cones” to find the appropriate node and follow subsequent links to child cones. Observations of Cone Tree use and of expert users’ problem solving in the Intelligent Zoom [BHD95] suggest that animating a path through a visualization of an information structure aids in information retrieval and navigation.

Ware and Franck’s investigations into motion cues in 3D visualization [WF96] confirmed the intuition from Cone Trees that simple rotation about an axis is effective in interpreting 3D information structures. They considered three kinds of rotation cues in both stereo and mono viewing conditions: passive (i.e., automatic, no user control); hand guided, and movement coupled to observer head position. They found that all three types of motion improved performance in using 3D graphs and all were more significant than the stereo cues alone.

Chang and Ungar [CU93] use techniques from film editing and cartoon animation in the **Self** user interface with the goal of enhancing cognitive comprehension and user engagement. *Motion blur* reduces temporal aliasing (this effect of an object “blinking out of existence” in one location and “blinking into existence” at another arises from large movements in a short interval and is due to the persistence of vision.) Filmic *dissolves* makes objects appear and disappear smoothly from view. More subtle exaggeration effects highlight the realism or credibility of the action. *Anticipatory action* is a small movement which occurs in the opposite direction of the subsequent animation and highlights the effect by subliminal prediction. *Slow-in-slow-out, follow-through* and *arc* rather than linear paths contribute to the perception of the animated objects as being “real”, that is, as having solidity and mass. The authors hypothesize that believable object motion makes the interface more enjoyable and offloads the burden of deciphering interface behaviour from “higher cognitive centers” to the perceptual system. Hudson and Stasko [HS94] have implemented support for similar character-animation based techniques in the **ArtKit** user interface toolkit.

Bharat and Sukiviriya propose an animation server architecture [BS93b] to emulate user interaction with a system with the goal of supporting script-driven animations for tutorial and groupware purposes.

Most recently Ware is experimenting with the use of *deictic* motion for narrative illustration [WFF97] by animating soft lines which persist for a while and then disappear. Deictic functions are communicative identification actions which directly show or point out referents, such as naming or describing the object, or specifying the target location [DAPG96]. In language, the words “this” and “that” have a deictic function. Physically referring to objects in space involves *deictic gesture*. Ware et. al. use linear stroking (“underlining”) movements for emphasis and highlighting; smooth, enclosing “ribbon” motions to introduce and encapsulate regions; and continuous, elongated motions to effect transitions to new areas.

5.1.2 Animation as Illustration

The most mature application of motion in displays is the animation of process behaviour over time for illustrative, analytical and explanatory purposes, e.g. in scientific, algorithm and program visualization.

Baecker and Small animated icons to identify and explain their function [BSM91]. The advantages of animation were particularly noticeable when the small size of icons meant a low resolution of information (i.e., intricate depiction was impossible). Ambiguity was reduced and users remembered the function of the particular icon better. Many operating systems animate process indicators. The process indicators can either require their own sub-displays (such as percent-done boxes [Mye85]) or overload the representation of other user interface symbols such as cursors [DFAB93]. Icons are sometimes used: the printer icon in the HP VUE [HP] user interface shows paper feeding through it when a print request is submitted and the desktop printer icon in the Macintosh [ACem] continually shows the degree of completion of the print job. Trend graphs, bar charts and other such analog displays in use in complex systems are instances of process indicators. A key feature of process indicators is that they must be true to the explicit *clock time* of the process; i.e., the representation must closely track the current state of the process and cannot flag behind nor surge ahead.

In contrast, animation used in algorithm, program and scientific visualization does not need to adhere to explicit clock time but rather to *relative time*: it can be replayed, slowed down or sped up, as long as the relative speeds of the changing elements remain constant with respect to each other (so that ways in which the changes occur are still meaningfully and faithfully represented.) Animation is a popular component of visual programming for prototyping and understanding (see [IJK90] for a discussion of visual programming and applications.) Bridgeland points out, however, that visual programming needs ways of expressing behaviour that implies an extension of the objects and their relationships rather than "just an arrangement" of the objects through time [Bri90]. Scientific visualization systems which rely on animation to explore complex events include meteorological, medical, geophysical and fluid flow dynamics (an NCSA CD-ROM sampler contains animations created in a multitude of scientific applications [NSC96]). Animation was first used in algorithm visualization in Baecker's well-known video of "Sorting out Sorting" [Bae81], which illustrates in parallel how different sorting algorithms re-arrange data and elicits an immediate perception of algorithm performance. Many algorithm visualization systems rely on animation to characterize the intended purpose and meaning of programs: examples include BALSA [Bro88], ACTION [HM90] and TANGO [Sta91]. The latter two enable users to design and implement animations by interactively specifying graphs of data and process structures. In the TANGO system [Sta91], users further specify the animation by laying out the path along which the objects will travel and defining the types of transitions they will undergo: for example, *move, resize, colour, fill, raise* and *alter visibility.* Limited control over rate of replay is supported by a *delay* transition.

The supposition by the creators of algorithm animation systems is that animation is beneficial to the learning and comprehension of how processes and systems work, since the viewer does not have to build up a mental image of the changes from static images and descriptions to understand the events [Won94], but there is some empirical evidence to shed doubt on this intuition. Palmiter and Elkerton [PE91] compared animated demonstrations, procedural textual instructions, and a combination of both in the tasks of learning and retaining interface procedures. They found that while the demonstration group was faster and more accurate in learning the tasks the text group had better retention, so that seven days later the text group was faster and as accurate as the demonstration group in performing the tasks. They propose that animation may in some cases hinder "deep" procedural learning since the simplicity of using demonstrations, even when text is provided, may encourage simple mimicry and discourage the use of the text. Wong [Won94] evaluated sets of animations in teaching abstract statistical concepts. She found some evidence that animations were more useful to students with low spatial ability. In general, while good animations had no positive effect on learning, poor animations were a hindrance. Students however, were more motivated to use animations in learning since they enjoyed them more. Wong concludes that the primary strengths of animation are in information *synthesis* and that static graphics are still preferable for information *analysis.*

5.1.3 Animation as Visualization

While the above systems use animation to illustrate the changes in simulated or abstract data objects over time, some effort has been directed to using animation as a direct representation of otherwise unrepresented system variables. Gronquist et. al. [GSAL96] and Alvarado [Alv93] model power system phenomena using a system of animated mass, spring and force vectors to show load flow, transmission capacity and transient stability. Vectors accelerate and decelerate in radial distance, width and angular displacement. Faults add to the kinetic energy of the concerned mass vectors, with the effect that as overall kinetic energy increases the system loses synchronism and goes out of balance. The authors state that "islanding" and coherency are particularly easy to visualize, even in large systems, since the diagram animation clearly shows groups of machines "swinging together".

Gobrecht and Ware [GWB96] use animated arrows travelling along links in a graph to show message passing activity in a multiprocessor system. Bursts and lulls in overall traffic can clearly be seen without extra representation or user computation. Fleet and Ware [FW96] use simple sinusoidal motion to show traffic on a link. A report on full graphical display design from CIGRE (Conférence Internationale des Grands Réseaux Electriques) proposes that power flow can be similarly represented by movement along a link, with the speed indicating proximity to threshold [CIG92].

