

High-Performance Knowledge Base Support for Monitoring, Analysis, and Interpretation Tasks

Jon Doyle, Isaac Kohane, William Long, and Peter Szolovits

Massachusetts Institute of Technology
Laboratory for Computer Science
and
The Childrens' Hospital (Boston)

December 4, 1996

This is the main text of a proposal submitted to the Defense Advanced Research Projects Agency in response to BAA #96-43, "High-Performance Knowledge Bases". The funded project will run for three years, beginning May 1, 1997 and ending April 30, 2000.

Contents

A Abstract	2
B Innovative claims	3
C Technical rationale	5
C.1 A look at the future	6
C.2 System Architecture	7
C.3 Testing and Evaluation	13
D Comparison with other ongoing research	14
E Proposer's previous accomplishments	17
References	19

A Abstract

Monitoring, analysis, and interpretation (MAI) tasks constitute central elements of many important applications, including plan execution sentinels, battlefield situation awareness, and intelligence analysis. These tasks vary in scope from narrow process-control monitoring, which require attending only to a fixed set of observable conditions, to national security intelligence analysis, which must attend to a broad range of observations, information, and actionable conditions that cannot be enumerated beforehand.

Although commercial knowledge-based system (KBS) technology has proven successful for the narrower process-control tasks, current technology does not support very well the broader, open-ended monitoring and interpretation processes. New information regularly leads analysts to attend to new classes of information and new potential threats, but current KBS technology depends on small groups of system developers hand-crafting carefully circumscribed knowledge bases, and does not permit distributed communities of users, even those highly knowledgeable in the domain, to exploit separately-developed large-scale knowledge bases or to add new analysis knowledge and methods to these knowledge bases in many small increments.

We propose to develop tools that, in the hands of system developers, will *enable rapid construction* of knowledge-based MAI systems, and in the hands of communities of users, will *enable ready extension* of such systems to cover new knowledge and methods. We will build on existing and under-development tools for rapid, distributed (web-mediated) construction of ontologies and knowledge bases.

- *We will design representations and procedures* for describing and recognizing conditions, trends, and patterns of causality, including but not limited to those needed for the challenge problem evaluations. The representations will combine logical, probabilistic, plausible, temporal, and causal patterns in a coherent manner, and will build on our established work on temporal and logical dependencies, probabilistic causal models, and trend templates. The procedures, organized into a library, will go far beyond the CommonKADS library [1], which at present provides only a few, and extremely limited, procedures for these tasks.
- *We will develop web-mediated tools* that permit the human analyst to construct representations of new conditions, trends, patterns, and monitoring procedures in terms natural to the analyst, and that then automatically integrate these representations into the libraries of representations and procedures. These tools will employ open interface technologies, such as CORBA, and will be integrated with ontology and knowledge-base development tools, such as the ISI Ontosaurus system.

Using the proposed tools, human analysts will be able, in a matter of minutes, to extend the monitoring system with elements that seek out new classes of complex events and that communicate such events to the appropriate consumers (human or machine).

B Innovative claims

Monitoring, analysis, and interpretation (MAI) tasks play important roles in most (and dominant roles in many) applications, but effective tools for constructing knowledge-based systems (KBSs) exist only for the narrowest instances, such as process-monitoring and sensor-fusion problems, such as autopilots or intensive care unit (ICU) monitoring systems. Though effective, these tools presuppose carefully circumscribed tasks, for example, a fixed set of sensors and control signals and a fixed range of user commands or queries, and require laborious hand-crafting by close-knit groups of KBS developers. The demand for limited scope of knowledge renders these tools inapplicable to the broader, open-ended tasks arising in battlefield situation awareness, intelligence analysis, and even in realistic plan execution sentinels. The focused, labor-intensive construction methods render these tools too expensive to use on many problems to which they apply. The labor needed to construct a second system of a given type is almost as great as that needed to construct the first one, as current tools do not facilitate much sharing of knowledge across applications. The tools also make it difficult for a distributed collection of application experts to collaborate in constructing the knowledge incrementally, forcing construction efforts to rely on dedicated expert attention that remove the subject experts from active practice for substantial periods.

We propose to develop MAITA—the Monitoring, Analysis, and Interpretation Tool Arsenal—a collection of representations, method libraries, and distributed, collaborative tools to provide effective support to both narrow and broad MAI tasks. The key contributions are as follows:

Representations tailored to MAI knowledge: To be truly effective, the language and representations for monitoring, analysis, and interpretation must be developed in their own right, independently of specific knowledge representation systems, but with careful attention paid to identifying terms the human analysts find natural. The language and representations must facilitate description of logical and probabilistic relationships, plausible arguments, causal chains and networks, and temporal constraints. No current knowledge representation system integrates all these elements in a form appropriate to the task here, though the trend template approach developed at MIT by Haimowitz and Kohane [18, 17, 28, 22, 14] provides many important elements. We will develop a fundamental revision of that approach that incorporates all the necessary types of knowledge and provides a suitable basis for a recognition system of general applicability.

Libraries of tailored methods for MAI tasks: Rapid construction of new MAI applications requires a well-stocked library of proven methods for all but the most esoteric variants of the applications. Such libraries do not exist at present. Commercial tools provide programming languages instead of method libraries, and even the otherwise admirable CommonKADS library of expertise modelling [1] gives monitoring and interpretation tasks short shrift, focussing attention on diagnostic and evaluative situations

in which the analyst has a complete model of normal and abnormal behaviors (e.g., an electronic circuit diagram, or a characterization of types of fraudulent loan applications) and a complete aggregated descriptions of a behavior in question (e.g., sample measurements of the voltages and currents in the circuit, or a completed loan application). These models apply poorly, if at all, to broader analysis tasks, in which the first problem is simply identifying whether some isolated data or event should be ignored as insignificant or should be considered part of a larger pattern of possible significance.

Exploiting broad knowledge bases and method libraries: Tools to aid in construction of MAI knowledge and methods both permit natural expression of such knowledge and methods and provide corresponding access to existing libraries of ontologies, knowledge and methods. The tools should handle the mechanics of translating between the various forms of library knowledge and procedures and the representations natural to MAI tasks.

