

Mark Derthick

Carnegie Mellon University
Human Computer Interaction Institute
Pittsburgh, PA 15213-3891
(412) 268-8812

5819 Kentucky Avenue
Pittsburgh, PA 15232
(412) 363-9859
mad@cs.cmu.edu
<http://www.cs.cmu.edu/~mad>

Research Interests

- Human-Information Interaction
- Interactive Information Visualization
- Knowledge Representation
- Exploratory Data Analysis
- Collection, Organization, and Summarization of Multi-Media Data Streams

Experience

1995 – Present

Carnegie Mellon University Visiting Scientist, Project Scientist, Research Scientist
Research on interactive visual environments for exploratory data analysis including:

- visual query languages
- integrated interface environments for exploration of multi-media data, including relational databases, free text, and images
- handling large data sets
- visualization of the history of users' interactions with these systems

1988 – 1994

MCC

Member of Technical Staff

Helped build the Cyc large-scale common-sense knowledge base by:

- implementing a smart spreadsheet application for intelligence data, featuring database integration and explanations generated by general-purpose translator from formal Cyc language to English. The spreadsheets performed equality reasoning, semantic contradiction detection, and automatic schema recognition.
- writing many smart interfaces and tools to support Cyc theory formulation, retrieval, browsing, testing, debugging, and statistical learning.
- formalizing the domains of commercial organizations, documents, and many others.

Summer 1984

IBM T. J. Watson Research Center

Intern

Developed a knowledge representation system for common sense reasoning, and utilized for mechanical design by prototype modification.

Summer 1982

Hewlett Packard Signal Analysis Division

Intern

Developed a test system for microwave oscillators.

Education

1983 – 1988

Carnegie Mellon University

PhD Computer Science

Thesis research, directed by Geoff Hinton and David Touretzky, developed a connectionist computational model similar to the Boltzmann Machine. Connectionist models use a densely connected network of very simple processors, each storing only one scalar value. Connectionist models are sometimes likened to the brain's network of neurons. Networks were built by compiling a high-level theory of a problem domain expressed in a formal language using two-valued logic. We used the language of KLONE, a pioneering knowledge representation and reasoning system. KLONE reasons sequentially using rules of deduction. When given an incomplete theory it returns "unknown" to some questions. Given an inconsistent theory, it may refuse to answer any questions at all. In contrast, the connectionist network degraded gracefully, returning a reasonable answer that minimized conflict.

Minor in Montague Semantics supervised by Dana Scott and Richmond Thomason.

1979 – 1983

Washington University in St. Louis

BS EE, BS CS *cum laude*

Areas of concentration included communication theory and digital system design. Photo editor of the *Washington Times* and an officer in the social fraternity Theta Xi.

Research Statement

The Research Statement has three parts:

- A description of analog models and why they're interesting
- A chronological survey of my past and present research, and how it all relates to analog models
- Strategic plan for expanding the scope and influence of my research

Driving Research Questions

Analog models fascinate me. An analog model is much like a piece of the real world. In the world, the laws of everyday physics ensure that no two objects occupy the same space, and it is never ambiguous whether an object is in New York or California. Analog computer models also represent complete and consistent information that is constantly updated as the model responds to external influences. At any point the answer to any question about what is true in the model can simply be read off. Spreadsheets are analog models, as long as the inputs are well behaved so that all formulas can be evaluated. Images are generally complete and consistent, as well. A picture of a person generally is usually unambiguous with respect to gender and hair color, for instance. By contrast, AI knowledge representation systems are language-based and can represent the abstraction person or disjunctions like the New York/California case.

As an example, an analog model can find the quickest rail route between two cities. The model is a network of beads for cities connected by strings whose length is proportional to the travel time. The fastest route can be found by grabbing the beads representing each city and pulling. Imagine hanging the network on a wall map so that each bead is on top of its city. Impediments like the the Rocky Mountains show up as regions of droopy string. Higher-level patterns like this are often called emergent representations. Humans, automated learning algorithms, or a combination can look for and characterize these representations. Analog models are the most effective starting point for discovering emergent patterns by humans, because their completeness allows them to be visualized as an image and people excel at image perception and interpretation.