5.2 Motion as Meaning

The previously reported work focuses on the use of animation as a technique to engage the user or to highlight certain semantic properties of objects without systematically investigating the salient properties of the motion itself. More interesting approaches consider which perceptual and interpretative characteristics of movement may convey meaning. We may roughly classify these as providing insight into *basic motion* (relating to perceptual properties); i*nterpretative motion* (the semantics assigned to different movement types); and *compound motion* (how the movement of objects relative to one another informs the perception of the relationships between them and of the object properties themselves.)

5.2.1 Basic Motion

Ware et. al. examined the basic motion parameters of velocity, amplitude, phase and frequency in both signalling [WBKC92] and data correlation tasks [WL94]. He used smooth linear motion as a "human interrupt" signal in the interests of seeing whether this would evoke the same direct pull of attention as blinking or flashing without causing the anecdotally reported associated irritation [Woo95a]. Subjects performed a primary task and were told to respond by hitting a key when they noticed movement of one of two small icons on either side of the top of the display. The icon was a small bar which grew and shrank vertically in an smooth, oscillatory fashion. Amplitude, side and velocity of the movement were varied. There was no effect for amplitude or side, but increases in **velocity** led to an increase in the number of quick responses and a decrease in the number of long ones. The good average response times (1.96 sec in the condition with the fastest velocities) indicated that subjects had no trouble noticing the interruption without any reported irritation factor. Even the slowest times were acceptable, suggesting that motion of this kind is a reasonable "attention getter". The experimenters posit that velocity in this case intuitively indicates urgency (*value*).

Relative **phase** proved an effective parameter for visually associating data (*integration*) [WL94]. Sinusoidal motion of data in a multidimensional scatter plot was compared with the static display dimensions of grayscale value, point size and position in a data correlation task. The three motion parameters were frequency, phase and amplitude, but only two were available simultaneously, since varying the relative phase is only meaningful if the objects are oscillating at the same frequency. The conventional scatter plot (x,y position) proved to be the most effective display technique, but the next most efficient was plotting vertical phase against X position, and it was "not significantly worse" than the scatter plot, and actually better than point size or grayscale value, two common techniques. Ware suggests an ecological basis for the percept of grouping objects in phase since natural phenomena such as schools and flocks move in similar phase. He expresses some doubt whether phase is effective with two or three groups moving simultaneously in different phases [War97]. Michotte found that *kinematic integration* took place with groups of objects moving simultaneously at the same **speed** and in **parallel directions**; i.e., the multiple movements were perceived as a single perceptual event [Mic63]. Bassili reports a similar result that elements are perceived to form a group when their motions share vector components [Bas76].

5.2.2 Interpretative Motion

Johannson [Joh73] and Bassili [Bas78] found that basic biological motion perception appears to differentiate between rigid (*mechanical*) and non-rigid (*animate*) movement. Bassili considered at a "moving lights" display of facial motion and established a relation between the rate of change in the facial motion and the degree of surprise. The duration of the motion indicated secondary emotional conditions (e.g., "fleeting disgust") [Bas78]. Lethbridge and Ware add two further conditions for the perception of animacy (from Marion et.al.): intentionality and a certain degree of randomness, in the sense of lack of repetition [LW90]. Thus, mechanical motion is perceived to be repetitive, automatic and constrained (in the sense that the pattern of motion does not deviate from a predictable path). Animacy is perceived from fluid, spontaneous, responsive, intentional and "free" (i.e., a certain erraticness and deviation from a strict pattern).

Kassin [Kas81] reports on studies by Tagiuiri who had subjects observe animated dots varying movement **angle**. He found that straight, linear paths gave the impression that the objects were "alert, well-reasoned, persevering, logical and ambitious"; erratic paths elicited the perception of "drunk, confused, immature, emotional, undependable and careless"; arched paths were "nonchalant, leisurely, relaxed and complacent". Michotte [Mic63] found that movement **speed** was important. Rapidity gave an impression of violence, slow movement an impression of gentleness, sudden reduction in speed was interpreted as hesitation and sudden, repeated variations in speed elicited impressions of nervousness and agitation.

Amaya et. al. [ABC96] used signal processing techniques to analyse emotion in motion. They captured movements from subjects performing two types of human activity, drinking from a cup and knocking on a door, in three "emotional contexts" (angry, sad and neutral). They identified two attributes which varied considerably over the different emotional movements: **speed** (frequency) and **spatial amplitude** (the range, or size of the motion). They divided the movements into "basic periods" (e.g "hand to cup", "cup to mouth", "cup down", "hand back"); determined the speed of the end effector along its trajectory in both the angry and neutral movements., and calculated speed transforms for both the neutral and angry movement by integrating the longest period along the trajectory and dividing it into frame "templates". Then the transform can be applied to a new, neutral movement by *time-warping* the frame distribution, interpolating between frames where needed. The spatial amplitude intensity for each joint is calculated for each movement period and nonlinear signal amplification is used to apply the amplitude transform to generate new joint positions from the new, neutral movement. They tested the approach by deriving angry and sad tranforms from the cup-drinking data. Then they applied the transforms to the neutral knocking data, and found a close match between the generated and the "real" (motion-captured) angry knocking motion data.

Traditional character animators have long relied on movement to convey the personalities and intentions of their characters. Techniques like *squash-and-stretch* effects, exaggerated object deformation and motion blur model the effects of forces. Anticipatory action, which as previously discussed is a small movement in the counter-direction of the "real" action, subliminally predicts intention and highlights the characters's subsequent move. Movement preparation and diminution (slow-in, slow-out), acceleration and deceleration fix the impression of the object's physical properties - weight, mass, and power. The smoothness of the movement (implemented by animators as transitions, or "inbetweens" to traditional animators and "keyframes" in the computer field [Las87]) affects how it is interpreted. Disney animators Thomas and Johnston give an example of this in describing head movement: a quick transition (no inbetweens) is seen as an abrupt hit to the head (even where no projectile is drawn); a few inbetweens portray a "nervous" subject, dodging something,; more inbetweens indicate a crisp nodding gesture (becoming "friendlier" as more transitions are added; and finally seven inbetweens give an impression of the subject easily craning his head to look at something, with no sense of urgency [TJ81].

Trajectory is also important in character animation. As [Las87] points out, in nature arcs are the most efficient paths by which a form can move from one place to another. Therefore, slightly arc'ed trajectories in animation are seen as more "natural" and less disconcerting. Our intuition is that purely "straight" paths are seen as abrupt spatial transitions, perhaps intensifying an impression of urgency but in the execution losing the essence of the action [Las87].

Carlson Vaughan's investigation into how people interpret emotion from movement in [Vau97] revealed four distinguishing characteristics: *path* (line the movement creates), *area* (use of space by the object), *direction* (direction of movement/animation) and speed (speed and tempo of object). Erect, open, slow movement was considered "beautiful"; narrow, cramped and jerky motion was considered "ugly" and mechanistic.