Exploiting distributed, collaborative expert and user contributions: Effective monitoring support will speed construction of monitoring systems by permitting humans expert at these tasks to provide many of the needed extensions themselves, without the aid of KBS system developers. In particular, users are well-positioned to identify areas in which existing KBs require extensions. Providing them with an entry point by which to add new specialized knowledge themselves (or to suggest new knowledge for formalization and integration by KB maintainers) serves to reduce the start-up system development costs and to simplify the evolution and maintenance of the system, as well as providing analysts with a tool to enhance their own effectiveness.

Tools to make these things convenient: Web-based technologies offer significant advantages in delivering services such as these to distributed user communities. They also form the delivery medium of KBS technologies under development, such as Stanford's Ontolingua and ISI's Ontosaurus systems. We will develop web-mediated tools that interoperate with and extend such existing systems.

Ensuring generality through sample applications: We will seek to ensure the generality, utility, and comprehensiveness of our representations, method libraries, and tools through integral development and testing on target applications both military and nonmilitary (e.g., battlefield situation assessment, intelligence analysis, logistic plan execution monitoring, epidemiology, medical diagnosis and therapy monitoring, general health maintenance monitoring).

C Technical rationale

Monitoring tasks constitute central elements of many important applications, both military (e.g., logistic plan execution, battlefield awareness, information warfare detection) and nonmilitary (e.g., job-shop plan execution, medical diagnosis, and epidemiology).

The difficulty of these tasks depends in part on the scope of the monitoring involved. Some process-control monitoring, for example, can be carefully circumscribed to require attending only to a fixed set of observable conditions, such as aircraft autopilots and intensive-care unit (ICU) alarm and patient-stabilization systems. At the other extreme, national security intelligence analysis and general medical diagnosis and health-maintenance monitoring must attend to a broad range of observations, information, and actionable conditions that cannot be enumerated beforehand, both because there are too many, and because not all are known (e.g., new terrorist groups may form, new technologies may provide new forms of attack, changing occupations may lead to new syndromes, and new diseases may arise). The first problem in the broader monitoring problems is simply determining whether there is in fact an interesting condition, not the (potentially simpler) problem of choosing how to address the condition. The scope of many important monitoring tasks, notably those involving plan execution, fall between these two extremes. The plans provide a focus and structure to the monitoring task, as they make obvious some natural conditions requiring attention (e.g., congestion at airports, achievement of mission tasks), but they do not fully circumscribe what needs looking after since indirectly related events may affect the plan operations (e.g., a sectarian riot may form suddenly and block traffic).

Monitoring tasks of any scope may benefit from a knowledge-based approach, but the benefits of this approach and the needs for further progress show up most clearly in the broader scope tasks. To see problems before they become imminent requires interpreting a broad stream of data to recognize many different patterns of facts which, in themselves, are only distantly related to the threat of interest through a tangled skein of causes and implications. Lists of indicators and warnings, currently used by analysts, include items both directly and indirectly relevant to the conditions of interest, but are very limited in the sorts of patterns they can detect, since the primary evaluation of such lists is simply to count the number of items which obtain. Improving the effectiveness of such techniques requires the ability to specify different ways the condition might arise, including both particular causes and general patterns of causes, and the ability to easily add new patterns of causes as they occur to the analyst. This ability in turn requires formalization of a broad range of knowledge about the world and means for specifying patterns of causation and interpretation.

Current knowledge-based system (KBS) technology does not support the broader monitoring and interpretation processes very well. It does not provide the broad bodies of formalized knowledge needed to interpret and correlate data automatically. It also does not provide tools to help human analysts cope with the open-ended nature of the intelligence task. New information regularly leads analysts to attend to new classes of

information and new potential threats, but current KBS technology depend on hand-crafting of representations and reasoning by system developers, and cannot be extended by users, even those highly knowledgeable in the domain.

We propose to develop methods for rapid construction and easy extension of automatic monitoring systems and the interpretational knowledge and procedures underlying them. These methods will give human analysts the ability to examine, refer to, tailor, and extend knowledge about causal patterns and procedures for recognizing important patterns in large-scale data streams. The bodies of interpretational knowledge constructed using these tools will then be usable to help human analysts understand ongoing events, first by filtering out irrelevancies in a more knowledgeable way than is currently possible with simple word-statistics methods, second by automatically highlighting patterns that are occurring in the data stream, and third by providing a backup to catch some human oversights.

We do not propose to develop systems for representing arbitrary knowledge, but expect to exploit existing and ongoing representational systems and libraries of formalized knowledge. Working with these as a base, we will construct systems that provide access to ontologies and libraries of formalized knowledge, but access in a form tailored to the monitoring, analysis, and interpretation (MAI) tasks. We will develop libraries of methods for describing causal and temporal structures useful in interpreting events, and systems and interfaces that provide convenient means for analysts to augment the knowledge of the monitoring system, and to change the behavior of the system in certain ways, especially through automating the construction of active sentinels to monitor specific conditions.

C.1 A look at the future

Colonel S., a staff intelligence officer during a major conflict, has his hands full analyzing the movements and intentions of the front-line forces. He remains concerned, however, about the possible upgrading of air defences around a key plateau and command post well to the rear of the fighting. He knows the enemy forces lack the new equipment, but believes groups in a neighboring "neutral" country may try to obtain the equipment and provide it to the enemy. Using his mouse to indicate the command post and relevant border regions with the neighboring country, he constructs a monitor which will automatically obtain and scan satellite photos to detect apparent meetings of vehicles from the neighbor with enemy forces in the border region and subsequent movement of those enemy forces in the direction of the command post. He indicates any possible activities along these lines be reported to him immediately for early targeting.

Jack R., an intelligence analyst focussing on international terrorism, reads a scientific report about using odor-free time-release buckyballs (TRBs) to encapsulate molecules. After ascertaining the types of equipment, materials,


and expertise needed to manufacture TRBs, he goes to his monitor library to examine his current methods monitoring terrorist attacks using chemicals. Using these methods as a base, he describes new patterns of attack involving manufacture of time-release buckyballs containing the toxin sarin, and sets up monitors that augment his current surveillance of Aum Shinrikyo remnants in Japan and the Russian Far East for contacts with TRB experts, travel to meetings reporting on TRBs, purchase of TRB production materials, sarin-consistent deaths and hospital admissions, etc. He indicates what data sources should stream through these various monitor elements and describes the points at which possibilities turn into probabilities. Knowing the breadth of possible activities fitting these criteria, he indicates that events matching some of these conditions be summarized each week, while events matching all be reported immediately.

