

Reasoning with Diagrammatic Representations

A Report on the Spring Symposium

B. Chandrasekaran, N. Hari Narayanan, and Yumi Iwasaki

■ We report on the spring 1992 symposium on diagrammatic representations in reasoning and problem solving sponsored by the American Association for Artificial Intelligence. The symposium brought together psychologists, computer scientists, and philosophers to discuss a range of issues covering both externally represented diagrams and mental images and both psychology- and AI-related issues. In this article, we develop a framework for thinking about the issues that were the focus of the symposium as well as report on the discussions that took place. We anticipate that traditional symbolic representations will increasingly be combined with iconic representations in future AI research and technology and that this symposium is simply the first of many that will be devoted to this topic.

The American Association for Artificial Intelligence (AAAI) held the Spring Symposium on Reasoning with Diagrammatic Representations from 25–27 March 1992. The emphasis of this symposium was diagrammatic (or pictorial) representations in problem solving and reasoning. The issues were not about how the raw sensory information in the visual modality is processed to form percepts; this process is the subject of image-processing and perception theories. Rather, the issues related to representations of diagrams and mental images and their functions in problem solving. The symposium attracted a large, diverse, and multidisciplinary audience: Philosophers, cognitive psychologists, design theorists, logicians, and AI researchers participated in 2-1/2 days of presentations and intense discussions.

Philosophers have been interested in the nature of mental imagery for a long time, and

debates about the reality and nature of mental images and visual representations have also raged in psychology. Design theorists have always been interested in the role of sketches and diagrams as design aids. However, logicians have traditionally disdained diagrams as merely heuristic aids to be discarded once the correct path to the real proof is obtained. Finally, although AI researchers flirted with diagrams in the early decades, especially in the work on geometric theorem proving, they have ridden on the wave of the so-called discrete-symbol-processing view or the logic view, and the representation and the use of diagrams have not been the center of attention for the last couple of decades.

Diagrammatic Reasoning: A New Emphasis in AI

Lately, there has been a ground swell of interest in AI on working with different types of representations and seeing how perceptual and motor components of intelligence can be integrated with the explicitly cognitive components. Correspondingly, there has also been an interest in treating the external world itself as a representation and in studying the multiplicity of types of representations that are available in it for an agent working in the world. Figure 1 shows the emerging picture of these concerns in AI. Because visual modality is one area that seems so rich in its participation in problem solving and reasoning, diagrammatic representation has been one subject of interest in this emerging milieu for research.

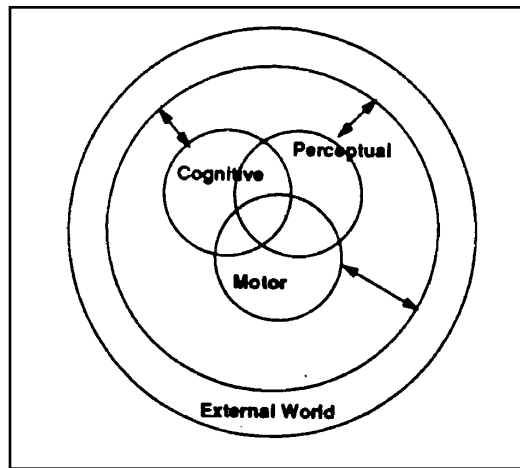


Figure 1. Multiple Types of Representations for a Situated Agent.

The representations include not only the traditional conceptual ones but also iconic and motor representations. The external world can also be used as an additional source of representations.

Organization

This symposium was organized by B. Chandrasekaran, Yumi Iwasaki, N. Hari Narayanan, and Herbert Simon. More than 70 researchers from 9 countries attended. The symposium consisted of five moderated sessions (moderators are given in parentheses): Imagery and Inference (Nancy Nersessian, Princeton University), Human Diagrammatic Reasoning—Analyses and Experimental Studies and Sketching (Irvin Katz, Educational Testing Service), Logic and Visual Reasoning (Pat Hayes, Stanford University), Computational and Cognitive Models of Diagrammatic Representation and Reasoning (Brian Funt, Simon Fraser University), and Qualitative Reasoning (Leo Joskowicz, IBM T. J. Watson Labs). Each session was structured as a series of presentations followed by moderator's comments and open discussion. The symposium concluded with a one-hour open discussion.