Laban and Lawrence [LL74] use a 4D space to classify human movement based on effort in which the axes are *exertion* (light - strong), *control* (fluent - bound), *effort* (flexible - direct) and *duration* (sustained - quick). Exertion is concerned with strength or weight (**W**); control with space (**S**), effort with flow (**F**) and duration with time (**T**). The eight basic (W,S,T) exertions are "slashing, gliding, pressing, flicking, wringing, dabbing, punching and floating". Figure 1 plots this space, where each corner represents a basic effort. Those connected by lines share two elements and differ in one element only. (Thus punching, for example, shares space and weight attributes with pressing and differs in its time signature.) The fourth observable phenomenon of effort is flow, in which one either "struggles against" or "indulges in" the effort.

They characterize the interpretation of movement as one of understanding and integrating combinations of these four factors. Thus, intense, emotional movement (such as punching) is a combination of strong, bound, direct, and fast; relaxed, calm, exploratory movement maps to light, fluid, flexible and sustained. We note that the Laban characterization of movement concentrates on the forces involved: so, for example, movements can be considered "direct" (controlled

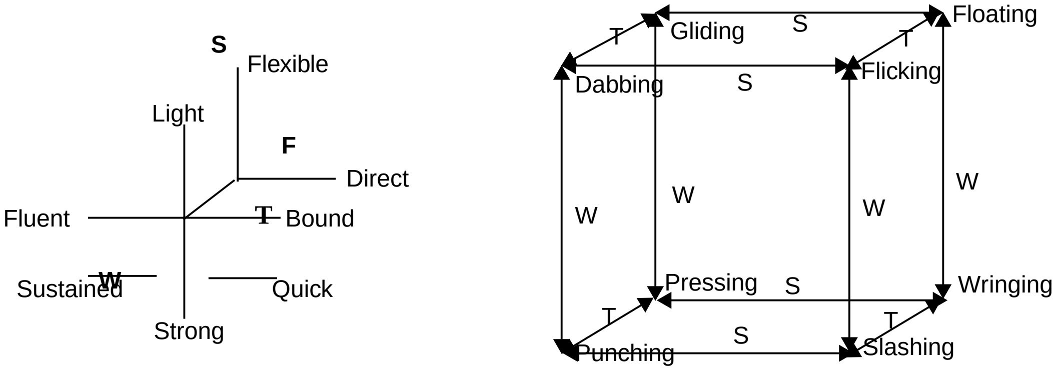

a. 4D “effort space”

0a.

Figure 1. The Laban Movement Space

use of space) and still differ in *absolute amplitude* (e.g. an angry knocking and a wagging finger). However, they may not differ in *relative amplitude*: that is, the extent to which they occupy the possible movement space.

An interesting aspect of Laban’s taxonomy is that one-sided exaggerations of effort may be perceived as having character or behavioural overtones. The authors give examples of laziness (exaggerated sustainment with little influence from other factors), hastiness (exaggerated quickness) and obstinacy (directness).However, Laban and Lawrence do not identify types of communicative motions per se; instead, they assess the effort characteristics of a person’s movement and use it to postulate character traits of the individual. They suggest people who move easily and rapidly seem to be freer as opposed to those who are “struggling against” time. Movements which demand less force or strength suggest relaxation and contentment. Flexibility and twisting in the movement space suggest happy exploration and security in the mover’s sense of space, where direct, constrained movements suggest stress. Controlling flow occurs when an individual tightly controls the progression of movement and indicates reticence. Having laid an initial groundwork for some truly sweeping psychological generations, the authors then conclude than the richness of combining these four factors is so complicated as to be unworkable and suggest abandoning any kind of psychological terms!

Finally, the authors make the point that efforts, and effort-rhythms, can be “transmitted more easily than thoughts”, so that in a dialogue between humans, movements which are independent of the content and object of the communication are still observed and internalized very quickly and powerfully.

Musical conducting is another field in which expression and direction are conveyed through movement. In this domain, the movement tempo is constrained to reflect the desired tempo of the music, and the physical area (where the hand is located when the motion takes place) is related to the organization and direction of the orchestra members and sometimes to signalling change in the dynamics and progression of the piece. In the grammar of musical conducting [Rud80], it appears that **movement amplitude**, **shape** (smooth and curved vs. straight and sharp-edged) and t**emporal continuity** are the expressive dimensions. Small, straight movements are considered

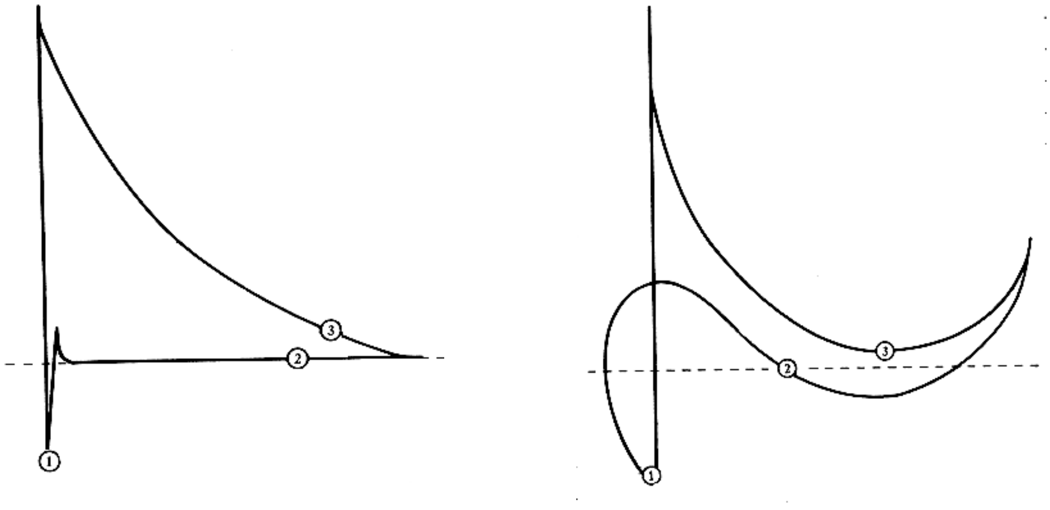

a. Nonespressivo

b. Espressivo

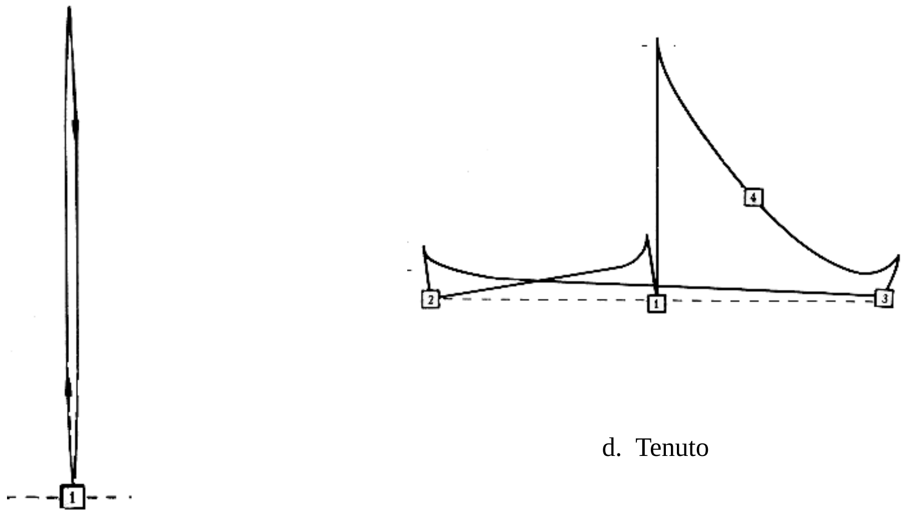

d. Tenuto

c. Marcato

Figure 2. Examples of Conducting Movements [Rud80]

