Proceedings of the 8th Conference on Computer Generated Forces and Behavioral Representation, 11-13 May 1999, Orlando, FL, Defense Modeling and Simulation Office & Defense Advanced Research Projects Agency, pp. 117-126.

Cognitive Modeling For Simulation Goals:
A Research Strategy for Computer-Generated Forces

B. Chandrasekaran
John R. Josephson
Laboratory for AI Research
Department of Computer & Information Science
The Ohio State University
Columbus, OH 43210
614-292-0923, 614-292-0208
Chandra@cis.ohio-state.edu, jj@cis.ohio-state.edu

Keywords:
Simulation goals, relation between fidelity and simulation goals, simulation task

ABSTRACT: *Simulation goals determine the fidelity requirements for cognitive models. We argue that the pursuit of high fidelity cognitive models, unfettered by detailed considerations of what we want the models for, is so unfocused as to be almost useless for practical purposes. What a cognitive model needs to contain is vitally affected by what kinds of questions one wishes to answer, i.e., the goals of the simulation. However, there is no reliable and systematic body of knowledge in the simulation community about the demands placed by the goals of simulation on models of cognitive (and other) entities. The mappings from simulation goals to fidelity and other requirements for cognitive models can be investigated empirically. We outline strategies for how this can be done.*

1. Fidelity Requirements and Simulation Goals

Fidelity, defined as faithfulness to the object or process being simulated, usually ranks high among the goals of simulation researchers and developers. Achieving fidelity requires, first of all, scientific understanding of underlying reality and, second, technologies that help turn scientific understanding into computer models to be used in simulation exercises. While the simulation community generally acknowledges that perfect fidelity is unattainable, because no model can exhaustively capture physical reality, still the goal has been to continually improve the potential for fidelity by continually improving scientific understanding of the phenomena involved, and by continually improving the related computational technologies.

Our scientific knowledge of the world of physics is sufficiently advanced that, arguably, the limitations in modeling military platforms and environments are now largely technological. This is not the case, however, with respect to modeling cognitive agents. We are still at a very early stage in our scientific understanding and computational modeling of human cognition. (We are even more primitive in our understanding and modeling of the cognitive behavior of groups of people, even highly structured groups like military organizations. Moreover, the most challenging of potential adversaries have computers, and we are very primitive in our understanding and ability to model computer-aided cognition, which is rapidly evolving.) Almost every problem of interest in CGF (Computer-Generated Forces) — mission planning, situation assessment, spatial reasoning, etc. — involves scientific issues being actively studied by cognitive modeling and artificial intelligence, two closely allied sciences concerned with building computational models of human and other forms of cognition. Thus, it seems that increasing the fidelity of cognitive agents in military simulation needs, more than anything else, scientific progress in these disciplines. However, military simulation technology does not simply need general progress in cognitive modeling, it needs progress of very specific kinds.

We argue that the pursuit of high fidelity cognitive models, unfettered by detailed considerations of what we want the models for, is so unfocused as to be almost useless for practical purposes. No cognitive model can be detailed enough, in every respect, to serve for every purpose we might conceivably use it for, and no practical purpose is such that an unbounded degree of fidelity is significantly useful for its practical achievement. We argue that what a cognitive model needs to contain is vitally affected by what kinds of questions one wishes to answer, i.e., the goals of the simulation. While we won't go so far as to say that fidelity is in the eye of the beholder, it is certainly something that must be judged by the needs of the model user.

However, there is no reliable and systematic body of knowledge in the simulation community about the demands placed by the goals of simulation on models of cognitive (and other) entities. We urge that these demands be studied empirically. We will describe how this can be done.