Knowledge representation languages are used to express theories formally. Usually the language is designed so that *computers* can reason about the theories efficiently. However, I think of information visualization as knowledge representation for humans. In both cases, finding the right representation is key. Thus knowledge, reasoning about knowledge, and presentation of knowledge all fit within the same framework. This provides a natural infrastructure for the nascent field of Human-Information Interaction, whose goal is to improve the way users find, interact with, and understand information in Internet and other information resources.

Most of my research relates to understanding the characteristics of analog models, their relationship to more expressive theories, and the applications for which they are most appropriate. In my current work, simple models are embedded in interfaces with a Model View Controller architecture. Further, in accordance with the Direct Manipulation principle that input=output, the view is the controller as well as the visualization. I draw on my training as a Computer Scientist to develop efficient algorithms so that the feedback appears continuous. I draw on my experience at MCC in knowledge representation to ensure that the computer's model, and the view of it, is appropriate for the domain and task.

These ideas almost constitute a formula for qualitatively changing the experience of using many current software applications: build an interface in Visage, and make it run fast even if you have to give up some functionality. Visage, developed in collaboration with the Maya Group and other members of the Sage group, embodies

the MV=C architecture. It is a flexible data visualization and exploration environment. For instance, AI scheduling systems currently run for hours. With Steve Smith, I developed an interface with 100ms feedback for looking at schedule tradeoffs, which are actually more useful than a detailed schedule, at least for the initial stages of planning. With Phil Gibbons and Andrew Moore I have developed algorithms to filter data with a million records with 100ms feedback, by restricting the range of visualizations. For the Informedia project, I developed an interactive image retrieval algorithm that can search 15,000 images in 500ms, while their previous algorithm required 30 seconds. Again, the interface design led to changes in the search algorithm. These interfaces are among those pictured below.

Research Chronology

CMU



My thesis work compiled a high-level symbolic domain theory into a sub-symbolic analog model. The questions it could answer were of the form “if these given conditions hold, what else will follow?” The conditions were expressed by forcing parts of the model into particular states. After some time for reaching equilibrium, the conditions that follow could be read off other parts of the model. The model is always in some state, so it will always give an unambiguous answer to any question. If the high-level theory is incomplete or inconsistent, the most likely state is one that minimizes conflicts and makes free choices that maximize the robustness of the solution to small changes in the theory.

MCC


```
(forall ?P
  (implies
    (isa ?P Person)
    (thereExstExactly
      2 ?LEG
      (and
        (isa ?LEG Leg)
        (anatomicalParts
          ?P ?LEG))))))
```

I built a system that used the Minimum Description Length (MDL) principle to 1) automatically discover emergent representations that are intelligible to people and 2) to find analogical mappings between domains. Knowledge of analogical mappings can help problem solving in unfamiliar domains. For instance if you already understand water flow, someone can explain heat flow to you by analogy. The analogy is helpful to the degree that the explanation is shorter than an explanation from first principles. This efficiency gain can be measured in bits. This turns it into an optimization problem, in which all contributing factors are weighed at once, as with any analog model.

CMU


Currently I build interfaces based on visualizations that embody analog models, with a primary goal of supporting Exploratory Data Analysis. Models that capture the important aspects of the domain must be developed in tandem with efficient algorithms, so that the experience of interaction is direct.

Scheduling Visualization




AI scheduling systems currently run for hours. With Steve Smith, I developed an interface with 100ms feedback for looking at high level schedule tradeoffs, which are actually *more* useful for understanding the constraints than a detailed schedule. The key is to show all possible sub-intervals of the schedule at once. Later intervals are farther to the right, and longer intervals are farther back. Then the set of requirements becomes a landscape of mountain ranges, where elevation shows the magnitude of resources required. The available resource capacity appears as a partially transparent inclined plane. Mountain peaks jutting above the capacity plane represent shortfalls, where requirements exceed capacity. The effect is similar to flying above a cloud layer, where it is very easy to identify mountain peaks that extend above the clouds.