“neutral”; quick, slightly curved motions which stop at the end of each iteration are “bouncy”, “energetic” and positive; and large, curved, swooping movements convey emotion and passion. The dynamics of the music and the volume of sound are expressed by the size (i.e, the spatial amplitude) of the gesture. A conducting *pattern* is a set of movements, the number of which corresponds to the number of beats in the music signature. Motions need not be symmetric in shape but the degree of curvature, spatial continuity (the extent to which the gestures are “sharp” with abrupt changes in trajectory) and temporal continuity (whether there are momentary stops in between beats) are important cues.

Figure 2 shows examples of several different conducting patterns. Figure 2a shows a *nonespressivo* pattern, which is described as a “plain, continuous, neutral” motion, which uses mainly straight lines and has no intensity information. Contrast this with the *espressivo* gesture in Figure 2b, which is curved and continuous, and whose intensity and extent of curvature increases with the “emotional quality” of the music. The *staccato* pattern (not shown) is quick and slightly curved with a stop on each count and bounce on the downbeat, and is characterized as “snappy”, “energetic” and “bouncy”: all terms which imply positive and enthusiastic feeling. The *marcato* pattern, in contrast, (Figure 2c) uses a heavy motion with a stop on each count and is interpreted as “forceful” and “aggressive”. The *tenuto* pattern in (2d) is related to the *marcato* in forcefulness but lacks its aggression. As the musical passion and dynamics intensify so would the curvature of the *marcato* and *tenuto* patterns.

5.2.3 Compound Motion

If we recall that our visual systems have motion-sensitive mechanisms for establishing grouping (section 4.1) then it would seem natural that we have some preattentive capabilities for assessing interactions and relationships between entities from their movement. Bassili [Bas76] and Berry [BS93a] studies of social perception and Michotte’s earlier work on causal attribution [Mic63] have found that simple existence of a temporal contingency is enough to indicate some interaction. i.e., as long as there was some perception of a temporal association, with object actions appearing to occur in some sequence within a small time interval (in the case of the Bassili experiments this was set at 7 frames in a film of 24 frames/sec, or ~290 msec.[Bas76].) Michotte’s experiments suggest that causality is perceived rather than interpreted: that is, a movement of object A (the *motor*) can be seen to cause the subsequent movement of B (the *projectile*) as a direct percept [Mic63]. Causality is perceived under appropriate conditions of time, space and speed of the two moving objects. The temporal interval must be small enough for the movements to be seen as a”whole”. If object speed is fast enough, there can be a fairly wide gap between the objects (50-70 mm). However, the objects must be seen to exist in the same plane. Relative speed acts as an integrating factor, determining, for example, whether A launched B (descending ratio of speed before and after impact) or A triggered B (ascending ratio). The key factor to causal perception is *movement ampliation*, where the movement of the motor object (A) extends into that of the projectile (B). To verify this Michotte considered the perceived effect of A’s movement on a qualitative change (appearance, disappearance, change in form or colour) of object B. When there was no opportunity for movement ampliation, no causality was remarked. Thus merely causing objects to appear to appear and disappear in temporal and spatial contiguity is insufficient for the impression of causality (as in flashing them on and off the screen in close coincidence); some form of kinematic integration must occur. Bassili [Bas76] further investigated temporal and spatial contingencies in social perception, and found that the simple existence of a temporal contingency is enough to indicate an *interaction*. However, the animations in which there was a spatial contingency, i.e., in which the objects were constrained to remain within certain relative distances, were perceived as more meaningful: subjects reported relationships such as chasing, following and hitting. Bassili speculates that the fact that animations in which the vector components of motion were similar were reported as most meaningful suggests that grouping percepts are important in specifying the nature of social interactions [Bas76].

Gibson established that when relative velocity is varied before and after an interaction observers perceive both causality and surface properties of the "objects" (e.g., hardness or softness) [Gib76].

While many social psychologists profess a profound belief that motion plays a fundamental role in social perception, (see [Kas81] for a review) the empirical evidence of which properties of motion are the effective information carriers has not yet been established. In their experiment with the animated triangles and the circle, Heider and Simmel were interested in how factors such as temporal and spatial proximity, range, velocity and direction of movement contributed to impressions about behaviour and its causes, but their experiment design was unsystematic and as a result the data were intractable [Kas81]. Nonetheless, the most interesting findings were a) that movement causes a perception of causality, later substantially explored and verified by Michotte, and b) subjects readily anthropomorphized the objects and movements, "attributing all kinds of emotional states, attitudes, motives and personality traits" to the objects [HS44]. Similar anthropomorphic interpretation was shown by Jetha's experiment in which subjects mapped the motion of individual abstract objects onto human dance sequences [Jet93].

Later, Braitenberg [Bra84] proposed that virtual "vehicles" directed by combinations of simple stimulus-response excitation-inhibition functions could manifest psychologically credible animistic behaviour. In an approach evocative of Braitenberg's "synthetic psychology", Lethbridge and Ware [LW90] used simple *behaviour functions* based on distance, velocity and direction to give the impression of complicated actions and interactions such as chasing, escaping, repulsion, collision and anticipation. They were concerned with achieving the impressions of intentionality and "randomness" to the extent required for the perception of animacy without incurring the computational cost of stochastic parameters. They modelled seemingly random and intentional behaviour of actors in an environment with stimulus-response functions calculating the velocity of an object i at time t as a function of the positions of all the objects in the environment at times $t-1$ and $t-2$. T1 functions determined object behaviour based only on positions of all objects at time $t-1$. The only time-dependent variables used in T1 responses are distance and inter-object direction; a multiplicative parameter is used to control speed. These simple (linear) functions can model situations where one object gravitates towards another, including collision, clinging, pushing, pulling, chasing, escaping, attraction and repulsion. T2 behaviours consider $t-1$ and $t-2$ states, and can model characteristic velocity, delay, momentum and anticipatory action.

Distance, velocity, inter-object direction, and temporal dependencies appear to be the major contributing factors to perception of the nature of the interaction. When velocity is too slow, "all impression of animacy is lost". All compound motions and interactions were reported in anthropomorphic, social terms, even when the motion itself exhibited "unnatural" behaviour (such as abrupt changes in velocity) [LW90].