In what follows, we give an overview of the field as we see it, interspersed with remarks about the contributions from presentations that were made. (The names of the researchers are given in parentheses, as is their affiliation the first time a name is mentioned.) These presentations constituted only a subset of the many excellent papers that were accepted. Space limitation precludes a discussion of all the papers here. The papers

appear in the working notes of the symposium,¹ which contain sections on diagram understanding and processing spatial expressions in addition to sections corresponding to symposium sessions.

Internal versus External Diagrammatic Representations

In discussing the role of diagrams in reasoning, we find it useful to make the following distinction about where the representations reside:

The *external world* is the world in all its detail and form as sensed in various modalities.

External diagrammatic representations are constructed by the agent in a medium in the external world (paper, and so on) but are meant as representations by the agent.

Internal diagrams or images make up the (controversial) internal representations that are posited to have some pictorial properties.

Some of the questions that arise in this context are, Are there any continuities in the mechanisms of processing and use regarding these representations? What are their common functional roles in problem solving and reasoning?

Mental Images and Their Status

There has been a tradition in psychology and philosophy that dismisses mental images as *epiphenomenal*; that is, they do not causally participate in reasoning or problem solving. Some even dismiss claims of mental imagery with the assertion that people who think they have mental images are just imagining things. No one at the symposium denied the phenomenal reality of mental images, but there were discussions about the sense in which they were pictures (P. Slezak, University of New South Wales, Australia). Obviously, this point is the heart of the imagery debate. However, the discussion in the field about this issue has been stymied by a lack of consensus on what various terms mean. We hope the following discussion is helpful in clarifying the relevant issues.

Let us refer to whatever it is that people have when they say they have a mental image as the *phenomenal image* and denote it by *I*. The agent can describe its content as propositions, but phenomenally, it is an image for the agent. Let *R* refer to the pattern

We anticipate that traditional symbolic representations will increasingly be combined with iconic representations in future AI research and technology

of activation in the neural structure in long-term memory that contained the information that gave rise to *I*. Let *P* refer to the pattern of activation in the neural structure when the agent reports having *I*. The neural pattern *P* occurs in the part of the cognitive architecture that corresponds to current awareness or deliberation.

What is the appropriate way to talk about *R*? We can talk abstractly about the content of *R*. We can certainly have a debate about which language type is the best to talk about or describe the contents of visual memory. We can draw pictures to describe the content; after all, we are talking about information that directly gave rise to the phenomenal image. Logicians in AI have argued that predicate logic can also describe the content well, and there is no need to assign a privileged and unique role to images as a language for describing this content. However, other logicians (J. Barwise, Indiana University at Bloomington, and J. Etchemendy, Stanford) argue that there is something fundamentally different about pictures as representational categories. Whatever the resolution of this debate, our point here is that the debate is about the appropriate descriptive language. It makes no sense to ask if *R* is a picture or a set of propositions: *R* is neural pattern. It does make sense to ask if there is a consistent interpretation of *R*'s content as an image, but such an interpretation is not in opposition to *R* having a consistent interpretation as a set of propositions. In this sense, image versus propositions is a false opposition about the contents of *R*. With advances in neuroscience, we might discover that the neural pattern *R* has a preferred interpretation as the content of a picture, but still talking about *R* as a picture would be a category error.

