Case Report ■

Trade-offs in Producing Patient-specific Recommendations from a Computer-based Clinical Guideline: A Case Study

PERRY L. MILLER, MD, PhD, SANDRA J. FRAWLEY, PhD

Abstract This case study explored 1) how much online clinical data is required to obtain patient-specific recommendations from a computer-based clinical practice guideline, 2) whether the availability of increasing amounts of online clinical data might allow a higher specificity of those recommendations, and 3) whether that increased specificity is necessarily desirable. The "quick reference guide" version of the guideline for acute postoperative pain management in adults, developed by the Agency for Health Care Policy and Research, was analyzed. Patient-specific data items that might be used to tailor the computer's output for a particular case were grouped into rough categories depending on how likely they were to be available online and how readily they might be determined from online clinical data. The patient-specific recommendations were analyzed to determine to what degree the amount of text produced depended on the online availability of different categories of data. An examination of example recommendations, however, illustrated that high specificity may not always be desirable. The study provides a concrete illustration of how the richness of online clinical data can affect patient-specific recommendations, and describes a number of related design trade-offs in converting a clinical guideline into an interactive, computer-based form.

■ *J Am Med Informatics Assoc.* 1995;2:238-242.

There is currently a major national emphasis on the development of clinical practice guidelines to help ensure that patient care is provided in a consistent, high-quality, and cost-effective fashion.¹⁻³ The computer is a potentially powerful vehicle to bring clinical guidelines to the clinician at the time of patient care. There are several ways in which a clinical guideline might be placed into an interactive, computer-based form. One general approach is for the computer to ask the clinician a set of questions describing the case at hand, and then to selectively produce those portions of the guideline relevant to the case. Alternatively,

the computer could obtain the clinical data from an electronic medical record.

In anticipating this scenario, several questions arise regarding the richness of the clinical data required in the online medical record. We use the term richness to refer to the extensiveness of the data items available and the ability of those data items to capture the nuances needed to assess the relevance of a particular recommendation to a patient.

Some clinical data, such as a patient's medications and current laboratory test results, are likely to be included in an online medical record. Other clinical data, such as the severity of a patient's pain measured using a standard assessment tool, are less likely to be stored online. There are a variety of reasons why a particular piece of clinical data might be difficult to have available online. Some data, such as the adequacy of a patient's home support system, involve subjective judgment. Other data, such as the exact character and location of a heart sound or abdominal pain, are linguistically difficult to characterize in a standardized fashion.

Affiliation of the authors: Center for Medical Informatics, Yale University School of Medicine, New Haven, CT.

Supported in part by NIH Contract N01 LM13537 from the National Library of Medicine.

Correspondence and reprints: Perry L. Miller, MD, PhD, Center for Medical Informatics, Yale University School of Medicine, 333 Cedar Street, P. O. Box 208009, New Haven, CT 06520-8009. e-mail: perry_miller@yale.edu

Received for publication: 7/11/94; accepted for publication: 2/13/95.

As a result, practical questions arise as to how much online clinical data is required to obtain patient-specific recommendations from a guideline, and whether the presence of additional online clinical data might usefully allow a greater specificity of those recommendations. This case study explores these questions in the context of a specific clinical guideline, the "short form" of the guideline for acute postoperative pain management in adults, developed by the Agency for Health Care Policy and Research (AHCPR).

The AHCPR Adult Postoperative Pain Guideline

The full form of the AHCPR guideline on acute pain management is a 145-page booklet containing a comprehensive review of the subject.⁴ The short form ("Quick Reference Guide for Clinicians") is a 22-page pamphlet that highlights the most clinically important adult aspects of the full guideline.⁵ This short version contains roughly ten full pages of text, two page-size figures, six pages of tables, including recommended drug dosages, and a one-page bibliography.

A number of pages contain discussions of a general nature, which are not patient-specific. The patient-specific textual material consists of roughly two and a half pages of the ten pages of text, discussing primarily the use of nonsteroidal anti-inflammatory drugs (NSAIDs) and opioids. This patient-specific textual material is discussed below.