Psychologists, then, give us evidence for the low-level perceptual capabilities to identify groups and certain interactions and relations; intuition and experience from these reported investigations, simulations and the long history of movement in the performing arts as communication confirm that movement, even of abstract entities, evokes powerful impressions of behaviour. What remains to be established is whether these impressions and perceptions can be usefully manipulated in a display environment to convey meaning about the information space.

6 Motion as a Display Dimension

6.1 Research Issues and Directions

My research is concerned with a principled approach to determining whether motion can be a useful “channel” in information visualization and generally in user interfaces in complex systems. To that end, I ask the following questions:

1. What are the salient perceptual features of motion? What are the emergent and behavioural properties? Can they be “tuned” to influence/alter its meaning?
2. What do motions “mean”? Is there any inherent tendency to assign any semantic association to types of motion? Can motion semantics be divorced from those of the moving object?
3. What is the coding granularity of motion? How many different motions can be used together for coding without interfering with each other? What other modalities reinforce/countermand the effects of motion?

Since we believe that effective decoding of a representation depends not only on mental economy but also its *ecological validity* we need to ask the additional question:

4. What can motion *afford* in the virtual ecology of the complex system interface, and how can we best exploit these affordances?

It seems reasonable that both the perceptual and the interpretative properties of motion need to be investigated together to understand what parameters relate to “meaningful” motion. We find it useful to taxonomize motion’s potential as a dimension for communication according to the three categories discussed in the previous sections: basic motion, interpretative motion and compound motion. Under *basic motion* we are interested in evaluating which basic parameters (or *kinetic primitives*) can be used to code information into simple motions. *Interpretative motion* studies should lead to a categorization of possible qualitative types and which types are most meaningful (can carry most information) under which conditions. Perhaps the most interesting area of investigation is that of *compound motion* to communicate the nature of relationships and linkages between data elements which are visually separated. My research will address whether we can derive useful constructs within this framework i.e., whether we can ascertain which basic characteristics are salient to which types of motion, and how such types of motion may be useful in ameliorating problems discussed earlier in this paper.

6.1.1 Basic Motion and Kinetic Primitives

Evidence from the reported areas indicates that the following may be considered as basic properties of motion: phase, velocity/frequency/rate/speed, periodicity (regular/ random?), trajectory (oscillatory vs. directional, curved vs. linear), position (anchored vs. floating), continuity/smoothness, size/spatial amplitude, transformation (rotation, translation, scale) and sustain/decay (the persistence of the motion’s visual presence along its path). Little is known about the proximity and similarity attributes of these properties. It will be fundamental to examine these to determine the emergent features of motion and moving objects: two possible such effects are grouping and causation.

6.1.2 Interpretative Motion

The type of motion pertains to its behaviour and affordances. (We note that in fact a complex motion may consist of a combination of several types.) Interesting types include autonomic, narrative/illustrative, expressive / intentional, inclusive, autonomous (the degree of autonomy such as passive, caused, active, reactive), locomotive, signal/alert, viewing (manipulating viewpoint), transitive/intransitive (i.e., having a direct or indirect effect on other objects), exertion (‘working”) and “jostling”.

6.1.3 Compound Motion

Compound motion involves a combination of two or more movement sequences which elicits the effect of a single perceptual and interpretative event, unlike simultaneous, similar motions which may have the emergent features of grouping (see above). The key issues in compound motion arise from both the characteristics of the individual movements and temporal and spatial constraints (i.e. how far apart can they be in both space and time before the effect is lost?)

6.1.4 Properties and Types

Table 1 suggests potentially interesting features of basic, interpretative and compoundcompound motion motion discussed earlier to be investigated in the context of this framework. One key aspect of investigation must be combination/exclusion contingencies: to what extent can properties of movement be used in more than one context: e.g., phase for both establishing groups (observer-related) and for identifying aspects of the grouping relationship (object-related).

Table 1: Motion Properties and Types

	Basic	Interpretative	Compound
Single object	**Basic** phase frequency/speed transformation/direction trajectory smoothness/ continuity duration / sustain position amplitude velocity shape/periodicity temporal continuity	signal active viewing “jostling” autonomy locomotive expressive exertion	
Groups of Objects	**Basic** phase frequency/velocity/speed direction duration position continuity trajectory size/amplitude	locomotive exertion expressive autonomy urgent signal	relative velocity relative trajectory sequence transition filmic techniques causation attraction repulsion

7 Potential Applications

Two great advantages of motion in a large and crowded display environment are its perceptual efficiency across a wide area and its compact use of screen real-estate and resources: it does not necessarily increase the density of the display. While extensive studies are required to establish the conditions of appropriate usage, we believe that motion holds great promise for improving visualization and user interface comprehension in the following six areas: most importantly, in annunciation, grouping and integration, and visualizing data relationships.

7.1 Annunciation and signalling

A key issue in supervisory control system interfaces is annunciation: how to ensure that users notice, comprehend and respond appropriately to alarms and system messages in a reasonable response time. In a dense display environment where multiple alarms and signals are concurrently active the user's visual system becomes overloaded and alarms are often missed. We have previously discussed what Woods calls cognitive tools [[woods84]] for selectively directing attention (see Section 2.2.1). We anticipate that motion will prove effective as such a cognitive tool. Since it appears that velocity and amplitude map somewhat intuitively to "urgency", we anticipate that motion can be tuned to represent alarm priority. Because smooth motion is much less disruptive than blinking (see [WBKC92]) the potential exists to gracefully draw the users' eyes to more important areas without disturbing the focus of attention on less important elements by coding the more important alarms with higher frequency motion.

7.2 Grouping and integration

Psychological evidence indicates that the perception of groups is a natural emergent feature of multiple similar motions. We predict that this use of motion may prove invaluable in fostering the immediate recognition of associated elements which may be widely scattered across the visual field. Other display dimensions such as colour and labelling are ineffective in such a broad area and tend to be already over-used.

7.3 Communicating data relationships

Current static graphical visualization techniques are ill-equipped to display dynamically shifting relationships between data elements. There is, for example, no seemingly intuitive way to represent causality, dependencies or even simple sequences without creating new, separate representations with additional abstractions and descriptions. The results from the Michotte [Mic63], Kassin [Kas81] and Heider et. al. [HS44] studies suggest that compound motion holds great promise for meaningfully and efficiently portraying certain relationships "in place"; that is, combining the movements of separate elements in their existing displays and representations in a way that elicits the immediate perception of how the data are related.

7.4 Data display and coding.

There is some evidence that aspects of motion such as phase can be used in a limited fashion to map data values in the same way that point size and colour saturation are used [WBKC92]. However, the obvious application is representing dynamic data such as traffic on a link or flow through a system. Examples include communications traffic in telecommunications, rate of production and flow in mining and petrochemical conduits and line load in power distribution.