TimeTravel Interface




Exploratory Data Analysis is an opportunistic process, where users try multiple avenues in search of understanding. Visage explicitly models these alternative pathways, and visualizes them as a tree. This makes it easy to switch back and forth, iteratively extending the most promising. New branches are created automatically when operations are performed in a past situation. In addition, visualization of the exploration path serves as documentation for how results were obtained. Navigation in the tree is done by dragging a cursor. This is more natural than current interfaces that use discrete undo/redo operations to restore previous states. User operations are shown in the tree as events. Selective undo and redo are accomplished by selecting subsets of events and dragging them out of the Time Travel interface (selective undo), or by dropping events on any point in the tree (selective redo).

Visual Query Environment




The Visual Query Environment (VQE) uses a navigation metaphor to capture the incremental process of Exploratory Data Analysis and to situate the analyst in the world of the data. It uses the concept of “aggregate” to bridge the gap between the expressive power of database query languages and the concreteness of data. The navigation path among aggregates resembles a graphical query language, but no “execute” step is needed to convert from specification to realization. Using efficient indexing, aggregate visualizations are updated continuously to maintain a sense of engagement with an analog model.

Navigation structures can be created in VQE, or a subset of a user’s normal Visage operations can be selected from the TimeTravel visualization and compiled into a visual query. Any visualizations that were involved in the user operations are copied into the VQE window as well. The navigation structure can be graphically edited, and data shown in the visualizations can be added or removed by drag-and-drop.



Sometimes abstract queries are desirable, because they can be reused on different data. Dragging all data out of VQE leaves a template for reproducing an analysis. A user can drop new data on the template, which will carry out the navigation operations and populate the visualizations. This declarative approach is more powerful than procedural macros, because navigation paths can be traversed either forward or backward. For instance, the original analysis may have begun with students and proceeded to the classes they take. On reuse, the user might drop classes on the query, which would look up the students.

Personal Information



An ongoing project is capturing and building manipulable models of my personal information stream. The diagram at left shows the dataflow from my email, calendar, contacts, tasks, web browsing history, phone calls, medical diary, Quicken, and Emacs and Microsoft Office visits and saves into Visage. Additional data comes from an Informedia Experience on Demand camera, microphone, and GPS receiver, and from a BodyMedia arm band that monitors heart rate, skin galvanic response, and posture. I want to build visualizations that can help me search and browse this collection for either particular sets of items or for emergent patterns. Current applications are tied to one or a few of these information sources, while people’s goals and interests are largely independent of modality or source application. As a starting point, Chris Neuwirth, Jim Morris, and I intend to generalize email programs to organize information by task rather than message.

Images



Current interfaces that search for similar images return an ordered list of matches and a match score for each. Researchers need insight into *why* an image is chosen in order to improve their retrieval algorithms. I developed an algorithm in which credit for a match can be assigned to individual pixels, and that runs in 500ms instead of the 30 seconds that the Informedia search engine previously took. Each pixel's contribution to the match is mapped to opacity, so what you see is what matched. The portion of the image to search, as well as the resulting color histogram, can be interactively updated. What was formerly a batch-oriented black box is now a manipulable analog model.

Strategic Plan I have led a number of proposals with a large number of collaborators from HCII, other CMU departments, and other institutions. With each one I have learned a little more about what it takes to have a successful HCI-oriented collaboration. I believe the pending NIH proposal "Exploratory Analysis and Visualization Software" contains all the needed ingredients:

- We have a large group of potential users who are unhappy with their current tools, and have written a letter of support promising access to their data and to be beta testers. These are public health analysts responsible for the National Survey of Child and Adolescent Well-Being (NSCAW), an ongoing nationally representative longitudinal study of children and families or other caregivers who have had contact with the child welfare system. The first wave included 5400 children. This gives us a clear target for user-centered design. Co-PI Kelly Kelleher, a physician with a Master of Public Health degree, understands the user community and the need for visualization to support Exploratory Data Analysis.
- Locally, we can bring in similar users for more frequent study. Many statisticians associated with Pitt and UPMC analyze similar survey data. Co-PI Howard Seltman, a physician and CMU Statistics Professor who has long collaborated with Dr. Kelleher, runs a summer program for graduate students from Historically Black Colleges and Universities that we can draw on for less experienced analysts.
- Jane Siegel and John Zimmerman complete the team with their expertise in user-centered design and evaluation.

