

# Building an Explanation Function for a Hypertension Decision-Support System

Ravi D. Shankar, MS<sup>1</sup>, Susana B. Martins, MD<sup>2</sup>, Samson W. Tu, MS<sup>1</sup>,  
Mary K. Goldstein, MD<sup>2,1</sup>, Mark A. Musen, MD, PhD<sup>1</sup>

<sup>1</sup>Stanford Medical Informatics, Stanford University School of Medicine, Stanford, California, USA

<sup>2</sup>VA Palo Alto Health Care System, Palo Alto, California, USA

## Abstract

*ATHENA DSS is a decision-support system that provides recommendations for managing hypertension in primary care. ATHENA DSS is built on a component-based architecture called EON. User acceptance of a system like this one depends partly on how well the system explains its reasoning and justifies its conclusions. We addressed this issue by adapting WOZ, a declarative explanation framework, to build an explanation function for ATHENA DSS. ATHENA DSS is built based on a component-based architecture called EON. The explanation function obtains its information by tapping into EON's components, as well as into other relevant sources such as the guideline document and medical literature. It uses an argument model to identify the pieces of information that constitute an explanation, and employs a set of visual clients to display that explanation. By incorporating varied information sources, by mirroring naturally occurring medical arguments and by utilizing graphic visualizations, ATHENA DSS's explanation function generates rich, evidence-based explanations.*

## Keywords:

Clinical Decision Support Systems; User-Computer Interface; Practice Guideline; Explanations

## Framing Coherent Explanations for a Decision-Support System

ATHENA DSS [1], a computer-based decision-support system for managing primary hypertension, was built as part of the ATHENA project at the VA Palo Alto Health Care System. ATHENA DSS is based on the EON architecture, which consists of a set of general components that work together to provide decision-support to clinicians for guideline-based medical therapy [2]. ATHENA DSS interprets the hypertension guideline model using specific patient data, and provides patient-specific recommendations for clinical management. It also generates coherent explanations that justify these recommendations. The need for explanations was

evident from the survey of 58 primary care health professionals participating in a guideline skills workshop. The survey was part of a separate guideline implementation project [3]. The participants included roughly equal numbers of physicians, nurse practitioners and registered nurses. As part of the survey, they were asked if they would want to see recommendations only or recommendations with explanations. A majority responded that they would like to see explanations for the recommendations. The participants were also given a multiple choice / multiple response question asking them about the type of explanation they preferred. The majority preferred seeing, at a minimum, the rules used to generate the recommendation. Nearly half wanted the relevant parts of the guideline document to be included in the explanation as well. One third preferred the rules, the guideline document and references to medical literature.

The survey gave us an understanding of what would be useful for clinicians using ATHENA DSS. We used it to develop a set of design criteria for the explanation function. The main driving points in developing the feature set were 1) the explanation should not only provide the rules that lead to the recommendation, but should also offer broader supporting material, such as pointers to the medical literature and the specific guideline document, and 2) the elements of an explanation should be organized and presented so that clinicians could navigate among them effectively. WOZ [4] is a multi-client framework that was developed to generate explanations for component-based systems. We adapted the WOZ framework to build the explanation function for ATHENA DSS.

## ATHENA DSS, a Hypertension Advisory System Based on the EON Architecture

EON is a knowledge-based system architecture that provides physicians with decision-support in guideline-

based care [2]. One of EON’s problem-solving components is the guideline-execution component. This component interprets an explicit model of clinical guideline using the patient data to generate guideline-based recommendations.

to display that evidence. With these ideas in mind, we used the following design criteria to develop it:

