When Is A Picture Worth A Thousand Words? — Allocation Of Modalities In Multimedia Communication

Eduard Hovy* and Yigal Arens
Information Sciences Institute
of the University of Southern California
Marina del Rey, CA 90292-6695
U.S.A.
Tel: 213-822-1511
Email: hovy@isi.edu, arens@isi.edu

May 3, 2002

Abstract

To be written.
Still need: References, comparison to previous work, conclusion.

Draft: Not to be quoted or fully believed — yet

*This author was supported in part by the Rome Air Development Center under RADC contract FQ7619-89-03326-0001.

Population	Location
8,000,000	(30N 31E)
2,200,000	(4S 15E)
1,060,000	(6N 3E)
818,000	(1S 36E)
697,000	(33S 18E)
475,000	(8S 13E)
61,000	(22S 17E)

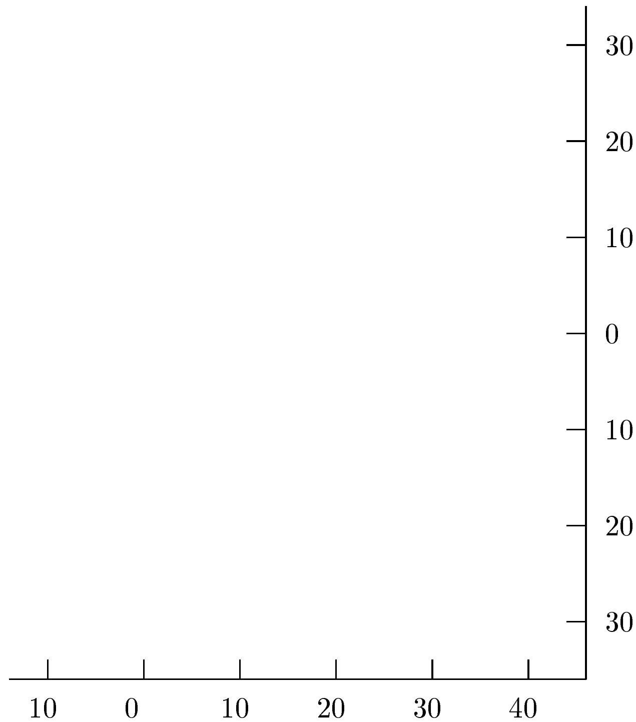

Figure 1: Populations of various cities in Africa.

1 Allocating Information to Modalities

1.1 Introduction

A picture is worth a thousand words. Figure 1 provides information about the populations and locations of various cities in Africa. But it is hard to use. Why? Certainly it provides all the information, contains a map of Africa, and even orders the cities by population size.

Clearly, though all the information is given, the presentation does not facilitate the reader's interpretation of it. In general, designating which information is to be displayed in which manner is a difficult problem. The solution depends on finding modalities with characteristics especially suited to pertinent characteristics of the information to be displayed and then using them to best advantage.

In this paper we describe a representational system for characterizing information and presentation modalities. This representational system provides the necessary basis for the automatic creation of presentations given data and knowledge about the available modalities. We have chosen to concentrate on the problem of creating output, as opposed to input, presentations.

We believe that a user interface's capacity for multi-media communication (henceforth, MMC) is essential for efficient and convenient human-computer communication. This conclusion is suggested to us both by our observation of human-human communication, and by practical considerations concerning the state of user interface technology.

Plainly, humans, when communicating among themselves, routinely utilize multiple media. It is a rare book that does not contain photographs, illustrations, or charts. Speakers in a typical scientific workshop would consider themselves severely handicapped in their ability to convey their thoughts to listeners, if denied use of the overhead and slide projectors. Even then, gestures would certainly be used to enhance the talk and convey some information visually. We believe it is safe to assume that humans select MMC in many situations due to the superior flexibility it provides over communication using any single medium.

From the practical standpoint, no single medium of human-computer communication available today is sufficiently developed to permit its exclusive use for the presentation of a very broad range of information. Natural language communication would appear to be most capable of the widest range of expression, but NLP technology is not yet developed enough to permit its comprehensive use. Given this state of affairs, it is imperative to provide for MMC simply as a means of assuring that all information which needs to be communicated between human and machine can in fact be transmitted.

