Cognitive Science Online

Vol 3, Issue 2, Fall 2005

http://cogsci-online.ucsd.edu

Letters

Articles

Department of Cognitive Science, University of California San Diego
9500 Gilman Drive, La Jolla, CA 92093-0515
cogsci-online@cogsci.ucsd.edu

Information

Cognitive Science Online is an online journal published by the UCSD Cognitive Science Department and seeks to provide a medium for the cognitive science community in which to exchange ideas, theories, information, advice and current research findings. This online publication is a peer-reviewed and highly interdisciplinary academic journal seeking contributions from all disciplines and methodologies investigating the mind, cognition and their manifestation in living, and possibly artificial, systems. For more information about this journal, submissions, back issues, please visit our website at http://cogsci-online.ucsd.edu

Contact Information

Department of Cognitive Science
University of California San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515
cogsci-online@cogsci.ucsd.edu

Editors

Advisory Editorial Board

Intersection Point

Sixteen years ago, the newly-formed UCSD Department of Cognitive Science, the first of its kind in the world, welcomed its first graduate students. The graduate program's mission has been to train researchers in the study of intelligent activity - whether carried out by humans, machines, or social groups - using a multidisciplinary approach informed simultaneously by current knowledge about the brain, behavior, and computation. While this mission remains steadfast, some practical aspects of doing research in cognitive science have changed since those formative days.

One such change has been a continued explosion of knowledge in each of cognitive science's contributing disciplines. This is the result of various advances including: the development of functional brain imaging; the sequencing of the human genome and the discovery of alleles associated with variation in cognition and behavior; the diffusion of digital videography, allowing fine-grained analyses of motor and social behaviours; and enormous increases in the power of computational systems to store and process information. This accumulation of new knowledge has accentuated more than ever one of the "special burdens" for researchers in such an interdisciplinary field - namely, "to be knowledgeable in and sympathetic to a large variety of fields and techniques" (as pointed out by the Introduction to the new Department in the 1989 UCSD Catalog). The challenge of keeping up to date with important findings is exacerbated by the sheer increase in volume of new literature - for example, by one estimate the number of scientific journals increased from 2,800 in 1960 to 6,800 in 1995.[1] Amidst this ballooning of scientific knowledge the challenge of integrating the contributions of neuroscience, psychology, anthropology, linguistics, artificial intelligence, and philosophy into one cohesive multidisciplinary field is more daunting than ever.

This burden has become manifest in a variety of forms. At recent meetings of the Cognitive Science Society, for instance, attendees may have noticed a strong division between computationally- and behaviorally-minded researchers, few venturing outside their respective symposia (neuroscientists are generally absent from CogSci meetings). One of CSO's editors was surprised to find his old friend and AI professor on the final day of CogSci 2004 - the meeting had been so compartmentalized that they had not seen each other in two full days of presentations.

Nowhere is the burden more apparent than in our graduate curricula. Cognitive science students face the challenge of integrating a mixture of techniques, theories, and findings, whose potential interrelationships may not be immediately apparent from their primary sources. Moreover, the skills necessary for successful research in any of cognitive science's sub-disciplines have limited applicability to others – for example, a student who is gifted at designing neuroimaging experiments may struggle with cognitive semantics.

The expanding breadth of our field is also apparent in the current issue of CSO. We

---

[1] Tenopir, C., & King, D. W. (2000). Towards Electronic Journals: Realities for Scientists, Librarians, and Publishers. Washington, D.C.: Special Libraries Association.

imagine that the intersection of readers who find both articles interesting and relevant may be small. Even as our journal seeks to "represent the diversity of ideas floating in our highly variegated field of cognitive science, as too often the lines that have traditionally partitioned its sub-disciplines begin to form impenetrable barriers" (as set forth in the inaugural Letter from the Editors),[2] we wonder whether we are truly succeeding in breaking these barriers.

The burdens of our multidisciplinary approach do not simply disappear by gathering Cognitive Scientists into a unifying department, organizing conferences, or showcasing diverse findings in an online journal. Rather, special efforts must be undertaken to make research in the various sub-disciplines more accessible to one another.

One such integrative effort, currently in development at UCSD, is the incorporation of "Datablasts" into departmental gatherings. These quick (<20min), refreshing talks are intended to inform a general audience about current ideas and practices within our far-ranging field. Since Datablasts are presented at department-wide functions, students and professors are exposed to research outside their regular circuit of lab meetings and lecture series, in an informal and interactive setting.

Another, more light-hearted channel that may facilitate communication across our sub-disciplines is introduced in this edition of CSO: Cognitive Science Movies. This index of over 100 popular films encompasses a wide variety of cognitive science themes. Movies, having universal appeal, may serve as starting grounds for discourse between traditionally segmented areas. Starting this quarter, the UCSD Cognitive Science Department will also be hosting Movie Nights, intended to foster more integration and community within our multidisciplinary program.

However, even the most valiant institutional efforts to unify cognitive science's burgeoning subject matter will not lift the burden from its individual researchers. We are each responsible for upholding the multidisciplinary torch, to seek and to appreciate research outside our familiar avenues, and to make special efforts at learning new experimental and observational techniques.

We hope that CSO, as a forum "representing the diversity of ideas" in cognitive science, reminds us of our "special burdens," but more importantly, provides our readership with the opportunity to become better-informed multidisciplinary thinkers.

Michael Kiang & Benjamin Motz

Department of Cognitive Science, UCSD

[2] Lovett, C., Saygin, A. P., & Yu. H. (2003). Letter from the editors. *Cognitive Science Online,* 1, p. i. http://cogsci-online.ucsd.edu/1/vol1_issue1.pdf

Cognitive Science Online, Vol.3.2, 2005 http://cogsci-online.ucsd.edu

Analysis of a Biologically-Inspired System for Real-time Object Recognition

Erik Murphy-Chutorian[1,*], Sarah Aboutalib[2] & Jochen Triesch[2,3]

[1]Dept. of Electrical and Computer Engineering, and
[2]Dept. of Cognitive Science
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0515

[3]Frankfurt Institute for Advanced Studies
Frankfurt, Germany

*Corresponding author. *E-mail address:* `erikmc@ucsd.edu`

Abstract

We present a biologically-inspired system for real-time, feed-forward object recognition in cluttered scenes. Our system utilizes a vocabulary of very sparse features that are shared between and within different object models. To detect objects in a novel scene, these features are located in the image, and each detected feature votes for all objects that are consistent with its presence. Due to the sharing of features between object models our approach is more scalable to large object databases than traditional methods. To demonstrate the utility of this approach, we train our system to recognize any of 50 objects in everyday cluttered scenes with substantial occlusion. Without further optimization we also demonstrate near-perfect recognition on a standard 3-D recognition problem. Our system has an interpretation as a sparsely connected feed-forward neural network, making it a viable model for fast, feed-forward object recognition in the primate visual system.

Introduction

Efficient detection of multiple objects in real-world scenes is a challenging problem for object recognition systems[3]. Natural scenes can contain background clutter, occlusion, and object transformations which make reliable recognition very difficult. In this work we develop a system that efficiently and accurately recognizes partially occluded objects despite position, scale, and lighting changes in cluttered real-world scenes.

Most modern recognition approaches represent specific views of objects as constellations of localized image features. The Scale Invariant Feature Transform, SIFT, is a well-known example (Lowe, 2004). In this approach, gradient histogram-

[3] Authors occasionally make a distinction between *recognition* (what is this object?), *detection* (e.g., is a face somewhere in this image?), and *multiple object detection* (is any of a set of known objects in this image?). In this paper, we use the generic term recognition to refer to all of these problems.

based SIFT descriptors are computed at Difference-of-Gaussian keypoints and stored along with a record of the key-point's 2D location, scale, and orientation relative to the training image. To detect an object in a new image, an approximate nearest neighbor search matches SIFT descriptors extracted from the image, and a Hough Transform detects and roughly localizes the object.

To improve performance for multiple objects, similar approaches have employed a quantized feature vocabulary[4], such that the set of features is shared across different object models (Murphy-Chutorian & Triesch, 2005). In this approach, every extracted local feature is compared to a much smaller set of vocabulary features by a fast nearest neighbor search, and a reference is stored with the key-point's 2D location relative to the location of the object. As the number of objects increases, the number of shared features need not grow proportionally. This benefit from shared features has been corroborated in a boosting framework (Torralba, Murphy, & Freeman, 2004). These authors demonstrated that by allowing only a fixed number of total features, using such shared features greatly outperforms a set of classifiers learned independently for each object class. Vocabulary-based recognition systems have also been proposed for single object recognition and image retrieval (Agarwal, Awan, & Roth, 2004; Leibe & Schiele, 2004; Sivic & Zisserman, 2003). This paper presents a novel framework for sharing multiple feature types, such as texture and color features, within and between different object representations. We learn probabilistic weights for the associations between features and objects so that any feature, regardless of type, can contribute to the recognition in a unified framework.

An interesting debate regarding the aforementioned recognition approaches is the question of how invariance to transformations (position, scale, rotation in plane, rotation in depth) should be achieved. On one end of the spectrum are approaches that try to hard-wire such invariance into the system by using invariant features. At the other end are approaches that try to learn certain invariance directly from training data. Our approach takes an intermediate stance, where position invariance is built into the system, and invariance to scale and pose are learned from training data.

System Overview

In brief, our system works as follows. During training, it creates a set of weighted associations between a learned set of vocabulary features and the set of objects to be recognized. During recognition, vocabulary features that are detected at interest points in the image cast weighted votes for the presence of all associated objects at corresponding locations, and the system detects objects whenever this consensus exceeds a learned threshold. In the following sections we describe these steps in more detail.

Feature Vocabulary

The recognition system uses a vocabulary of local features that quantize a potentially high-dimensional feature space. Our implementation uses color and texture feature

[4] The term *vocabulary* is analogous to the *feature dictionary* used in previous work (Murphy-Chutorian & Triesch, 2005).

vocabularies. The color features are represented as 2D hue saturation vectors, corresponding to the local average of 5x5 pixel windows. The Euclidean distance in polar hue-saturation coordinate space provides the basis for comparing color features. To learn the color feature vocabulary, we extract color features at the locations of objects in a large training set of images and cluster them with a standard K-means algorithm to arrive at our 500 entry color feature vocabulary.

The texture features are 40-dimensional Gabor jets (Lades et al., 1993) comprised of the magnitude responses of Gabor wavelets with 5 scales and 8 orientations, for details see (Murphy-Chutorian & Triesch, 2005). For vertically or horizontally oriented Gabor jets, the necessary convolutions can be efficiently calculated with separable filters. For all other orientations, the image can be first rotated and then processed with the same filters. Our implementation processes all 40 convolutions in approximately 200ms on a 2.8Ghz computer. To compare two Gabor jets, x and y, we use the normalized inner product,

$$S(\mathbf{x},\mathbf{y}) = \frac{\mathbf{x}^T\mathbf{y}}{||\mathbf{x}||\,||\mathbf{y}||}, \tag{1}$$

which is robust to changes in brightness and contrast. By normalizing the vectors and computing only the inner product at runtime, the calculations are reduced. To learn the Gabor feature vocabulary we extract many Gabor jets at interest point locations from around the objects in a large set of training images. As an interest point operator we choose the Harris corner point detector which is highly stable over multiple views of an object (Harris & Stephens, 1988). We use a modified K-means clustering to compute a 4000 entry Gabor jet vocabulary. The modification of the K-means clustering consists of normalizing the jets to unit magnitude following each iteration of the algorithm.

Given either feature type, finding the nearest vocabulary features that best represent it requires a nearest neighbor search in a 2-, or 40-dimensional space, respectively. An approximate kd-tree algorithm accomplishes this efficiently (Mount & Arya, 2005). We have found that the system performs optimally if we use the six nearest Gabor jets and the single nearest color-jet for each respective vocabulary query. As a consequence, our initial encoding of the image in terms of its features is extremely sparse with only 6 out of 4000 Gabor features or 1 out of 500 color features being activated at a given interest point location.

Transform Space

A 2D-Hough transform space (Ballard, 1981; Lowe, 2004) ~partitions the image space into a set of regions or bins for each object. During recognition, the detected vocabulary features cast weighted votes for the presence of an object in a specific bin, storing the consensus for classification[5]. The optimal size of the bins will be discussed in its own section.

---

5 To avoid the problem of boundary effects from the discrete Hough bins, each feature votes for a bin and its 8 neighboring bins.

Feature Associations

Initially, we develop a sparse set of associations between the features and objects. If an object and feature are both present in a training image, the system creates an association between the two. This association is labeled with the distance vector between the location of the feature and the center of a manually drawn bounding box around the object, discretized at the level of the bin spacing. Duplicate associations, (i.e. same feature, same object, same displacement) are disallowed. Once all of the training images have been processed in this way, the system begins a second pass through the training images to learn a weight for each of the associations. Assuming conditional independence between the inputs given the outputs, Bayesian probability theory dictates the optimum weights are given by the log-likelihood ratios,

$$w_{fm\vec{d}} \equiv \ln \mathrm{P}(X_f = 1|Y_{m\vec{d}} = 1) - \ln \mathrm{P}(X_f = 1|Y_{m\vec{d}} = 0), \qquad (2)$$

where X_f is a Bernoulli random variable describing the presence ($X_f = 1$) or absence ($X_f = 0$) of feature f in the scene, and $Y_{m\vec{d}}$ is another Bernoulli random variable indicating the presence or absence of object m at a discretized spatial offset $\vec{d}$ from feature f[6]. Figure 1 shows the distribution of the log-likelihood weights for the color and the texture features. Not surprisingly, the higher-dimensional texture features tend to be more discriminative as they have a higher average log-likelihood weight.