2. Research Strategy

Here is the strategy in a nutshell:

> *Conduct a set of case studies wherein concrete examples of military uses of simulation are studied empirically. For each case, analyze the demands on simulated entities and characterize these demands relative to the purposes for which the simulation is conducted. Map out the tasks to be performed, both by the simulation user and by simulated agents, in terms of goals and subgoals. Develop concepts and vocabulary for describing these task structures and for characterizing the similarities and differences among the different uses of simulation. Seek generalizations that aid description and analysis and help to predict the characteristics of further examples.*

All cognitive simulation is done in the twin context of some tasks performed in the world by some agents that we wish to simulate, and some purpose for which the simulation is used. While instances of such tasks and purposes may be too numerous to list, nevertheless there are certain generic types of them that can be identified. Generic military simulation purposes include training, mission rehearsal, acquisition decision-making, development of strategy and tactics, planning, and evaluation of doctrine. Of course, the trainees, acquisition decision-makers, and so on, are also performing cognitive tasks when they use simulation exercises for their respective purposes. Thus, there are two sets of cognitive tasks that need to be analyzed: tasks of the agents who are being simulated, and tasks of the agents who are using the results of simulation for some purpose. We will use the terms “subject-agent task,” and “simulation-user task” to be clear about which kind of task we intend when there is danger of confusion.

Let us consider two examples, a use of simulation for training a company commander in tactics and a use of simulation in helping to assess an acquisition proposal. In the training example, the subject-agent tasks include the cognitive tasks performed by the trainee's subordinates and the opposing units. In the acquisition example, similarly, the subject-agent tasks are the cognitive tasks performed by the various agents in the military situation. Let us first consider how the simulation-user task would proceed in the absence of computer-based simulation models. In the training of a company commander, the trainer would likely create an exercise in which the trainee is exposed to real command problems. That is, the trainee is placed in a military exercise in which the subject agents are given certain goals and scripts. The exercise would create field situations to which the trainee would have to respond. In the acquisition example, similarly, the exercise would involve military actions in which the relevant agents use specified equipment, and the agents are instructed to achieve various goals and act in various ways. In both cases, the user controls the exercise in such away that the subject agents produce behavior that is necessary for the user's goals to be accomplished. When some of the subject agents are replaced by CGFs, the requirements on the behaviors of the CGFs are identical to the requirements on real agents in the exercises.

The kind of case study that we are proposing is aimed at identifying the requirements on agent behaviors, whether or not the agents are to be humans or CGFs. In each study, analysis would identify which entities need to be modeled for achieving the simulation goals, what kinds of information each entity needs, what kinds of behavior and information each entity produces, and what kinds and degrees of fidelity are needed. The outcome of such a study would be a kind of task analysis [1], a hierarchical decomposition of the simulation goals into subgoals, and a specification of the requirements for the various subgoals. We are especially interested in the behavioral information that the simulation user needs from cognitive agents in the domain. For example, a subgoal in training a company commander might be to expose the trainee to situations where the opposing side adopts a certain tactic, or the opposing side has certain weapons. For this subgoal to be accomplished, the simulation users – in this case, the trainer and the trainee – would need, among other information, information about the behavior of the opposing units. On analyzing the information need, one might discover that, for this particular subgoal, the opposing agent's behavior need only be modeled as making action choices probabilistically. That is, as long as the trainee is exposed to a variety of situations that cover the full range of opponent's behavioral choices, the training subgoal would be achieved. Alternatively, the subgoal might require that the opponent generate actions that satisfy a specified optimality criterion, say, minimax. Identifying the requirements for agent behavior helps the analyst to determine what agent models will be adequate and appropriate for the simulation exercise.

In order to do the analysis properly, the analyst needs to have suitable vocabulary for describing the behaviors of the subject agents. In cases where the analyst is not in possession of this vocabulary, an analysis of the tasks and behaviors of the subject agents may be required. The results of this investigation would also take the form of a task analysis, in which the goals and subgoals of the subject agents are made explicit and information requirements are identified. This analysis only needs to be performed to the degree needed for the simulation purposes. For some of the purposes, it may be necessary to understand how to achieve the goals in the context of the human architecture, while for other purposes only the logical requirements of the task need to be modeled without concern about how humans actually achieve a certain goal. For example, for a particular simulation purpose, the task of the agent may be simply defined as generating a minimax-optimal action, while for another purpose, the task structure might include subtasks in which decisions make explicit demands on the agent's capacity for short-term memory.