The representational system we describe in this paper is preliminary. It still requires additions and refinement. We encourage discussions, suggestions, and criticism of our system at the workshop.

The rest of the paper has the following structure: it first outlines the information-to-modality allocation scheme, then defines some basic terminology, then develops characteristics of presentation modalities and characteristics of information. Next, it provides some examples, and finally outlines some problems and further issues to be considered.

1.2 The Basic Idea

The simplest MMC systems operate by assigning a particular modality to each type of data to be presented. A typical rule is:

□ *Rule 1: Ships' locations are presented on maps.*

The approach presented here is based on a two-step generalization of the basic scheme.

The first generalization is to assign a modality not to each data type, but instead to each feature that characterizes data types. Thus instead of rule 1, one writes the rule:

- ☐ *Rule 1': Data duples (of which ships' locations are an example) are presented on maps, graphs, or tables.*

Of course, when considering subsets of features, one invariably gets underspecific rules. To provide more specificity one has to formulate such additional rules as:

- ☐ *Rule 2: Data with spatial denotations (such as locations) are presented on modalities with spatial denotations (such as maps).*

However, note that this rule deals not with the map modality but instead with a characteristic of the map modality. It suggests the second step of the generalization.

The second generalization is to assign characteristics of data not to modalities, but instead to characteristics of modalities. The two example rules now become:

- ☐ *Rule 1: Data duples (of which locations are an example) are presented on planar modalities (such as graphs, tables, and maps).*
- ☐ *Rule 2: Data with spatial denotations (such as locations) are presented on modalities with spatial denotations (such as maps).*

In this example, the two rules together suffice to specify maps uniquely as the appropriate modality for location coordinates. Of course, though, one can present the same information using natural language, as in *"the ship is at 15N 79E"*. Thus one is led to rephrase rule 2 to arrive at a more general (but very powerful) formulation:

- ☐ *Rule 2': Data with specific denotations are presented on modalities which can convey the same denotations.*

Since language, pictures, and maps can carry spatial denotations (while, say, graphs or histograms usually do not), we once again require additional rules in order to specify a unique modality. However, since each of the three mentioned modalities can be perfectly suitable in the right context, the rules we formulate now may not absolutely prohibit modalities; rather, the rules should be context-dependent in ways which enable the selection of the most appropriate modality. Thus we are led to rules such as:

- ☐ *Rule 3: If more than one modality can be used, and there is an existing presentation, prefer the modality/ies that is/are already present as exhibits in the presentation.*
- ☐ *Rule 4: If more than one modality can be used, and there is additional information to be presented as well, prefer modality/ies that can accommodate the other information as well.*

Rule 4 has important consequences. If one is to present not only the location of a ship, but also its heading, then both language and a map would do, since both modalities have facilities for indicating direction (in the case of language, an appositive phrase with the value *"heading SSW"*; and in the case of a map, an icon with an elongation or an arrow). If in addition to this now one adds the requirement to present the nationality of the ship, natural language has such a capability (the adjective *"Swiss"*, say) but due to limitations of the map modality, one of the icon's independent characteristics must be allocated to convey nationality. Of course, this requires the addition of a description of the meaning of the different characteristic values the icon can have and their denotations (for example, a table of color vs. nationality). The extra presentation overhead makes a map a less attractive modality than natural language for presenting a single ship's location-heading-nationality data (but not, perhaps, for a number of ships).

The algorithm for modality allocation can be described as a constraint satisfaction problem, where the constraints arise from rules requiring data characteristics to be matched up optimally with modality characteristics, given data to present, and using one or more modalities from a given collection. In somewhat more detail:

The algorithm derives its power from the two-step generalization away from fixed data-to-modality allocation rules. Our rules select *all* modalities that are capable of satisfactorily presenting the desired data. The final choice of modality can then be left to other criteria unrelated to the data, allowing more freedom than in a scheme forcing a single modality for any type of data. This added freedom can be used tailor the presentation to suit the specific context in which it is being created, the purpose for which it is being made, or particular preferences of the user.