Dr. K, an experienced knowledge-engineer, directs an effort at constructing an intelligence-monitoring system for a CERT-headed Internet security warning system aimed at early detection of new threats to the nation's information infrastructure. Rather than starting from scratch, she goes to the library of monitoring structures she recently developed for the CDC/DOD epidemiology tracking system, and has a basic system up and running in the space of a week, with a fairly substantial system in place within a couple of months.

C.2 System Architecture

We propose to construct a system, MAITA (Monitoring, Analysis, and Interpretation Tool Arsenal) built around several libraries of MAI information and a collection of editing, query, and reference tools. Library/tool combinations already exist in various systems, such as the ISI Ontosaurus and Stanford Ontolingua and Protege systems. Our system architecture will seek to build on these, and we expect to implement our system as an extension of one of these systems if this is compatible with the composition of the overall HPKB effort. For concreteness, we will describe the system in terms of an extension of ISI's Ontosaurus, but other implementations should be possible.

The core of a monitoring system consists of a set of data sources, which vary from application to application together with a set of monitors operating over these sources.




We augment this core with an environment consisting of editing, query, and information tools plus libraries of signal transducers, methods for correlating signals and recognizing

trends or patterns, models of monitoring control and interpretation processes, models of protocols for alerting people or processes of results, methods for display and presentation of results, and knowledge bases formalizing both general knowledge of the world and specific domain knowledge. The following describes each of these elements in turn.

Data Sources: The types of data sources of interest in MAI tasks varies with the task, and most such tasks will involve only a subset of the possible types of data sources. We anticipate providing for data sources including continuous signals, sampled continuous signals, text or message streams, propositional information, graded propositional information, i.e., with uncertainty (probability, evidence) or imprecision (fuzzy) measures, and databases, including relational or object databases, maps and geographical databases, and databases of images and sounds. The architecture permits chaining of monitoring processes, so that the results of one process can serve as a data source for other processes.

Monitors: The main operational element of the architecture is a set of monitors operating over the data sources.



Each of these monitors takes a subset of the data sources as inputs, processes these inputs with one or more signal transducers and correlators in accordance with a monitoring model, and reports the results in accordance with an alerting and display model. The set of monitors is dynamic, including, for example, sentinels posted by the execution systems of automated campaign and logistics planners, and new identifiers of indicators and warnings constructed by human analysts in response to a dynamic environment.

Signal transducers: Signal transducers transform a signal into one or more new signals. The most familiar variety of signal transducers all concern continuous or time-series signals. These include linear extrapolation and interpolation, trend line fitting, wavelet decomposition, fourier transforms, summary statistics, outlier detection, threshold detection, and others. The range of useful signal transducers appears to be more limited for propositional signals, including, for example, translation into new propositions, and measures of change statistics (e.g., when last changed, how frequently changing). Sometimes propositions encode data that may be usefully viewed as a sampled signal; for example, sequential reports on the location of a person or piece of equipment may usefully be aggregated into a map-based data series and analyzed with corresponding techniques. The most common signal transducers are provided in commercial products and libraries, and we plan to acquire or imitate these existing transducers wherever possible.

Signal correlators and trend templates: Signal correlators take several streams of data as inputs and provide one or more new signals (or propositions) as output—one can think of them as multi-signal transducers—and constitute one of the most important elements of the MAI system. Normally it is coordinated changes in different aspects of a situation that signal an event or trend of interest. A discrepancy between rates of use and procurement of a resource can signal a problem needing attention. Determining a new possible motivation for an agent may cast the activity reflected in other data streams in a new light.

The building blocks of signal correlators include standard continuous-signal operations such as differencing, modulating, and demodulating, but the most interesting building blocks for knowledge-based applications are those correlating propositional and graded information, such as rules, reasons and argument structures, bayesian probabilistic networks, causal networks, and temporal constraints. Existing monitoring systems show how to combine some of these. Determining how to combine these types of correlations will constitute one of the principal foci of this research.

We will use these signal-correlating building blocks to construct a library of abstract and special signal correlators called *trend templates*, after the representation by that name developed at MIT by Haimowitz and Kohane in the $TrenD_x$ system [18, 17, 28, 22, 14]. A trend template (TT) is an archetypal pattern of data variation in a related collection of data. Each TT has a temporal component and a value component. The temporal component includes landmark time points and intervals. Landmark points represent significant events in the lifetime of the monitored process. They may be uncertain in time, and so are represented with time ranges (min max) expressing the minimal and maximal times between them. Intervals represent periods of the process that are significant interpretation. Intervals consist of begin and end points whose times are declared either as: offsets of the form (min max) from a landmark point, or offsets of the form (min max) from another interval's begin or end point. The representation is supported by a temporal utility package (TUP) that propagates temporal bound inferences among related points and intervals [24, 23]. The value component characterizes constraints on individual data values and propositions and on computed trends in time-ordered data, and specifies constraints that must hold among different data streams.

In matching a trend template to data, two tasks are carried out simultaneously. First, the bounds on time intervals mentioned in the TT are refined so that the data best fits the TT. For example, a TT that looks for a linear rise in a numeric parameter followed by its holding steady while another parameter decays exponentially must find the (approximate) time boundary between these two conditions. Its best estimate will minimize deviations from the constraints. Second, an overall measure of the quality of fit is computed from the deviations. The most appropriate language of trends and constraints will vary from domain to domain, and we expect to build a rich set of capabilities to populate the ontology of trends. For the constraint language, we have so far explored mainly linear and quadratic regression models for numeric data, absolute and relative numerical constraints on functions of the data, and logical combinations of such descriptions and propositions.

We plan to develop the ability to build other TTs using descriptions that characterize any outputs of signal transducers and additional models of correlation among signals. The template library will be expanded over the life of the effort, with research and new applications leading to new additions. Moreover, augmenting the library with new templates will form one of the key operations in practical use of the system, allowing analysts to codify new indicators and warnings as they are identified.

The measures of quality that tells how well various TTs fit the monitored data become either time-varying signal or propositional outputs of the signal correlators and trend detectors, and provide the appropriately processed inputs for making monitoring decisions.