The research strategy we propose thus requires that we study real instances of uses of simulation to generate an understanding of how simulation purposes impact on the kinds of models that are needed. We believe that by performing such studies over a series of examples, increasingly generic understanding will be achieved about the relationship between purpose types and agent model types. Generalizations will emerge that will be useful in staging exercises and assessing their significance. To give a suggestive example, it is plausible that, for a broad range of uses, individual agents must be represented for one or two echelons below the level of primary focus, but that aggregated agents will suffice for yet lower echelons. We note that, in the unlikely event that no generalizations emerge, and nothing at all can be carried from case to case, still it will have been useful to have studied in detail the requirements on agent models to help guide development for the specific cases.

We propose that the community start with a few — on the order of three or four — combinations of generic cognitive tasks and generic simulation purposes. Situation assessment in the context of training, and mission planning in the context of acquisition decision making, are examples of such combinations. For each such combination, we propose that a specific real instance be identified. The next step will be to perform empirical studies and task analyses of both the subject-agent tasks and the simulation-user tasks. An analysis of the subject-agent tasks will describe how each task determines the type and fidelity of the information the agent needs and how the agent uses the information. The analysis will identify what an agent needs to know, what the needs from other entities that he/she interacts with, what problem-solving strategies are used, the potential impact of so-called performance modulators on the decisions made, etc. The parallel analysis of the simulation-user tasks will describe how the simulation user makes use of the information about the subject agents performance of their tasks. In the training example, one might characterize the nature of models of the opposing agents. Is it necessary for them to generate: optimal behaviors, likely behaviors, a range of plausible behaviors, or a wide range of behaviors including unlikely and irrational behaviors? When can subordinate echelons be modeled as aggregate agents? When is it important to model the effects of fatigue on decision making?

Going beyond a small initial set of selected generic task/purpose combinations, as the community gathers experience in task analysis and cognitive modeling, additional examples can be studied to obtain and validate better generalizations and in order to achieve a better general understanding of the mapping from simulation goals to fidelity and other requirements for cognitive models. The initial research will also suggest additional concrete studies that can be done.

How to set up the collection of data and perform the analyses require additional thinking, but the central point is that case studies should be performed to analyze the actual requirements for cognitive models in real-world examples of military uses of simulation. In addition to well-established techniques of cognitive modeling from psychology, task analysis techniques from artificial intelligence can be effective in analyzing the structure of knowledge and the processing strategies that subject agents must have to satisfy simulation goals. Knowledge acquisition techniques from knowledge engineering can be used to determine and document what builders of simulation exercises already know about requirements for agents. In addition, the sensitivity of simulation outcomes to various aspects of agent fidelity can be assessed by directed experiments with simulation models.

2. 1. Research Strategy and Environmental Modeling

The proposed research strategy can also help in characterizing the requirements on models of environments and platforms. The issues that gave rise to the proposed strategy apply not just to fidelity issues regarding CGFs. Very similar issues arise about requirements for environments and physical platforms. The empirical studies that we propose may similarly be applied to develop an understanding of the relationship between simulation goals and fidelity requirements for models of any of the entities that need to be simulated. In fact, there is no need to conduct separate studies, one for understanding requirements on cognitive agents, one for requirements on environments, and so on. One carefully conducted study can identify the constraints on all entity models that are needed for the instance of simulation use, and generalizations over a number of studies, similar to what we suggested for CGFs, should be possible for models of other entities as well.