Patient-specific Clinical Data Items

The first step in our analysis was to list all the patient-specific clinical data items that *might* be used to tailor the guideline's recommendation to a particular case, and to group these data items into three categories:

1. Data items that are reasonably straightforward to have online in an electronic patient record
2. Data items that might be moderately difficult to have available online
3. Data items that could be very difficult to have online without asking a very specific question of the patient or the clinician

We have divided the data items into these classes because we are particularly interested in how the presence or absence online of data items in categories 2 and 3 might affect the specificity of patient-specific recommendations from the guideline.

Category 1: Reasonably Straightforward Data Items

Patient age
Drug allergies

NSAIDs and opioid drugs used (including dose, schedule, and route)

Concurrent monoamine oxidase inhibitor (MAOI) use
Thrombocytopenia

A patient's age, drug allergies, medications, and laboratory test results are among the most basic types of clinical data that one would typically expect to be available online.

Category 2: Data Items Moderately Difficult to Have Online

Patient is postoperative
Postoperative day
Impaired renal function
Coagulopathy
Patient is at risk for bleeding or gastric ulceration
History of opioid use before surgery
Respiratory depression

The data items in this category may in one way or another be difficult to ascertain with full confidence from the data items available in an online record. Many of these data items are actually inferences that need to be deduced from other data items in the medical record.

The presence of certain items in this list may be arguable. For example, a patient's postoperative status is often likely to be known because a major operation is likely to be recorded. On the other hand, all such procedures may not be reliably recorded and it may be difficult to distinguish with certainty which procedures involve postoperative pain. Impaired renal function may often be readily determined from online test results, but sometimes may involve clinical judgment. Similarly, exactly what is meant by respiratory depression may be difficult to quantify. The guideline itself arbitrarily defines respiratory depression as a respiratory rate below 10 breaths per minute, but this definition is clearly simplistic and does not even apply (as a reason not to give a narcotic) when the patient is on a ventilator.

Category 3: Data Items Very Difficult to Have Online (Without Asking a Specific Question of the Patient or the Clinician)

Pain severity
Pain is adequately controllable with medications
Patient can tolerate oral intake

The data items in this list involve subjective judgments and would usually require that a focused question be answered by the patient (e.g., pain severity) or the clinician. In fact, the severity of pain is best measured using a standard pain assessment tool, an example of which is included in the guideline. Al-

though some postoperative wards do record pain severity information online, this is uncommon and requires the use of a standard tool.

Example Trade-offs: Availability of Data versus Specificity of Recommendations

We now address two general topics: 1) the question of how necessary the data items in categories 2 and 3 above are to allow patient-specific recommendations, and 2) the effect on the specificity of the advice produced if they are not available. A good way to illustrate the various issues concretely is to examine specific recommendations in the pain guideline from this perspective.

In the examples that follow, the italicized statements are recommendations that could be produced if only data items in category 1 were available online. We then discuss how these recommendations might be made more highly specific if additional clinical data were online. We also discuss potential advantages and disadvantages of higher specificity.

Example 1

Pharmacologic management of mild to moderate postoperative pain should begin, unless there is a contraindication, with an NSAID. However, moderately severe pain should normally be treated initially with an opioid analgesic, with or without an NSAID.

This statement introduces the patient-specific advice, and helps place the statements that follow into perspective. Notice that the first sentence applies to mild and moderate pain, and the second sentence to severe pain. If the severity of a patient's pain is stored online, one could design the guideline so that only the appropriate one of these two sentences is produced. This example illustrates two interesting issues.

1. This is the only place in the guideline where the severity of the patient's pain might affect the recommendations produced. As a result, the only effect on the guideline of going to all the work to quantify a patient's pain and record it in the computer would be to save at most a single sentence of output.