A similar argument holds for data type features: rules refer to only those features of the data type that are pertinent to presentation, and are not complicated by references to irrelevant considerations.

In both cases, leeway is created that can be exploited when multiple pieces of information must be presented simultaneously. The appropriate mechanism is some type of constraint satisfaction or planning/negotiation process.

To help make the algorithm clear, examples are given in Section 5.

2 A Vocabulary of Features

2.1 The Need for a Vocabulary

One of the central problems of multimodal human-computer communication is the allocation of display resources for presenting various pieces of information.[1] We take the

[1] We assume here that the task of finding precisely what information to be displayed can be separated from the task of planning which information to present in which modality. While this may not be true,

approach that each type of information is characterized by the features salient to its display, and each modality is characterized by the features it is suited to display. The task facing us is to develop a vocabulary of features of information and of modalities which will enable some planning/negotiation process to determine the allocation of information to its most suitable modality for each presentation instance.

What should the vocabulary contain? How can it be simultaneously specific enough the support and describe the abilities and features of quite diverse modalities and still be general enough to allow comparisons to be made between them? The vocabulary should:

- describe all features of the information that are salient for presentation purposes,
- describe all features of presentation modalities that can be utilized to convey information,
- be general enough to allow comparisons and specific enough to differentiate between different modalities and information,
- be easily extended to encompass new capabilities,
- support comparisons and decisions.

Before developing the vocabulary, we define some useful terms.

2.2 Terminology

The following terms are used to describe presentation-related concepts. We take the point of view of the process that generates the modality; where the information consumer's subjective experience may differ, we indicate this.

1. Consumer: The person interpreting a presentation.

2. Presentation: The collection of presentation instances required to deliver a complete collection of data to the consumer.

3. Presentation Instance: A single isolated performance by a computational system during which one coherent portion of the given data is delivered to the information consumer. In the absence of a modality which produces a temporally protracted exhibit (e.g., an animation facility), a presentation instance is static. It usually constitutes one screenful of data.

4. Modality: A single mechanism by which to express information stored in a computer. Examples: tables, forms, maps, beeps, written natural language, spoken natural language. Though modalities may be decomposable into components, we treat

we believe it is a reasonable starting hypothesis.

each modality as an indivisible process ('subroutine') that can be invoked by one call of the presentation manager.

5. Medium: A hardware facility utilized by a modality to realize the expression of the information. For example, the CRT monitor is the medium for displaying information in tables.

6. Exhibit:

- **6.1 Complex Exhibit:** A collection, or composition, of several simple exhibits.
- **6.2 Simple Exhibit:** That which is produced by one invocation of one modality. Examples of simple exhibits would be: a map, a paragraph of text, a table, a beep. Stated differently (using terms defined below), a simple exhibit is created when, given a substrate, a modality is used to generate an information carrier that expresses the given information.

A presentation instance is an exhibit, usually a complex one.

The following terms describe the roles played by the two parts of each simple exhibit:

7. Substrate: The background in a simple exhibit. That which establishes, to the consumer, physical or temporal location, and often the semantic context, within which new information is presented to the information consumer. The new information will often derive its meaning, at least in part, from its relation to the substrate. Examples: a screen (on which information may be presented); a grid (on which a marker might indicate the position of an entity); a page of text (on which certain words may be emphasized in red); a noun phrase (to which a prepositional phrase may be appended). An empty substrate is possible.

8. Information Carrier: That part of the simple exhibit which, to the consumer, communicates the principal piece of information requested or relevant in the current communicative context. Examples: a marker on a map substrate; a prepositional phrase within a sentence predicate substrate. A degenerate carrier is one which cannot be distinguished from its background (in the discussion below the degenerate carrier is a special case, but we do not bother explicitly to except it where necessary. Please assume it excepted).