- The explanation function should identify the types of

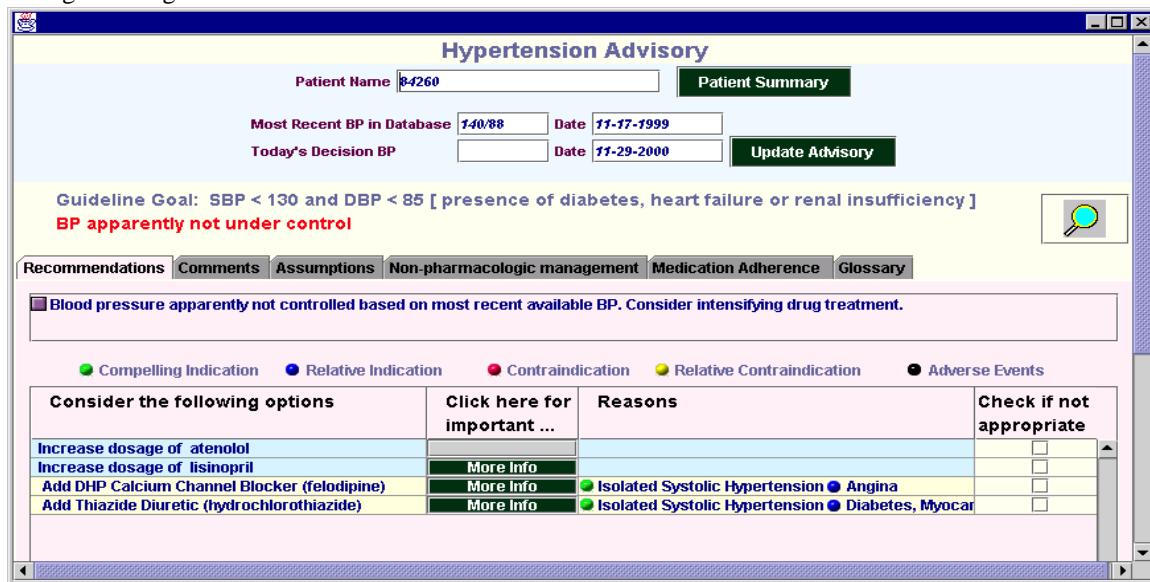


Figure 1: ATHENA DSS recommendations screen. The screen presents patient-specific recommendations. Clinicians can ask the system to explain recommendations such as “Add Thiazide Diuretic.”

Another component is the database-mediator component, which handles all the patient-data requirements of the problem-solving components. The domain-model component provides relevant medical terms and relations used by the other components. The guideline-knowledge-base component provides computer interpretable representations of clinical care guidelines that are used by the problem-solving components.

We used the EON components to build the ATHENA DSS application. Using the EON guideline model, we created a computer interpretable representation of major portions of the hypertension guideline described in the Sixth Report of the Joint National Committee on Prevention, Detection, Evaluation and Treatment of High Blood Pressure (NIH 1997), and of the VA hypertension guidelines. ATHENA DSS applies relevant patient data to the guideline model, determines if a patient is eligible to be treated under the guideline. It determines the patient’s target blood pressure, decides if his blood pressure is under control, identifies his risk group, and provides drug recommendations and general management messages. It displays the patient-specific treatment information at the point of care, and allows a clinician to interact with the system to modify patient data and fetch updated advisories, or to request explanation on any of the recommendations (Figure 1).

### Design Features of the Explanation Function

We wanted our explanation function to cover two broad issues: (1) the location of evidence underlying a recommendation, and its content; (2) the most effective way

sources that could provide information for generating explanations.

We wanted ATHENA DSS to generate an explanation that not only shows the rules used to arrive at the recommendation, but also other relevant information. For example, in explaining why the system recommended thiazide diuretic in Figure 1, we wanted to show relevant patient data, including medical evidence that supports using this drug, and links to the appropriate text in the guideline document.

- The explanation function should implement a mechanism for specifying what information should be part of an explanation.

The type of information that should be included in an explanation depends on the type of recommendation being explained. For example, when explaining a drug recommendation, the explanation function should show drug indications for a patient’s conditions and the evidence for the recommendation, and when explaining a patient’s risk group it should provide a description of the risk group and relevant patient data. Therefore, for each type of recommendation, it needs a template identifying the elements that constitute the explanation.

- The explanation function should integrate the different pieces of information coherently.

It should use a mechanism for specifying the relationships among the different elements of an

explanation. These relationships aid ATHENA DSS in generating a cohesive explanation out of varied pieces of data. For example, by using the