Monitoring control and interpretation models: Models of monitoring control and interpretation processes describe procedures for conducting monitoring activities. These range procedures may vary across several dimensions, including level of abstraction, types of information used, degree of passivity.

The most abstract control and interpretation procedures serve as a base for more specific ones, but will be rarely used directly. The real strength of the library of monitoring models will be in identifying specific combinations of representations, procedures, and domain characteristics that offer significant power compared to the more abstract procedures.

Variations in the types of information used provide much of the richness of the library of monitoring models. The TrenD_x [18, 17, 28, 22, 14] trend monitoring system developed at MIT uses a partial-match strategy operating over a trend templates consisting primarily of temporal constraints. More refined monitoring models would amend this procedure to take probabilistic or default information into account, or to embed background knowledge of the domain in the matching strategy (e.g., always try matching location first before bothering with other information).

Degree of passivity forms another important dimension of variation. At one extreme, simple lists of indicators and warnings may just monitor a set of propositional inputs to detect the presence or absence of a set of specific conditions, and the output of the procedure is to simply report the set of present conditions, or perhaps just the number of conditions present at a given time. At the other extreme, active monitors may start with such a list of indicating conditions and continuously actively seek out new information to determine the presence or absence of these conditions, as opposed to simply waiting for notifications of presence to enter as inputs. Intermediate monitoring procedures might simply filter inputs passively until some threshold is reached, and then switch to an active mode to confirm or deny remaining conditions.

Degree of passivity is closely tied to notions of the utility of information. Once an active search is underway, the best strategy is to seek first the information most useful to answering the question in the time allowed, but utility considerations also arise in formulating the thresholds at which monitors “go active”. For example, with only a few pieces of information, learning an additional item on an indicators list may not change

the quality of the match significantly. But at some point, learning an additional item makes each of the remaining items very significant, and “going active” at that point may well be the appropriate path. Because of this, the library of monitoring models represents in part models of the utility of information. This information is also used in the library of alerting models.

Another dimension of variation is whether the monitoring represents the activity of a single agent or is distributed across multiple agents. For many target applications, such as battlefield situation awareness and intelligence analysis, the distributed monitoring model arises naturally. The monitoring procedure library will therefore include procedures that distribute the effort in various ways, including fixed distributional arrangements, hierarchical tasking (as is done in military planning at different echelons), and economic models in which analyst processes distribute tasks through a market in intelligence information. We will draw here on both our ongoing research on market-guided planning and computation [8, 9], and on our web-mediated mechanisms for distributed medical record retrieval and analysis.

Combination and cascading of monitoring procedures leads to additional library elements, since one may sometimes combine synergistic but separate monitoring procedures into more effective ones. Some of these combinations on monitoring procedures mirror combinations of trend templates, thus reflecting portions of the trend template library, but the dataflow connections among monitoring procedures mean that the monitoring procedure library must be treated on its own rather than as a derivative of the trend template library.

Procedures for a small number of fairly abstract monitoring procedures have been codified already in the CommonKADS library of problem solving methods [1], but most of these concern fairly active procedures for diagnosing devices for which complete structural and functional information is available. Expanding on this basis to cover the broader tasks not addressed by the very restrictive CommonKADS assumptions will constitute one of the important contributions of this research.

As part of the formalization of the library of monitoring models, we will develop an ontology for monitoring processes, including concepts such as causal structures (chains constitute only the simplest such structures), partial matches, evaluations of significance and likelihood, and focus of attention. We will seek to build on the ongoing ARPI work on planning ontologies in formalizing this ontology. In particular, plans represent a specific type of causal and procedural structure, and plan monitoring (sentinel) tasks constitute a very important class of monitoring applications.

Alerting models: Alerting models describe criteria for deciding what to do with conditions detected by monitoring procedures; who to notify, when to notify them, and how to notify them. Alerting models are essential since analysts have priorities among the conditions of interest to them, and normally wish to hear about the most urgent and important items right away, with the lesser items deferred for consideration later. Most of the work of alerting models occurs in describing the utility of different results to dif-

ferent agents at different times. These utilities often can be grouped into classes, and the library of alerting models provides templates for specifying utility ascriptions specific to particular alert consumers (human or machine).

We will build the library of alerting models on both extant procedures for making alerting decisions and on methods for convenient specification of utility information. The medical informatics literature contains an unsystematic variety of alerting procedures, but few tied to explicit notions of utility (see, for example, [21, 20, 34, 36, 37]). One element of this research will be to use explicit utility models to develop a systematic collection of alerting procedures that includes the ones already reported in the literature. We will also build on our past work [48, 11, 49, 13] on qualitative representation of utility information, which has developed logical languages that can express generic preferences (“prefer air campaign plans that maintain a center of gravity over those that distribute forces more widely”), and that relate this notion of preference to the notion of problem-solving or planning goals (interpreting goals as conditions preferred to their opposites, other things being equal). We will develop utility models that combine both qualitative preference information with approximate numerical models of common utility structures (e.g., utility models that increase up to some time and then drop off to model deadline goals, as in [15]), along with automatic procedures for combining such information into qualitative decision procedures and numerical multiattribute utility functions suitable for quick evaluation of alternatives.

We expect utility models account for most of the variation in alerting models, though some variation can arise through the the sets of possible recipients and media used to communicate alerts. Selecting and tailoring utility models will be a key facility in making the monitoring system responsive to individual analysts, since the desired behavior will depend strongly on the utility of the particular conditions being monitored and on the context of other conditions and tasks faced by the analyst.

Display and presentation models: Human analysts require the results of analyses to be presented in intelligible forms, such as graphs, charts, etc. The library of display and presentation models describes standard forms of useful forms. As with the signal transducer library, the main forms of display and presentation methods are available in commercial and standard products. We will make use of these existing resources to provide this necessary capability, but do not view this as a research element of this effort (though it may well be a research area of continuing interest to HCI researchers).

Monitor construction tools: To permit easy augmentation and refinement of the set of monitors and bodies of monitoring knowledge and procedures, the system will provide a set of editing tools for creating, copying, removing, filling out, and revising trend templates, monitoring models, and alerting models, as well as a set of informational tools for querying existing ontologies, knowledge bases, and reference materials.