How about the quality for *P* of being picturelike? *P* is the neural activation that corresponds to the agent having a phenomenal image. All the points that we made about the content of *R* are applicable to the content of *P* as well, but we can also ask what it is about *P* that gives the agent the experience of the phenomenal image. One hypothesis is that the activation pattern of *P* has commonalities in both the shape and the location of the activation patterns with some part of the architecture that is involved in seeing. Specifically, if the architecture under discussion is the part that is involved in the final stages of perception, having the activation pattern *P* would correspond to the experience of image *I*. So far, there is nothing specifically picturelike about *P* other than it gave rise to *I*. However,

as we discuss in the next section, some image-processing theories propose that the neural-activation patterns corresponding to seeing have a systematic spatial array character. To the extent that *P* shared the patterns with those that occur in perception, there would be a sense in which *P* would be picturelike. At the symposium, a proposal was made (B. Chandrasekaran and N. H. Narayanan, The Ohio State University) about *P* and what features it shared with the pattern that arises during perception.

Another dimension to how mental images arise is the degree to which typical mental images are the result of composition and construction operations at run time as opposed to the retrieval process of complete images. When people are asked to recall images of objects that they are familiar with (for example, the map of South America), they often produce images in which different parts were moved and rotated to standard positions in relation to each other (B. Tversky, Stanford). Thus, support is given to the constructivist hypothesis. However, the issue of how the elementary pieces are represented is left open (that is, how far down does the constructivist hypothesis go? Are there any visual primitives in representation?).

The World, Diagrams, and Mental Images: Some Distinctions

We use the term *visual information* to refer to information that humans can extract by inspection from an image or from the world by directing visual attention to it. What can be extracted visually by inspection varies somewhat from person to person because of training and other personal factors; hence, an exhaustive list cannot be given of what this information consists of, but shapes, certain simple spatial relations, color, texture, and so on, are the kind of information we have in mind for this term.

It is useful to make a number of distinctions so that the issue of mental images and their representation and use is not conflated with the issue of diagrammatic reasoning in general.

Seeing: Seeing is perceiving the three-dimensional (3D) reality of the visual world. As we mentioned, most of this is in the realm of image processing, from sensory information in the retina to the construction of, say, 3D shape descriptions.

It is possible that scanning a mental image is scanning only in a metaphoric sense.

What kinds of representations can support the range of behaviors associated with the use of images and diagrams?

Problem solving by using vision on the world: Here, as the problem-solving process sets up subgoals that require visually obtainable information from the world, visual attention can be shifted to the relevant part of the world, and information can be extracted.

Drawing: An external representation is constructed, often from memory, of relevant aspects of some visual domain, possibly imaginary. There are interesting issues about how the pieces of the drawing are conceived, related, annotated or otherwise modified, and so on.

Using drawings for problem solving: Using drawings for problem solving involves many of the same processes as using vision on the world for this task (scanning by attention-driven mechanisms, using perception for extracting needed information, and so on), but the drawings are almost always much more simplified, though they preserve some of the visual information about the world that is necessary for problem solving. Also, many drawings are abstract and are not meant as veridical visual representations of the real world, but they are used to represent information about some problem in a form that can visually be extracted easily.

Imaging: *Imaging* is making a mental image of some aspect of a real or imaginary visual world. Here, some of the issues are similar to drawing, but there are open issues about how the images are generated and processed (for example, what mechanisms are shared between processing an external diagram and an internal one).

Using mental images for problem solving: Some of the issues here are similar to those for drawings in problem solving, but it is not clear what scanning really means here or what role perception plays.

A point that repeatedly comes up in discussions on imaging is the degree to which the hardware is shared between seeing and imaging. In the previous section, we talked about the location and shape of activation patterns being shared between mental images and perception. It is unlikely that early visual processing (that is, to use David Marr's [vision researcher and author of *Vision*] language, extraction of the primal sketch, 2-1/2-dimensional sketch, and so on, from the retinal

array) is being performed on a mental image. (To suppose that the full range of image-processing operations that are performed during the seeing of objects in the world or external diagrams is also done on mental images is truly to fall into the vortex of infinite regress!) However, a repeatedly proposed hypothesis is that there is a visual buffer in which the final stage of visual perception resides, and this visual buffer is shared between imaging and seeing. In Marr's view, one can think of image processing as taking place in layers. Each layer preserves some aspects of the pictorial quality of the visual world, the lower layers corresponding to the results of early image-processing operations and the final layer corresponding to the visual buffer whose elements are perceptually interpreted elements of the scene. The buffer preserves the spatial relations between the elements, each of which is already, in some sense, interpreted, not requiring low-level visual processing. Only the elements of the visual buffer are available to the agent for access (in perception) or construction (in imaging).