9. Carried Item: That piece of information represented by the carrier; the 'denotation' of the carrier.

For purposes of rigor, it is important to note that a substrate is simply one or more information carrier(s) superimposed. This is because the substrate carries information as well[2]. In addition, in many cases the substrate provides an internal system of semantics

[2]Note that from the information consumer's point of view, Carrier and Substrate are subjective terms; two people looking at the same exhibit can interpret its components as carrier and substrate in different ways, depending on what they already know. For example, different people may interpret a

which may be utilized by the carrier to convey information. Thus, despite its name, not all information is transmitted by the carrier itself alone; its positioning (temporal or spatial) in relation to the substrate may encode information as well. This is discussed further below.

10. Channel: An independent dimension of variation of a particular information carrier in a particular substrate. The total number of channels gives the total number of independent bits of information the carrier can convey. For example, one icon can convey information by its shape, color, intensity level, and by its position and orientation in relation to a background map. The number and nature of the channels depend on the type of the carrier and on the exhibit's substrate.

3 Describing Modalities

3.1 Modality Characteristics

In this section we describe how some modalities exhibit an internal structure and then list a number of useful characteristics of modalities in general.

3.1.1 Internal Semantic Systems

Some information carriers exhibit an internal structure that can be assigned a 'real-world' denotation, enabling them subsequently to be used as substrates against which other carriers can acquire information by virtue of being interpreted within the substrate. For example, a map used to carry a region of the world possesses an internal structure — points on it correspond to points in the region it charts. When used as a background for a ship icon, one may indicate the location of the ship in the world by placing its icon in the corresponding location on the map substrate. Examples of such carriers and their internal semantic systems are:

graph tracking the daily value of some index differently as follows: someone who is familiar with the history of the index may call only the last point of the graph, that is, its most recent addition, the information carrier, and call all the rest of the graph the substrate. Someone who is unfamiliar with the history of the index may interpret the whole line plotted out as the information carrier, and the graph's axes and title, etc., as substrate. Someone who is completely unfamiliar with the index may interpret the whole graph, including its title and axis titles, as information carrier, and interpret the screen on which it is displayed as substrate.

Carrier	Internal Semantic System
Table	categorization according to row and column
Graph	coordinate values on graph axes
Map	'real-world' spatial location based on map denotation
Form	'real-world' denotation according to form field
Picture	'real-world' spatial location based on picture denotation
NL sentence	'real-world' sentence denotation
Ordered list	ordinal sequentiality

Other information carriers exhibit no internal structure. Examples: icon, beep, and unordered list.

An internal semantic system of the type described is always intrinsic to the item carried.

3.1.2 General Useful Characteristics of Modalities

In addition to the internal semantics listed above, modalities differ in a number of other ways which can be exploited by a presentation manager to communicate effectively and efficiently. The values of these characteristics for various modalities are described in Table 1.

Carrier Dimension: Values: *0D, 1D, 2D*. Definition: A measure of the number of dimensions usually required to exhibit the information presented by the modality. That is, whether the modality presents information as a single quantity (0D), presents it linearly in space or time (1D), or presents different portions of the information in different orders.

Internal Semantic Dimension: Values: *0D, 1D, 2D, >2D, 3D, #D, ∞D*. Definition: A measure of the number of dimensions present in the internal semantic system of the carrier or substrate; that is, the dimensionality of the information presented. A single binary bit is *0D*, a pair of values *2D*, a list of entities *#D*, a complex but undecomposable entity (such as a picture) *∞D*.

Temporal Endurance: Values: *permanent, transient*. Definition: An indication whether the created exhibit varies during the lifetime of the presentation instance.

Granularity: Values: *continuous, discrete*. Definition: An indication of whether arbitrarily small variations along any dimension of presentation have meaning in the denotation or not. If so, the modality is *continuous* (e.g., graphs); if not, it is *discrete* (e.g., tables).

Medium Type: Values: *aural, visual*. Definition: What type of medium is necessary for presenting the created exhibit.