3. Discussion

Simulation is used in the military for a number of different purposes: training, mission rehearsal, doctrine evaluation, and acquisition decisions, to name a few. Programmatically, technology for interactive simulation has been pursued in three separate but parallel streams: physical platforms (tanks, helicopters, etc.), physical environments (terrains, oceans, skies, etc.), and computer-generated forces (CGFs -- cognitive agents at different echelon levels). But this programmatic division has resulted in each stream of development pursuing ever-increasing fidelity, relatively unconcerned with how the models will be actually used for different simulation goals, and unconcerned with the relationship to the models of entities in the other streams. Excellent renderings have been created of the crests of waves metamorphosing into spray and of individual leaves shaking in the wind. Similarly concerned with fidelity, the cognitive modeling community can easily point to the myriad ways in which current CGF models fall short of their human counterparts. But how, and when, and what kinds of fidelity are required, have not been major concerns.

Fidelity is certainly an attractive goal in the abstract, but there are some limitations on the pursuit of fidelity, especially in CGFs:

- Fidelity in cognitive agents is an open-ended enterprise. In contrast to the case of physical platforms and environments, the scientific basis of AI/Cognitive Science is nowhere near maturity. Thus, even if cost were no object, we simply can't produce models of the fidelity required for some purposes. However, the simulation effort cannot wait until the day that the scientific theory of cognitive agents is mature. We need to know what simulation goals can be accomplished with current knowledge, and what simulation goals require advances in what aspects of the underlying science.

- There are different types of cognitive models with different properties in different dimensions. There are aggregate and individual models, decision-theoretic and game-theoretic models, models where behavior is generated by constraint-satisfaction, by problem-space search, by probability distributions over possible behaviors, etc. These models have varying costs for development, instantiation, and deployment, differing run-time demands on computational resources, and offer different fidelities in different dimensions of behavior. An aggregate model may have higher fidelity in certain dimensions than a model that consists of a large number of interacting constituent agents and would probably be cheaper and faster to compute. If we wish to train someone for the worst case, a game-theoretic mini-max model of the opponent might give the right kind of fidelity, while for some other training goals it would be more useful to have an opponent that non-deterministically displays a range of plausible behaviors. Sometimes fidelity about the effect of sleeplessness on performance will be needed, sometimes not. Little is known about how to choose appropriate models, and how to interoperate them.

The general recommendation from the cognitive modeling community is that research in a range of basic scientific issues needs to be undertaken in order to help make progress in CGFs. The problem is that the list of issues that need attention for better CGFs is almost exactly co-extensive with the list of issues for cognitive modeling in general. From natural language understanding, to planning, to situation assessment, to the role of stress and emotions in decision-making, the scientific agenda of cognitive modeling is virtually the same as the list of areas for progress in the fidelity of CGFs.

While there is indeed much that the CGF program needs from research in cognitive modeling, we argue in this report for a specific strategy for conducting this research that we believe will help give us greater insights about the requirements for cognitive models for simulation goals. Cognitive modeling research can then focus more effectively on how to improve CGFs incrementally, and in a way that is motivated by the needs of simulation, rather than by an abstract pursuit of ever-increasing fidelity.

3. 1 Tasks of Military Agents and Simulation Goals

As a first approximation, we might identify the tasks performed by subject agents at different echelon levels as generic types. There is general agreement in the community that the tasks of agents in the lower echelons take place in scales of time and space that are much smaller than the tasks at higher echelons, and that the behavior is more reactive. This seems to be the case independent of the branch of the military. Also, in the case of lower echelons, the doctrinally appropriate behaviors of the agents can be prespecified as a set of reactive rules – rules that specify what actions to take under what conditions, without the need to engage in problem-solving. On the other hand, the behaviors of agents in the upper echelons cannot be completely specified in this reactive framework. As we move up the echelon level, the spatio-temporal scope gets increasingly wider. These agents, even when following doctrine, often need to implement doctrine for the concrete situation by additional problem solving. These kinds of regularities provide an initial basis for characterizing generic types of agents in the military.