Through this project I believe I will break through my Computer Science biases and become a much more rounded HCI researcher. Sage, Visage, and Informedia have not yet embodied good HCI in this sense.

In my statement of career goals three years ago I said clinical medicine was going to be my primary focus. With the NIH proposal, as well as possible projects with Highmark and the Pittsburgh Regional Healthcare Initiative, I am now poised to make this happen. With other HCII faculty interested in medical applications and socially responsible computing, I would like to make CMU a leader in this area.

I am eager to share ideas with students as well as colleagues. I hope to reach a steady state of supporting two or three PhD students, and to teach a graduate visualization seminar often enough that every PhD student has the opportunity to take it, starting with Spring 2004.

- Grants**
1. NIH, "Improving Bayesian Phylogeny." Joseph Kadane, Mark Derthick. 6/15/2003 – 6/14/2007. \$730,642. Subcontract from the University of Wisconsin, Madison.
 2. DoD/STTR, "An Electronic Workspace for the Commander." Mark Derthick. Phase I (2002-2003, \$29,700) & Phase II (2003-2004, \$140,000). Subcontract from Maya Design.
 3. Advanced Research and Development Activity (ARDA) AQUAINT Program "Question Answering from Errorful Multi-Media Data Streams." 07/03/02-01/02/04. Howard Wactlar, Alex Hauptmann, Mark Derthick, John Lafferty, Steven Roth, Laurie Waisel. \$1,115,833.
 4. NSF Award #0121641 AWSFL008-DS3 ITR/IM: "Capturing, Coordinating and Remembering Human Experience." Howard D. Wactlar, Takeo Kanade, Michael G. Christel, Alexander G. Hauptmann, Mark Derthick. 10/1/01 – 9/30/04. \$2,000,000.
 5. CMU Seed Grant, "Visual Discovery: Exploring the Universe at Interactive Speeds." 1999. Robert Nichol, Mark Derthick, Andrew Connolly, Andrew Moore, Jeff Shneider, Larry Wasserman, Nathan Stone, Joel Welling, Mike Levine, Steven Roth. \$96,000.
 6. MCC Internal Grant, "Using the Minimum Description Length Principle to Learn Features and Analogical Mappings." 1990. Mark Derthick. \$110,000.
- Pending Proposals**
7. NIH, "Exploratory Analysis and Visualization Software." Mark Derthick, Kelly Kelleher, Howard Seltman, Jane Siegel, John Zimmerman. \$1,588,968.
 8. NSF ITR: "Detecting Video Perspectives from Global Sources in support of Intelligence Analysis and Contextual Awareness." Howard Wactlar, Mike Christel, Mark Derthick, Alex Hauptmann, Dorbin Ng. \$3,951,018.
- Books**
1. Mark Derthick. *Mundane Reasoning by Parallel Constraint Satisfaction*. Research Notes in Artificial Intelligence. Pitman, London, 1990. Reprint of PhD thesis, also available as CMU Technical Report CMU-CS-88-182.
- Book Chapters**
2. Mark Derthick. Finding a maximally plausible model of an inconsistent theory. In John A. Barnden and Jordan B. Pollack, editors, *Advances in Connectionist and Neural Computation Theory*, Volume 1: High-Level Connectionist Models, pages 241-258. Ablex, Norwood, NJ, 1991.
- Journal Articles**
3. Mark Derthick and Stephen F. Smith. [An Interactive 3D Visualization for Requirements Analysis](#). To appear in *Journal of Scheduling*.
 4. Mark Derthick. [Mundane Reasoning by settling on a plausible model](#). *Artificial Intelligence*. 46(1-2):107-157, November 1990.
 5. Mark Derthick and Steven F. Roth. [Enhancing Data Exploration with a Branching History of User Operations](#). *Knowledge Based Systems*, 14(1-2):65-74, March 2001..
- Journal Articles (submitted)**
6. Mark Derthick and Steven F. Roth. A Navigation Semantics for Visual Queries. *ACM Transactions on Information Systems*..