Generic Modality	Carrier Dimension	Int. Semantic Dim.	Temporal Endurance	Granularity	Medium Type	Default Detectability	Baggage
Beep	0D		transient	N/A	aural	high	
Icon	0D		permanent	N/A	visual	low	
Map	2D	2D	permanent	continuous	visual	low	high
Picture	2D	3D	permanent	continuous	visual	low	high
Table	2D	2D	permanent	discrete	visual	low	high
Form	2D	>2D	permanent	discrete	visual	low	high
Graph	2D	2D	permanent	continuous	visual	low	high
Ordered list	1D	#D	permanent	discrete	visual	low	low
Unordered list	0D	#D	permanent	N/A	visual	low	low
Written sentence	1D	∞D	permanent	discrete	visual	low	low
Spoken sentence	1D	∞D	transient	discrete	aural	medhigh	low
Animated material	2D	3D	transient	continuous	visual	high	high
Music	1D	?	transient	continuous	aural	med	low

Table 1: Modality characteristics.

Default Detectability: Values: *low, medlow, medhigh, high.* Definition: A default measure of how intrusive the exhibit created by the modality will be. This obviously depends on the size/loudness of the exhibit, and on the substrate, but the value given here characterizes a baseline case.

Baggage: Values: *low, high.* Definition: A gross measure of the amount of extra information a consumer must process in order to become familiar enough with the substrate to correctly interpret a carrier on it. Substrates with *high* baggage value are distracting unless the amount of new data presented on them justifies their use.

3.2 How Carriers May Convey Information

As part of an exhibit, a carrier can convey information along one or more **channels**. For example, with an icon carrier, one may convey information by the icon's shape, color, and possibly through its position in relation to a background map. The number and nature of the channels depends on the type of carrier and the substrate.

The semantics of a channel may be *derived* from the carrier's spatial or temporal relation to a substrate which possesses an internal semantic structure; e.g., placement on a map of a carrier representing an object which exists in the charted area. Otherwise we say the channels is *free*.

Among *free* channels we distinguish between those whose interpretation is *independent* of the carried item (e.g., color, if the carrier does not represent an object for which color is relevant); and those whose interpretation is *dependent* on the carried item (e.g., shape, if the carrier represents an object which has some shape).

In the rest of this subsection we discuss these channel classes, ending with a discussion of how temporal variation along certain channels can be used by carriers to convey additional information.

3.2.1 Channels Derived from the Substrate's Internal Semantics

By being used with substrates with internal semantic systems as background, information carriers 'inherit' channels that can be altered to convey meaning. Of course, these channels convey meaning only when interpreted against the particular substrate. For example, interpreted against the neutral visual substrate (an infinite featureless flat plane), an icon can have no orientation, position, or size. But interpreted against a particular visual substrate such as a terminal screen, an icon has size, position, and orientation, all of which are channels through which information may be expressed — although it need not be. The MacIntosh computer uses an icon that varies in shape (presenting a Bomb or Trash Can), but assigns no meaning to the icon's channels of position or of orientation. Some substrate-derived channels are:

Channel	**Characteristic Values**
Position	[coordinates]
Orientation	N/E/S/W...
Size	[relative size]

3.2.2 Free Channels Independent of the Carried Item's Features

Information carriers, whether containing internal semantic systems or not, can be augmented by the addition of one of a set of further properties that can provide additional information channels. The interpretation of an augmented carrier is usually additive: its interpretation against the substrate and the information conveyed via the additional channels are added together. For example, icons on a map may represent the locations of ships. Coloring the icons differently (an independent channel, here) adds information of an arbitrarily designatable kind (indicating ships' nationalities, say). These properties are unrelated of the substrate upon which the carrier appears (assuming that the property's value keeps the carrier distinct from the substrate; as stated above, we are not bothering to treat degenerate cases in this paper). Some independent channels are:

Channel	Characteristic Values
Texture	smooth/hatched/...
VisualIntensity	high/med/low
AudioIntensity	high/med/low

Note that free channels can be used to good effect *across substrates and even modalities*, a use which is unnatural with substrate-derived properties and impossible, by their nature, with substrate-generated internal semantic systems. For example, in a presentation involving written text and a map with various icons, each icon can be assigned a specific color and the text dealing with that icon's denotation can be written in the same color in order to signal the underlying correspondence.