As we have described, all cognitive simulation is done in the twin context of some cognitive tasks performed by some agents (or groups of agents) and some purpose for which the simulation is used. Agent tasks include such things as formulating air tasking orders, assessing enemy battalion strength and location, choosing avenue of approach, selecting target, etc. Simulation purposes include such things as training a company commander in tactics, assessing some set of doctrine, evaluating a proposed acquisition to assess its potential impact on the outcome of engagements, etc. Every instance of simulation can be located in this two-dimensional space. Together, the task and the purpose determine the kind of information and knowledge needed by the agent models, the required fidelity of the models in different dimensions, and the kinds of processing strategies that are appropriate or needed to produce the desired behaviors.

Certain tasks/purposes may require close attention to human processing limitations, while other tasks/purposes may be handled by ignoring how humans do it, but instead developing a suitable algorithm for the purpose that may even better suited than the model of the corresponding human. Some tasks/purposes may need modeling of the detailed surface of the wing of an aircraft, while other tasks/purposes might simply call for the aircraft to be modeled as a point in space. For some situations it may be necessary to make the agents as smart as the best human performing the task, while for other situations it may be necessary to ensure that the model produces irrational as well as rational behavior to ensure coverage of real world behavior.

In our view, the requirements placed on models of military agents are poorly understood in the simulation community. This is not to say that those in the field charged with using simulation for various purposes have not developed useful heuristics and intuitions about such requirements. On the contrary, it is quite likely that those involved in using simulations for specific problems have a good sense of the sensitivity of results to different dimensions of fidelity of the cognitive models available to them. However, any such practical knowledge does not seem to have moved to the community that is charged with producing the simulation technology.

3. 2 Can one agent model answer all questions?

A view seems to be current that, with the improvements in computing power, there is no longer any need to be concerned about the different requirements on cognitive models due to different simulation purposes. In this view, the only difference between different simulation purposes is how we analyze, filter and abstract the simulation results, not in what models we use or how we actually run the simulations. The cost of computing and the quality of the models are supposed to have gotten to the point that we can have one model of agents that can answer all questions. This view regards the concern with the relationship between simulation purposes and the types of cognitive models as an artifact of the funding sources for simulation, and not an intrinsic issue in simulation. We believe that this view is the result of overgeneralization of a useful insight: that indeed we don't need to handcraft from scratch cognitive models for each simulation setup or exercise, and that a large part of what is in cognitive models can be shared across exercises. It is true that we can formally treat the union of all the current models for an agent as the "one model" that will answer all the current questions about the agent. Even if such a universal model of an agent or agent type could be implemented, it would still be good to understand which details and dimensions must be represented in such a model to answer all current needs, which are of marginal value, and which are superfluous. As scientific progress results in models of more and more aspects of cognitive agents, it will become computationally unwieldy to exercise aspects of the agent that are not relevant to the simulation question of current interest. It is computationally demanding to simulate the detailed behavior of a single neuron, much more so to simulate the human brain in detail! Thus, there need to be limits to representing and invoking detail, whether or not the best implementational strategy turns out to be the use of a single all-purpose model for each individual agent or agent type.

4. More on Task Analysis

A task analysis is a functional decomposition of an information processing task. We have written elsewhere about the idea of a task structure [1], [2]. which is a representation of a task in terms of its inputs and outputs, the methods that are applicable for it, and the conditions under which each method is applicable. Such an analysis is intended to answer questions about how the task is accomplished or may be accomplished. The task is considered to be comprised of coordinated subtasks and those subtasks are comprised of sub-subtasks, and so on. The subtasks are those accomplishments that may be used to accomplish the task, and the sub-subtasks are accomplishments that may be used to accomplish the subtasks. For each task (subtask, sub-subtask, etc.) there will be one or more possible methods that can be used to accomplish it. Methods are ways to accomplish a task, as using an abacus, or using ones fingers, or using a calculator are alternative methods for performing the task of adding integers. In a task structure, associated with a method is a description of how it uses knowledge and inference to achieve its goals, and a specification of what subgoals (subtasks) it sets up and requires to be achieved before it can succeed. Alternative methods for a task may make use of a common subtask. This kind of decomposition can be done recursively, bottoming out in methods that achieve subgoals by “direct action” rather than by setting up additional subgoals. A method may be a procedure, where the sequencing of steps is pre-specified, or it might specify how subtasks are controlled and coordinated in some other way.