2. It is not completely clear that it is desirable to produce only one of these sentences. If a user is interested in learning about pain management in the context of a particular case, seeing both recommendations might be useful. To introduce pain management, an interactive guideline designer might therefore decide to include both sentences independent

of the patient's pain severity. As a result, because these two comments are so central to pain management, it is a design judgment whether to make this recommendation more specific.

In other situations, of course, the availability of a clinical data item could allow the guideline to omit an obscure, esoteric fact that was irrelevant to the case at hand. Thus the clinical content of a recommendation is important in determining whether higher specificity is desirable in a particular recommendation.

Example 2

NSAIDs must be used with care in patients with thrombocytopenia or coagulopathies and in patients who are at risk for bleeding or gastric ulceration. However, acetaminophen does not affect platelet function, and some evidence exists that two salicylates (salsalate and choline magnesium trisalicylate) do not profoundly affect platelet aggregation.

If only category 1 data were available, this recommendation would always be produced when the clinician was proposing to use an NSAID. Notice, however, that the recommendation applies only "in patients with thrombocytopenia or coagulopathies and in patients who are at risk for bleeding or gastric ulceration." Here again, the comment contains fundamental information about NSAIDs, and a guideline designer might decide that the comment is worthwhile to include even if none of these conditions applies.

Furthermore, one would presumably want to be quite certain that none of these conditions applies before allowing the interactive guideline to omit this comment. Certain of these conditions may be difficult to infer from an online medical record, particularly in a patient who has just undergone a surgical procedure that might predispose to bleeding on that basis alone. Here again, the only effect of the availability of these several pieces of clinical data would be to allow the guideline to omit these two sentences (and only when none of the conditions applies).

Example 3

If the patient cannot tolerate oral medication, alternative routes such as rectal administration can be used. At present, one NSAID (ketorolac) is approved by the Food and Drug Administration for parenteral use.

This is an example of a recommendation that depends on one of the data items in category 3 (whether the patient can tolerate oral intake). If the computer could determine this clinical fact, it could omit this

comment in those cases where the patient could tolerate oral intake.

The only sure way to determine this fact with certainty would be to query the clinician. On the other hand, it might be easier for the clinician to just read (or quickly skim over) these two sentences than to read and respond to a question before getting the advice. This is a particular issue in this example because even the clinician might not be sure of the answer, in which case this recommendation would be relevant to include. This example illustrates that the decision of how specific to make an interactive guideline can require a fairly subtle clinical understanding of the domain.

Potential Effect on the Amount of Text Produced

The examples above illustrate how increased specificity in patient-specific recommendations can be achieved when additional clinical facts about the patient are available online. A next logical question is how this might affect the entire set of patient-specific recommendations produced by the interactive guideline. How much shorter would the recommendations be if all data were online, compared with the length of the recommendations if only category 1 data items were available online?

To help answer this question, Table 1 shows how much text might be omitted if data items in categories 2 and 3 were available online. As seen in this table, there are a total of 10½ sentences in the patient-specific recommendations that might not have to be produced (in certain clinical circumstances) if all the data items in categories 2 and 3 are available online. To put this information into perspective, one needs to know how much text is typically produced by the recommendations of the guideline produced using only category 1 data.

1. *When only NSAIDs are used*, the less-specific recommendations (based on only category 1 data items) typically contain nine sentences, of which seven *might* be omitted if more data items were online (but *might not* be omitted depending on the values of those data items). In this case, therefore, the recommendations might be reduced by roughly 78%.

2. *When only opioids are used*, the less-specific recommendations (based on only category 1 data items) typically contain 27 sentences, of which six *might* be omitted if more data items were online. In this case, therefore, the recommendations might be reduced by roughly 22%.