Similarly, some of the motor and perceptual mechanisms are shared between seeing the world, seeing a drawing (which is a representation of some possible world), and using an image during problem solving. The extraction of certain classes of visual predicates seems to be shared by all these processes. Scanning is also generally thought of as an activity that can be performed on the world, the drawing, and the mental image, although it is doubtful that an agent has to engage the same muscular actions in scanning a mental image as he/she would in scanning an external image of the world. It is possible that scanning a mental image is scanning only in a metaphoric sense, but instead, what really happens is the reconstruction of the image to correspond to the results of scanning.

In the constructivist view, images are constructed in the visual buffer from pieces that are set in a certain relation. The pieces themselves tend to be stereotypical in character, resulting in relations between the elements of the image that violate verisimilitude. In this sense, imaging is partly like drawing, with all the errors and biases that are features of most people's drawings. The constructivist view still leaves open the issue of how the elementary pieces of the image are stored.

A persistent issue about images is the degree to which additional direct work is done on the images by visual modality-specific operations to yield further information.

This issue can be illustrated by considering two examples, one owed to R. Lindsay (University of Michigan) and the other attributed to Z. Pylyshyn (who was not present at the symposium):

Take one step north, one step east, and one step south. How do you get back to the starting point?

Imagine a vat of blue paint, and imagine pouring yellow paint into it and stirring it. What do you see happening to the paint in the blue vat?

To the extent that articulations of cognitive strategies that people consciously adopt for solving a given problem can be relied on, it seems that a common strategy for the first problem is to construct a mental map and extract from it, by inspection, the relation of the starting point to the ending point. It is also possible to directly use prior factual knowledge to conclude that the starting point is to the west of the ending point. A deliberative syllogistic reasoning account, distinct from both the image-based strategy and the direct application of factual knowledge, is theoretically possible but rarely reported. However, in the paint example, a strong argument can be made that the conclusion that one sees the paint in the blue vat turning green is not given by the image, but rather, it is retrieved from one's store of factual knowledge and imparted to the image.

Representational Issues

What kinds of representations can support the range of behaviors associated with the use of images and diagrams? At the symposium, there were proposals based on perceptual primitives (Chandrasekaran and Narayanan) for mental images and on a distinction between spatial and visual representations as the basis for computational imagery (J. Glasgow and D. Papadias, Queens University, Kingston, Ontario). Two-dimensional pixel representations and pixel-pattern-to-pixel-pattern inference rules were presented as components of a deductive system that uses only picturelike representations (G. Furnas, Bellcore).

Diagrams and Images in Problem Solving

The symposium gave rich evidence that images and diagrams are used extensively in reasoning and problem solving in all sorts of domains: scientific discovery (Y. Qin and H.

Simon, Carnegie Mellon University [CMU], and P. Cheng, CMU), economic reasoning (H. Tabachneck and Simon, CMU), architecture (J. G. Wickham, University of Toronto), creative and conceptual design (V. Goel, University of California at Berkeley, and B. Faltings, Swiss Federal Institute of Technology), mathematical proofs (Barwise; Etchemendy; S-J. Shin, University of Notre Dame; K. Stenning and J. Oberlander, University of Edinburgh; D. Barker-Plummer, Swarthmore College; and S. C. Bailin, CTA Inc.), biological reasoning (A. Kindfield, UC Berkeley), molecular scene analysis (Glasgow and Papadias), geometric theorem proving (K. Koedinger, CMU), planning (Faltings and P. Pu, University of Connecticut), college physics (G. Novak, University of Texas at Austin), and procedure instructions (B. Pedell and H. Kuwahara, Michigan Technological University). What properties of images make them so useful for problem solving?