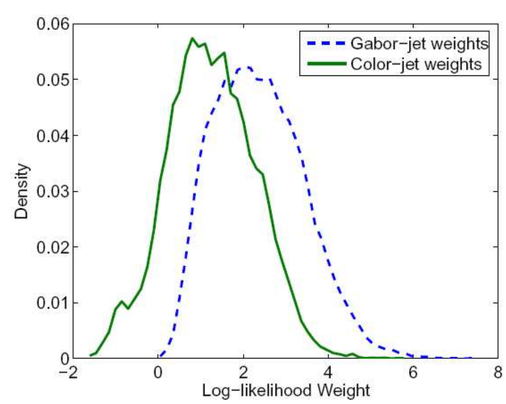

Figure 1. Distribution of Log-Likelihood Weights for each feature type

Optimum Detection Thresholds

During recognition, all of the detected features cast weighted votes to determine the presence of the objects. If any Hough transform bin receives enough activation, this

6 It may seem at this point that a naive Bayes rule expansion could be applied with these log-likelihood ratios and known priors to obtain the posterior probability than an object is present, but the underlying conditional independence assumption is highly erroneous in our case and leads to rather poor performance.

suggests the presence of the object. To determine a detection criterion, we develop optimum thresholds from the maximum a posteriori (MAP) estimator under Gaussian assumptions. Let Y_m be a Bernoulli random variable describing the presence of the object m and let tm be a continuous random variable corresponding to the maximum bin value in the Hough parameter space of m. The MAP estimator, $\hat{y}_m$, describes the most likely value for y_m given the value of t_m:

$$\hat{y}_m = \arg\max_{y_m} \mathrm{p}(t_m|Y_m{=}y_m)\mathrm{P}(Y_m{=}y_m) \tag{3}$$

$$= \begin{cases} 0 & : \quad \mathrm{p}(t_m|0)\mathrm{P}(0) \geq \mathrm{p}(t_m|1)\mathrm{P}(1) \\ 1 & : \quad \mathrm{p}(t_m|0)\mathrm{P}(0) < \mathrm{p}(t_m|1)\mathrm{P}(1) \end{cases}, \tag{4}$$

where $P(y_m)$ is the prior probability that $Y_m = y_m$, and $p(t_m| y_m)$ is the conditional pdf of t_m given $Y_m = y_m$. We then define the optimum threshold, θ_m, as the value of t_m which satisfies

$$\mathrm{p}(t_m|Y_m{=}0)\mathrm{P}(Y_m{=}1) = \mathrm{p}(t_m|Y_m{=}1)\mathrm{P}(Y_m{=}1) \tag{5}$$

For $t_m > \theta_m$ it is more probable that the object is present in the scene, and for $t_m < \theta_m$ it is more probable that the object is absent. Assuming that $p(t_m \mid y_m)$ is a Gaussian distribution, we can fully determine $p(t_m \mid Y_m = y_m)$ knowing only the first and second order moments, μ_{ml} and σ^2_{ml}, where $l = 1$ if the object is present. We estimate the moments from the training data and find θ_m by solving the quadratic equation:

$$\frac{1-p}{\sqrt{2\pi\sigma^2_{m0}}}\exp^{-\frac{(\theta_m-\mu_{m0})^2}{2\sigma^2_{m0}}} = \frac{p}{\sqrt{2\pi\sigma^2_{m1}}}\exp^{-\frac{(\theta_m-\mu_{m1})^2}{2\sigma^2_{m1}}}, \tag{6}$$

where $p = P(Y_m{=}1)$. Assuming $\mu_{m1} > \mu_{m0}$ and $\sigma^2_{ml} > \sigma^2_{m0}$ as is always the case for our data, the solution is given as

$$\theta_m = \frac{-b - \sqrt{b^2 - 4ac}}{2a}, \tag{7}$$

with

$$a = \sigma^2_{m0} - \sigma^2_{m1}$$

$$b = 2(\mu_{m0}\sigma^2_{m1} - \mu_{m1}\sigma^2_{m0})$$

$$c = \mu^2_{m1}\sigma^2_{m0} - \mu^2_{m0}\sigma^2_{m1} + 2\sigma^2_{m0}\sigma^2_{m1}\ln\left(\frac{(1-p)\sigma_{m1}}{p\sigma_{m0}}\right).$$

Experiments and Results

The CSCLAB cluttered scenes database was used to test the performance of our system (Murphy-Chutorian & Triesch, 2005). It consists of 500 scenes of 50 everyday objects against cluttered, real-world backgrounds with significant occlusion. Each scene contains 3 to 7 objects as shown in Figure 2. The objects are presented at roughly the same viewpoint in every scene, but there remains differences in depth, position, rotation, and lighting. The depth changes cause considerable scale variation among the object classes, which vary by a factor of two on the average. The system learns scale-invariant representations by building a conglomerate set of associations from training images of objects at representative scales. Alternatively, it could be trained with fewer scenes, explicitly presented at multiple scales (Burt & Adelson, 1983). In addition, the database contains scenes of all ten backgrounds without objects, as well as scenes of every background with each object by itself. All of the scenes have associated XML files that store the manually-labeled bounding boxes and names of the objects for supervised training and evaluation.

The dataset was split into three sets. The first set contained 100 multiple object scenes which were used to create the feature dictionary. The second set contained 100 additional multiple-object scenes and all of the individual object scenes. This set provided the training data for learning associations between vocabulary features and objects and the corresponding weights. The third set, containing the remaining 200 multiple-object scenes, was presented to the system for recognition.

Figure 2. Labeled Example Scene from the CSCLAB dataset

Feature Sharing

Figure 3 demonstrates the amount of feature sharing in the learned representations for the 50 objects from the CSCLAB data base. In Figure 3(a) we show histograms of how frequently a feature is shared between representations of different objects. Interestingly, there is a sizable fraction of features that are shared by many objects, and only few features are not shared at all, i.e. they are specific to one object only. Figure 3(b) shows how often a feature is shared within one or multiple views of a single object. Noting that there are no duplicate associations, this denotes the number of associations between this feature and the object with different discretized displacements. One can see that this intra-object sharing is happening less often than the inter-object sharing, but this is a meaningless ratio, since it directly depends on our choice of the Hough bin size and number of objects.

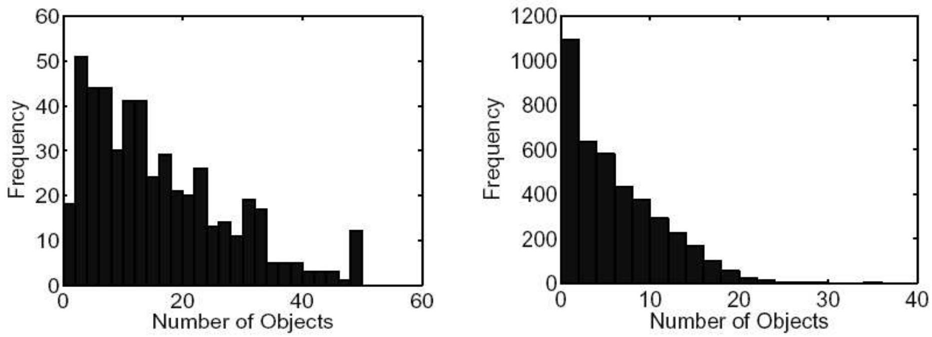

(a) Histogram of inter-object sharing, showing the number of objects that connect to each feature for color-jets (left) and Gabor jets (right).

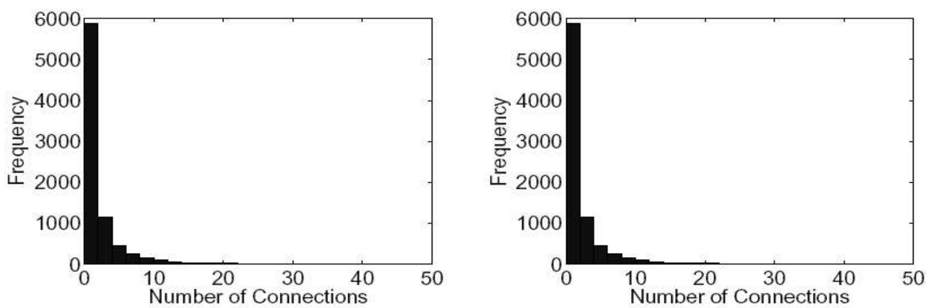

(b) Histogram of intra-object sharing, showing the number of times a feature connects to the same object for color-jets (left) and Gabor jets (right).

Figure 3. Feature Sharing

Optimal Bin Size

The optimal size of the Hough transform bins is determined by a trade-off between two competing factors. If the bin size is too small, votes from the same object may fall into different bins because of variations in object appearance such as scale or rotation. Larger bins, however, increase the risk of a spurious accumulation of votes from background clutter or unrelated objects into a single bin, which can lead to a false positive detection. Because of this trade-off, there exists an intermediate bin size that yields optimal performance (Aboutalib, 2005). We investigated this effect by systematically varying the bin size[7]. Figure 4 shows the result. The tradeoff favoring

---

7 In this experiment we kept the bin size fixed for every object, but an object specific selection of the bin size may further improve performance.

intermediate bin sizes is clearly visible. Based on this result, we use 16x16 pixel Hough transform bins to maximize recognition.

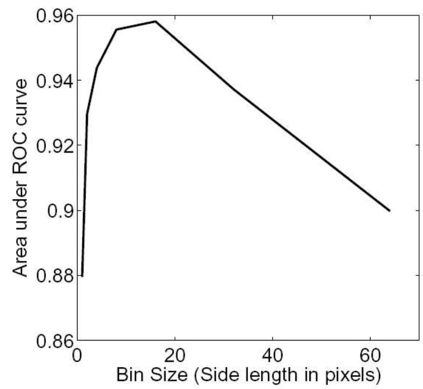

Figure 4. Area under the averaged ROC curves for various bin sizes

Recognition Performance

Figure 5 and Figure 6 show histograms of the detection rates and false positive rates for the 50 objects in the CSCLAB dataset. The detection rate is defined as the fraction of objects that were successfully detected, and the false positive rate is the fraction of images in which an object is incorrectly detected. In this application, the system is able to detect most of the objects more than 80% of the time while maintaining less than a 5% false positive rate. The system has the most difficulty with the objects that lack sufficient texture, or have significant transparencies. Performance examples are shown as ROC curves for the best, median, and worst individual ROC curves are given in Figure 7. Figure 8 shows an average of the spline-interpolated ROC curves for all of the objects. In the course of the experiment, 10 of the 50 objects were perfectly recognized with a 100% detection rate and no false positives. Figure 9 provides examples of the system's recognition ability. On a 2.8Ghz personal computer, our system requires approximately one second to recognize all of the objects in a 640x480 pixel image.

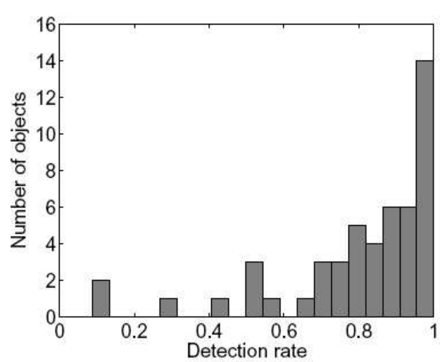

Figure 5. Histogram of the individual object detection rates at optimum thresholds

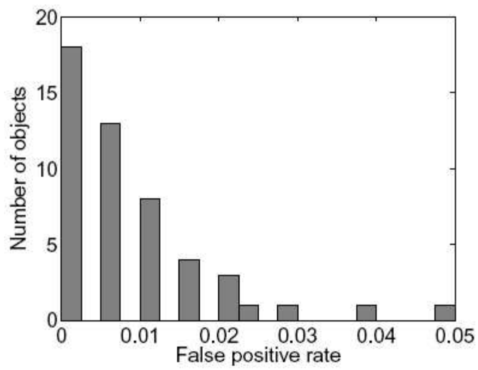

Figure 6. Histogram of the individual object per-image false positive rates at optimum thresholds

Neural Network Interpretation and Relation to Models of Biological Object Recognition

It is frequently argued that the remarkable speed of primate object recognition suggests a processing architecture that is essentially feed-forward in nature, and prominent models of biological object recognition are feed-forward processes (Fukushima, Miyake, & Ito, 1983; Riesenhuber & Poggio, 1999). Feed-forward models are unlikely to be able to account for all aspects of primate object recognition, but they may be a reasonable approximation in many situations.

We can interpret our system as a simple feed-forward neural network. In this case, the input layer consists of the vocabulary features at every possible discretized location. The output layer consists of the objects at every possible Hough bin. The activation of an output node, $y_j = y(m_j, q_j)$, is given by the linear summation,

$$
\begin{aligned}
y_j &= \sum_i w_{ij} x_i, \\
x_i &\equiv x_{(f_i, p_i)}, \\
w_{ij} &\equiv w_{f_i m_j (q_j - p_i)},
\end{aligned}
\tag{9}
$$

and the weights of the network are the log-likelihood ratios of the features mentioned earlier. In this context, x_i is the *ith* binary input node that “fires” whenever the shared vocabulary feature, f_i, is found anywhere inside the unit’s “receptive field,” and *w*ij is the weighted connection between y_j and *x*i. p_i is the location of input x_i, and q_j is the location of y_j . We assume $w_{ij} = 0$ whenever *y*j and x_i are not connected. Although the weighted connections are learned from the relative displacement between the input and output nodes, this can be interpreted as *weight sharing* in a neural network with connections based on absolute displacements.

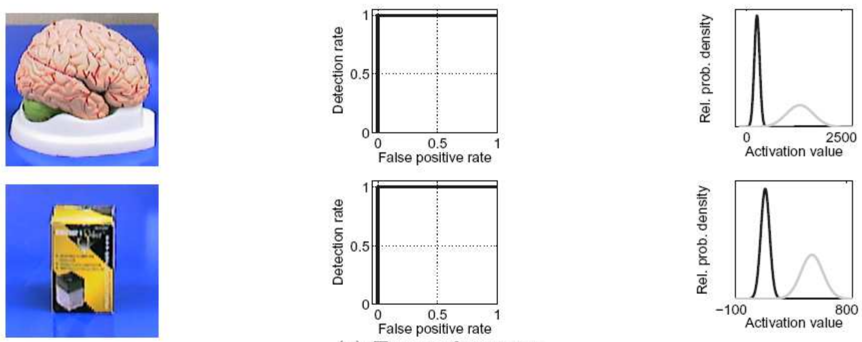

(a) Top performance

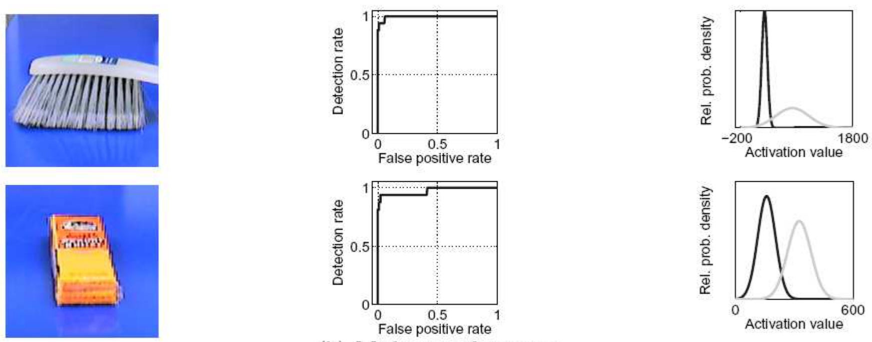

(b) Median performance

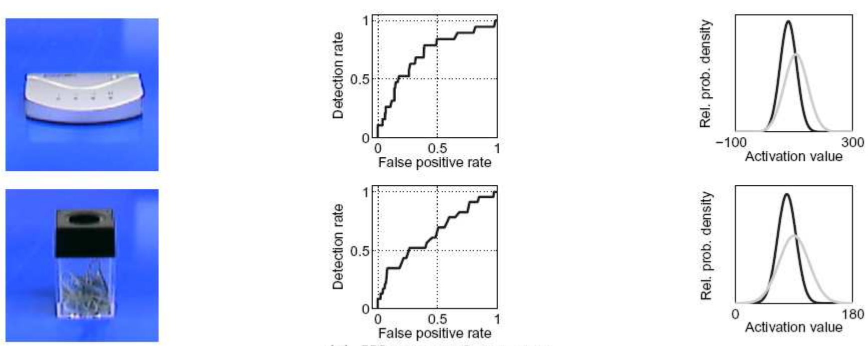

(c) Worst performance

Figure 7. ROC curves and estimated conditional pdfs (black: object absent, gray: object present) for individual object examples

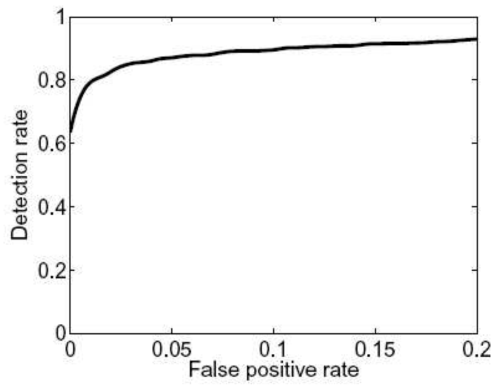

Figure 8. Averaged ROC curve for all 50 objects

The feed-forward neural network interpretation of our system suggests that one could view it as an abstract model of primate object recognition. In fact the introduction of Gabor wavelet features into computer vision systems was inspired by biological findings. In this context, our shared Gabor jet features loosely correspond to shape selective cells in area V4. Compared to the other models mentioned above, the binning operation inherent in the Hough transform mechanism corresponds to non-linear operations that introduce a degree of shift invariance in the above models. The sparseness of connections from these features to object detectors (corresponding to populations of cells in inferotemporal cortex) is also in line with biological considerations, but in stark contrast to many previous models of biological object recognition, we obtain excellent performance on a difficult real-world recognition problem. To do so in real-time paves the way for the development of more elaborate models of visual cognition that model object recognition and learning in the context of ongoing behavior.