Conference Papers

7. Mark Derthick, Michael G. Christel , Alexander G. Hauptmann, and Howard D. Wactlar. [Constant Density Displays Using Diversity Sampling](#). To appear in *Proceedings of the IEEE Information Visualization Conference (InfoVis'03)*.
8. Mark Derthick. Interfaces for Palmtop Image Search. In *Proceedings of the Joint ACM/IEEE Conference on Digital Libraries*, Portland, OR, July, 2002, pp. 340-341. ACM Press, 2002.
9. Mark Derthick and Steven F. Roth. [Example-based generation of custom data analysis appliances](#). *Proceedings of Intelligent User Interfaces (IUI '01)*, Santa Fe, NM, January, 2001, pp. 57-64.
10. Joel Welling and Mark Derthick. [Visualization of Large Multi-Dimensional Datasets](#). in *Proceedings of Virtual Observatories of the Future 2000*. Pasadena, CA. June, 2000. Ed. R. J. Brunner, S. G. Djorgovski, and A. Szalay.
11. Mark Derthick and Steven F. Roth. [Data exploration across temporal contexts](#). In *Proceedings of Intelligent User Interfaces (IUI '00)*, pages 60-67, 2000. (Nominee for Best Paper.)
12. Mark Derthick, James Harrison, Andrew Moore, and Steven F. Roth. [Efficient multi-object dynamic query histograms](#). In *Proceedings of the IEEE Information Visualization Conference (InfoVis'99)*, pages 84-91, 1999.
13. Mark Derthick, John A. Kolojechick, and Steven F. Roth. [An interactive visualization environment for data exploration](#). In *Proceedings of Knowledge Discovery in Databases (KDD'97)*, pages 2-9, 1997.
14. Mark Derthick, John A. Kolojechick, and Steven F. Roth. [An interactive visual query environment for exploring data](#). In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '97)*, pages 189-198, 1997.
15. Mark Derthick. A minimal encoding approach to feature discovery. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pages 565-571, 1991. Also published as [MCC tech report ACT-CYC-212-91](#), June 1991.
16. Mark Derthick and Joe Tebelskis. 'Ensemble' Boltzmann units have collective computational properties like those of Hopfield and Tank neurons. In *Proceedings of the IEEE Conference on Neural Information Processing Systems*, pages 223-232. American Institute of Physics, 1988.
17. David S. Touretzky and Mark A. Derthick. Symbol processing in connectionist networks: Five properties and two architectures. In *Proceedings of IEEE Spring COMPCON87*. IEEE, February 1987.
18. Mark Derthick. A connectionist architecture for representing and reasoning about structured knowledge. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, pages 131-142. Lawrence Erlbaum, 1987.
19. Mark Derthick. Counterfactual reasoning with direct models. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 346-351, Seattle, Washington, July 1987.
20. Mark Derthick and David C. Plaut. Is distributed connectionism compatible with the physical symbol system hypothesis? In *Proceedings of the 1986 Cognitive Science Conference*, pages 639-644. Lawrence Erlbaum, 1986.

**Unrefereed
Conference
Papers**

21. A. Hauptmann, R. Yan, Y. Qi, R. Jin, M. Christel, M. Derthick, M.-Y. Chen, R. Baron, W.-H. Lin, T.D. Ng. Video Classification and Retrieval with the Informedia Digital Video Library System. In *Proceedings of The Eleventh Text Retrieval Conference (TREC 2002)*, Gaithersburg, MD, November 2002.
22. Mark Derthick. Introduction to mundane reasoning. In David Touretzky, Geoffrey Hinton, and Terrence Sejnowski, editors, *Proceedings of the 1988 Connectionist Models Summer School*, pages 291-300. Morgan Kaufmann, 1988.