To give some flesh to this description of task analysis, we present in the Appendix a sketch of a task analysis for the generic military cognitive task of battlefield situation assessment. Battlefield situation assessment may be required of certain agents in a simulation exercise. Depending on the purposes of the exercise, situation assessment might be performed infallibly (e.g., by “cheating” and providing subject agents with direct access to high-level information about opponents), or it might be performed fallibly by providing agents with the kinds of information that would realistically be available in an actual military situation. The situation assessment task is a type of what a type of inference called *abductive inference.* The analysis in the Appendix is based on the task analysis of abduction in general and situation assessment in particular given in [3].

The task analysis presented in the Appendix will necessarily be very schematic, both because the kind of case studies we have called for have not yet been conducted, and because we give it here with more generality and imprecision than will characterize concrete examples of types of military agents and the situation-assessment tasks they face. Moreover, as we have suggested, task analysis should be performed taking account of the significance of the cognitive task in the context of simulations performed for specific purposes, such as concrete instances of training or of planning of defenses.

5. Concluding Remarks

We agree with the general consensus that, for making rapid progress in computer-generated forces, a program of basic research in cognitive modeling is necessary, since there are far too many things we don't understand about cognitive tasks performed by military agents. However, all forms of cognitive modeling are not equally useful for the purposes for which military simulations are conducted, and all avenues of research in cognitive modeling are not equally timely and significant for the development of useful CGFs. We propose a research strategy to build a body of knowledge that will help simulation exercise builders to map from their goals to the cognitive models that are most appropriate. Besides the importance of such knowledge for the construction of simulation exercises, it can also serve to direct modeling research toward actual and significant modeling needs.

We expect that researchers whose main interest is in cognitive modeling will be supportive of the focus in this report on analysis of the cognitive tasks of the agents involved in real military decision making. However, extending the analysis to include the user dimension, as we propose in this document, may strike some as bringing application considerations prematurely into basic research. Yet, simulation users are performing cognitive tasks as well; they are using information from the subject agent models to solve problems and improve performance. What is an application concern for general cognitive modeling research is a basic research concern for the science of simulation. *Understanding the relationship between simulation goals and the requirements on models is a basic research issue in simulation science*.

In summary, we have argued that the goal of increasing the fidelity of the models of entities in simulations leads to a mistaken tendency to conduct undirected research on cognitive and environmental modeling. Instead, we suggest that there is a need to develop a body of knowledge that helps us to understand the relationships between the real-world tasks of agents that we wish to simulate, the purposes of the simulation, and the requirements on the models. We need to develop a body of knowledge that tells us more about when aggregate models would be sufficient, when models of performance moderators are needed, when agent models need to produce optimal behavior, when tracing of the rationale behind behavior is important, when perceptual capabilities must be modeled in detail, and so on. We have outlined a strategy for research to develop a body of knowledge that relates the requirements on models to the purposes for which the models are to be used.

6. Acknowledgments

This work was performed under a contract with SAIC to consult for DARPA Advanced Simulation Technology Thrust. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of DARPA or of the United States Government. We acknowledge useful conversations with Dell Lunceford and Lashon Booker, but the authors take full responsibility for the contents of the paper.

7. References

[1] B. Chandrasekaran, "Design problem solving: A task analysis," *AI Magazine*, Vol. 11, no. 4, 1990, p. 59-71.

[2] B. Chandrasekaran, Todd Johnson, and Jack W. Smith, "Task Structure Analysis for Knowledge Modeling," *Communications of the ACM*, 33-9, Sep, 1992, pp. 124-136.

[3] John R. Josephson and Susan G. Josephson, (Editors), *Abductive Inference: Computation, Philosophy, Technology*, Cambridge University Press, 1994.