3.2.3 Free Channels Dependent on the Carried Item's Features

When the carried item (the 'real-world denotation' of the carrier) exhibits some feature which matches a channel of the carrier (i.e., when the channel type is the same as that feature which we wish to present), an appropriate value should be chosen for that channel. Thus, for example, if the carried item is a ship, and the carrier an icon, and we wish to present the ship's shape, then the icon's shape channel should be preferred for this purpose. And if due, say, to unavailability of the appropriate icons, another channel must be used to convey the ship's shape, it would still be appropriate to exercise care with the shape channel. If given an identifiable shape, the icon should look like a ship and not like a pineapple or a squirrel.

Some potentially dependent channels are:

Channel	Characteristic Values
Color	red/blue/...
Shape	plane/ship...

3.2.4 Temporal Variation in Carriers

Most of the carrier channels can be made to vary their presented value in time. Time variation can be seen as an additional channel which provides yet another degree of freedom of presentation to most of the other channels. The most basic variation is the alternation between two states, in other words, a flip-flop, because this guarantees the continued (though intermittent) presentation of the original basic channel value. Some values for the temporal variation of channels are:

Channel	**Temporally Varied**
Position	hopping
Orientation	wobbling
Size	pulsating
Shape	throbbing
VisualIntensity	flashing
AudioIntensity	ululating

4 Describing Data

In this section we develop a vocabulary of presentation-related characteristics of data.

Broadly speaking, there are three types of information concerning data that must be considered when designing a presentation for a data item: Properties of the specific item of data which we wish to present; properties associated with the class to which an item of data belongs; and properties of the collection of items that will eventually be presented, and which the current data item is only one member of. A summary of the terms appears in Table 2.

Type	**Characteristic**	**Values**
Data-Item Property	Dimensionality	0D, 1D, 2D, >2D, ∞D
	Transience	live, dead
	Urgency	urgent, routine
Data-Class Property	Order	ordered, nominal
	Density	dense, discrete, N/A
Data-Batch Property	Volume	singular, little, much

Table 2: Data characteristics by type.

In the rest of this section we will go through the various data characteristics and present rules for modality choice that depend on those characteristics.

4.1 Dimensionality

A single data item will often be decomposable as a vector of simple components. For example, data concerning a ship might be a vector of length 3 — the components being the ship's name, and its longitude and latitude. Other data, such as that contained in a photograph, has a complex internal structure which is not decomposable. We define the *dimensionality* of the latter as *complex*, and of the former as the dimension of the vector.

Since all the data must be represented in some fashion, the following must hold (where *simple* dimensionality has a value of 0, *single* the value 1, and so on, and *complex* the value ∞):

The Basic Dimensionality Rule of Presentations

□ *Rule: Dim(Data) $\leq$ Dim(Carrier) + Free Channels(Carrier) + Internal Semantic Dim(Substrate)*

In addition, we have separate rules that apply to data of differing dimensions. For the purpose of presentation rules we divide data into four classes with respect to dimensionality as follows:

- *Simple:* Simple atomic data entities, such as an indication of the presence or absence of email.

  □ *Rule: As carrier, use a modality with a dimension value of 0D.*

  □ *Rule: No special restrictions on substrate.*

- *Single:* The value of some meter such as the amount of free disk space left. Associated rule is:

  □ *Rule: No special restrictions on substrate.*

- *Double:* Pairs of data components, such as coordinates (graphs, map locations), or domain-range pairs in relations (automobile $\times$ satisfaction rating, etc.).

  □ *Rule: As substrate, use modalities with internal semantic dimension of 2D.*

  □ *Rule: As substrate, use modalities with discrete granularity (e.g., forms and tables) if data-class of both components is discrete.*

  □ *Rule: As substrate, use modalities with continuous granularity (e.g., graphs and maps) if data-class of either component is dense.*

  □ *Rule: As carrier, use a modality with a dimension value of 0D.*

- *Multiple:* More complex data structures of higher dimension, such as home addresses. It is assumed that data of this type requires more time to consume (hence the last rule in this group).