## The Explanation Function Using WOZ

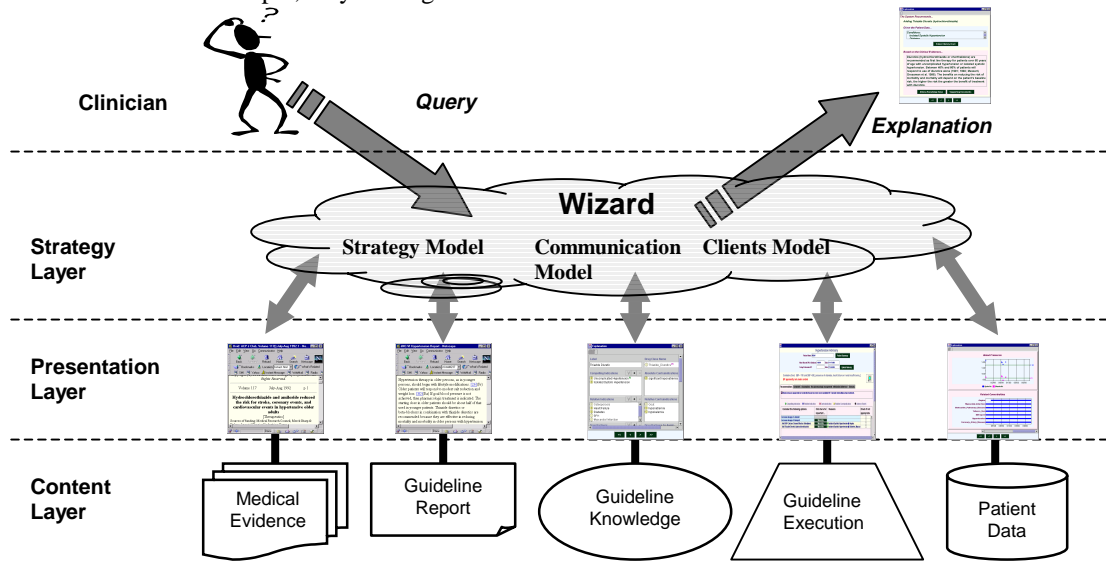


Figure 2: The WOZ Explanation Framework. When a clinician makes a query, the Wizard consults the Strategy Model to identify the information required to generate an explanation. It then looks up the Clients Model to find out which visual clients can provide the information. Using the visual clients, the Wizard generates an explanation.

relationships among disparate pieces such as patient data, the recommendation and the guideline knowledge, the system can generate a cohesive explanation: *Given that the patient has Isolated Systolic Hypertension (ISH), the system recommends adding thiazide diuretic or a calcium channel antagonist because ISH is a compelling indication for thiazide diuretic or calcium channel antagonist.*

- The explanation function should present the information at an appropriate level of abstraction.

Displaying all the pieces of the explanation along with the recommendation could create an information overload. We avoided this problem by clustering the explanation pieces into different levels of abstraction, and presenting them to users one level at a time. This approach allows users to start from a high-level explanation and drill-down to get details. For example, the explanation illustrated in the previous bullet text would be at the top level, and the evidence showing that ISH is a compelling indication for thiazide diuretic would be at the next level down.

- The explanation function should facilitate recursive explanations.

It should chain its explanations together so that some of the pieces of the explanation could, in turn, be explained. For example, a clinician might query why the patient has ISH in the previous examples.

We built the explanation function for ATHENA DSS based on the WOZ explanation framework [4]. WOZ is broadly made of three layers: **the content layer**, which constitutes the different sources of information that contribute to the explanation, **the presentation layer**, which consists of the set of visual clients that display the information in the content layer, and **the strategy layer**, which describes the information that should be presented to explain a specific recommendation (Figure 2).

### The Content Layer

The information that is used for generating an explanation in ATHENA DSS can be classified as **reasoning information**, which is used to compute the recommendation and **supporting information**, which provides evidence for the rules that are used in the computation. The reasoning information includes the patient data, the guideline rules, and the medical domain knowledge (such as drug information). The supporting information includes the guideline document, abstracted peer reviewed summaries of trials, the trial reports, and links to the guideline knowledge model. The reasoning information is provided by the underlying EON components. The guideline-execution component provides reasoning information to support its recommendations. The information includes references to relevant objects such as drug-indication objects in the guideline model, patient data used in the reasoning, appropriate static explanatory messages stored in the guideline model, and dynamically created messages that explain the evaluation of complex criteria. For the support information, we compiled medical literature that provides evidence for the drug recommendations. We also annotated parts of the JNC VI hypertension guideline report, which

states medical evidence. We then linked the rules in the guideline model knowledge base to the appropriate evidence literature and to the appropriate parts of the guideline report.

### The Presentation Layer

ATHENA DSS uses a set of visual clients that display the explanation information in the content layer graphically. The visual clients can receive messages that contain instructions on what to display. For example, an instruction to the patient-data visual client can be: *Display the blood pressures of Joe Smith from June 1999 to May 2000.* We built a communication model with a set of terms that will be used in encoding these messages. We built new visual clients such as the one that displays the main recommendations screen (Figure 1). We also created clients by wrapping existing applications with a message-handling layer. For example, Protégé is the knowledge-acquisition tool that we used to create and store the hypertension guideline model [5]. We built a visual client around Protégé to display specific guideline model objects. Another example is utilizing applications such as Netscape and Acrobat Reader to display medical literature. Each source of information in the content layer is associated with a visual client, and, and we built a clients model that describes the type of information each visual client can display. The communication model and the clients model are stored in the WOZ knowledge base.