Discussion

We presented a new framework for multiple-object detection with a vocabulary of shared features. Using multiple feature types and sparse, weighted associations between vocabulary features and objects, we demonstrated object detection in cluttered real-world scenes despite significant scale variation and occlusion in real-time. Since the system can be interpreted as a feed-forward neural network, it may be viewed as an abstract model of object recognition in the primate visual system, although this was not the main focus of this research.

In a full 3-D recognition task on the otherwise much simpler COIL database, our system showed excellent performance. Evaluating our system on a full 3-D recognition problem that also includes clutter, occlusions, and lighting variations remains a topic for future research. At present, there are no available benchmark databases of this kind. Performance gains could be achieved by the addition of other feature types. Transparent objects and objects lacking unique texture and color were the most difficult to detect, and this could be remedied by the addition of features that could detect these objects by their characteristic shape. The framework presented in this paper easily accommodates additional features. A further avenue for future research is the incorporation of stereo information and the explicit modeling of object occlusions (Eckes, Triesch, & Malsburg, 2005).

We would also like to investigate the ability to learn objects with only minimal supervision, since hand-labeled training data as we have used here is tedious to create. Recent pilot work has demonstrated this system's potential for learning object representations in a semi-autonomous fashion through online demonstration, where objects are simply shown to the system for an extended period of time as they undergo scale and pose changes and the system detects, tracks, segments, and learns to recognize these objects without additional human intervention (Murphy-Chutorian, Kim, Chen, & Triesch, 2005).

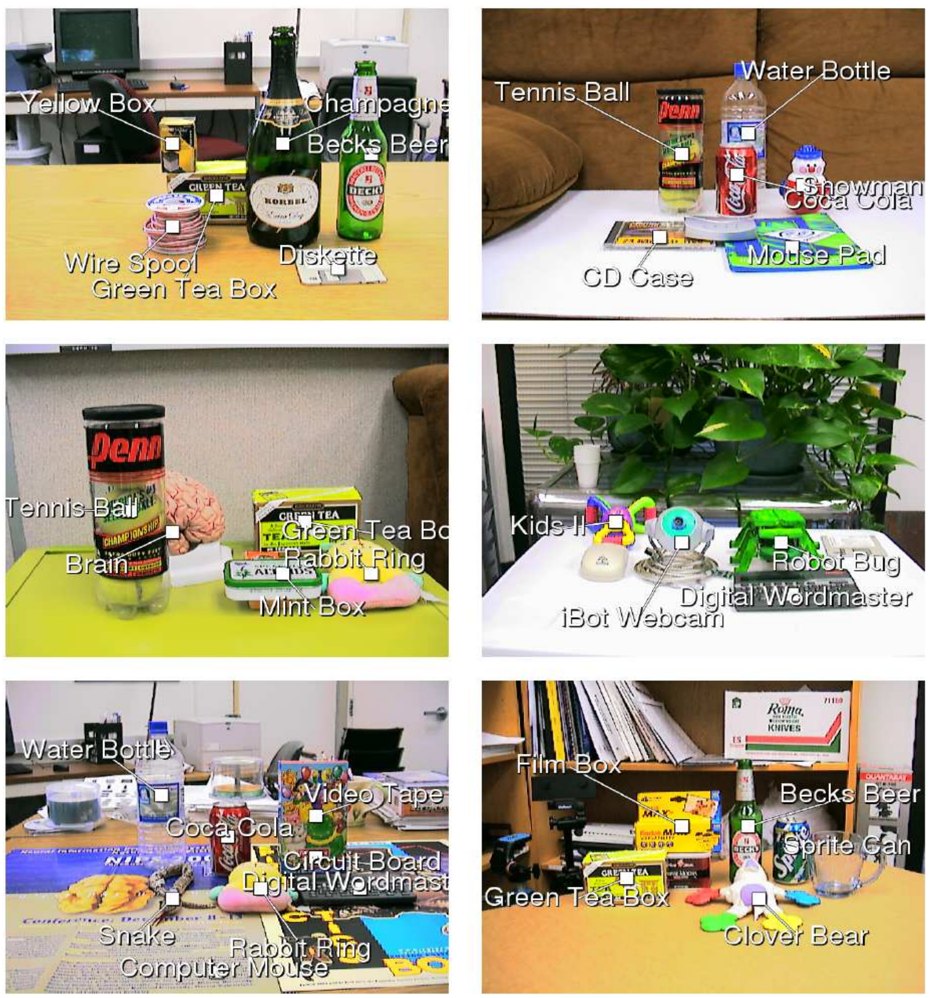

Figure 9. Example Recognition Results (squares indicate the estimated object center)

Acknowledgments

This work was supported by the National Science Foundation under grants IIS-0208451 and IIS-0233899. We would like to thank Hyundo Kim, Huei-Ju Chen, Emily Ashbaugh and Nathan Rosencrans for their help in creating the cluttered scenes database. We would also like to thank Boris Lau, Alan Robinson and Tom Sullivan for their helpful suggestions and programming contributions.

References

Aboutalib, S. (2005). *Position invariance in a view-based object recognition system.* Unpublished honor's thesis, University of California, San Diego.

Agarwal, S., Awan, A., & Roth, D. (2004, November). Learning to detect objects in images via a sparse, part-based representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 26* (11), 1475–1490.

Ballard, D. (1981). Generalizing the Hough transform to detect arbitrary shapes. *Pattern Recognition, 13* (2), 111–122.

Burt, P. J., & Adelson, E. H. (1983). The laplacian pyramid as a compact image code. *IEEE Transactions on Communications, COM-31*,4, 532–540.

Eckes, C., Triesch, J., & Malsburg, C. von der. (in press). Analysis of cluttered scenes using an elastic matching approach for stereo images. *Neural Computation*.

Fukushima, K., Miyake, S., & Ito, T. (1983). Neocognitron: A neural network model for a mechanism of visual pattern recognition. *IEEE Transactions on Systems, Man, and Cybernetics, SMC–13* (5), 826–834.

Harris, C., & Stephens, M. (1988). A combined corner and edge detector. *In Proceedings of the Alvey Vision Conference* (pp. 147–151). Manchester.

Lades, M., Vorbr¨uggen, J. C., Buhmann, J., Lange, J., Malsburg, C. von der, W¨urtz, R. P., et al. (1993). Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers, 42*, 300–311.

Leibe, B., & Schiele, B. (2004). Scale invariant object categorization using a scale-adaptive mean-shift search. *In Proc. deutsche arbeitsge- meinschaft fr mustererkennung pattern recognition symposium*. Tuebingen, Germany.

Lowe, D. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, *60* (2), 91–110.

Mount, D., & Arya, S. (2005, May). ANN: A library for approximate near- 12 est neighbor searching, version 1.1. (http://www.cs.umd.edu/_mount/ ANN/).

Murphy-Chutorian, E., Kim, H., Chen, H.-J., & Triesch, J. (2005). Scalable object recognition and object learning with an anthropomorphic robot head. *In*

Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. San Diego, CA.

Murphy-Chutorian, E., & Triesch, J. (2005, January). Shared features for scalable appearance-based object recognition. *In Proceedings of the IEEE Workshop on Applications of Computer Vision*. Breckenridge, CO, USA.

Nene, S., Nayar, S., & Murase, H. (1996, February). Columbia object image library (COIL-100) (Tech. Rep.). Columbia University.

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience, 2*, 1019–1025.

Sivic, J., & Zisserman, A. (2003, October). Video google: A text retrieval approach to object matching in videos. *In Proceedings of the IEEE International Conference on Computer Vision*. Nice, France.

Torralba, A., Murphy, K., & Freeman, W. (2004). Sharing features: efficient boosting procedures for multiclass object detection. *In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.

Cognitive Science Online, Vol 3.2, pp.1-14, 2005 http://cogsci-online.ucsd.edu

Theory Pictures as Trails: Diagrams and the Navigation of Theoretical Narratives

J.R. Osborn
Department of Communication
University of California, San Diego
jrosborn@ucsd.edu

Abstract

This paper examines diagrams as academic and theoretical tools. Drawing upon the work of Gilles Deleuze and Felix Guattari (1987), a diagram is defined as an abstract machine for constructing arguments. The theoretical diagram provides neither a direct representation of the natural world nor a representation of a natural data set, but a suggested theoretical walk through a landscape of data. It is a tool for learning how to see, how to reason, and how to narrate. The paper begins with a closer examination of diagrammatic thought and the ways in which diagrams differ from other visual representations. It then introduces Vannevar Bush (1945) and follows his idea of associative trails through more recent attempts at modeling semantic associations (Semantica Inc., 2005) and the use of "trails" as narrative markers in the sequential art of comics (McCloud, 1993). These trails, in turn, lead to a discussion of academic work practices, trajectory (Strauss, 1993), and the means of navigating information ecologies (Hutchins, 1996; Bowker & Star, 1999). Finally, the path returns to visualization practices, where it uncovers diagrams as a distinct strategy which scholars may employ as a method of analysis. Along the way, diagrams are offered as both examples and theoretical models. For, among their other benefits, diagrammatic models construct a visual language and represent what is difficult to express in prose.

Introduction

Perhaps out of a desire for intelligibility, we can imagine that, in order to follow a complex trajectory, the human mind begins with simple elements and constructs a cultural object, which outlines both constraints to which it must submit and choices it is able to make

– adapted from A.J. Greimas (1987, p. 48)[8]

[8] The opening quote has been reorganized and adapted from the explanatory note of Greimas' (1968) essay "The Interaction of Semiotic Constraints," which introduces the structure of the now famous semiotic square. In adapting the quote for the purposes of this paper, I have switched the position of two phrases: "construct a cultural object" and "follow a complex trajectory." The original quote, which functions to explain a quite different purpose, is presented below with the altered sections highlighted in italics.

> "Perhaps out of a desire for intelligibility, we can imagine that, in order to *achieve the construction of cultural objects (literary, mythical, pictorial, etc.)*, the human mind begins with

By altering Greimas' note as my opening quote, I also wish to reorganize a corner of discussion regarding visualization strategies. The "cultural objects" to be examined are theoretical diagrams in the social sciences, and the "complex trajectories" are the methodologies of study which these diagrams represent. These diagram objects—such as Greimas' (1987) own semiotic square, Fauconnier and Turner's (2002) basic "blending" diagram, and models of communication transfer—provide trajectories which scholars may utilize in analyzing data. But this is not to say that diagrams force data into a particular form. Much criticism has been directed toward structuralism and its attempts to force data into pre-constructed molds. But rather than viewing the diagram as a closed structure, I ask what the diagram opens up. A diagram is not a stamp placed upon the data. A diagram, as the opening quote suggests, offers a series of choices and constraints, a roadmap of choices for navigating through data. And like geographic maps, diagrams only provide a possible outline or itinerary; they do not determine the specifics of how a journey will unfold.

More generally, I wish to disengage diagrams from the burgeoning field of information and data visualization. Visualization research revels in producing new pictures of large data sets. These images map collected datasets and present a new view of the evidence. But this experimentation of imagery depends upon a data collection, which can be isolated and quantified. The graphic diagram examined in the following pages, however, offers something quite different: the opportunity to present theoretical models in a visual format beyond the formality of written language. Diagrams contain language, but they break the grammar of language. They replace the relations of words and concepts with lines, arrows, and shapes. Decisions, regarding what language to include and what language to replace, rest upon the qualitative judgment and critical choices of those drawing the diagram. The positions of diagram terms are critically chosen from the beginning rather than mapped by a computer for later manipulation. Indeed, the critical positioning of terms, the spatial topology of the diagram, imbues the diagrammatic image with a sense of coherence and meaning.

What, then, is a diagram? In the following discussion, theoretical diagram maps an argument such that it can be approached and contemplated as an image. All images, including evidential photographs and visualized datasets, provide an argument: a series of choices as to what will and what will not be included. But the diagram entails not only a choice of framing but the additional choices of layout and relation as well. A drawn diagram offers a narrative argument, a story of what moving across the image entails. Lines and arrows display a functional relation between terms: this path can be followed in this way. One aspect of academic work, I suggest, is the practice of building methodological tools for navigating ecologies of information. And a diagram is a visual representation of these navigational trails. As the opening quote suggests, diagrams are cultural objects composed of simple elements, and these simple elements allow human cognition to follow a complex trajectory. The diagram is neither a direct representation of the natural world nor a natural data set, but a suggested theoretical walk through the landscape of data.