  □ *Rule: As substrate, use modalities with discrete granularity if data-class of all components is discrete.*

  □ *Rule: As substrate, use modalities with continuous granularity if the data-class of some component is dense.*

- ☐ *Rule: As carrier, use a modality with a dimension value of at least 1D.*

- ☐ *Rule: As substrate and carrier, do not use modalities with the temporal endurance value transient.*

- *Complex:* Data with internal structure that is not decomposable, such as photographs.

  - ☐ *Rule: Check for the existence of specialized modalities for this class of data.*

4.2 Transience

Transience refers to whether the information to be presented expresses some current (and presumably changing) state or not. Presentations differ accordingly.

- *Live:* The information presented consists of a single conceptual data item (that is, one carried item) that varies with time (or in general, along some linear ordered dimension), and for which the history of values is not important. Examples are the amount of money owed while pumping gasoline or the load average on a computer. Most appropriate for *live* data is a single exhibit.

  - ☐ *Rule: As carrier, use a modality with the temporal endurance characteristic transient if the update rate is comparable to the lifetime of the carrier signal.*

  - ☐ *Rule: As carrier, use a modality with the temporal endurance characteristic permanent if update rate is much longer.*

  - ☐ *Rule: As substrate, unless the data is already part of an existing exhibit, use the neutral substrate.*

- *Dead:* The other case, in which information does not reflect some current state, or in which it does but the history of values is important. An example is the history of some stock on the stock market; though only the current price may be important to a trader, the history of the stock is of import to the buyer.

  - ☐ *Rule: As carrier, use ones that are marked with the value permanent temporal endurance.*

4.3 Urgency

Some data, or some values of data, may be designated *urgent*, requiring presentation in such a way that the consumer's attention is drawn. This characteristic takes the values *urgent* and *routine*:

- *Urgent:* This situation is exemplified in emergencies, whether they be imminent meltdowns or dangerously little disk space remaining. Rules of modality allocation are:[3]
  - ☐ *Rule: If the data is not yet part of a presentation instance, use a modality whose default detectability has the value high (such as an aural modality) either for the substrate or the carrier.*
  - ☐ *Rule: If the data is already displayed as part of a presentation instance, use the present modality but switch one or more of its channels from fixed to the corresponding temporally varying state (such as flashing, pulsating, or hopping).*
- *Routine:* The normal case.
  - ☐ *Rule: Choose a modality with low default detectability and a channel with no temporal variance.*

4.4 Density

The difference between data that is presented equally well on a graph and a histogram and data that is not well presented on a histogram is a matter of the density of the class to which the data belongs. In the former case we have a *discrete* data class; an example is the various types of car made in Japan. In the latter — a *dense* data class; an example is the prices of cars made in Japan.

- *Dense:* A class in which arbitrary small variations along a dimension of interest carry meaning. Data in such a class is best presented by a modality that supports continuous change:
  - ☐ *Rule: As substrate, use a modality with granularity characteristic continuous (e.g., graphs, maps, animations).*
- *Discrete:* A class in which there exists a lower limit to variations on the dimension of interest. Appropriate modalities are as follws:
  - ☐ *Rule: As substrate, use a modality with granularity characteristic discrete (e.g., tables, histograms, lists).*

[3] The rules should take into consideration all exhibits currently on the screen to make sure the new exhibit actually stands out. We have not addressed that problem yet, and we refer only to *default* detectability of carriers.

4.5 Volume

A batch of data may contain various amounts of information to be presented. If it is a single fact, we call it *singular*; if more than one fact but still little relative to some some task- and user-specific threshold, we call it *little*; and if not, we call it *much*. This distinction is useful because not all modalities are suited to present *much* data.

- *Much:* The relatively permanent modalities such as written text or graphics leave a trace to which the consumer can refer if he or she gets lost doing the task or forgets, while transient modalities such as spoken sentences and beeps do not. Thus the former should be preferred in this case.