7.5 Representing change.

It is important to convey not only the fact that change has occurred in a timely fashion but also the nature and rate of that change, Our visual sensitivity to the continuity, or “smoothness”, of movement and to changes in its frequency and amplitude suggest that we can efficiently and intuitively communicate change by simply varying these parameters, thereby changing the qualitative nature of the motion. For example, we could animate a data representation such as a text element to convey that it had recently changed and change the nature of the movement to indicate to what degree it had done so.

7.6 General visibility concerns

Consider for data that there are two levels of “representational state”: one to do with meaning and direct content (qualitative, quantitative value or state, relative to other data), and one to do with the “display” or “mediated” state of the information. For example, is it currently visible? Is it occluded, or outside the screen “portal”? Is it currently undisplayed but available (i.e., not explicitly excluded by the user) or has it been filtered out? Display state is a combination of size, colour, location, neighbourhood density/display density, and shape/form.

We believe that motion may prove useful in manipulating the display configuration (viewpoint) to draw attention or at least perception to the desired area. In the ecological sense, the user is “walked around” to see information of note. Motion which is used to disambiguate 3D structures, smooth transitions and perceptually group data is related to the display state rather than to basic meaning. “Jostling” windows or otherwise animating icons to indicate where information may be hidden is one potential example.

8 Implementation Issues

8.1 Perceptual Artifacts

Sekuler [STK81] identifies two confusing and contradictory artifacts of motion perception which must be avoided in temporal displays.

- The *Motion After-Effect* (*MAE*) occurs after constant observation of a moving pattern for 15 seconds or more. If the motion suddenly stops, the pattern will seem to move slowly in the opposite direction. The MAE arises only if the stationary pattern is on the same part of the retina as was the moving pattern; an eccentrically viewed MAE appears faster and lasts longer than one in the center of the visual field. It can be cancelled by a small counter-movement prior to the “real” cessation of motion.
- *Induced motion* is the effect which one set of moving objects exerts on the perceived velocity of another set of moving objects (e.g., a standing train which appears to be moving when the train beside it moves). The perceived speed of an object is reduced if surrounded by faster moving ones and increased if surrounded by more slowly moving ones. A moving frame containing a stationary dot gives the impression of the dot moving. In cases where the relative velocity of objects is a meaningful dimension mechanisms must be investigated to annul this effect.

Braunstein reports on the related effect of *motion parallax* in which objects which move faster appear to be closer, even in a two-dimensional display [Bra90].

8.2 Real-Time Display Requirements

Apparent motion is the illusion of continuous motion resulting from the momentary presentation of objects in orderly locations on the visual field, which is, of course, how discrete frames are perceived to link into movement sequences. The key issue is the required temporal resolution. Differences of 200 msec or more cancel apparent motion [STK81]. "Smooth" motion, however, requires 12-14 frames a second, which we will use as our base requirement of "real-time" performance. (Video resolution is 30 frames/sec: film, 24/sec). Most computer workstations have 60 Mhz interleaved displays: alternating halves of a frame, or "fields", are refreshed every 16.67 msec. The challenge in using motion as a display dimension is to guarantee that sufficient temporal resolution can be maintained to reliably elicit the desired perceptual impression, and to ensure the correct synchronization of movements. Operating systems such as UNIX and NT offer only weak support for real-time programming but we anticipate that the combination of current graphics display capabilities and proper use of multitasking will satisfy the temporal resolution requirement. Of potential interest is the approach taken in the ArtKit animation toolkit [HS94], which uses an *animation dispatch agent* to synchronize animation steps with the redraw cycles.

Animation research has concentrated on veridical motion: realistic, believable simulations of the motion of complex articulated figures which are based on either dynamics, motion capture or procedural models of movement. (The reader is directed to [FvDFH90] for a discussion of computer animation.) Dynamics are based on the physical principles of forces and mass, and while the motion produced reliably creates physically credible sequences, the calculation is extremely computationally expensive, involving numerical solutions of large sets of equations which cannot be executed in real time. Forward kinematics, on the other hand, in which only the geometric and movement properties are specified, can be solved in real-time. The evidence that humans employ kinematic rather than dynamic principles in perceptual operations [CM93] encourages the feasibility of this approach.

Lethbridge and Ware's experience with their system of behaviour functions is an effective example of how believable motion and interactions between sets of objects can be simulated using simple sets of partial response equations. They report that calculating the state of the environment at each time step is of complexity $O(n^2)$ where n is the number of objects in the environment. Real-time performance with simple geometrical objects is easily obtainable, although they do not know whether the approach is extensible to animating articulated figures with many more degrees of freedom.

Film techniques such as dissolves, motion blur and slow-in, slow-out involve 2nd and 3rd order continuity of motion [Las87]. Chang and Ungar [CU93] achieve these effects in real time using colour table animation, in which the colour space is divided into separate colour maps which can be used to cycle through representations. The resulting reduction in available colours was not a drawback to display resolution since they were portraying user interface objects rather than continuously varying data values. In systems where continuously varying colour is used as a display dimension, colour table animation may not be a reasonable technique, and we must consider other options.

9 Conclusion

We believe that motion holds great promise as a dimension for displaying information in user interfaces to complex systems because it is perceptually efficient, interpretatively powerful and currently under-used. Of great interest is its potential to intuitively represent data relationships and higher-order system behaviour which static graphical methods cannot. However, little knowledge exists to guide its application in information display. This paper summarizes types of movement characterization from diverse disciplines and proposes an initial taxonomy of motion properties and application to serve as a framework for further empirical investigation into motion as a useful display dimension. Potentially “codable” properties of motion are identified from perception research. A higher-level categorization of motion types is drawn from a review of interpretative movement. There is evidence of qualitative differences in the motion of a single object as opposed to the combined motions of several objects. An ecological perspective is used to anticipate the applicability of different motion types to different uses Finally, implementation issues are discussed with respect to perceptual artifacts which must be avoided and to minimum requirements for temporal resolution, perceptual synchronization and animation techniques.

10 References

[ABC96] K. Amaya, A. Bruderlin, and T. Calvert. Emotion from motion. In *Proceedings of Graphics Interface 1996*, pages 222–229. Morgan Kaufman, 1996.

[ACem] Inc. Apple Computer, The Macintosh Operating System.

[Alv93] F.L. Alvarado. Visualization of power system security with matlab. Workshop on Visualization Methods for Electric Power Industry Applications, EPRI, October 1993.

[ATP95] M. Adams, Y. Tenney, and R. Pew. Situation awareness and the cognitive management of complex systems. *Human Factors*, 37(1):85–104, 1995.

[Bae81] R. Baecker. Sorting out sorting. Morgan Kaufman, 1981.

[Bai91] L. Bainbridge. Multiplexed vdt display systems. In G. Weir and J. Alty, editors, *Human-Computer Interaction and Complex Systems*, pages 189–210. Academic Press, 1991.

[Bai97] L. Bainbridge. The change in concepts needed to account for human behavior in complex dynamic tasks. *IEEE Transactions on Systems, Man and Cybernetics*, 27(3):351–359, 1997.

[Bas76] J. Bassili. Temporal and spatial contingencies in the perception of social events. *Journal of Personality and Social Psychology*, 33(6):680–685, 1976.