Appendix

Example: Task Analysis of Battlefield Situation Assessment

Goal: Monitor the external environment for entities of interest. Entities are interesting if they have the potential to impact the current mission. They potentially include threat forces, friendly forces, and other entities.

When an interesting entity is detected:

- determine its spatial location and movement
- recognize or classify it
  - friend or foe
  - type, e.g., military unit type
- infer high-level attributes about it
  - threats
  - alertness
    - has it detected us?
  - readiness status
  - recognized behavior, e.g., maneuver
  - goals, e.g., mission objectives
  - plans

The output (main product) of the situation-assessment task is a continually maintained composite description of the situational state, including meta-descriptions annotating aspects of the situational-state description with confidence estimates, specific reservations, and alternative interpretations. This description of situational state is used to provide specific reports, and to support decision making.

Inputs to the situation-assessment task include sensor data, intelligence reports, and current position, as well as information from mission briefings, weather reports, and terrain maps.

(Note: The situation-assessment process can provide valuable feedback to information-acquisition systems, for example, by guiding sensors to focus on areas and objects likely to provide information to discriminate between alternative interpretations. Such feedback may also be considered to be an output of situation assessment, although it is not the main product, since it is intended to achieve information-gathering subgoals that arise from subgoals of maintaining the description of situational state.)

In order for an agent to perform the inferences required to maintain an assessment of the current situation, it needs extensive knowledge about individuals and types of entities of interest, including knowledge about their capabilities and knowledge about probable goals and plans. Specific pieces of this knowledge will be used in specific ways for specific subtasks, e.g., knowledge of size may be used to confirm or rule out a potential classification hypothesis.

The situation-assessment process forms and maintains a composite interpretation of data. High confidence in the various components of that interpretation depends on chains and networks of inference that are well-justified. A typical pattern of justification, which may only be implicit in an inference process, is given by:

> D is a collection of data (facts, observations, givens).
> H explains D (would, if true, explain D).
> No other hypothesis can explain D as well as H does.
>
> ---
>
> Therefore, H is probably true.

Such a pattern is commonly called *abduction*, or *inference to the best explanation.*

The justification strength of an abductive conclusion depends generically on several considerations, including:

- how plausible H is in itself
- how much better H is than alternative explanations
- how thorough has been the search for alternative explanations

Thus, the task of assessing the strength of a situation-assessment conclusion will typically set up subtasks concerned with evaluating the conclusion according to these considerations. (We have written about abductive inferences extensively in [3].)

Various methods may be used for the situation-assessment subtasks concerned with locating and identifying entities and with inferring high-level characteristics of them. However, explicit use of abductive inferences to accomplish these subtasks has the advantage that justifications may be read off and reported as well as critically assessed to determine appropriate confidence levels. Alternative methods for recognition and classification include simple pattern recognition based on prestored patterns engineered by hand, and recognition based on patterns created by training, as in neural nets. Each method has its advantages and disadvantages, which can be spelled out explicitly.

In general, interpretations of situations should be coherent and consistent, which will require coordinating the inferences that are concerned with different entities of interest. Moreover, the composite interpretation of data formed and maintained by situation assessment may usefully be considered to comprise several levels or layers with, for example, identification of entity features at a lower layer than identification of entity types, which is in turn at a lower layer than determination of entity goals. Thus, the overall composite interpretation of data may have a complex justification structure with the overall form of layered abduction, where the interpretations formed at lower layers become the data that are explained at higher layers. Again, if the processes that form and maintain interpretations of data use abductive inference explicitly, the justifications will be accessible for reporting and critical oversight.