---

simple elements and *follows a complex trajectory, encountering on its way* both constraints to which it must submit and choices it is able to make." (p. 48)

To guide this walk, a model is drawn. This model cannot predict the events and encounters of a specific stroll, but it can guide the trajectory. It outlines a method for recreating the path at a latter date, if not the specific expressions. This ability to recreate a type of experience parallels the abstract machine of Deleuze and Guattari (1987): "The abstract machine is pure Matter-Function—a *diagram* [italics added] independent of the forms and substances, expressions and contents it will distribute" (Deleuze & Guattari, 1987, p. 141).[9] A diagram is a function of matter, a model for shaping matter. Diagrammatic machines shape matter into a form of expression, and the contents of expression are inextricably tied to the form of their expression (Deleuze &Guattari, 1987). The operations involved in forming an expression are distinct from the contents that form makes possible. These operations are diagrammatic, and their image constitutes the diagram. By way of example, imagine the patent process. In order to patent a machine, a required drafting diagram presents an outline of its construction and functioning. Patents rely upon drawn mechanical arguments because these drawings model the consistent creation, repair, and replacement of a type of machine. Each machine created from a patent diagram is a specific object, a specific content and form of expression. But a single diagram provides the model by which these machines are built. The diagram outlines the operations which bring this form into being. Similarly, social science diagrams are operational models for the construction of a narrative argument. Like patent drawings, they show each component of the argument form, and how it fits together with other components. By modeling a stable form, they allow it to hold content. But, unlike patent drawings, the content is not a physical machine; the content is a series of thoughts: "Diagrams are simple drawings or figures that we think with or through" (Knoespel, 2001, p. 146). Reading a diagram, the viewer asks: What does this line mean in terms of my argument? What part of my argument does this shape represent? By answering these questions, theorists think through a diagram and build expressions in the form of diagram models. A diagrammatic form outlines a model, but the specific expression arises through the act of building. Thus, diagrams both formalize thought and provide a means of discovery. Indeed, the thinking through of a diagram is precisely what formalizes the discovery.

Every method of discovery is an abstract machine. But I also locate the diagram as a concrete type of visual object. In doing so, I borrow a distinction between *sentential* and *diagrammatic representations* (Larkin & Simon, 1987). Sentential representations model expression as a single sequence of characters, a spoken string, or block of written text. Diagrammatic representation, in contrast, indexes information by spatial location. Examining these representations as tools for problem solving highlights the differences of their forms:

> In a diagrammatic representation, the expressions correspond, on a oneto- one basis, to the components of a diagram describing the problem. Each expression contains the information that is stored at one particular locus in the diagram, including information about relations with adjacent loci.

[9] Deleuze and Guattari (1987) dedicate a section of their essay "On Several Regimes of Signs" in *A Thousand Plateaus* to a discussion of diagrammatic thought and the abstract machine. Deleuze also discusses the diagram in his books *Foucault* (1988) and *The Fold* (1989). See Knoespel (2001) for an accessible general introduction to the Deleuzean theory of diagrams and Massumi (1992) for a meditation upon the implications of diagrammatic thinking in the ontology of Deleuze and Guattari (1987).

> The fundamental difference between our diagrammatic and sentential representations is that the diagrammatic representation preserves explicitly the information about *the topological and geometric relations* [italics added] among the components of the problem (Larkin & Simon 1987, p. 66).

A problem solving approach assumes that representations are task oriented, and that representations are created in order to examine a specific problem. But oriented tasks may not drive representation, and diagrams may offer general theoretical models rather than specific solutions. This is especially true of social science diagrams, where the image presents abstract material as a spatial ordering. Rather than preserve an existing spatial topology, such diagrams *apply* spatial and geometric relations to components of a more abstract issue. Diagrams represent the reasoning and thought processes of their authors upon the plane of the page, but this mapping need not reflect a concrete distribution of objects beyond the page. One method for solving problem is to offer a better representation of the problem (Larkin & Simon, 1987; Norman, 1993; Hutchins, 1996),[10] and diagrammatic thought may take problems with no direct spatial relations and represent them as a spatial argument. Interacting with the diagrammatic representation will provoke new insight and suggest alternative solutions. This method of discovery is the process of thinking through the diagram as an abstract machine.

In the following pages, I seek to develop strategies for examining diagrams as theoretical tools. In doing so, I first address the definition of the diagram in more detail, asking how diagrams are different. What does a diagram seek to display? How does this relate to other visual signs? Where do diagrams fit into a typology of graphics, and how might this highlight their possible uses and differences? Secondly, I conduct an initial archeology of graphic representation, and how it came to be understood as a tool for modeling abstract thoughts. Beginning with Vannevar Bush's (1945) idea of associative trails, I ask how trails of argumentation are constructed. Following the trail further, I explore more recent attempts to model associations (Semantica software), and the use of "trails" as narrative markers in the sequential art of comics (McCloud, 1993). In the third section, the trope of the trail leads back to academic work practice, with a discussion of trajectory (Strauss, 1993) and the means of navigating information ecologies (Hutchins, 1996; Bowker & Star, 1999). Finally, I return to visualization strategies and uncover narrative diagrammatic models (Greimas, 1987) as a distinct type of representation. These models, I suggest, offer theoretical narratives which scholars employ as methods of analysis. Throughout the essay, diagrams are offered as both examples and theoretical models. For, among their other benefits, diagrams construct a visual language and represent what is difficult to express in prose.

Seeing How Diagrams Are Different

---

[10] Hutchins especially wishes to steer cognitive science away from the model of cognition as problem solver. In its place, he offers a model of distributed cognition, in which cognitive activity arises from the interaction of individuals with their environment. Environmental interaction relies heavily upon cultural models and the available representations for describing that environment. A focus upon representations and the ability to translate between them resonates with discussions in the sociology of knowledge as well. Meaning and cognition arise through the translation of representational forms into other forms and models, both external and internal. For a useful introduction and overview of these issues within a sociology of knowledge framework, see Jules-Rosette (2004).

In order to demonstrate how diagrams utilize spatial organization as an abstract machine, I offer Stuart Hall's famous image of "Encoding/decoding" (1990, see Figure 1). The diagram summarizes the first half of Hall's (1990) article in a simple image, which can then be referred to as Hall later suggests three possible positions of the encoding/ decoding relation.[11] Hall's image introduces a structured trajectory with five distinct moments of communication.[12] The moments of Encoding and Decoding are "determinate" moments in comparison to the privileged position of the discursive form of the message (labeled in the figure by "Programme as 'meaningful' discourse") (Hall, 1990, p. 129). The graphic diagram reflects this priority by situating the determinate moments beneath the privileged position. More importantly, however, isolating the encoding and decoding moments highlights that their respective meaning structures (labeled "meaning structures 1" and "meaning structures 2") do not constitute a direct identity. Rather, the degree of symmetry between these distinct moments relates the degree of understanding between sender, who occupies the knowledge frameworks in the initial position, and receiver, who constructs the knowledge frameworks of the final position. The model graphically challenges the study of communication with a new research agenda: compare degrees of symmetry and asymmetry between the encoding and decoding positions. In doing so, suggests Hall, scholars may better approximate exactly what is being communicated by a specific meaningful program.

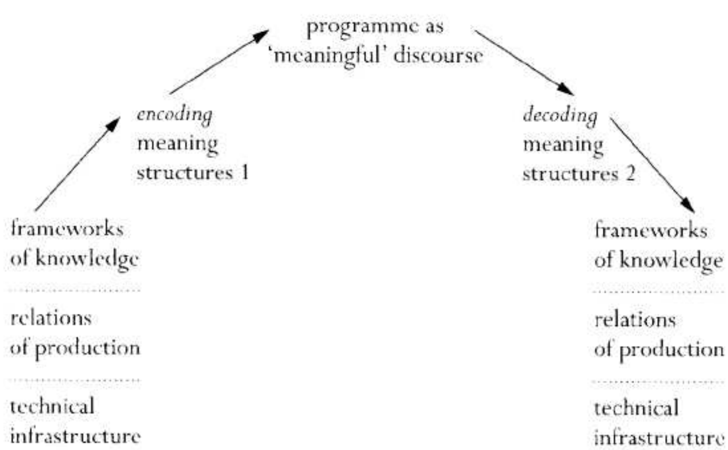

Figure 1. Stuart Hall's (1990) Encoding/Decoding relation

11 Although I focus upon the image of the diagram in this essay, I do not seek to privilege this representational form over written prose or spoken language. Rather, I support a descriptive model in which multiple modes of communication are used to approach a single topic. The diagram, in this regard, accompanies rather than replaces the text. The specifics of this imagetext relationship deserve further study. Although unexamined in the current essay, Roland Barthes' (1977) reflections on the relation between a photograph and its caption provide fruitful ground for beginning such an analysis.

12 Hall offers his five-step model as a direct challenge to the "mathematical model" of a simple sender-transmission-receiver loop, as proposed by the cybernetics of Claude Shannon (1948) and Norbert Weiner (1948).

By parceling the communicative event into a series of five moments, the diagram outlines Hall's suggestion of relevant topology for the problem at hand. The construct translates his theory of communication into a spatial graphic mode, and Hall has chosen a graphic representation in order to imbue his argument with a spatial typology. This typology can then be preserved as suggestions of the model are translated back into the sentential representations of a written text. In the second half of his brief article, Hall does just this. Thinking through the diagram, he uncovers a series of theoretical positions relating the moments of encoding and decoding. Differences of position arise as differences of symmetry across the two moments. Thinking through the symmetries, Hall follows the lines of connection, asks what each moment implies, and finally compares two of these moments (the moment of encoding, and the moment of decoding) to uncover three possible positions. Thus, three different contents are created through a similar set of operations. As Deleuze and Guattari (1987) suggest, the model provided a diagrammatic abstract machine for producing content. The abstract machine of Hall's diagram, which outlines these operations, does not circumscribe a single position of understanding. Rather, it provides a machine for outlining how alternate understandings may arise from a single message.

Hall's image offers both an example of diagrammatic representation and a model of communication. The priority accorded the message form arises from the encoding of an event as a story: the event must become a story before it can become a *communicative* event. Likewise, we can ask what kind of a story the diagram must offer before it becomes a useful analytic tool. In claiming that diagrams become stories, I position them as graphic images of narrative representation. Narrative representations portray unfolding actions and processes of change through the presence of a vector (Kress & Van Leeuwen, 1996).[13] Vectors lead the viewer to perceive the image as a process rather than a timeless description. In diagrammatic thought, the following of this vector is the thinking through of forming an expression. The vectors of Hall's model are easy to notice: the arrows indicating movement from the left to right. These conspicuous vectors offer a deceivingly straight-forward argument: the terms function like nouns and the arrows connecting them function like verbs. Relations among the text fragments may then form clauses, such as "Frameworks of knowledge are encoded by meaning structures 1 in order to become programmes of meaningful discourse." By reading Hall's article, however, we find that translation is not so simple. Hall presents at least five pages of written discussion to explain what his graphic image entails, and none of this discussion is tied to a specific instance of communication. Thus, a single arrow may indicate the need for a verb of relation, but an abundance of verbs and multiple explanations can replace it; "The meaning potential of diagrammatic vectors is broad, abstract, and difficult to put into words" (Hall, 1990, p. 59). The strength of the diagram rests with the numerous ways its vector connections can be explained. Diagram vectors represent more than a single sentential representation, and following the narrative vector of a diagram offers an explanation of what these connections represent in specific circumstances. As a

[13] Kress and Van Leeuwen (1996) differentiate between two types of images: narrative and conceptual. Narrative images present stories in the form of vectors, whereas conceptual images present static qualities or classification schemes. The category of narrative image covers a range of representations apart from diagrams with defined arrows. Lines of sight, depicted roads, or suggested movement of actors are just as likely to provide an image with narrative as the clearly marked arrows of diagrammatic representations.

matter-function (Deleuze & Guattari, 1987), the diagrammatic image condenses the vector (function) into a specific form of expression (matter). These expressions may vary, even as the operations of their expression arise from a single diagram. The variability of relations between encoding and decoding, not its structural determinism, allows Hall to draw three distinct codes from a single model.

But does a diagram's lack of linguistic specificity also uncover a weakness of its representational form? Does the ambiguity of translating vectors into words allow the diagram to promote a set of relations without adequate description? Can diagrams provide a crutch for weak arguments? And do the "simplest cases" of diagrammatic figures betray the complexity of the written text they claim to represent (Lynch, 1991)? "Simplest case" diagrams do not perform an independent representational function. Like Hall's image, they simply restate the written language of an article in graphic form. Converting written text into a graphic display, the diagram makes an argument "look" consistent without furthering the discussion (Lynch, 1991).[14] And by tricking the viewer with the appearance of logical visual consistency, simple diagrams provide the article with a greater weight of authority. Lynch labels this extra weight "rhetorical mathematics" because it cloaks the argument within an image of logical formality:

> Although theory pictures are neither naturalistic nor mathematical representations, they evoke an impression on mathematicity. . . . In an important way, these usages are metaphorical, not mathematical, because often it is difficult to imagine how numerical coefficients ever could be assigned to the structural axes and and causal pathways (Lynch, 1991, p. 12-13).

But why should we wish to replace an image with numerical coefficients? We may wish instead to replace the simplistic image with a new body of text, a text equally consistent with the diagram yet distinct from the original text. Lynch criticizes the openness of the simple diagram, but in order to do so, he returns to quantification.[15] But the benefits of diagrammatic images rest in the ambiguity of their vectors, not their quantification. Diagrams borrow from both written language and mathematics, while breaking the rules of both. The diagram resides halfway between mathematics and something yet to be explained (Knoespel, 149). Mathematically, it isolates variables, but it fails to precisely define or explain these variables. Thus, the openness of the diagram is both its challenge and its gift. The openness challenges the viewer to think through the image, to produce thoughts via the abstract machine. But the

---

[14] Another line of argument suggests that diagrams function as mnemonic devices rather than sources of additional insight. John Law (1986) labels this movement of *interessement*, a method for interesting or enrolling readers in the text. The image also offers a handle for remembrance which readers may hold as they leave, and the diagram may serve to mobilize resources in support of an argument (Latour, 1986, Lynch, 1995), even as it offers little additional information.

[15] Joseph Gougen's work with Algebraic Semiotics offers an interesting contrast with rhetorical mathematics. Algebraic semiotics presents semiotic transformations and "morphisms" within a formal system of algebra, attempting to flatten the diagram, along with all semiotic systems, to the realm of formal logic. Gougen outlines a series of equivalences and algebraic axioms, but the application of these rules to concrete "semiotic morphisms" remains problematic. His work, therefore, appears to swallow the decoy of rhetorical mathematics. But I suggest Gougen's work is itself diagrammatic, offering a series of problems to think through "mathematically," even as these problems cannot be formally defined in the language of mathematics. Once again, the insight arises from the process of translation, not formal transformation.

openness also offers numerous results from the process of thinking through. But this openness is also a weakness, because the diagram itself cannot validate its arguments. That openness must be filled, and each of these results examined separately. To assume that "simplistic" images uselessly restate the written contents of a text assumes that the written contents are themselves easily graspable. But if the text is complex, the image may provide a scaffold for understanding that complexity. As an alternate route for understanding a text, diagrams translate the text in spatial terms. But translation is always partial, and the diagrammatic representation can never replace the specifics of sentential representations.