### The Strategy Layer

The central component of the explanation function is the Strategy Layer, which consists of a module called the *Wizard*. The Wizard consults the strategy model to decide what information is required to justify a conclusion, and uses the clients model to identify which visual clients can display that information. The strategy model adapts Toulmin’s argument structure to identify the different elements of an explanation and organize them [6]. The elements might include the conclusion itself, the strength of the conclusion, the circumstances that contradict it, the data used in arriving at it, and the general body of evidence supporting it. The Wizard fashions an explanation for a conclusion by using argument structures as templates. Because specifying argument structures can be a daunting task for each individual conclusion, we grouped the recommendations into classes. For example, all drug recommendations are grouped into one class and all general comments are grouped into another class. Then, for each class we specified an argument structure that identifies the information required to explain all the claims in that class. The argument structure for explaining drug recommendations is shown in Figure 3. All the argument structures are encoded using the terms in the communication model and are stored in the WOZ knowledge base.

### The Generation of Explanations

When a clinician requests an explanation, the Wizard identifies the class to which the recommendation belongs. It then selects the argument structure associated with that class from the strategy model. Next, it uses the abstract

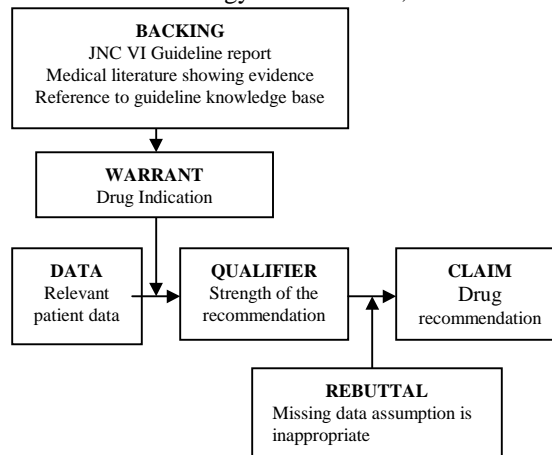


Figure 3: An argument structure for drug recommendations. An explanation can be generated using the elements of the structure. Its presentation can be organized using the relationships among the elements

descriptions in the argument structure to get the specific pieces of information relevant to the computation of the recommendation. It compiles the information using the visual clients, and organizes it for presentation. A clinician can navigate through this organization to explore the different levels of the explanation (Figure 4). A detailed description of how the WOZ explanation framework generates an explanation can be found in Shankar and Musen [4].

### Discussion

The explanation function of ATHENA DSS is unique in terms of displaying the evidence underlying the recommendation generated [7]. We believe that explicitly displaying evidence empowers clinicians to make decisions based on it, rather than simply on recommendations generated from a practice guideline. The system gives clinicians access to updated evidence, which educates them, and it also offers information that they can use to discuss recommendations with their patients. As a result, the risks of bad outcomes are reduced and adherence to treatment is enhanced.

An important role of the explanation function is to help developers of the decision-support system to debug the system. Our explanation framework helps to integrate the development of the guideline model with the guideline interpretation and explanation. We used Protégé to develop our hypertension guideline model and the medical domain model. Protégé allows embedding applications such as ATHENA DSS. After developing the domain model and the guideline model, we can use ATHENA DSS within the

Protégé environment to compute recommendations for a test patient, and to get the explanation for a recommendation. Since the explanation function uses the guideline model as one of the information sources, the explanation will include a reference to the relevant guideline model object that describes the rule used to compute the recommendation. The developer can inspect the rule and the recommendation,