In the following discussion, we use the word *diagram* to stand for both mental images and externally drawn diagrams. (We already provided some arguments for the hypothesis that the processes that are involved in extracting visual information from diagrams and images share many elements. At any rate, this hypothesis is a working hypothesis for the rest of the discussion in this section.) The following argument suggests why and under what conditions diagrammatic representations might be helpful.

Diagrams preserve or directly represent locality information. A number of visual predicates are efficiently computed by the visual architecture from this information, for example, neighborhood relations, relative size, and intersections. This ability makes certain types of inferences easy or relatively direct. (It should be emphasized that only some visual predicates are easily computed by the visual architecture. The visual system is not good at making many inferences for which information is directly available in the diagram. For example, given a large circle and a smaller circle, the visual system can directly tell that one is smaller than the other, but given two complicated shapes, where one of the shapes has a smaller area than the other, the visual architecture cannot compare them directly or easily without additional measurements and calculation.)

This ability to get some of the needed information visually explains the role of diagrams in problems that are essentially spatial, such as geometry problems. Even here, there are interesting strategies by which the dia-

What properties of images make them so useful for problem solving?

*diagrams
aid in the
organization
of cognitive
activity.*

gram is additionally manipulated. Additional constructions are made on the diagram (Lindsay), which enable the detection of new visual information (or emergent properties, according to Koedinger) in the next cycles of inspection. Also, symbolic annotations can be made on the diagram, which enable a new round of inferences to be made, not by the visual architecture but by use of information in the conceptual modality. The result of this inference can be represented in the diagram, which can then enable the extraction of additional visual information. This highly interlaced sequence of extraction of visual information and symbolic inference making is what gives this whole approach its power: Each modality obtains information that it is best suited for and then sets up additional information that makes it possible for the other modality to arrive at additional information for which it is best suited (K. Myers and K. Konolige, SRI International; Narayanan and Chandrasekaran; and Barwise and Etchemendy). Cognitive models for such visual interactions were proposed and discussed (E. Rogers, Georgia Institute of Technology).

Now, the previous discussion explains the role that diagrams play in problems that have an intrinsic spatial content. What about problems that are not spatial? What explains the ubiquity of diagrams in such problems as well? The key is the existence of mappings with the property shown in figure 2.




Figure 2. Mapping from a Nonvisual to a Visual Problem.

R is a nonvisual representation from which it is desired to compute a predicate P . V_R is a visual representation of R such that P_R and, hence, P can quickly be computed by the visual architecture.

Suppose that R is some (nonvisual) representation and that P is a predicate that we are interested in computing about this representation. Suppose a mapping is available such that R goes to V_R , a visual representation. Suppose also that some visual predicate P_R can be extracted from V_R efficiently by the visual architecture and that P can be obtained from P_R by the reverse of the mapping from R to V_R . If these conditions are satisfied, then

one can effectively use the visual representation for the original problem.

All the assumptions in this scenario might make it sound as if it is a rare event for such mappings to be found, and it would only be infrequently that a nonvisual problem might be aided by such visual analogs. In fact, however, for commonsense reasoning phenomena, over a period of time, we learn a whole repertoire of such analogs constructed by the culture. Think of how common mappings from temporal phenomena to spatial phenomena are. We often represent time as a line and reason with lengths of lines when we want to compare durations.

In general, mapping to visual representations is a possible option for nonvisual problems whose properties can be mapped to spatial properties that are easy to visually recognize. The phrase *easily recognizable visually* is important here because many properties are spatial but not easily recognizable visually. Properties that are easily recognizable visually include such things as larger and smaller, longer and shorter, more to the right (left, above, below), thinner and thicker, and brighter and darker. In these types of properties, what is important seems to be the notion of (total) ordering in the one or two dimensions in which the properties of interest vary. Another type of spatial property that is often used to represent nonvisual properties is containment, as in the case of Venn diagrams. Spatial containment of the sort used in Venn diagrams is also easily recognized visually. However, containment does not necessarily impose a total ordering among represented objects.