Criticizing the diagram for failings in the realms of mathematical potential or linguistic content judges the diagram in accordance with rules of a foreign system. The spatial typology of diagrammatic representation presents a system of rational imagery, which differs from both figurative representations and linear sequences of musical, verbal, or mathematical notation. In his fascinating 1967 study *Sémiologie graphique* (*The Semiology of Graphics*), Jacques Bertin outlines eight variables of the graphical system: two dimensions of the plane, plus differences in size, value, texture, color, orientation, and shape.[16] Within these limits, Bertin creates a tripartite classification of graphic types: diagrams, networks, and maps. But in relation to the current essay, an important distinction needs to be drawn with his definition of diagram: "a graphic is a diagram when correspondences on the plane can be established among all elements of *another* component" (Bertin 1983, p. 193). Bertin's diagram involves a graph of axes, and points on the plane relate the variable of one axis with another. In such a model, the diagram is mathematically specific: it begins by attributing meaning to the two planar dimensions and then plots the correspondences. Bertin (1983) limits the diagram to three dimensions, because his typology only addresses "classic graphics" involving the fixed image upon a page, but his logic of the diagram is not limited in number of dimensions. Computer visualizations, which allow users to map data correspondences across numerous dimensions or move among a series of multivariable representations also fulfill this definition.

A diagram such as Stuart Hall's (1990), in contrast, does not plot data along a set of axes. Rather, it models narrative vectors as a visual argument. Thus, to examine the differences between Hall's theoretical model and the correspondence-based diagram of Bertin, we must disentangle what types of information each presents. The plotted diagram of Bertin's typology presents a relational argument drawn from a collected data set. It provides a picture of the data, and is guided by the following questions: What type of graphic should be used? And what graphic image best relates the visual variables to indexed components of the information? The encoding/decoding model, on the other hand, operates in the reverse direction: its spatial argument provides a guide for isolating the components of information. This process is the practice of diagrammatic thought, utilizing the diagram as an abstract machine:

---

16 Bertin's system offers a chapter detailing each of these variables, explaining their possibilities and constraints, and is therefore an incredibly useful guide for graphic design. But the theory also assumes problem-solving model of representation, assuming that graphics merely strive to represent a best view of the data. A similar set of assumptions drives much of computer-aided information visualization (Card, Mackinlay, & Shneiderman 1999; Ware 2000; Wilkonson 1999), as well as guidelines for graphic design (Tufte 1990, 1997; Tonfoni 1998; Berryman 1984)

> A diagram has a function analogous to constructing a *plot for a narrative* argument [italics added]. Once a diagram has completed its prephilosophical task of mapping a conceptual space, the diagrammatic nodes must be animated with figures who speak in coherent and consistent dialogue (Knoespel 2001, p. 150,).

A diagrammatic model provides the narrative plot, and the work of theorizing adds the figures who speak in "coherent and consistent" dialogue. These figures are translated into the frame of the diagram, such that the model may "speak for itself." Hall's three positions offer separate views of the encoding/decoding relation, yet all three result from a single diagrammatic image. The image does not graph these positions; it offers a set of operations for discovering the multiple voices.

The two models differ in their level of abstraction.[17] Bertin's diagrams operate at a level of empirical and evidential representation, where a change in the image displays a change in evidence. But like Lynch (1991), I wish to move away from discussions of images presented as evidence and toward the examination of visual aides in theoretical arguments. "Theory pictures" operate at higher levels of abstraction, where a change in the diagram indicates a change in the *type* of evidence to be collected. Altering terms in diagrammatic representation alters the abstract machine, and altering a machine will produce a new type of object. Hall's analysis of communication would produce a new set of relations had he isolated six moments rather than five. Thus, Hall's model of encoding and decoding is itself a code, a code for parceling an event of communication into a sequence of five moments. As a mode of representation distinct from sentential language, diagrams parcel experience in new and different ways, and challenge us to consider connections of thought, which are difficult to model through written language alone. Diagrams offer the possibility of theoretical representations beyond the realm of spoken and written language: representations built upon a spatial typology rather than the rules of linguistic grammar.[18] But why consider the modeling of thought as a spatial typology? What metaphors allow us to imagine thought as a spatial layout beyond the realms of sentential representations cannot?

Visions of Information Architecture

Vannevar Bush's essay "As We May Think" was published twice during 1945 and immediately hailed as a groundbreaking vision of the future comparable to Ralph Waldo Emerson's address "The American Scholar" (1837). The essay first appeared in the July issue of *The Atlantic Monthly*, and a shortened, illustrated version followed in the September edition of *Life*. During World War II, Bush rose to prominence as a highranking military engineer and chief organizer of the Manhattan Project, and his bold essay formulates a new direction for science and engineering as research shifts

---

[17] The concept of levels of abstraction is borrowed from Gregory Bateson (1972) and the stimulating insights of *Steps to an Ecology of Mind*. I use the concept generally and do not relate presented models to specific levels of Bateson's discussion. As a preliminary suggestion, Bertin's (1983) diagram may be said to offer models of proto-learning, whereas the theoretical models (such as Hall's (1990) image and those throughouth this paper) operate at the level of deuetro-learning, or learning to learn (Bateson, 1972).

[18] The diagram's ability to break with linguistic grammar through the creation of a new spatial grammar may offer a clue to the diagram's appeal in structuralist circles. In structuralism's attempt to condense all thought to language, diagrams offer the ability to comment on language from beyond its borders.

away from the victorious war effort. But the July and September publication dates, which announced a new research agenda in times of peace, ironically bracketed the extreme violence of the war's end, a violence Bush himself was instrumental in achieving: the atomic bombings of Hiroshima and Nagasaki during August of the same year.

"As We May Think" predicts a series of inventions, which Bush believes will revolutionize the practices of knowledge and memory, including miniature personal cameras, the growth of microfilm storage, a "vocoder" speech to type translator, and a powerful calculator dubbed the "thinking machine" by the editors of *Life.* The article's centerpiece, and source of its lasting influence, however, is a device labeled "Memex." The memex is a mechanical aid to extend memory through personalized filing and rapid information selection.[19] Although the record of collected information continues to grow, consultation of this record and its subsequent translation into useful knowledge remains mired in outdated methods: "Selection [of texts] is a stone adze in the hands of a cabinetmaker" (Bush, 1945, p. 99). Sixty years later, with an increasing cascade of information, Bush's concerns and suggestions remain starkly contemporary:

> This is the essential feature of the memex. The process of tying two items together is the important thing. . . .
>
> Thereafter, at any time, when one of these items is in view, the other can be instantly recalled merely by tapping a button below the corresponding code space. Moreover, when numerous items have been thus joined together to form a *trail*, they can be reviewed in turn, rapidly or slowly, by deflecting a lever like that used for turning the pages of a book. It is exactly as though the physical items had been gathered together from widely separated sources and bound together to form a new book. It is more than this, for any item can be joined into numerous trails (Bush, 1945, p. 103-104).

The solution to poor indexing and information selection rested upon the creation of associative trails, which, when stored in a Memex, could efficiently and easily retrieve information at a later date. Today, Bush's (1945) charge for a "selection by association, rather than by indexing" (p. 102) remains unanswered. Web-based hypertext links numerous documents, but user-defined trails of association cannot blaze across unconnected texts. Rather, they can only follow those links already embedded in the text.[20]

What interests me about Bush's (1945) suggestion of associative trails, however, is not the intricacies of hypertextual navigation, but their proclaimed "analogy" with a theory of cognition. The memex, and Bush's reflections on its possibilities, were

---

[19] Bush's vision of the memex changed over time in response to new technology and theories of cognition. The discussion in this paper, however, focuses upon the original idea presented in 1945. *From Memex to Hypertexr: Vannevar Bush and the Mind's Machine* (1991), edited by James M. Nyce and Paul Kahn collects Bush's writings regarding the memex, along with commentary, reflections, and supporting documents.

[20] Randall Trigg (1991) offers a thoughtful, although now somewhat dated, comparison between Bush's trailblazing and hypertext construction. Ted Nelson, who coined the word hypertext in his influential *Computer Lib/Dream Machines* also argues that hypertext has failed to live up to its potential for modeling narrative.

firmly rooted in an environment of utopian thinking imagined through the lens of analog technology (Nyce & Kahn 1991; Owens, 1991). The potential memex manifests itself as a direct modeling of the brain, and the promise this holds for personalizing the storage and retrieval of information:

> When items are thus tied together in a chain, when an item in the chain can be caused to be followed by the next, instantly and automatically, wherever it may be, there is formed an associative trail through the material. It is closely analagous [sic] to the trails formed in the brain, and it may be similarly employed (Nyce and Kahn, 1991, p. 58).

The trails of the memex diagram a process of thought. And by following these trails, one can recreate earlier associative trails. Physically, the memex provides a personal memory prosthesis containing all the books, facts, letters, records, and communications with which an individual came in contact (Bush 1945): a library demarcated with personal trails of association. But the function of memex is an abstract machine: a set of operations for rebuilding the thoughts diagrammed by its trails. The memex allows individuals to perfectly recreate expressions of thought, just as patent diagrams provide a means for recreating physical machines. Moreover, by recording personal trails, the memex makes those trails visible, and, once visible, trails may be followed by others. Bush (1945) foresaw a profession of trailblazers who delighted in finding new and useful trails through the enormous mass of the common record. The vision of memex is much more than a call for diagrams; it is a call to institutionalize, share, and mechanize diagrammatic thought. In the form of the memex, diagrammatic trails of interpretation guide information selection, overcoming the "stone adze" of standardized indexing.[21]

Although the memex champions the lofty ideals of diagrammatic thought and abstract machines, Bush's interest in trailblazing may have much humbler origins. For his Master's Thesis from Tufts College, Bush invented the Profile Tracer, a machine for measuring the distance traveled by surveyors over uneven ground (Owens, 1991). During these years of study, the intellectual atmosphere of Tufts engineering school was dominated by Gardner Anthony, who advocated the graphic language as "an exercise in writing straight and thinking straight" (Owens, 1991, p. 26). Anthony's short book (with lengthy title) *An Introduction to the Graphic Language: the Vocabulary, Grammatical Construction, Idiomatic Use, and Historical Development with Special Reference to the Reading of Drawings* (1922) proclaims the uniqueness of technical graphic drawing as a system of orthography, vocabulary, and grammar.[22] Understanding graphic language allows the engineer to "express ideas in the most concise manner with absolute accuracy of detail, using the greatest care to avoid ambiguity" (Anthony, 1922, p. 81). Building upon architectural and technical patent drawing, Anthony champions the diagram for its specificity. Here, the diagram's

---

21 The memex records more than links between information. It also records the interpretative act of associating two distinct texts. Associative trails interpret information rather than index information, and the abstract machine operates as interpretation (Massumi 1992, 17).

22 Surprisingly, Anthony's text is rarely mentioned by later writers, such as Bertin (1983) and Wilkonson (1999), who pursue the similar aim of outlining a grammar of graphic language.

ability to escape language makes it powerful, contrasting sharply with the unspecific weakness of Lynch's (1991) simplest case. Whereas Lynch's criticism begins from a diagram's poverty of theoretical complexity, Anthony's praise arises from their usefulness in the building of structures and objects. But what prevents bridging this strength of manufacturing into the realm of the abstract? The abstract machine suggests just this: that theories and interpretations, like physical machines, can be reproduced diagrammatically. The only difficulty rests with how. When Bush (1945) ends his visionary essay by asking if the connection between the human senses and knowledge absorption may be established more directly, he echoes the Gardner Anthony's Graphic Language. With the suggestion of associative trails, he also offers a partial answer of how that may be accomplished.

Similar suggestions for modeling the trails of association continue in academic discussion, and San Diego based company Semantica Research, Inc. has recently reinvigorated the prospect of visually displaying association. Without directly citing Bush, Semantica's byline "I see what your thinking" renews his call to make thought process directly accessible through the senses and reiterates the virtues of Gardner' Anthony's graphic language. Semantica's product line provides software for the creation, viewing, and sharing of semantic networks. Networks consist of three hierarchical levels: 1. Concepts, 2. the Relations connecting concepts, and 3. Instances, which encompass at least two related concepts (Analyst, 5; Network, 3; see Figure 2).[23] Concepts are entered into a network and associative trails connect them with other concepts. Naming these associations transforms them into relation, and links two concepts as a single Instance. Like memex, Semantica's model relies upon a proclaimed analogy with the practices of memory:

> Our Semantica products quickly and easily capture what experts know, organize it, and visually *represent it the way that humans store information in long-term memory* [italics added]. Unlike traditional databases, which try to fit knowledge into rigid structures of tables and rows . . . , [Semantica] allows the expert to model their internal mental structure and expose it to others within and outside of the organization (http://www.semanticresearch.com).

Semantica visually represents the collective mental structures of an organization. The rhetoric of Semantica's papers and press releases emphasize this visuality, hoping to reintroduce experimental visualizations to the center of intellectual discourse. Like Anthony's graphic language, Semantica champions the detail and specificity of visual representation. But like Lynch's theoretical pictures, the company examines visual artifacts at the levels of "metacognition," and theoretical sophistication: "[Semantica] reflects our thoughts back to us as concretized, visible things instead of momentary, fleeting entities" (Semantica in Education, 2003, p. 11).[24] Although Semantica offers

---

23 Semantica's three levels fulfill a similar function as Charles Pierce's concepts of firstness, secondness, and thirdness. Although Semantica white papers present the product in terms of education and semantic network theory, they do not draw upon the field of semiotics.

24 Can semantic networks function as abstract machines? Geneviéve Teil and Bruno Latour (1995) ask a related question in their essay "the hume machine: can association networks do more than formal rules?," which explores the possibilities of association networks to model computerized data analysis. Like Bush, the authors emphasize the personalization of networks: "the possibility *for the actors themselves* to define their own reference frames as well as the metalanguages used within them" (1). Teil and Latour share many of the conclusions of this essay, including their final suggestion that association networks are "moving closer and closer to techniques of narrative" (*ibid.,* 9).

a personal environment for the construction of associative trails, it does not allow users to solidify a geometry above the level of relation. Instances simply connect two concepts, they cannot build shapes or patterns.

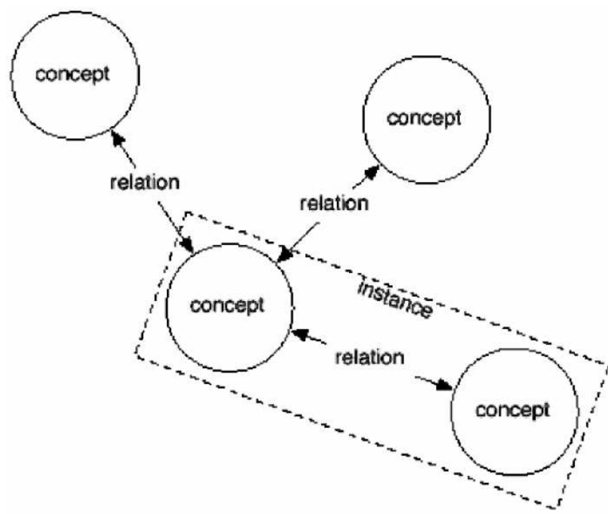

Figure 2. The three levels of a Semantica Knowledge structure: 1) *Concept*, 2) *Relation* and 3)*Instance*.

But the spatial topologies of what I have been calling theoretical diagrams exist as a completed image, not just a collection of relations. The connections of diagrams are positional, not symmetrical (Jameson, 1987, p. xv), and understanding the placement of terms plays an important role in grasping the operations of an abstract machine. The pieces of a patented machine do not connect haphazardly; they adhere to an order and placement demanded by the accompanying diagram. Likewise, diagrammatic theories build narrative arguments through spatial relations of conscious and critical placement. As a spatial argument, the positioning of terms trumps their specific definitions. In response to this provision, I returned to the metaphor of the trail and found, in the writings of Scott McCloud as definition of "trails" as spatial narratives.