Constructing trend templates involving propositional conditions requires a system for representing these conditions. For practical use, this means having on tap one or more

formalized knowledge bases and ontologies to provide the vocabulary and background information needed to express the monitoring conditions. Development of the content of such knowledge bases will not be a focus of this research, except as needed to provide test cases and to handle the challenge problems. The need to interoperate with knowledge bases does, however, provide strong motivation for embedding the monitor-editing operations in an existing knowledge-base editing system, as this will avoid needless duplication of effort. We will seek to exploit existing tools for editing knowledge bases as much as possible, for example, extending the Ontosaurus web-based knowledge base editor to revising trend templates etc. The focus of our effort on editing capabilities will be on providing new operations and operations tailored to the specific representations of trend templates and monitoring and alerting models, in order to maximize the convenience and simplicity of augmenting these bodies of knowledge to the analyst.

Simply editing the libraries to contain new trend templates and monitoring and alerting models does not suffice to put new monitors into operation. Construction of monitors proper will be done using the same sort of editing tools to specify input data sources, signal transducers, recognition templates, and monitoring, alerting, and presentation models. With these elements specified, the system will “compile” the combination into an active process operating over the data sources and add this process to the set of active monitors. Part of our research will address this compilation process. Additional research opportunities are possible to optimize the operation of the set of active monitors (e.g., restructuring monitors with overlapping interests to avoid redundant analyses), but we will pursue these only if time permits.

C.3 Testing and Evaluation

The principal aim for the MAI toolkit is to facilitate rapid construction and augmentation of monitoring systems and subsystems across a wide range of monitoring tasks. To evaluate our progress, we will test the system throughout its development on at least three classes of applications. One class of tests will be the annual challenge problems developed by the HPKB effort. Though the exact nature of these challenge problems are to be determined, we know of numerous MAI tasks involving plan sentinels, battlefield awareness, and intelligence analysis directed toward crisis recognition. Other classes of tests will be conducted on a suite of sample problems already available in-house that were developed with older technology. These include the $Trend_x$ system for monitoring pediatric growth problems, several systems for monitoring intensive-care unit (ICU) alarms and stabilization systems, and general health monitoring in the Guardian Angel project [44]. Recoding of the $Trend_x$ library of trend templates and the ICU monitoring systems would provide rapid tests of the convenience, expressiveness, and power of the language and libraries developed in this research. The main form of evaluation will be performance on the challenge problem test suites.

D Comparison with other ongoing research

The bodies of ongoing work most relevant to the proposed research are work on editors for ontologies and knowledge bases, process model libraries, and monitoring systems. We discuss each of these in turn.

Editors for ontologies and knowledge bases clearly have much to offer our plans to construct editors for trend templates and monitoring and alerting models, and we intend to exploit and extend rather than reinvent such editors as much as possible. Three of the main current efforts of interest are the Stanford Ontolingua system, the SRI Graphical Knowledge Base Editor (GKB Editor), and the ISI Ontosaurus system. Each of these systems offers basic facilities for browsing and editing knowledge bases, and for translating knowledge among different representational forms. They mainly differ in the expressiveness of the underlying representational system, their web-accessibility, and their graphical editing capabilities. Ontolingua and Ontosaurus offer the expressiveness of first-order logic, while GKB Editor handles a generic frame-based representation of somewhat lesser expressivity. Ontolingua and Ontosaurus also operate via the web, while GKB Editor operates from a workstation. Only GKB Editor, however, offers true graphical editing capabilities, permitting the user to manipulate the structure of concepts directly from pictorial presentations of the conceptual hierarchy and interrelationships. These three systems have been evolving towards each other, and we are confident they would provide a good basis for extension for use in the MAITA system, obviating the need to implement yet another web-based editing system. These systems do not suffice without extension, however, since the level of convenience we wish to provide human analysts requires supplementing these generic knowledge-base editing tools with editing operations (including drag-and-drop operations for constructing monitoring processes) specialized to the key representational elements of MAI knowledge and processes.

The primary body of work on process models for MAI tasks consists of the problem solving methods present in the CommonKADS library [1], which are also used in the Stanford Protege project and the ISI Expect project. The CommonKADS library contains a good range of abstract procedures for a number of generic tasks, but its coverage of methods for MAI tasks is very limited.

The CommonKADS project puts forward methods for “assessment” and diagnosis as the main methods for MAI tasks. While the methods included for these tasks are all important, they simply do not cover many of the important classes of MAI tasks. The CommonKADS notion of “assessment” consists of taking a “case” and “system description” as inputs and giving a “decision” as outputs. For example, a loan-fraud detection task would involve taking a completed application for a loan and deciding whether the application was legitimate or fraudulent. The CommonKADS library provides an array of different methods for such tasks, but none of these fit general MAI tasks very well. For example, to shoehorn intelligence analysis tasks into this framework requires interpreting the “case” as the current sum of knowledge, so there is no good sense in which one gets different cases, only the same case at different times. The KADS abstractions

cover this, but the level of abstraction is much too high, and the CommonKADS library does not include specializations appropriate to the analysis task. Worse still, the “system description” is taken to be static, where in the analysis setting, what is considered to be abnormal or dangerous changes over time and with the new information coming in. That is, the case is the same as the system, and both change together. In addition, there is little or no structure to what is considered a case.

The CommonKADS library also includes more detailed procedures for diagnosis and prediction, but the procedures concern model-based diagnosis and other settings in which the monitor possesses complete or nearly complete information about the structure and intended behavior of the system being monitored. These assumptions are highly inappropriate for the more open-ended range of MAI tasks.

Essentially all of the important structure of MAI tasks lies outside of the extant CommonKADS assessment library, hence the attention of the proposed research to identifying and formalizing the needed extensions to this library.

Specific monitoring systems, as opposed to codifications of libraries of knowledge for constructing monitoring systems, are well represented in the literature and in commercial products. The most relevant work, other than our own, on monitoring knowledge and methods appears in the literature on trend detection and “temporal abstraction”, especially in the work of Shahar [43] and Das [2] at Stanford. These efforts focus on representing temporal relationships and on methods for identifying patterns of temporal relationships as instances of more abstract events. This work provides a good foundation for MAI activities, but intelligent monitoring and analysis involve more than just temporal information. Structuring relevant sorts of non-temporal information, especially information about logical implication, statistical correlations, and causation, is crucial, but lacking in most abstraction-based treatments. Statistical trend detection, on the other hand, does not adequately exploit the constraints and structuring information that templates provide. We plan to design representations for monitoring conditions that integrate the best representations devised for each of these separate types of knowledge.