[Bas78] J. Bassili. Facial motion in the perception of faces and of emotional expression. *Journal of Experimental Psychology*, 4(3):373–379, 1978.

[Ber83] J. Bertin. *Semiology of Graphics*. University of Wisconsin Press, 1983.

[BHD95] L. Bartram, F. Henigman, and J. Dill. The intelligent zoom as metaphor and navigation tool in a multi-screen interface for network control systems. In *Proceedings of IEEE Conference on Systems, Man and Cybernetics ’95*, 1995.

[BHDH95] L. Bartram, A. Ho, J. Dill, and F. Henigman. The continuous zoom: A constrained fisheye technique for viewing and navigating large information spaces. In *Proceedings of ACM UIST ’95*, 1995.

[BODH94] L. Bartram, R. Ovans, J. Dill, and W. Havens. Contextual assistance in user interfaces to complex, time-critical systems: The intelligent zoom. In *Proceedings of Graphics Interface ’94*, pages 216–234. Morgan Kaufman, 1994.

[Bra84] V. Braitenberg. *Vehicles, Experiments in Synthetic Psychology.* The MIT Press, Cambridge, Massachusetts, 1984.

[Bra90] M. Braunstein. Human performance models for computer-aided engineering. pages 89–105. Academic Press, 1990.

[Bri90] D. Bridgeland. Simulacrum: A system behavior example editor. In T. Ichikawa, E. Jungert, and R. Korfhage, editors, *Visual Languages and Applications*, pages 191–202. Plenum Press, New York, 1990. behavior functions and meaningful motion.

[Bro88] M.H. Brown. Perspectives on algorithm animation. In *Proceedings of CHI ’88*, pages 33–38. ACM SIGCHI, 1988.

[BS90] R. Baecker and I. Small. Animation at the interface. In B. Laurel, editor, *The Art of Human-Computer Interface Design*, pages 251–267. Addison-Wesley, 1990. taxonomy of animation uses in user interface.

[BS93a] D. Berry and K. Springer. Structure, motion, and preschoolers' perceptions of social causality. *Ecological Psychology*, 5(4):272–283, 1993.

[BS93b] K. Bharat and P. Sukaviriya. Animating user interfaces using animation servers. In *Proceedings of UIST '93*, pages 69–79, November 1993. UI animation implementation, for demonstration.

[BSM91] R. Baecker, I. Small, and R. Mander. Bringing icons to life. In *Proceedings of CHI '91*, pages 1–6. ACM SIGCHI, April 1991.

[Byb92] L. Byblow. Bc tel pscc information processing report. Technical Report IGI92-0039, Intelligent Graphic Interface Project, 1992.

[Cas91] S.M. Casner. A task-analytic approach to the automated design of information graphic presentations. *ACM Transactions on Graphics*, 10(2):111–151,, April 1991.

[CCMG97] A. Chappell, E. Crowther, C. Mitchell, and T. Govindaraj. The vnav tutor: Addressing a mode awareness difficulty for pilots of glass cockpit aircraft. *IEEE Transactions on Systems, Man and Cybernetics*, 27(3):372–385, 1997.

[CIG92] CIGRE. Man-machine interface using full graphics (draft). Technical Report SC35 WG01, ConfÈrence Internationale des Grands RÈseaux Electriques, 1992.

[Cle93] W. Cleveland. *Visualizing Data*. AT&T Bell Labs, New Jersey, 1993.

[CM93] L. A. Cooper and M. P. Munger. Extrapolating and remembering positions along cognitive trajectories: Uses and limitations of analogies to physical motion. In Eilan et al. [EMB93], pages 112–131.

[Cut86] J.E. Cutting. *Perception with an Eye for Motion*. MIT Press, 1986.

[DAPG96] A. De Angeli, D. Petrelli, and W. Gerbino. Interface features affecting deixis production: A simulation study. Technical report, Department of Psychology, University of Trieste, 1996.

[DBB91] J. Dickinson, L. Byblow, and W. Byblow. Transalta human factors report. Technical Report IGI91-0016, Intelligent Graphic Interface Project, 1991.

[DdKCV91] F. Decortis, V. de Keyser, P. Caccabiue, and G. Volta. The temporal dimension of man-machine interaction. In G. Weir and J. Alty, editors, *Human-Computer Interaction and Complex Systems*, pages 51–72. Academic Press, 1991.

[DFAB93] A. Dix, J. Finlay, G. Abowd, and R. Beale. *Human-Computer Interaction*. Prentice-Hall, 1993.

[EMB93] N. Eilan, R. McCarty, and B. Brewer, editors. *Spatial Representation: Problems in Philosophy and Psychology*. Blackwell, 1993.

[FvDFH90] J.D. Foley, A. van Dam, S.K. Feiner, and J.F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, 1990.

[FW96] D. Fleet and C. Ware. Visualizing communications: Motion perception experiments. In *Proceedings of IEEE Visualization '96*, 1996.

[Gav93] W.W. Gaver. What in the world do we hear? an ecological approach to auditory event perception. *Ecological Psychology*, 5(1):1–29, 1993.

[GGB89] W. Gilmore, D. Gertman, and H. Blackman. *User-Computer Interfaces in Process Control: A Human Factors Engineering Handbook*. Academic Press, San Diego, 1989.

[Gib76] J.J. Gibson. *The Ecological Approach to Perception*. Houghton Miflin Co., Boston, 1976.

[GSAL96] J. Gronquist, W. Stehares, F.L. Alvarado, and R. Lasseter. Animated vectors for visualization of power system phenomena. *IEEE Transactions on Power Systems*, 11(1):267–273, February 1996.

[GWB96] C. Gobrecht, C. Ware, and V. Bhavsar. 3-d visualization of message passing in distributed programs. Technical Report TR96-103, Faculty of Computer Science, University of New Brunswick, January 1996.

[Hea96] C. Healey. Choosing effective colours for data visualization. In *Proceedings of IEEE Visualization ’96,* 1996.

[HM90] C. Hughes and J. M. Moshell. Action graphics: A spreadsheet-based language for animated simulation. In E. Jungert T. Ichikawa and R. Korfhage, editors, *Visual Languages and Applications,* pages 203–235. Plenum Press, New York, 1990. behavior functions and meaningful motion.

[HP] Inc. Hewlett-Packard. The vue user interface.

[HS44] F. Heider and M. Simmel. An experimental study of apparent behavior. *American Journal of Psychology*, 57:243–259, 1944.

[HS94] S. Hudson and J. Stasko. Animation support in a user interface toolkit: Flexible, robust and reusable abstractions. In *Proceedings of UIST ’94,* pages 57–67, November 1994.

[IJK90] T. Ichikawa, E. Jungert, and R. Korfhage, editors. *Visual Languages and Applications*. Plenum Press, New York, 1990. behavior functions and meaningful motion.

[Jet93] Z. Jetha. On the edge of the creative process: An analysis of human figure animation as a complex synthesis task. Master’s thesis, School of Engineering Science, Simon Fraser University, 1993.