Building upon Will Eisner's definition of comics as "sequential art", McCloud (1993) examines the possibilities of comics as a distinctive art form. In doing so, he dedicates significant amounts of discussion to the representation of time (See Chapters 3 and 4). McCloud (1993) highlights that, although individual comic frames are static, action occurs as the reader moves from one image to the next. The action does not occur within the marked frames, but in the unmarked "gutter" separating frames. From the emptiness of the gutter, the reader creates closure and imagines the movement from one comic frame to the next. Or, to utilize the language of Semantica, the gutter allows the reader to connect two comic terms and relate them as a single instance. All readers of a single comic follow the shared narrative of the frames, but the specific details of narrative action rest upon the individual reader (see Figure 3). By demonstrating how comic narratives share a story with guided actions but ambiguous details, the analogy of the gutter offers a clue to solving the puzzle of diagrammatic representations. Terms of a diagram, like comic frames, choose the essential elements of a story. But the axe of a diagrammatic argument strikes in the gutter, when an individual reader elaborates the details as a specific expression.

In a subsequent book, McCloud (2000) redefines the sequential art of comics as "an artist's map of time itself" (p. 206). Comics translate temporal relations into a spatial layout, and reading this layout provides the vector for their narrative. But the direction of graphic narrative vectors need not follow the right to left, top to bottom arrangement of classic comics on a printed page. Rather, the narrative path from comic frame to comic frame can adopt an infinity of forms.[25] But how will readers know which path to follow from frame to frame? McCloud (2001) suggests connecting frames with a simple line and in an online series, he labels these lines "trails". Although McCloud mentions Vannevar Bush as a predecessor to the digital publishing revolution, he does not specifically cite Bush's use of the term "trail." However, McCloud's vision of the comic artist carving narrative trails across an infinite canvas reflects the "professional trailblazer" of Bush. The comic artist creates a map of time in space, and outlines the temporal trail across that space. Trails connect frames depending upon the associations of the artist. But the specific details and interpretation of those associations are left to the reader. The result is a visual artifact, both entertaining and operational. Comic artists share their stories, but only as readers think through the connections. McCloud revels in the possibility of comic art as a dialogue between artist and reader. And describing these possibilities, he unwittingly recalls Bush's work with the Manhattan project: "Comics is a *powerful idea* . . . like an atom waiting to be *split*" (McCloud, 2000, p. 238-241).

Figure 3. McCloud's "graphic" example of blood in the gutter. McCloud provides this example to demonstrate how the reader of comics produces closure by imagining the details which connect two frames.

---

[25] See McCloud's discussion in chapter ? of *Reinventing Comics.* A number of comic examples with non traditional trail structure are also available on his website (http://scottmccloud.com)

Watching Academic Work

The effects of splitting McCloud's atom were indeed powerful, and propelled his work beyond the realm of comics. His name is now cited in relation to graphic design, film studies, and reflexive use of alternative text formats. I too continue the discussion. What might the concept of trails, as outlined by both Bush and McCloud, contribute to an analysis of diagrammatic representation? I suggest academic practice can be likened to the professional trailblazing of Bush. Like McCloud's comic artists of the future, academics carve narrative trails across an infinite canvas. In the realms of information retrieval a "search" requires following a trail. Likewise, ethnographic and ethnomethodological researchers often "trail" participants in order to understand local practices. Both search and research trails, however, operate within a larger project: the collection and analysis of information, which is then written and presented either in conference proceedings or journal articles. Do the narratives of these journals form trails of their own? And how might the narrative representations of diagrams relate to these wider narratives? Can diagrams represent these narratives as maps of intellectual space, just as comics offer spatial maps of narrative time? And if so, should academics be writing diagrammatic comic books?

The trails of academic argumentation rely on more than just association. They require the analysis and clarification of these associations such that another scholar may discern their information and argument structure. The subtlety of descriptive social theory is one of its greatest strengths. Trails of academic research are rarely obvious, and walking them entails a careful following, a careful reading. Careful explanation of theoretical connections combines the associative trails of Bush with the narrative trails of McCloud. Associations become lucid, shared, and discussed once they are situated within a narrative.[26] Academic research parallels the trailblazing of Vannevar Bush (1945) by linking distinct sources of specialized information in order to further a larger claim. But academics go one step further: they analyze and explain the wider significance of their trail. Information is richest when it offers multiple meanings and a wealth of possible interpretations.[27] These meanings gather in webs of information ecology, which, like biological ecologies, are densely interwoven, messy, redundant, and complex (Bowker & Star, 1999). As a result, the utility of associative trails across an information landscape is not self-evident. Interpretation reduces ecological richness, grounding analysis in specific contexts or local practices.

In order to examine these local work processes, Anselm Strauss (1993) employs the term *trajectory*. A trajectory involves both the emergence and persistence of phenomena, as well as the multiple actions contributing to the phenomena. The important dual meaning asks both *what* trail does an object of analysis follow and *how* is that trail shaped by interactions with other objects and actors. Trajectories result from interaction, and this interaction puts the trailblazer back on the trail.

---

[26] The importance of narrative for theoretical exposition and argumentation is also a major theme of Francois Lyotard's *The Postmodern Condition* (1979)

[27] In an insightful section dealing with the semiotic theory of codes, Umberto Eco (1976) bridges both Stuart Hall's (1990) encoding/decoding model and the multiplicity of information: "Information is a value depending on the richness of choices. . . . This richness of the message is only *reduced* by the addressee when he [sic] selects a definitive interpretation" (p. 141).

Trails, both the associative trails of memory and the dusty trails of the countryside, alter as they are traversed. As action continues, trajectories merge, diverge, interlace, and change direction. In academic circles, the interpretation, conceptualization, and projection of phenomena redirects the unfolding trajectory of an analyzed object. Analysts gather strings of information from the information ecology, and knot these strings along lines of association and interpretation. Knots make associations explicit, the work of tying applies the critical judgment of an expert, and academic writing publishes the knots as a finished essay, a series of sentential representations.

As an alternative strategy, diagrammatic representations explicitly highlight aspects of analysis (through a clear display of terms and relations) while leaving other aspects ambiguous (the specificity of these connections). A diagram offers a narrative trail waiting to be completed. Like a comic book, moments of shared narrative are clearly marked, but most of the argument occurs in the "gutter." Approaching a diagrammatic model, the reader must provide closure, relating cross-term connections as a single instance. The researcher may imagine these connections with as little or as much detail as they wish, but like the readers of McCloud's comic, it is they who drop the axe. How one drops the axe betrays a theoretical commitment, and communities of researchers develop around these shared sets of commitments. *Communities of practice* wield similar axes, and swing them in similar arcs. Consequently, they share more detailed narratives and chop similar trails across the information landscape. Research trajectories blaze trails across the information ecology. These trails are then shared, so that others with similar axes may run along them, rather than chop a new path.

But merely suggesting research and publishing as trailblazing and the sharing of trails does not answer the charge of how trailblazing benefits from the use of diagrams. Having moved from trails to trajectory, I was not surprised, therefore, when the next piece of the puzzle accompanied a discussion of navigation. Throughout *Cognition in the Wild* (2001), Edwin Hutchins utilizes diagrams as situating devices in order to share the cognitive strategies of navigators. Thinking through Hutchins' diagrams places the reader in relation to the navigational markers being described. Diagrammatic thought, like navigation, is a method for getting from point A to point B, from one location of understanding to another. In the process, diagrams play a "piloting role" (Deleuze and Guattari, 1987, p. 142): they suggest a new way of seeing from a new perspective and present a new visible landscape of the information ecology. Likewise, Hutchins' diagrams pilot the reader to an understanding of Micronesian navigation. Micronesian navigators do not direct a moving canoe among stationary islands. Rather, they maintain a stationary canoe as the islands move by on either side. For those of us familiar with geographic maps, the Micronesian model is difficult to grasp. We are too firmly positioned in a community of practice which imagines the geographic landscape as a stationary set of markers. But the model of mapping stationary locations is equally difficult from the perspective of the Micronesian navigator (Hutchins, 1979). In a useful explanation of this confusion, Hutchins offers a thought experiment:

> Go at dawn to a high place and point directly at the center of the rising sun. That defines a line in space. Return to the same place at noon and point again to the center of the sun. That defines another line in space. I assert that the sun is located where the two lines cross. Does that seem wrong? Do you feel that the two lines meet where you stand and nowhere else? (Hutchins, 1979, p. 81)

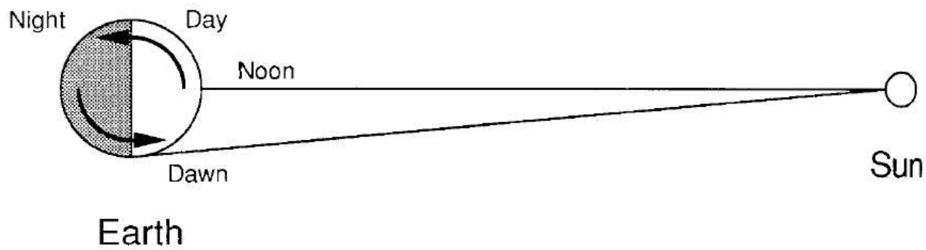

Figure 4. Hutchins' diagrammatic example of how two lines, which appear to cross at an individual standing on earth, can be shown to meet at the sun. Thinking through the diagram situates the viewer beyond the solar system.

Intuitively, the two lines appear to cross at the point where the individual stands, but a diagram displays how they may meet at the sun (See Figure 4). The diagram is drawn from a perspective beyond the solar system. Thinking through the diagram places the viewer beyond the solar system as well. The diagram is both a situating device and a coding scheme, allowing the viewer to see the world from the perspective it establishes (Goodwin, 1994). Geographic maps, like the solar system diagram, place the viewer above the landscape. Micronesian navigation, in contrast, systematically organizes its representations around the position of the canoe. The two representational systems lend themselves to distinct sets of inferences, and calculations (Hutchins, 2000), but both provide useful navigational models. The diagrams of each model provide a means for locating markers within a landscape, and therefore play a piloting role. Navigation, like cognition, occurs as a system of interaction between individuals, the environment, and the markers highlighted within that environment.[28]

The strength of Hutchins's thought experiment arises through the reader's notice of shifting reference frames. In the experiment, the reader begins, like a Micronesian navigator, from the frame of their body as it observes the sun. The diagram, however, draws them out of this frame and positions them beyond the solar system. Thinking through the diagram moves the reader from one position to the next: from the frame of the body to an external viewpoint. But most of Hutchins' navigation diagrams operate in the reverse direction. In order to situate readers on Micronesian canoes, he translates diagrams of geographical positioning into images of horizon lines and positional *etak* islands. Via these diagrams, the reader is removed from his/her position above the ocean and placed within the Micronesian canoe. Shifting diagrams move the horizon around the individual, just as the stars move about a canoe. Diagrams move readers from one island of thought to another, and the abstract machine formalizes a new interpretation by reconstructing the perspective of the canoe.

---

28 Hutchins works these discussions into his argument within the larger category of distributed cognition. Distributed cognition offers an alternative model of cognition from the "official history of cognitive science," (356-259) in which cognition occurs as much outside the head as within it. Rather, cognition occurs as a system, in which individuals interact with the built environment and the tools within it.

In a later paper, Hutchins discusses this practice of thinking-through diagrams as the use of material anchors for conceptual blending. The theory of conceptual blending (Fauconier & Turner, 2002) outlines a cognitive trajectory in which two mental spaces become blended to create a new mental space (See Figure 5). The blending trajectory is represented by a narrative diagram consisting of four spaces: 1) a generic space, which holds the structure that the input spaces share; 2) two input spaces; and 3) the blend. Selective elements of the input spaces project into the blend, where they give rise to new emergent structure. In the diagram, the square in the blend space represents emergent structure. Elaborating this structure—a process known as "running the blend" (Fauconnier & Turner, 2002)—gives rise to new elements, which are indicated by the small white circles of the blend space. By running the blend, individuals discover new properties, relations, and elements. Stable input models facilitate the process, and one method for achieving stability is the creation of physical models or "material anchors" (Hutchins, 2000). Material anchors represent an input space ready at hand, which may then be blended with another mental space. The analog clock, which presents a cyclical model of time divided into two series of twelve-hour segments, provides a good example. Familiarity with reading clocks results from apprenticeship in a community of practice, and, once gained, an individual can blend the structure of the clock face with knowledge of day and night in order to specify the time (Hutchins, 2000). Material anchors provide ready-made mental models, which, when blended with specific circumstances, allow individuals to navigate their surroundings and produce local meaning.

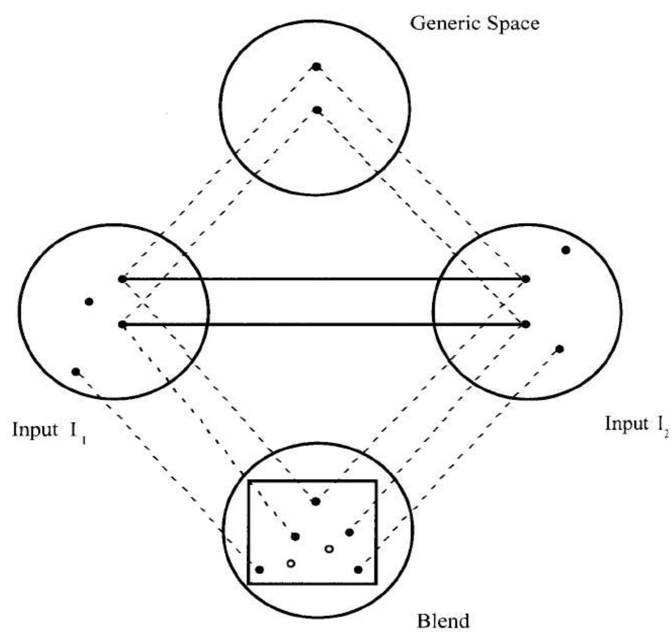

Figure 5. The basic blending diagram. The images consist of 1) a *generic space*, 2) two *input spaces*; and 3) the *blended space*. The square in the blended space represents emergent structure, and the small white circles are new discoveries.

The spatial typologies of diagrammatic representations fulfill a similar role in the analysis of data. Diagrams provide a ready-made structure with which to interpret and reduce the richness of information. Just as Fauconnier and Turner (2002) utilize a diagram to explain conceptual blending, conceptual blending can also be employed to explain the usefulness of diagrammatic representations. The diagram provides an image of mental structure, and sharing this representation attempts to create a common understanding. Blending this spatial typology with collected data provides an opportunity for running the blend, and uncovering new discoveries. Thinking through diagrams produces new discoveries, and formalize a specific expression of the abstract machine. Externalizing operations of discovery as diagrammatic representations provide images of shared mental structure such that others can “see what [the expert] is thinking” (Semantica). Interactively thinking through these representations situates individuals and introduces the shared psychology of a community of practice. As navigational tools for reasoning, diagrams operate as a form of “professional vision”:

> Inscription practices are accomplished through appropriate use of artifacts [such as diagrams]. Supporting such tool use are sets of perceptual structures, the ability to see what and where to measure. Moreover, we are able to glimpse how these structures are passed from one generation to the next through apprenticeship (Goodwin, 1994, p. 615).