The formation of a composite interpretation by abductive inference can be accomplished by accomplishing two main subtasks at each layer of interpretation:

- subtask 1: *generating elementary explanatory hypotheses*,
- subtask 2: *synthesizing a composite best explanation*

The subtask 1, consists of a set of specific subtasks of generating an elementary hypothesis to explain each data item of interest. Each of these subtasks may be accomplished using two main sub-subtasks:

- sub-subtask 1.1: *evoke a relevant hypothesis-forming concept*
- sub-subtask 1.2: *instantiate an elementary hypothesis* (adapted to the case)

The sub-subtask 1.2, instantiate an elementary hypothesis, has two main sub-sub-subtasks:

- ▪ sub-sub-subtask 1.2.1: *score the elementary hypothesis for initial confidence*
- ▪ sub-sub-subtask 1.2.2: *determine the explanatory coverage of the hypothesis*

This decomposition is very general. A typical situation interpretation is a composite description, and somehow, explicitly or implicitly, there must be some method of choosing which elementary hypothesis to consider, some way of making them specific to the case, and some way of combining and deciding what to accept.

Evocation (sub-subtask 1.1) can occur bottom-up, a hypothesis being stimulated for considerations that are cued by the data presented at a layer below. That is, the presence of a certain data item suggests that certain hypotheses are appropriate to consider. More than one hypothesis may be suggested by a given data item. Evocation can also occur top-down, either as the result of priming (an expectation from a level above), or as a consequence of data-seeking activity from above, which can arise from the need for scoring (as when you look for a characteristic feature of an entity type under consideration, e.g., look for the gun on a possible tank). Evocations can generally be performed in parallel, and need not be synchronized.

Instantiation (sub-subtask 1.2) occurs when each stimulated hypothesis is independently scored for confidence (sub-sub-subtask 1.2.1), and a determination is made about what part or aspect of the data is accounted for by the hypothesis (sub-sub-subtask 1.2.2). Instantiation is in general top-down from a stimulated hypothesis toward lower-level data. During instantiation, data may be sought that was not part of the original stimulus for evoking a hypothesis. Each hypothesis is given a confidence value on some scale, which can be taken to be a “local match” or *prima facie* likelihood, a likelihood of being true based only on consideration of the match between the hypothesis and the data, with no consideration of interactions between potentially rival or otherwise related hypotheses. Typically, many evoked hypotheses will get very low scores, and can be tentatively eliminated from further consideration. The data that are accounted for by a hypothesis may or may not be identical to the data on the basis of which the hypothesis was scored, or the data that did the evoking.

In the course of instantiation the hypothesis set may be expanded, for example by including subtypes and supertypes of high-confidence hypotheses, if the space of potential hypotheses is organized hierarchically by level of specificity (e.g., a hypothesis might be generated identifying a platform subtype – not just a tank but a specific type of tank).

Instantiation is done most efficiently when it is based on matching against prestored patterns of features, but slower processes of instantiation are also possible whereby the features to match are generated at run time.

The result of instantiation activity is a set of hypotheses, each with some measure of confidence, and each offering to account for some portion of the data. Since within a particular wave of instantiation activity, hypotheses may be considered independently of each other, the process of instantiation can go on in parallel.

Composition (subtask 2) occurs when the instantiated hypotheses interact with each other and (under favorable conditions) a coherent best interpretation emerges. At the beginning many hypotheses will probably have intermediate scores, representing hypotheses that can neither be taken as practically certain, nor as being of such a low confidence as to be ignorable. So knowledge of interactions between the hypotheses is brought to bear to reduce the degree of uncertainty, increasing confidence in some of them, and decreasing confidence in others.

Author Biographies

B. CHANDRASEKARAN is professor Emeritus, a senior research scientist, and the director of the Laboratory for AI Research (LAIR) at Ohio State University. His research focuses on knowledge-based systems, causal understanding, diagrammatic reasoning, and cognitive architectures. He was editor in chief of IEEE Expert from 1990-1994. He is a Fellow of IEEE, ACM and AAAI.

JOHN R. JOSEPHSON is a research scientist and associate director of LAIR at Ohio State University. His primary research interests are knowledge-based systems, abductive inference, causal reasoning, theory formation, speech recognition, perception, diagnosis, the logic of investigations and the foundations of science. He is co-editor of *Abductive Inference: Computation, Philosophy, Technology*, Cambridge University Press, 1994.