[JH70] B. Julesz and R.I. Hess. Inability to perceive the direction of rotation movement of line segments. *Nature,* pages 243–244, 225 1970.

[Joh73] G. Johannson. Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics,* 14(2):201–211, 1973.

[Joh75] G. Johannson. Perception of motion and changing form. *Scandinavian Journal of Psychology,* 5(3):181–298, 1975.

[Kas81] S. Kassin. Heider and simmel revisited: Causal attribution and the animated film technique. *Review of Personality and Social Psychology,* (3):145–169, 1981.

[KT92] D. Kahneman and A. Treisman. The reviewing of object files: Object-specific integration of information. *Cognitive Psychology,* 24(2):175–219, 1992.

[Las87] J. Lasseter. Principles of traditional animation applied to 3d computer animation. *Computer Graphics*, 21(4):35–44, August 1987.

[LL74] R. Laban and F.C. Lawrence. *Effort: Economy of Human Movement*. Macdonald and Evans, 1974.

[LW90] T. Lethbridge and C. Ware. Animation using behavior functions. In Ichikawa et al. [IJK90], pages 237–252. behavior functions and meaningful motion.

[Mac86] J.D. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141, April 1986.

[Mic63] C. Michotte. *The Perception of Causality*. Methuen, 1963.

[MOW97] D.H. Merwin, J.V. O'Brien, and C.D. Wickens. Perspective and coplanar representation of air traffic implications for conflict and weather avoidance. In *Proceedings of the 9th International Symposium on Aviation Psychology*, 1997.

[MRC91] J. Mackinlay, G. Robertson, and S. Card. The perspective wall: Detail and context smoothly integrated. In *Proceedings of ACM CHI '91*, pages 173–181. ACM SIGCHI, April 1991.

[MS97] C. Mitchell and G. Sundstrom. Human interaction with complex systems: Design issues and research approaches. *IEEE Transactions on Systems, Man and Cybernetics*, 27(3):265–273, May 1997.

[Mye85] B. Myers. The importance of percent-done process indicators for computer-human interfaces. In *Proceedings of ACM CHI '85*, pages 11–17. ACM SIGCHI, 1985.

[NSC96] NSCA. Ncsa digital gallery cd-rom. CD-ROM, 1996.

[oA86] Instrument Society of America. Graphic symbols for process displays: American national standard. Technical Report ANSI/ISA-S5.5-1985, Instrument Society of America, 1986.

[PBF$^+$93] Z. Pylyshyn, J. Burkell, B. Fisher, C. Sears, W. Schmidt, and L. Trick. Multiple parallel access in visual attention. *Canadian Journal of Experimental Psychology*, 1993.

[PE91] S. Palmiter and J. Elkerton. An evaluation of animated demonstrations for learning computer-based tasks. In *Proceedings of ACM CHI '91*, pages 257–263. ACM SIGCHI, April 1991.

[Rey87] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, August 1987.

[RMC91] G. Robertson, J. Mackinlay, and S. Card. Cone trees: Animated 3d visualizations of hierarchical information. In *Proceedings of ACM CHI '91*, pages 189–195. ACM SIGCHI, April 1991.

[Rud80] M. Rudolf. *The Grammar of Conducting*. Schirmer, 1980.

[Sch89] K. Schlueter. Perceptual synchronization in window systems. M. math, Department of Computer Science, University of Waterloo, 1989.

[SM] Inc. Sun Microsystems. The openwindows user interface.

[Sta91] J. Stasko. Using direct manipulation to build algorithm animations by demonstration. In *Proceedings of ACM CHI '91*, pages 307–314. ACM SIGCHI, April 1991.

[Sta93] J. Stasko. Animation in user interfaces: Principles and techniques. In L. Bass, editor, *User Interface Software*, pages 81–101. Addison-Wesley, 1993.

[STK81] R. Sekuler, D. Tynan, and S. Kennedy, R. A sourcebook of temporal factors affecting information transfer from visual displays. Technical Report JUN-TR540, U.S. Army Research Institute for the Behavioral and Social Sciences, June 1981.

[SW95] N. Sarter and D. Woods. How in the world did we ever get into that mode? mode error and awareness in supervisory control. *Human Factors*, 37(1):5–19, 1995.

[TJ81] F. Thomas and O. Johnston. *Disney Animation: The Illusion of Life*. Abbeville Press, 1981.

[Tuf90] E. Tufte. *Envisioning Information*. Graphics Press, 1990.

[Vau97] L. Carlson Vaughan. Understanding movement. In *Proceedings of ACM SIGCHI '97*, pages 548–549, 1997.

[VCP95] K.J. Vicente, K. Christofferson, and A. Pereklita. Supporting operator problem-solving through ecological interface design. *IEEE Transactions on Systems, Man and Cybernetics*, 25(4):529–545, April 1995.

[War97] C. Ware. Private communication. 1997.

[WAS86] J.A. Wise and N.C. Abi-Samra. The design of display systems for electrical control centers. In *Proceedings of the International Conference on Power System Monitoring and Control*, pages 205–210, 1986.

[WBKC92] C. Ware, J. Bonner, W. Knight, and R. Cater. Moving icons as a human interrupt. *International Journal of Human-Computer Interaction*, 4(4):341–348, 1992.

[WC95] C. Wickens and C. Carswell. The proximity compatibility principle: its psychological foundation and relevance to display design. *Human Factors*, 37:473–494, Sept. 1995.

[WF96] C. Ware and G. Franck. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15(2), April 1996.

[WFF97] C. Ware, G. Franck, and D. Fleet. Narrative vision. Video., 1997.

[WL94] Colin Ware and Serge Limoges. Perceiving data displayed through oscillatory motion. Technical Report TR94-089, Faculty of Computer Science, University of New Brunswick, 1994.

[WML94] C. Wickens, D. Merwin, and E. Lin. Implications of graphics enhancements for the visualization of scientific data: Dimensional integrality, stereopsis, motion and mesh. *Human Factors*, 36:47–61, March 1994.

[WMT96] C.D. Wickens, S. Miller, and M. Tham. The implications of data-link for representing pilot request information on 2d and 3d air traffic control displays. *International Journal of Industrial Ergonomics*, 18:283–293, 1996.

[Won94] A.Y.K. Wong. *The Use of Animation in Computer Assisted Instruction*. PhD thesis, Department of Psychology, Carleton University, January 1994.

[Woo84] D. Woods. Visual momentum: a concept to improve the cognitive coupling of person and computer. *International Journal of Man-Machine Studies*, pages 229–244, 1984.

[Woo95a] D. Woods. The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 1995.

[Woo95b] D. Woods. Toward a theoretical base for representation design in the computer medium: Ecological perception and aiding human cognition. In Flach, Hancock, Caird, and K. Vicente, editors, *An Ecological Approack to Human-Machine Systems I: A Global Perspective*, pages 157–188. Lawrence Erlbaum Associates, 1995.

[WS91] N.J Wade and M. Swanston. *An Introduction to Visual Perception*. Routledge, 1991.