The narrative outlines of diagrammatic representation contribute material anchors for navigating new trails and running new blends. Diagrams clearly label the categories, and terms relevant to a specific community (Goodwin, 1994), condensing the richness of the information landscape into bounded frames. The trails of narrative may then be traced from frame to frame, with the resulting stories shared across a community

Diagrammatic representations construct spatial typologies in an attempt to share mental structure and arrive at a collective psychology. By displaying the associative trails of “experts,” they stabilize “professional vision” for communities of practice. If the perspective of professional vision is difficult to grasp, a diagram situates the novice by reorganizing the information landscape. Through the lens of the diagram, islands of thought swim by the viewer, and the abstract machine reproduces a shared perspective. Diagrammatic markers offer signposts for navigating shared narratives, but the closure of filling the gutter with detail provides individuals with a unique trail of personal associations. The critical analysis and descriptive sharing of this personal trail is the work of theorizing. A collective structure, the shared functions of the diagram, help plan the journey, but the detail lies in the traveling. [29]

Looking at Visualizations

Returning to the cascade of visualization artifacts, how well do they answer this call for representations of shared psychology? Do information visualizations provide

---

[29] Hutchins (2001) wishes to steer cognitive science away from the model of cognition as problem solver. In its place, he offers a model of distributed cognition, in which cognitive activity arises from the interaction of individuals with their environment. Environmental interaction relies heavily upon cultural models and the available representations for describing that environment. The distributed cognition resonates with discussions in the sociology of knowledge. In both fields, meaning is understood not as a solution to tasks but as a translation from one setting to another (See Jules-Rosette, 2004).

adequate tools for navigating the trails of an information landscape? Do they produce new realities, new insights, and new interpretations? Diagrams present markers for navigating an ecology of information, and theorists navigate the blend of diagrammatic space and specific data. The theorist, following the comic-style narrative of a diagram, colors the gutter with precision and description. The diagram outlines a trajectory, and introduces the landscape. It provides a narrative structure for traversing the landscape along a vectored trail. But the individual researcher must still walk the trail, and the discoveries of that trail arise from critical, careful, and conscientious marking of space they discover.

This contrasts sharply with the current flood of information visualizations, which more closely resemble the diagrams of Jacques Bertin (1983). Computer-aided visualization strategies map the correspondences of massive data tables, in order to produce new views of data. Researchers then interpret and analyze these views to explain patterns, discrepancies, or interesting points of convergence. Such visualizations provide the matter upon which interpretation functions, rather than a Matter-Function (Deleuze & Guattari, 1987): a diagrammatic model for guiding interpretation. Building upon Bertin's aphorism that "graphics is the visual means of solving logical problems," Card, Mackinlay, and Schneiderman (1999) define *Information Visualization* as "the use of computer-supported, interactive, visual representations of data to amplify cognition" (6). The goal of visualization, they continue, is insight, not pictures. But, unlike both Bush's (1945) trails of associated information and theory pictures of diagrammatic thought, "scientific visualizations tend to be based on physical data" (Card, Mackinlay & Schneiderman, 1999). Visually transforming a data set may highlight patterns of interest, but the image *only* highlights; it does not offer an explanatory narrative. Reading a computer-generated visualization image rests within the narrative of a predefined task. It is an event of the story, not the arch of the story. Not surprisingly, therefore, the authors provide a diagram displaying the narrative process of using visualizations in the service of a task (See Figure 6). First, *Data Transformations* map *Raw Data* into *Data Tables*; next, *Visual Mappings* transform data tables into *Visual Structures*; and, finally, *View Transformations* complete the process by creating new *Views* (Card, Mackinlay & Schneiderman, 1999). Every step of the process benefits from information visualization, but the trajectory of work remains the same. The diagrammatic figure displays the dominant narrative of an abstract machine, in which each stage of visualization is merely a cog.

Thus, the narrative representation, in which visualization contributes to a task, provides a better example of diagrammatic thought than the visualization strategies themselves. The diagram is not in service of an external task; the diagram explains the task. Or, more precisely, *the diagram is the task*. The diagrammatic function shapes matter into a form of expression, and "the diagrammatic or abstract machine does not function to represent even something real, but rather constructs a real that is yet to come, a new type of reality" (Deleuze and Guattari, 1987). Diagrams offer new interpretations of reality, the movement from one island of thought to another. Theory diagrams function like the navigational images of Hutchins: they situate the viewer in the canoe. In Micronesian navigation, the canoe remains stationary as a new location arrives for the navigators. Similarly, the diagram remains stable, as a new information-scape arrives for the theorist who is riding it. Diagrammatic thought does not task itself with moving to an already known and geographically mapped location.

Rather, the image explains how a new location, a new reality, and a new interpretation may gather around the diagrammatic theory. The new locations of diagrams result from coding schemes, which teach the viewer how to see. These ready-made models parcel events via a spatial topology which escapes sentential representation; they produce notational systems for representing new methods of seeing (Norman 1993). Like visualizations, they offer stable artifacts. But they are also instructions for producing new cognitive artifacts; they are abstract machines for the reproduction of an expressive form. As abstract meta-representations, diagrams represent methods of representation (Norman 1993).

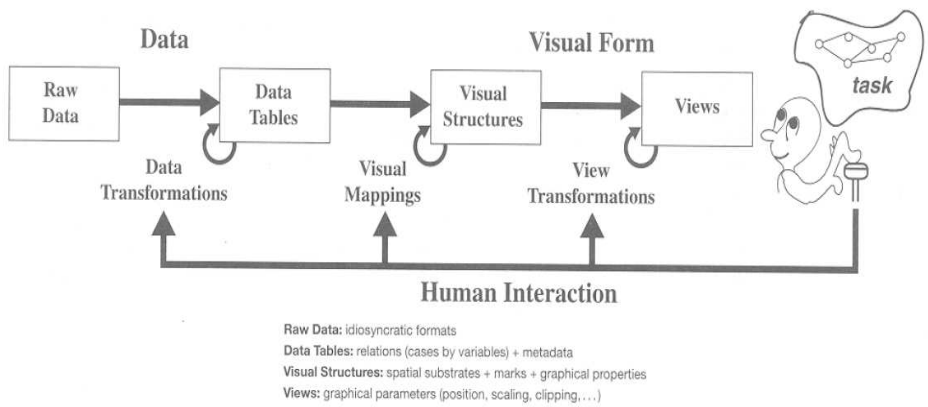

Figure 6. "The Uses of Computer Visualization." Presented by Card, Mackinlay and Schneiderman, (1999), the diagram portrays how *data transformations*, *visual mappings*, and *view transformations* of computer information visualization are situated within the narrative of a task oriented work trajectory.

The diagrammatic symbol of the semiotic square, from which the opening quote was adapted, exemplifies this logic (See Figure 7).[30] Greimassian semiotics specifically aspires to create meta-representations for translating between levels of language: "the investigation of meaning is by definition a metalinguistic activity that paraphrases and translates words and utterances by other words and utterances" (Peron, 1987). The semiotic square is one such attempt. As an analytic framework, it translates the language of narrative into a spatial construct, and numerous theorists, including Bennetta Jules-Rosette (2004), James Clifford (1988), Frederic Jameson (1987), and Katherine Hayles (1999), have utilized the square. Through a shared set of operations, these theorists have expressed a variety of contents. Since all these utterances share a common form, they function as "immutable mobiles" (Latour, 1986), and may be shared, compared, studied, and exchanged.[31] The semiotic square creates a notational

---

30 Figure 7 presents a diagram of the simple semiotic square. The full set of relations and terms which the square makes possible are outlined in Appendix A.

31 See Latour (1986, 2221), Lynch (1985), and Norman (1993) for discussions of comparison across representations, and the mobilization of immutable mobiles for furthering academic claims.

system with which to compare numerous domains through a shared spatial topology. More importantly, however, this shared topology outlines a trajectory for producing yet more descriptive forms. The terms of the semiotic square are connected by narrative vectors. These vectors must be animated to speak, and they have spoken with many voices. Expounding upon these connections—filling their "gutters" and giving them voice—retranslate the model into prose and produce additional insight.

Thus, the semiotic square operates as narrative representation on multiple levels. On a primary level, it isolates the structure of narrative. On a higher level, it returns this structure as a process for constructing other narratives. But how does the system work? What operations does the square suggest? What trail does it blaze? First, the square requests the isolation of two oppositional terms, or semes. These semes become the characters, whose subsequent actions will complete the story. Each character-term generates another character, its simple negative. Beginning with the opposition of student and teacher, for example, we generate the terms of non-student and non-teacher. Notice that these negatives are not synonyms of the original terms. The qualities of a non-student differ greatly from the qualities of a teacher. As the story continues, relations develop between characters, and the square unfolds into its full structure:

> The entire mechanism is capable of generating at least ten conceivable positions out of a rudimentary binary opposition (which may have originally been no more than a single term) . . . The square [offers] a kind of "*discovery principle*" [italics added] . . . One can, in other words, very properly use this visual device to map out and to articulate a set of relationships that is much more confusing, and much less economical, to convey in expository prose (Jameson, 1987, p. xiv-xv)

The pedagogical function of the square translates a complexity of relations into a single image, and produces a narrative argument from a simple opposition of terms. Limited to prose, scholars may overlook these structures due to their difficult expression. But translating the same material via the square offers a new form of expression, and this new form may simplify aspects that were previously difficult. The square outlines a model for thinking through the very connections it represents. And it diagrammatically produces a new landscape around any navigator who holds a steady course between two opposing terms.

Navigating the semiotic square, theorists recreate the form of expression, but the contents of expression vary greatly. Greimas' (1987) own example begins with the binary of permissible and unacceptable sexual relations. He then offers nine dense pages of description, outlining the meaning of each position and each relation. But these outlines in prose are themselves incredibly technical and abstract. Applying them to specific sexual encounters or martial relations would require yet more extensive translation work. The terms of the square present essential elements of the narrative, but the "gutter" separating these terms can be continuously filled with ever-increasing detail.[32] The structure of the semiotic square orders spatial typology, but its

---

32 A counter example can be found in the book *Mapping the Dynamics of Science and Technology* edited by Ari Rip. The book contains a diagram surprisingly similar to the semiotic square, however, the diagram itself is not analyzed as a semiotic square. The square is given as an example of an article as a network of connections. This presentation is given as an example of a network, which can be drawn from an article. But the other than that, the figure is given a mere three sentences of attention in the written text. The figure,

"discovery principle" flourishes through interaction with that structure. And its insight arises from the clearing of details from that structure's gutter.

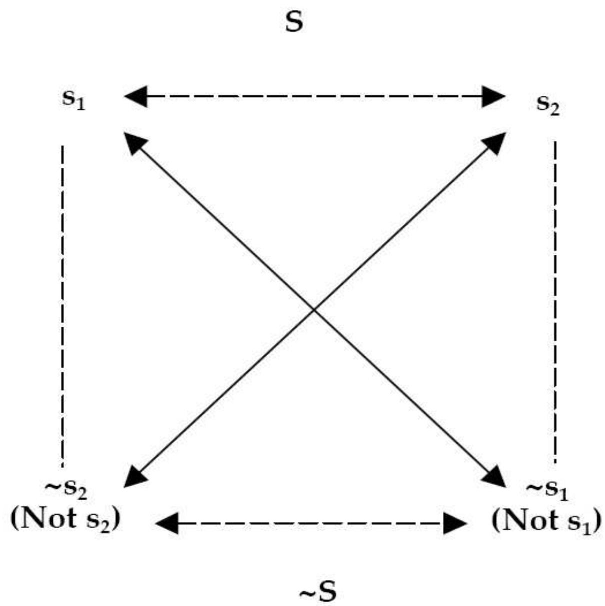

Figure 7: The basic semiotic square of Algirdas Greimas. The square begins with the terms *s1* and *s2* and unfolds into its complete structure and set of relations. For an explanation of each position and relation, see Appendix A. s1 ~s2 (Not s2) s2 ~s1 (Not s1) S ~S

The weakness of the diagram is the collapse of multiplicity into a homogenous structure, but the strength of the diagram rests with the emergence of meaning between its fractures (Massumi 1992). The fractures of the gutter offer obstacles in the work of navigating a diagrammatic trail. Narration helps smooth these fractures, recreating them in line with a shared story. One function of narrative is to mitigate deviations from a pattern such that they once again conform to collective standards (Bruner, 1990).[33] Communities of practice tell stories to maintain coherence, and these stories are often based upon shared forms and collective understandings. The diagram offers a material anchor, a narrative representation, of these shared forms. Just as Hutchins' diagrams explain and teach us how to see with alternative

although resembling the semiotic square, fulfills quite another role: that of an evidential marker. However, for those familiar with the semiotic square, the image suggests a narrative which could be thought through. The connections need not be arbitrary connections, but a map for navigating and narrating the ecology of the information in the article discussed by Rip.

33 The analysis of how narratives mitigate deviations and smooth an information landscape requires further study. Here, I simply offer one of Bruner's narrative characteristics. How this process unfolds as a narrative process is only hinted at, and needs development along the lines of Genette (1972), Lotman (1977), and Eco (1976).

navigation models, a theorist can refer to the shared model of the diagram as a way to legitimize and share their narrative. Thinking through the diagram creates a story of the collected data. Narrating the through the diagram, a community locates shared navigational markers, such that they may then fill the fractures between them. The diagram situates the viewer, and in doing so, is also able to transfer the viewer. Riding on a theoretical canoe, the information landscape passes by the steady image.

Conclusion: Diagrams as Trailheads

The question is not, Is it true? But, Does it work? What new thoughts does it make possible? With these questions, Brian Massumi (1992) begins his discussion of the work of Deleuze and Guattari (1987). And with the same questions, I end my discussion of the diagram. The diagram is an operational graphic, a model for situating a theorist and constructing an argument. It is a tool for learning how to see, how to reason, and how to narrate. Narrative representations are stories waiting to be told, forms of matterfunction ready to hold content. As schematic models, they offer input structures for running a blend and formalizing discovery. But just as crucial as diagrams are to the building process, they also mark the regulation of construction and maintenance after completion (Knoespel, 2001). The patent drawing standardizes a machine; and the diagrammatic image standardizes a formal argument. The image of a diagram is not only a method, but a picture of that method as well. The first aspect fulfills its role as a meta-representation: an abstract machine for producing representations of a certain form. The second aspect is that certain form, a material representation which can be held, compared, shared, and combined with other representations. The picture of a diagram may even fall victim of its own method, such as the semiotic square did when Greimas' students built it into the semiotic frieze.[34] As representations of representations, diagrammatic images contribute to the process of *unlimited semiosis*, the process by which new meaning arises from the transformation of meaning (Eco, 1976).