Diagrammatic representations play another important role (Novak, Koedinger, Simon, and co-workers). In this role, diagrams help in selecting methods to solve a problem, not so much in making immediate inferences. That is, diagrams aid in the organization of cognitive activity.

However, there are also limitations to the diagrammatic form of reasoning. Because diagrams are concrete representations (that is, they are models in the sense of Johnson-Laird's mental models), when inferences can be made, they are fast and direct. This concreteness enables such representations to avoid some of the problems associated with the frame problem in inference (S. Huffman and J. Laird, University of Michigan). However, unless special techniques are employed, they cannot be used as such to perform reasoning involving universal quantification. Disjunctive reasoning (especially when the

numbers of disjunctions are large) becomes hard in a purely visual mode. These limitations of a purely diagrammatic reasoning are why most nontrivial problems involve the integration of both visual and nonvisual modes of reasoning (Barwise and Etchemendy).

Quite a bit of discussion was devoted to the use of sketches in the task of design (Goel and Faltings). Sketches are often intended to be vague in some aspects; that is, the person who draws one is not committed to all the dimensions and relations as drawn. Several speakers argued that this vagueness plays a functional role: It helps the designer avoid overcommitment to those aspects of the design to which he/she is not yet ready to make a precise commitment but, at the same time, still take advantage of the visual mode for organizing problem-solving activity and inference making. In fact, what makes sketches especially useful is the fact that a sketch is not so much vague as it is something that stands for a family of precise models.

The role of generic visual forms in facilitating creative discoveries, demonstrated by experiments in which subjects constructed mental images of such forms and then explored various ways of interpreting them, was a topic of discussion (R. Finke, Texas A&M University). A pilot study of cognitive difficulty in display-based problem solving, in which subjects solved Tower of Hanoi isomorphs by directly manipulating computer displays of problem objects, was also presented (M. Lewis and J. Toth, University of Pittsburgh).

Some of the discussion related to the role of images and diagrams in qualitative reasoning about the physical world. It was suggested (Narayanan and Chandrasekaran) that much of our commonsense knowledge about how objects in the world behave under various forces and collisions is actually in the form of abstract perceptual chunks that directed the reasoning activity. Representation of object-configuration diagrams and predictive knowledge indexed by shapes, as well as visual events, were emphasized as central problems. This work contrasts with much of the current work in qualitative physics, which emphasizes symbolic and axiomatic reasoning in this task. The use of visually based analog simulations for predicting the behavior of liquids to augment traditional symbolic methods was also proposed (J. Decuyper, D. Keymeulen, and L. Steels, Free University, Brussels). Using visual representation for analog simulation was, in fact, a general theme

(echoed by Lindsay).

As mentioned earlier, *normative logic* (that is, logic as an approach for deriving sound conclusions) has long treated diagrams as mere heuristic aids in the organization of thought and as objects that should not appear in the proofs that are finally produced. A number of papers that were presented at the symposium challenged this tradition in logic. The general point that united all of them was that it is possible to devise logical systems that take diagrams seriously and that have their inference rules organized partly as diagrams with symbolic annotations. Some interesting systems based on this approach were discussed or demonstrated at the symposium (Barwise and Etchemendy; Shin, Stenning, and Oberlander; and Barker-Plummer and Bailin).

Concluding Remarks

The symposium was successful in many dimensions. It fostered many interdisciplinary interactions: Psychologists, philosophers, and computer scientists listened to each other about problems of common interest. The discussions were intense. The symposium revealed an emerging consensus on the need to expand AI's representational repertoire to include diagrammatic representations and integrate their use with traditional representations. Furthermore, it highlighted contributions that researchers in other disciplines can make to this endeavor. Interestingly, the fact that researchers in different disciplines bring different assumptions to bear on their work was brought into sharp focus.