**Technical
Reports**

23. Mark Derthick. [*The minimum description length principle applied to feature learning and analogical mapping*](#). Technical Report ACT-AI-234-90, MCC, June 1990.
24. Mark Derthick. [*An epistemological level interface for CYC*](#). Technical Report ACT-CYC-084-90, MCC, February 1990.
25. Mark Derthick. *Variations on the Boltzmann Machine learning algorithm*. Technical Report CMU-CS- 84-120, Carnegie-Mellon University, Pittsburgh PA, August 1984.

Book Reviews

26. Mark Derthick. Book review of Connections and Symbols. *Artificial Intelligence*, 43(2):251-265, May 1990.

Videos

27. Mark Derthick and Stephen F. Smith. [*An Interactive 3D Visualization for Requirements Analysis*](#). Accompanies Paper 3.
28. Mark Derthick, Michael G. Christel , Alexander G. Hauptmann, and Howard D. Wactlar. [*Constant Density Displays Using Diversity Sampling*](#). Accompanies Paper 7.
29. Mark Derthick, Kelleher, Howard Seltman, Jane Siegel, John Zimmerman. [*Exploratory Analysis and Visualization Software*](#). Submitted with Proposal 7.
30. Mark Derthick. [*Interactive Visualization of Video Metadata*](#). (Video Demo). In *Proceedings of the Joint ACM/IEEE Conference on Digital Libraries*, Roanoke, VA, 2001, page 453. ACM Press, 2001.
31. Mark Derthick and Steven F. Roth. [*Data exploration across temporal contexts*](#). 2000. Accompanies Paper 11.
32. Mark Derthick, John A. Kolojejchick, and Steven F. Roth. [*An interactive visual query environment for exploring data*](#). 1997. Accompanies Paper 14.

Reviewer NSF, IJCAI, FLAIRS, AAI, IEEE Information Visualization program committee, CHI, UIST, East-West International Conference on Human-Computer Interaction, ACM Transactions on Graphics, IEEE Transactions on Visualization and Computer Graphics, IEEE Internet Computing, Journal of Visual Languages, Graphics Interface, IEEE Transactions on Information Systems

Societies ACM, AAI, IEEE Computer Society, Tau Beta Pi, Sigma Xi, Eta Kappa Nu

Consulting
2000 Cycorp. Visualization of a large AI common-sense knowledge base.

**Invited
Presentations**
2000

Microsoft Research. *Interactive Visualizations for Exploratory Data Analysis*.
Lotus Research. *Interactive Visualizations for Exploratory Data Analysis*.
Andersen Consulting (CSTaR). *Interactive Visualizations for Exploratory Data Analysis*.
USC Information Sciences Institute. *Interactive Visualizations for Exploratory Data Analysis*.

1994 Salk Institute. *Learning hidden unit representations with the Minimum Description Length principle*.

1992 AAI Spring Symposium. *The minimum description length principle applied to feature learning and analogical mapping*.

1991 Siemens Corporate Research Center, Princeton, NJ. *The minimum description length principle applied to feature learning and analogical mapping*.

1989 Stanford University. *Interfaces for Browsing and Editing the Cyc Knowledge Base*.

1988 IBM T. J. Watson Research Center. *Mundane Reasoning by Parallel Constraint Satisfaction*.

AT&T Bell Labs, Murray Hill, NJ. *Mundane Reasoning by Parallel Constraint Satisfaction*.

MCC. *Mundane Reasoning by Parallel Constraint Satisfaction*.

SRI. *Mundane Reasoning by Parallel Constraint Satisfaction*.

ICSI. *Mundane Reasoning by Parallel Constraint Satisfaction*.

High-Level Connectionism Workshop, Las Cruces, NM. *Finding a maximally plausible model of an inconsistent theory*.

1987 MIT. *Mundane Reasoning by Parallel Constraint Satisfaction*.

1985 St. Michaels Workshop on Connectionist Symbol Processing, St. Michaels, MD. *Analysis of Learning in Boltzmann Machines*.