The Guardian project [21, 20] at Stanford has developed a highly dynamic programming environment for the construction of very flexible monitoring systems. It puts very strong emphasis on giving the system the ability to reason, during the monitoring process, about the most appropriate data collection, interpretation and integration strategies. It places correspondingly less emphasis on the ease of constructing relatively simpler monitoring strategies beforehand, and has not developed detailed libraries of monitoring modules to support easy assembly. In our proposed work, we intend the background knowledge about monitoring and about the domain and monitoring task to be used more at the time a monitoring process is assembled and configured, not dynamically during its execution. We believe that this approach will lead to more efficient monitoring systems and greater ease of their development and configuration.

Commercial technology for monitoring and control offers good models of some of the capabilities we seek, but does not offer the flexibility, modularity, or construction tools of interest here. The G2 system offered by Gensym Corporation provides a very

good example. This system provides a good base of the “object-level” monitoring capabilities, namely the ability to accept inputs from several types of sources, a library of single-signal filters (linear extrapolation, fourier transforms, etc.), and a knowledge-based reasoning component for constructing multisignal analysis systems. While the library of single-signal filters and the primitives of the multisignal analysis language provide good starting points, they fail to cover some important types of knowledge (probabilities, causality). More importantly, G2 provides only a programming language, and not a structured library of procedures at various levels of specificity. Finally, G2 is structured as a heavyweight, stand-alone application, and does not provide the environment needed to support distributed efforts by multiple collaborating analysts. In G2, adding a new process to monitor some additional threat requires programming the new recognition procedure (without library support) and then recompiling and reinstalling the resulting overall monitoring process. For distributed, collaborative MAI efforts, what is needed instead is the ability to toss a new monitoring element into an ongoing process.

E Proposer’s previous accomplishments

Our research group has a long history of work in diagnostic and monitoring reasoning in the medical domain, and has made many contributions to complex probabilistic and heuristic reasoning, model-based reasoning, reasoning at multiple levels of abstraction, explanation generation, learning from experience, dealing with time-dependencies systematically, and modeling and using preferences in decision-making. We have also contributed to medical knowledge representation, knowledge representation in general, truth-maintenance, qualitative reasoning, and modeling repetitive decision-making.

Our first diagnostic program that attempted to model reasoning processes related to the problems faced in MAI tasks was the “Present Illness Program” (PIP), reported in *American J. Med.* in 1976 [35]. The present illness problem requires dealing with a stream of often unrelated conditions or reports and trying to determine if there is a problem, and if so, what the problem is. The PIP adopted a hypothetico-deductive framework for diagnostic reasoning, using strong cues from the patient presentation to trigger hypotheses, both logical criteria and a pseudo-probabilistic scoring scheme to confirm or eliminate hypotheses, and explicit differential links to revise hypotheses when discrepant information arose. Later versions introduced a simple model of time, categorizing both patient data and a hypothesis-oriented time line along the dimension: past, recent-past, now, near-future, future. Our interest in temporal reasoning has continued through the doctoral work of Kohane [24, 23], exploring temporal constraints in diagnostic reasoning and Temporal Utility Package (TUP); Russ [38, 39, 40, 41], who designed a control structure that supports reasoning about unreliable streams of time-oriented data and applied it to diabetic ketoacidosis; and Haimowitz [18, 17], who studied trend detection in pediatric growth data and in ICU monitoring in the $TrenD_x$ system [28, 22, 14].

We have substantial experience in implementing monitoring and analysis environments. In 1991, Dr. Kohane completed the implementation of an on-line medical chart (the Clinician’s Workstation–CWS) [30, 25, 29, 26] for the Division of Endocrinology at Children’s Hospital. This system has now been in full operation for 5 years and provides on-line access to clinic notes, clinic measurements, demographics, pharmacy data, laboratory results, problem lists and reports from ancillary departmental systems (e.g., radiology) to several clinical divisions at Children’s Hospital. Dr. Kohane also designed and led the implementation of a data integration and display system for the Multidisciplinary Intensive Care Unit, and more recently has led development of the W3-ICU web-based ICU monitoring system.

We have a long-term participation in knowledge representation efforts. Hawkinson and Szolovits worked in the mid-1970’s on the OWL [45] and BrandX [46] representation schemes that provided great flexibility and opportunities to exploit linguistic analogies but suffered from a lack of semantic rigor. When current more restrictive KR systems were built in the 1980’s, we tried to use KL/ONE to represent medical knowledge and found that too much expressive ability had been sacrificed for semantic cleanliness and computational efficiency [16, 19]. Doyle and Patil produced a major and influential

critique of this trend for the KR community [10].

Doyle's continuing work on truth maintenance and nonmonotonic reasoning [3, 33, 12, 5, 7, 6, 9] has been complemented in recent years by studies with Michael Wellman (now on the faculty at University of Michigan) of qualitative representations of preference information [48, 11, 49, 13], by studies of the use of economic mechanisms in controlling distributed reasoning and activities [4, 8, 9], and by work on constructing ontologies for plans and the process of planning. The ontology research has been conducted in conjunction with the ARPI Planning Ontology Construction Group.

We are engaged in a number of projects that exploit the revolutionary capabilities of the World Wide Web (W3) in innovative ways. Our W3-EMRS project [31, 27] re-engineers electronic medical record systems to use the distributed, multi-platform capabilities of the W3 to build more effective, more flexible, more secure and cheaper to implement record systems. In addition, this project is building virtual records that integrate health information from multiple institutions to reconstruct a patient's longitudinal health history from fragments stored at different hospitals, health centers, doctors' offices, etc. A related project [47] uses similar mechanisms to distribute real-time data via W3, to allow remote monitoring of patients in intensive care from any authorized remote site. Our Guardian Angel project [44] is developing personal health information systems that help patients at home manage significant aspects of their own health care, maintain records on their condition, treatments and responses, communicate with health care providers, and access educational resources that help them understand their conditions, all via the W3.