For example, what many AI researchers considered to be a given (S. M. Kosslyn's imagery theory), researchers in psychology considered to be problematic. A number of ideas and systems were placed on the table for use in diagrammatic representations by both humans and computers and for a wide variety of problem-solving tasks. An electronic mailing list was set up for continued discussions and information exchange.²

Following the symposium, an informal survey of its impact was conducted among the participants. An interesting picture emerged from the responses. Most respondents said that their main goals in attending the symposium were to learn about cognitive and computational research concerning the representation and use of diagrams, meet other researchers, and publicize their own work, and they agreed that the meeting definitely facilitated these goals. Many were

*what many
AI researchers
considered to
be a
given,
researchers
in psychology
considered
to be
problematic*

pleasantly surprised by the diversity of the work presented, but some lamented the lack of specific applications and connectionist models. Although the program devoted almost a third of the total time to discussion periods, respondents felt that more time was needed. This comment is an indicator of the enthusiasm generated, and the mailing list is, indeed, serving as a forum for continued dialogues. This enthusiasm foretells of a bright future for research in this area.

The survey also showed that a precise characterization of diagrammatic representation and reasoning processes and a shared view on their role and advantages did not emerge from the symposium. Different groups used these terms with different meanings in mind, thereby conflating issues that ought to be kept distinct. However, this fact is not surprising given the ground-breaking nature of this gathering and its scope. Indeed, we would have been surprised to find a consensus view emerging! Instead, the symposium served as a catalyst to initiate discussions on fundamental issues, which still continue through electronic mail.

Thus, the symposium was successful in meeting the following goals: initiating interdisciplinary contacts and discussions, providing an overview of current research (albeit spotty at places), and promoting increased research interest in the topic. It did not result in a consensual characterization of the area or a list of future research problems. Instead, it brought out important issues and initiated discussions regarding these issues. Most importantly, it was agreed that the topic of the symposium certainly deserves more research attention and that future meetings of this kind would be of great value.

Acknowledgments

First and foremost, we acknowledge the contributions of the participants: They made the symposium so successful. We also thank AAAI for facilitating the meeting, survey respondents for sending us feedback, and Dave Barker-Plummer of Swarthmore College for setting up the diagrams mailing list. Chandrasekaran and Narayanan's work in organizing the symposium and preparing this article was supported by the Defense Advanced Research Projects Agency under Air Force Office of Scientific Research contract F49620-89-c-0110. Narayanan also acknowledges support from BP America in the form of a fellowship. Peter Patel-Schneider made useful editorial comments on the draft. Finally, we owe Pat Hayes a great debt for reading

through an earlier draft of this article and making his customary friendly but thorough critiques. The final version has benefited enormously from his critique, but we do not imagine that it resolves all his criticisms satisfactorily.

Notes

1. The working notes were intended only for distribution to the attendees. However, interested readers can contact Hari Narayanan at <narayan@harl.hitachi.co.jp> for the list of papers and subsequently get in touch with the authors directly for reprints.

2. Please send requests to <diagrams-request@cs.swarthmore.edu>.



B. Chandrasekaran directs the Laboratory for AI Research and is a professor of computer and information science at The Ohio State University. His research interests include knowledge-based systems, the use of images in problem solving, and the foundations of cognitive science and AI. Chandrasekaran received his Ph.D. in electrical engineering from the University of Pennsylvania in 1967. He is editor in chief of *IEEE Expert* and is a fellow of the Institute of Electrical and Electronic Engineers and the American Association for Artificial Intelligence.



N. Hari Narayanan is a visiting research scientist at the Hitachi Advanced Research Laboratory, Hatoyama, Saitama 350-03, Japan. He is interested in how perceptual representations aid intelligent behavior, with his current research focusing on

problem solving with diagrams and diagrammatic representations. He is currently editing a special issue of *Computational Intelligence* on reasoning with diagrammatic representations.



Yumi Iwasaki is a researcher at Stanford University's Knowledge Systems Laboratory. Her research interests include model-based reasoning and reasoning with pictorial representations. She received her B.A. in mathematics from Oberlin College, her

M.S. in AI from Stanford University, and her Ph.D. in computer science from Carnegie Mellon University. She is a member of the American Association for Artificial Intelligence, the Japanese Society for AI, Computer Professionals for Social Responsibility, and Sigma Xi.