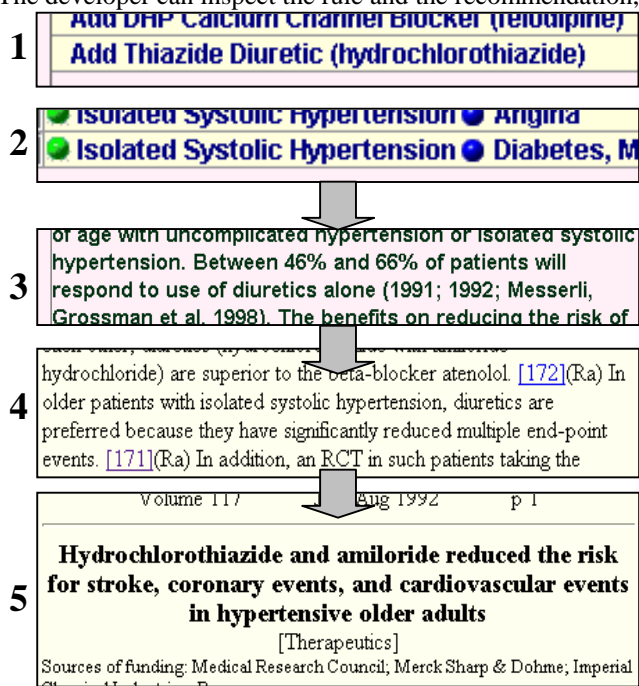


Figure 4: Drilling Down an Explanation Path. 1) A drug recommendation; 2) Reason for the recommendation; 3) A brief summary supporting the reason; 4) What JNC VI report says about the reason; 5) An abstract of one of the reference articles.

make changes to the rule if necessary, get updated recommendations, and iterate through the debugging process.

An evaluation of the explanation function in the actual clinic setting will provide a measure of the usefulness of the explanations in supporting clinicians' therapeutic decisions.

One way to measure the explanation function is through the program logs, which contain information about frequency of use for each patient and the proportion of patients per clinician. Time spent viewing the explanation function can be measured, as can the depth of searching. This information can provide correlations between accessing explanatory functions and therapeutic decisions.

Measures of usefulness of explanations could be captured by the clinician's response to a set of questions that have a range of semantically and logically distinct response options [8]. In order to avoid recall bias, these questions would appear in a box, just before closing the explanatory function.

Clinician satisfaction with the explanatory function in terms of ease of access, quality of information, utility of the information, content and visual display can be measured in a survey.

## Acknowledgments

This work has been supported, in part, by grant LM05708 from the National Library of Medicine, and VA grants HSR&D CPG-97-006, CPI 99-275, and RCD-96-301. The views expressed in this article are those of the authors and do not necessarily represent the views of the Department of Veterans Affairs. We thank Valerie Natele for her valuable editorial comments.

## References

- [1] Goldstein MK, Hoffman BB, Coleman RW, Musen MA, Tu SW, Advani A, Shankar R, and O'Connor M. Implementing clinical practice guidelines while taking account of changing evidence: ATHENA DSS, An easily modifiable decision-support system for managing hypertension in primary care. *Proc AMIA Symposium*, 2000(20 Suppl): 300-304.
- [2] Musen MA, Tu SW, Das AK, Shahar Y. EON: A component-based approach to automation of protocol-directed therapy, *Journal of the American Medical Informatics Association*, 1996: 3(6), 367-388.
- [3] Goldstein M, Coleman R, Advani A, Munroe P, Gholami P, Lavori P, and Hoffman BB. Implementing practice guidelines for hypertension: effect of computer-generated patient-specific recommendations for drug therapy. *Medical Decision Making*, 1999. 19: 529.
- [4] Shankar RD, Musen MA. Justification of Automated Decision-Making: Medical Explanations as Medical Arguments, *Proc AMIA Symposium*, 1999; supplement, 395-399.
- [5] Musen MA, Gennari JH, Eriksson H, Tu SW, Puerta AR. PROTEGE II: Computer support for development of intelligent systems from libraries of components, *MEDINFO 95, The Eighth World Congress on Medical Informatics*, Vancouver, B.C., Canada, 1995; 766-770.
- [6] Toulmin S. The uses of argument, *Cambridge University Press*, Cambridge MA, 1958.
- [7] Schiffmann RN, Liaw Y, Brandt CA, Corb GJ. Computer-based guideline implementation systems: A systematic review of functionality and effectiveness, *Journal of the American Medical Informatics Association*, 1999: 6(2), 104-114.
- [8] Friedman CP, Wyatt JC. Developing a measurement technique. In: *Evaluation Methods in Medical Informatics*. New York: Springer Verlag, 1997: 119-153.

## Address for correspondence

Ravi D. Shankar

MSOB Room X215, 251 Campus Dr.  
Stanford University School of Medicine

Stanford, CA 94305-5479, USA  
shankar@smi.stanford.edu