Long's Heart Disease program (HD) [Long92, Long92a, Long94], addresses the complex treatment of patients with heart failure, providing both a diagnostic and therapy planning component. Diagnosis is based on an approximate probabilistic method that works over a network of clinically-significant causal concepts, and therapy prediction is based on predicting the influence of possible interventions in a complex feedback system by using signal-flow analysis techniques. HF has proven to be quite effective at diagnosis in certain subdomains, and remains under active development to augment its diagnostic acumen and to further develop and test its therapeutic side. Current work includes the creation of W3-based interfaces that allow cases to be entered anywhere in the world and analytical results returned to the widespread community of users.

Former and current students have developed modeling and analysis methods for time-oriented data that are directly relevant to the proposed project. Dr. Tze-Yun Leong, now a professor at Singapore's National University, developed methods of modeling recurring decisions using semi-Markov decision processes [32]. Milos Hauskrecht is currently completing a doctoral thesis on the efficient analysis of partially-observable Markov decision processes. Yao Sun, MD, is pursuing his PhD studies and has implemented closed-loop controllers for ventilating infants based on fuzzy control algorithms. Alex Yeh [50] (now at MITRE Corp.) and Elisha Sacks [42] (now a professor at Purdue University) have both developed methods for the analysis of dynamic systems, especially those with repetitive behavior.

References

- [1] J. A. Breuker and W. Van de Velde, editors. *The CommonKADS Library for Expertise Modelling*. IOS Press, Amsterdam, 1994.
- [2] A. K. Das and M. A. Musen. A comparison of the temporal expressiveness of three database query methods. In *Nineteenth Annual Symposium on Computer Applications in Medical Care*, pages 331–337, New Orleans, LA, 1995.
- [3] J. Doyle. A truth maintenance system. *Artificial Intelligence*, 12(2):231–272, 1979.
- [4] J. Doyle. Rationality and its roles in reasoning. *Computational Intelligence*, 8(2):376–409, 1992.
- [5] J. Doyle. Reason maintenance and belief revision: Foundations vs. coherence theories. In P. Gärdenfors, editor, *Belief Revision*, pages 29–51. Cambridge University Press, Cambridge, 1992.
- [6] J. Doyle. Final report on rational distributed reason maintenance for planning and replanning of large-scale activities. Submitted to Rome Laboratory, available via <http://www.medg.lcs.mit.edu/doyle>, Oct. 1994.
- [7] J. Doyle. Reasoned assumptions and rational psychology. *Fundamenta Informaticae*, 20(1-3):35–73, 1994.
- [8] J. Doyle. A reasoning economy for planning and replanning. In M. H. Burstein, editor, *Proceedings of the 1994 ARPA/Rome Laboratory Knowledge-Based Planning and Scheduling Initiative Workshop*, pages 35–43, San Francisco, CA, 1994. Morgan Kaufmann.
- [9] J. Doyle. Toward rational planning and replanning: rational reason maintenance, reasoning economies, and qualitative preferences. In A. Tate, editor, *Advanced Planning Technology: Technological Achievements of the ARPA/Rome Laboratory Planning Initiative*, pages 130–135. AAAI Press, Menlo Park, California, 1996.
- [10] J. Doyle and R. S. Patil. Two theses of knowledge representation: Language restrictions, taxonomic classification, and the utility of representation services. *Artificial Intelligence*, 48(3):261–297, Apr. 1991.
- [11] J. Doyle, Y. Shoham, and M. P. Wellman. A logic of relative desire (preliminary report). In Z. W. Ras and M. Zemankova, editors, *Methodologies for Intelligent Systems, 6*, volume 542 of *Lecture Notes in Artificial Intelligence*, pages 16–31, Berlin, Oct. 1991. Springer-Verlag.

- [12] J. Doyle and M. P. Wellman. Rational distributed reason maintenance for planning and replanning of large-scale activities. In K. Sycara, editor, *Proceedings of the DARPA Workshop on Planning and Scheduling*, pages 28–36, San Mateo, CA, Nov. 1990. Morgan Kaufmann.
- [13] J. Doyle and M. P. Wellman. Representing preferences as *ceteris paribus* comparatives. In S. Hanks, S. Russell, and M. P. Wellman, editors, *Proceedings of the AAAI Spring Symposium on Decision-Theoretic Planning*, 1994.
- [14] J. Fackler, I. J. Haimowitz, and I. S. Kohane. Knowledge-based data display using trendx. In *AAAI Spring Symposium: Interpreting Clinical Data*, Palo Alto, 1994. AAAI Press.
- [15] P. Haddawy and S. Hanks. Representations for decision-theoretic planning: Utility functions for deadline goals. In B. Nebel, C. Rich, and W. Swartout, editors, *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning*, pages 71–82, San Mateo, CA, 1992. Morgan Kaufmann.
- [16] I. J. Haimowitz. Using NIKL in a large medical knowledge base. TM 348, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, Jan. 1988.
- [17] I. J. Haimowitz and I. S. Kohane. Automated trend detection with alternate temporal hypotheses. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pages 146–151, Chambery, France, 1993.
- [18] I. J. Haimowitz and I. S. Kohane. An epistemology for clinically significant trends. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 176–181, Washington, DC, 1993.
- [19] I. J. Haimowitz, R. S. Patil, and P. Szolovits. Representing medical knowledge in a terminological language is difficult. In *Symposium on Computer Applications in Medical Care*, pages 101–105, 1988.
- [20] B. Hayes-Roth, S. Uckun, J. E. Larsson, J. Drakopoulos, D. Gaba, J. Barr, and J. Chien. Guardian: An experimental system for intelligent ICU monitoring. In *Symposium on Computer Applications in Medical Care*, Washington, DC, 1994.
- [21] B. Hayes-Roth, R. Washington, D. Ash, R. Hewett, A. Collinot, A. Vina, and A. Seiver. Guardian: A prototype intelligent agent for intensive-care monitoring. *Journal of AI in Medicine*, 4:165–185, 1992.
- [22] I. Kohane and I. Haimowitz. Hypothesis-driven data abstraction. In *Symposium on Computer Applications in Medical Care*, Washington, DC, 1993.