  □ *Rule: As carrier, do not use a modality the temporal endurance value transient.*

  □ *Rule: As substrate, do not use a modality the temporal endurance value transient.*

- *Little:* There is no need to avoid the more transient modalities when the amount of information to present is *little*.

- *Singular:* A single atomic item of data. A transient modality can be used. However, one should not overwhelm the consumer with irrelevant information. For example, to display information about a single ship, one should not draw a map.

  □ *Rule: As substrate, if possible use a modality whose internal semantic system has* low *baggage.*

5 Examples

5.1 Example 1

We first present three simple tasks (in parallel, in order to illustrate various aspects of the algorithm).

Given: the task of presenting Paris (as the destination of a flight, say).
Available data (three separate examples): the coordinates of the city, the name Paris, and a photograph of the Eiffel Tower.
Available modalities: maps, spoken and written language, pictures, tables, graphs, ordered lists.

The data characteristics can be tabulated as follows:

	Coordinates	**Name**	**Photograph**
Data	48N 2E	Paris	Eiffel Tower
Dimensionality	*double*	*single*	*single*
Volume	*singular*	*singular*	*singular*
Density	*dense*	*discrete*	*discrete*
Transience	*dead*	*dead*	*dead*
Urgency	*routine*	*routine*	*routine*

The modality characteristics are listed among those in Table 1.

The allocation algorithm classifies data characteristics with respect to characteristics of modalities, according to the rules outlined in Section 4. The modality with the most desired characteristics is then chosen to form the exhibit.

Handling the coordinates data: As given by the rules mentioned in Section4, data with a *dimensionality* value of *double* is best presented in a substrate with a *dimension* value of *2D*. This means that candidate substrates for the exhibit are maps, pictures, tables, and graphs. Since the *volume* is *little*, *transient* modalities are not ruled out. The value *dense* for the characteristic *density* rules out tables. The values for *transience* and *urgency* have no further effect. This leaves maps, and graphs as possible modalities. Next, taking into account the rules dealing with the internal semantics of modalities, immediately only maps remain (maps' internal semantics denote spatial locations, which matches up with the denotation of the coordinates). Given that no prior presentations exist, and that no other information is present, a map modality is selected to display the location of Paris.

Handling the name: The name Paris, being an atomic entity, has the value *single* for the *dimensionality* characteristic. By the appropriate rule (see Section 4.1), the substrate should be the neutral substrate or natural language and the carrier one with *dimension* of *0D*. Since the *volume* is *singular*, a *transient* modality is not ruled out. None of the other characteristics have any effect, leaving the possibility of presenting the single word `Paris` on a blank screen or of making a sentence such as *"The destination is Paris"*, either spoken or written.

Handling the photograph: The photograph has a *dimensionality* value *complex*, for which appropriate rules specify modalities with *internal semantic dimension* of *3D*, and with *granularity* of *continuous* (see Section 4.1) — animation or pictures. Since no other characteristic plays a role, the photograph is simply to be presented using the picture modality.

5.2 Example 2

DRAFT NOTE:
music example

6 Open Issues

DRAFT NOTE:
orientation of lists; axis allocation on graphs and tables; relative size and scale; screen outlay and occlusion; harmonious management of all exhibits together; styles

7 Conclusion

References

[Bertin 83] Bertin, J. *Semiology of Graphics*, trans. by J. Berg. The University of Wisconsin Press, 1983.

[Feiner 88] Feiner, S. An Architecture for Knowledge-Based Graphical Interfaces. ACM/SIGCHI workshop on *Architectures for Intelligent Interfaces: Elements and Prototypes*, Monterey, CA, May 1988.

[Hovy and Arens 90] Hovy, E. A. and Y. Arens. When is a Picture Worth a Thousand Words? — Allocation of Modalities in Multimedia Communication. *AAAI Symposium on Human-Computer Interaction*, Stanford University, May 1990.

[Mackinlay 86] Mackinlay, J. *Automatic Design of Graphical Presentations*. Ph.D. dissertation, Department of Computer Science, Stanford University, 1986.