What new thoughts, then, does the diagram make possible? The question differs slightly from of the goal of Vannevar Bush, but it shares his hope of personalizing our relation to information retrieval. Bush sought the ability to navigate collected records by running along the trails of association. But we no longer need to run from text to text. Rather, with the aid of powerful search engines, we can instantly jump to those texts relevant to our queries. What is needed, therefore, are not trails for locating information, but markers for navigating what is returned. Diagrams present tools for personalizing this process. As representations of narrative processes and discovery principles, we may borrow their spatial reasoning and run new blends. The spatial topologies of diagrammatic representations place us upon a trail. But these are not records of trails that have already been walked, interpreted, and associated. They are, rather, trailheads opening into unexplored territory. Equipped with navigational aids of diagrammatic thought, we venture off, in search of new directions and new stories.

[34] For a brief discussion of the semiotic square's transformation into an expanding semiotic freize, see Jules-Rosette (2004, 17-20).

Appendix A: The Semiotic Square of Algirdas Greimas (1987)

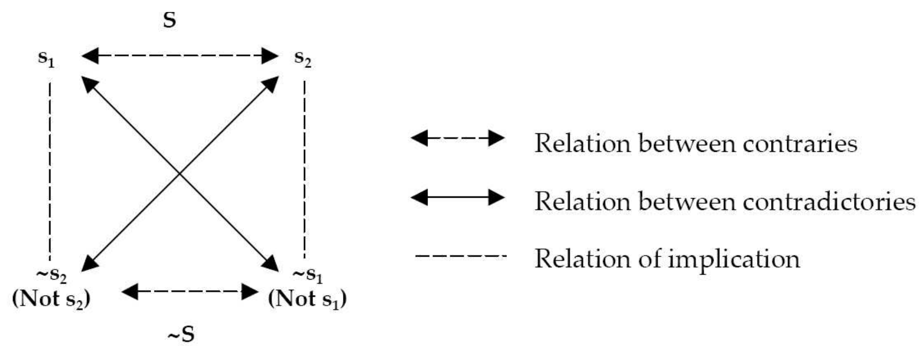

Relations (50-51)

1. **Hierachical:** *hyponymic* relations are established between s1, s2 and S; ~s1, ~s2 and ~S
2. **Categorical:**
    a. A relation of *contradiction* is established between S and ~S; and at the hierarchically inferior level between s1 and ~s1, between s2 and ~s2
    b. A relation of *contrariety* articulates s1 and s2 on the one hand, and ~s1 and ~s2 on the other.
    c. A relation of *implication* is established between s1 and ~s2 on the one hand, and s2 and ~s1 on the other.

Six Systematic Dimensions (51)

1. Two ***axes*,** S and ~S: their relation is one of contradiction. S may be termed the axis of the complex: It subsumes s1 and s2. ~S is the axis of the contradictories ~s1 and not ~s2
2. Two ***schemata***: s1+~s1 define schema 1; s2+~s2 define schema 2. Each of the schemata is constituted by the relation of contradiction
3. Two ***deixes***: The first is defined by s1 and the relation of implication between s1 and ~s2; the second by the implication between s2 and ~s1

The Status of Manifested Contents (61-62)

1. The ***disjunctive*** mode a. Disjoined from the other three terms; it is then isolated in the manifestation. For example, we have s1 vs. (s2, ~s1, ~s2). Thus, there is one manifestation possible for each of the four terms. b. Disjoined from another term; it becomes part of a distinctive opposition. There are six possible manifestations: s1 vs. s2; s1 vs. ~s1; s1 vs. ~s2; s2 vs. ~s1; s2 vs. ~s2; ~s1 vs. ~s2
2. The ***conjunctive*** mode: Six binary oppositions that define what are called complex terms can correspond to the six immanent manifestations of the constitutional structure. **s1 ~s2 (Not s2) s2 ~s1 (Not s1) S ~S** Relation between contraries Relation between contradictories Relation of implication

References

Anthony, Gardner C. (1922). *An Introduction to the Graphic Language: The Vocabulary, Grammatical Construction, Idiomatic Use, and Historical Development With Special Reference to the Reading of Drawings*. Boston: D.C. Heath and Co. Publishers.

Barlow, H., Blakemore, C., & Weston-Smith. M. (1990). *Images and Understanding: Thoughts about Images, Ideas about Understanding*. Cambridge: Cambridge University Press.

Barthes, R. (1977). *Image, Music, Text*. (S. Heath., Trans.) New York: Hill and Wang.

Bateson, G. (2000). *Steps to an Ecology of Mind* (Rev. ed.). Chicago: The University of Chicago Press.

Belew, R. K. (2000). *Finding Out About: A Cognitive Perspective on Search Engine Technology and the WWW*. Cambridge: Cambridge University Press.

Berryman, G. (1984). *Notes on Graphic Design and Visual Communication*. Los Altos, CA: William Kaufmann, Inc.

Bertin, J. (1983). *Semiology of Graphics: Diagrams, Networks, Maps*. (W. J. Berg., Trans.) Madison: University of Wisconsin Press. (Original work published 1967)

Bolter, J. D. (1991). *Writing Space: The Computer, Hypertext, and the History of Writing*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Bowker, G. C., &Leigh Star, S. (1999). *Sorting Things Out: Classification and Its Consequences*. Cambridge: The MIT Press.

Bruner, J. (1990). *Acts of Meaning*. Cambridge: Harvard University Press.

Bush, V. (1945). As We May Think. In J. M. a. P. K. Nyce (Eds.) *From Memex to Hypertext: Vannevar Bush and the Mind's Machine* (pp. 85-113). San Diego: Academic Press, Inc.

Callon, M., Law, J., & Rip, A. (1986). *Mapping the Dynamics of Science of Technology: Sociology of Science in the Real World*. London: The Macmillan Press Ltd.

Card, S., Mackinlay, J., & Schneiderman, B. (1999). *Readings in Information Visualization: Using Vision to Think*. New York: Morgan Kauffmann.

Clifford, J. (1988). *The Predicament of Culture: Twentieth-Century Ethnographer, Literature, and Art.* Cambridge: Harvard University Press.

Deleuze, G. (1988). *Foucault*. (S. Hand Trans.) Minneapolis: University of Minnesota Press. (Original work published 1986)

Deleuze, G. (1993). *The Fold: Leibniz and the Baroque*. (T. Conley Trans.) Minneapolis: University of Minnesota Press. (Original work published 1988).

Deleuze, G. & Guattari, F. (1987). *A Thousand Plateaus: Capitalism and Schizophrenia*. (B. Massumi. Trans.) Minneapolis: University of Minnesota Press. (Original work published 1980)

Eco, U. (1976). *A Theory of Semiotics*. (T. A. Sebeok, Ed.). Bloomington: Indiana University Press.

Eisenstein, S. (1974). *Film Sense.* Harcourt, Brace, Jovanovich.

Fauconnier, G. & Turner, M. (2002). *The Way We Think: Conceptual Blending and Mind's Hidden Complexities*. New York: Basic Books.

Fisher, K. M. & Hoffman, R. (2003) Knowledge and Semantic Network Theory. in *Semantic Research, Inc. White Papers*. San Diego: Semantic Research, Inc.

Genette, G. (1972). *Narrative Discourse: An Essay in Method*. Tr. J. E. Lewin. Ithaca: Cornell University Press.

Goldfarb, B. (2002). *Visual Pedagogy: Media Cultures in and beyond the Classroom*. Durham: Duke University Press.

Goodwin, C. (1994). Professional Vision. *American Anthropologist,* 96(3), 606-633.

Greimas, A. J. (1987). *On Meaning: Selected Writings in Semiotic Theory*. (W. a. J. S.-S. Godzich Ed.) . (P. J. a. F. H. C. Perron Trans.) Minneapolis: University of Minnesota Press.

Hall, S. (1990). Encoding/decoding. In S. During (Ed.) *The Cultural Studies Reader*, (pp. 97-113). New York: Routledge.

Haraway, D. (1991). *Simians, Cyborgs, and Women: The Reinvention of Nature*. London: Free Association.

Hayles, K. (1999). *How We Became Posthuman: Virtual Bodies in Cyberneics, Literature, and Informatics*. Chicago: University of Chicago Press.

Hutchins, E. (1996). *Cognition in the Wild*. Cambridge: MIT Press.

Hutchins, E. (2000). *Material Anchors for Conceptual Blends*. University of California, Distributed Cognition and HCI Laboratory.

Jacobson, R. (1999). *Information Design*. Cambridge: MIT Press.

Jameson, F. (1987). Foreword. In W. a. J. S.-S. Godzich Ed. *On Meaning: Selected Writings in Semiotic Theory*. Minneapolis: University of Minnesota Press.

Jules-Rosette, B. (1990). *Terminal Signs: Computers and Social Change in Africa*. (T. A. Sebeok, Roland Posner, and Alain Rey, Eds.) New York: Mouton de Gruyter.

Jules-Rosette, B. (2004). Sociology of Knowledge in *Sociology G/232, Graduate Seminar*. University of California San Diego.

Kirsh, D. (1990). When is Information Explicitly Represented? In P. P. Hanson , Ed. *Information, Language, and Cognition*, Vancouver: University of British Columbia Press.

Kirsh, D. (1993). *The Intelligent Use of Space.* Department of Cognitive Science, University of California, San Diego.

Knoespel, K. J. (2001). Diagrams as Piloting Devices in the Philosophy of Gilles Deleuze. *Theorie – Litterature – Enseignement: Deleuze-chantier* 19, 145-165.

Kress, G. & van Leeuwen, T. (1996). *Reading Images: The Grammar of Visual Images*. New York: Routledge.

Larkin, J. H. & Simon. H. A. (1987). Why a Diagram is (Sometimes) Worth Ten Thousand Words. *Cognitive Science* 11(1), 65-99.

Latour, Bruno. (1986). Visualization and Cognition: Thinking With Eyes and Hands. *Knowledge and Society: Studies in the Sociology of Culture Past and Present* 6, 1-40.

Latour, B., & Woolgar, S. (1979). *Laboratory Life: The Social Construction of Scientific Facts*. London: Sage Publications.

Law, J. (1986). The Heterogeneity of Texts. In M. Callon, J. Law & A. Rip, Eds. *Mapping the Dynamics of Science and Technology: Sociology of Science in the Real World*, (pp. 67-83). London: The MacMillan Press Ltd.

Lotman, J. M. (1977). The Structure of Narrative Text. In D. P. Lucid, Ed. *Soviet Semiotics: An Anthology*, (pp. 193-197). Baltimore: The Johns Hopkins University Press.

Lucid, D. P. (1977). *Soviet Semiotics: An Anthology*. Baltimore: The Johns Hopkins University Press.

Lynch, M. (1985). *Art and Artifact in Laboratory Science*. H. Garfinkel, Ed. Boston: Routledge & Keegan Paul.

Lynch, M. (1991). Pictures of Nothing? Visual Construals in Social Theory. *Sociological Theory* 9(1), 1-21.

Lyotard, J.-F. (1984). *The Postmodern Condition: A Report on Knowledge*. G. Bennington. & B. Massumi, Trans. Minneapolis: University of Minnesota Press. (Original work published 1979).

Massumi, B. (1992). *A User's Guide to Capitalism and Schizophrenia: Deviations from Deleuze and Guattari*. Cambridge: MIT Press.

McCloud, S. (1993). *Understanding Comics: The Invisible Art*. New York: HarperCollins Publishers, Inc.

McCloud, S. (2000). *Reinventing Comics*. J. Kahn, Ed. New York: Paradox Press.

McCloud, S. (2001). Follow that Trail! in *I Can't Stop Thinking*.. http://scottmccloud.com/comics/ icst/icst-4/icst-4.html.

Nelson, T. H. (1987). *Literary machines:Tthe report on, and of, Project Xanadu concerning word processing, electronic publishing, hypertext, thinkertoys, tomorrow's intellectual revolution, and certain other topics including knowledge, education and freedom*. Swarthmore, PA: Theodore Holm Nelson.

Norman, D. A. (1993). *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*. Reading, MA: Addison-Wesley Publishing Company.

Nyce, J. M. & Kahn, P. (1991). *From Memex to Hypertext: Vannevar Bush and the Mind's Machine*. San Diego: Academic Press, Inc.

Owens, L. (1991). Vannevar Bush and the Differential Analyzer: The Text and Context of an Early Computer. In ,J. M. a. P. K. Nyce, Eds. *From Memex to Hypertext: Vannevar Bush and the Mind's Machine* (pp. 3-39) San Diego: Academic Press, Inc.

Perron, P. J. (1987). Introduction. *On Meaning: Selected Writings in Semiotic Theory*. Minneapolis: University of Minnesota Press.

Semantic Research, Inc. Position Paper: Semantica as an Analyst's Tool. in *Semantic Research, Inc. White Papers*. San Diego: Semantic Research, Inc.

Semantic Network Picture Book. in *Semantic Research, Inc. White Papers*. San Diego: Semantic Research, Inc.

Semantica In Education. *Semantic Research, Inc. White Papers*. San Diego: Semantic Research, Inc.

Shannon, C. (1949). *The Mathematical Theory of Communication*. Urbana: University of Illinois Press.

Strauss, A. (1993). *Continual Permutations of Action*. New York: Aldine De Gruyter.

Sturken, M. & Cartwright, L. (2001). *Practices of Looking: An Introduction to Visual Culture*. Oxford: Oxford University Press.

Teil, G. & Latour, B. (1995). The Hume Machine: Can association networks do more than formal rules? *Stanford Humanities Review,* 4(2).

Tonfoni, G. (1998). *Information Design: The Knowledge Architect's Toolkit*. Lanham, MD: The Scarecrow Press, Inc.

Trigg, R. H. (1991). From Trailblazing to Guided Tours: The Legacy of Vannevar Bush's vision of Hypertext Use. In J. M. a. P. K. Nyce, Eds., *From Memex to Hypertext: Vannevar Bush and the Mind's Machine*, (pp. 353-367). San Diego: Academic Press, Inc.

Tufte, E. R. (1990). *Envisioning Information*. Cheshire, CT: Graphics Press.

Tufte, E. R. (1997). *Visual Explanations: Images and Quantities, Evidence and Narrative*. Cheshire, CT: Graphics Press.

van Leeuwen, T. & Jewitt, C. (2001). *The Handbook of Visual Analysis*. London: SAGE.

Ware, C. (2000). *Information Visualization: Perception for Design*. New York: Morgan Kauffman.

Watzlawick, P., Helmick Beavin, J., & Jackson, D. D. (1967). *Pragmatics of Human Communication: A Study of Interactional Patterns, Pathologies, and Paradoxes*. New York: W.W. Norton & Company, Inc.

Wiener, N. (1948). *Cybernetics, or Control and Communication in the Animal Mind and the Machine*. Cambridge: MIT Press.

Wilkonson, L. (1999). *The Grammar of Graphics*. New York: Springer-Verlag, Inc.