- [23] I. S. Kohane. Temporal reasoning in medical expert systems. In R. Salamon, B. Blum, and M. Jørgensen, editors, *MEDINFO 86: Proceedings of the Fifth Conference on Medical Informatics*, pages 170–174, Washington, Oct. 1986. North-Holland.
- [24] I. S. Kohane. Temporal reasoning in medical expert systems. TR 389, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, Apr. 1987.
- [25] I. S. Kohane. Maintaining alternate interpretations of data from multiple sources in a clinical event monitoring system. In *MEDINFO 92: Proceedings of the Seventh Conference on Medical Informatics*, pages 483–489, Geneva, Switzerland, 1992.
- [26] I. S. Kohane. Getting the data in: Three-year experience with a pediatric electronic medical record system. In J. G. Ozbolts, editor, *Symposium on Computer Applications in Medical Care*, Washington, DC, 1994. To appear.
- [27] I. S. Kohane, P. Greenspun, J. C. Fackler, C. Cimino, and P. Szolovits. Building national electronic medical record systems via the world wide web. *Journal of the American Medical Informatics Association*, 3(3), 1996.
- [28] I. S. Kohane and I. J. Haimowitz. Encoding patterns of growth to automate detection and diagnosis of abnormal growth patterns. *Pediatric Research*, 33:119A, 1993.
- [29] I. S. Kohane, R. Lauerman, L. R. First, and J. A. Majzoub. Knowledge-based automated growth monitoring. *Pediatric Research*, 31:124A, 1992.
- [30] I. S. Kohane and D. P. McCallie. A dynamically reconfigurable clinician’s workstation with transparent access ot remote and local databases. In *AMIA*, 1990.
- [31] I. S. Kohane, F. J. van Wingerde, J. C. Fackler, C. Cimino, P. Kilbridge, S. Murphy, H. Chueh, D. Rind, C. Safran, O. Barnett, and P. Szolovits. Sharing electronic medical records across multiple heterogeneous and competing institutions. In J. J. Cimino, editor, *Proceedings 1996 AMIA Annual Fall Symposium*, pages 608–612, Philadelphia, PA, Oct. 1996. American Medical Informatics Association, Hanley and Belfus.
- [32] T.-Y. Leong. *An Integrated Approach to Dynamic Decision Making under Uncertainty*. PhD thesis, MIT, Cambridge, MA, May 1994. Laboratory for Computer Science Technical Report MIT/LCS/TR-631.
- [33] D. McDermott and J. Doyle. Non-monotonic logic—I. *Artificial Intelligence*, 13:41–72, 1980.
- [34] S. M. Ornstein, D. R. Garr, R. G. Jenkins, P. F. Rust, and A. Arnon. Computer-generated physician and patient reminders. *Journal of Family Practice*, 32:82–90, 1991.

- [35] S. G. Pauker, G. A. Gorry, J. P. Kassirer, and W. B. Schwartz. Towards the simulation of clinical cognition: Taking a present illness by computer. *American Journal of Medicine*, 60:981–996, 1976.
- [36] D. M. Rind, C. Safran, R. S. Phillips, W. V. Slack, D. R. Calkins, T. L. Delbanco, and H. L. Bleich. The effect of computer-based reminders on the management of hospitalized patients with worsening renal function. In P. Claytons, editor, *Proceedings Symposium Computer Applications in Medical Care*, pages 28–32, Washington, DC, 1991. McGraw-Hill.
- [37] D. M. Rind, C. Safran, R. S. Phillips, Q. Wang, D. R. Calkins, T. L. Delbanco, H. L. Bleich, and W. V. Slack. Effect of computer-based alerts on the treatment and outcomes of hospitalized patients. *Archives of Internal Medicine*, 154:1511–1517, 1994.
- [38] T. A. Russ. Temporal control structure reference manual. TM 331, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, June 1987.
- [39] T. A. Russ. Using hindsight in medical decision making. In *Symposium on Computer Applications in Medical Care*, 1989. Also published in *Computer Methods and Programs in Biomedicine*, 32(1): 81–90, May 1990.
- [40] T. A. Russ. Using hindsight in medical decision making. *Computer Methods and Programs in Biomedicine*, 32(1):81–90, May 1990. Also published in *Proceedings of the Symposium on Computer Applications in Medical Care*, pp. 38–44, 1989.
- [41] T. A. Russ. *Reasoning with Time Dependent Data*. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, Cambridge, MA, Sept. 1991. To appear as MIT/LCS/TR-545.
- [42] E. P. Sacks. Automatic analysis of one-parameter planar ordinary differential equations by intelligent numerical simulation. *Artificial Intelligence*, 48(1):27–56, 1991.
- [43] Y. Shahar. *A Knowledge-Based Method for Temporal Abstraction of Clinical Data*. PhD thesis, Stanford University, Stanford, CA, 1994. Available as Computer Science Department report CS-TR-94-1529.
- [44] P. Szolovits, J. Doyle, W. J. Long, I. Kohane, and S. G. Pauker. Guardian Angel: Patient-centered health information systems. Technical Report MIT/LCS/TR-604, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, May 1994.
- [45] P. Szolovits, L. Hawkinson, and W. A. Martin. An overview of the OWL language for knowledge representation. In *Proceedings of the Workshop on Natural Language*

Interaction with Databases, Schloss Laxenburg, Austria, Jan. 1977. Also appeared as MIT/LCS/TM-86(1977).

- [46] P. Szolovits and W. A. Martin. BRAND X: Lisp support for semantic networks. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 940–946, 1981.
- [47] K. Wang, I. Kohane, K. L. Bradshaw, and J. Fackler. A real time patient monitoring system on the world wide web. In J. J. Cimino, editor, *Proceedings 1996 AMIA Annual Fall Symposium*, pages 729–732, Philadelphia, PA, Oct. 1996. American Medical Informatics Association, Hanley and Belfus.
- [48] M. P. Wellman and J. Doyle. Preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, 1991.
- [49] M. P. Wellman and J. Doyle. Modular utility representation for decision-theoretic planning. In *Proceedings of the First International Conference on AI Planning Systems*, 1992.
- [50] A. S. Yeh. PLY: A system of plausibility inference with a probabilistic basis. Master’s thesis, Massachusetts Institute of Technology, Cambridge, MA, June 1983.