Table 1 ■

Numbers of Sentences of the Patient-specific Recommendations That Would Not Have to Be Produced if Data Items in Categories 2 and 3* Were Available Online and Had Values That Did Not Require Comment

Data Item	Extra Output Necessary If Data Item Is Not Available Online
Category 2	
Impaired renal function	½ sentence if meperidine used
Thrombocytopenia or coagulopathy or	2 sentences if NSAID used
Patient is at risk for bleeding or gastric ulceration	
History of opioid use before surgery	1 sentence if opioid used
Respiratory depression	1 sentence if opioid used
Category 3	
Pain severity	1 sentence
Pain is adequately controllable with medications	1 sentence if opioid used
Patient can tolerate oral intake	2 sentences if route is not oral 2 sentences if NSAID used

*Categories 2 and 3 are explained in the text.
NSAID = nonsteroidal anti-inflammatory drug.

Discussion

The previous sections clearly show that there are trade-offs between the richness of online clinical data and the degree of specificity possible in patient-specific recommendations produced from a computer-based guideline. In the context of one guideline, we have illustrated the problem with concrete examples and have provided a rough assessment of how the amount of overall advice produced can be affected. We now discuss a number of issues that arise.

1. *The clinician can recognize whether a recommendation applies to a patient.* Implicit in the discussion above is the assumption that the clinician can distinguish whether a particular part of a recommendation applies to a patient. This is possible because the recommendations themselves describe the conditions in which they are relevant. Thus the goal in designing an interactive guideline is *not* to ensure absolutely that no material is produced that is not relevant to the patient, but rather to minimize the amount of such material and to let the clinician decide what is specifically relevant.

2. *High specificity may not always be necessary or desirable.* As illustrated in the examples, the mere fact that a data item may allow one to make a guideline more specific does not mean that this is automatically desirable. Particularly when a comment is fairly central

to the topic under discussion, it may be helpful to include it even though it may not apply directly to the case at hand. As a result, when converting a guideline to computer-based form, a fairly subtle design decision may be involved to determine the degree of specificity that is clinically most appropriate. Different levels of specificity may be appropriate for different clinical users. A medical student might benefit from less specific recommendations containing more material that helps put the topic into broader context, while an experienced clinician might appreciate more focused, highly specific recommendations. Also, a less specific guideline involves less complex computer logic and may be considerably easier to build, test, and maintain than a highly specific guideline. From this standpoint alone, a less-specific guideline may sometimes be desirable.

3. *A great deal of "work" may produce relatively little gain.* In certain cases, a considerable amount of work may be required to make certain data items available online, yet may have relatively little payoff in terms of reducing the amount of text output. In fact, certain data may be available to the computer only by asking specific questions of the clinician. It may be considerably more work to answer these questions than to read a small amount of extra text.

4. *Different clinical events have different frequencies.* Some are very frequent, some are very rare. The frequency of a particular clinical event may affect how important it is to have a data item online. If it is very likely that the data item will have a value that allows text to be suppressed, then it will be more useful to ensure that that data item is available online.

5. *Different clinical recommendations have different characteristics.* From a clinical standpoint, certain recommendations are relatively central to a problem as a whole, while some comments are relatively esoteric and may be distracting to the clinician if produced in cases where they are not relevant. In the latter situation, it might be particularly desirable to suppress a comment when it does not apply.

6. *Different amounts of text depend on different data items.* Another observation is that certain data items may affect the relevance of a considerable amount of text, while other data items may have relatively little effect on the total amount of text produced.

7. *A computer-based guideline could adapt its specificity automatically to the available data.* An interesting further observation is that a computer-based guideline could be designed so that it automatically adapted the spec-

ificity of its recommendations to the amount of data available online. This could be accomplished by writing a highly specific interactive guideline, and then adding conditions that also test the availability of data items online in determining whether a comment should be produced. When used with a patient information system with relatively little clinical data, the guideline's recommendations could be less specific. When used with a system with more robust clinical data, the same guideline could produce more highly specific recommendations.

8. *Some guidelines have few, if any, patient-specific differences.* Another consideration is that some clinical guidelines have relatively little patient-specific differences in the recommendations that are made, or none at all. For example, a guideline might recommend that a particular procedure (e.g., a sickle cell test) be performed for all infants. In such a case, if the guideline applies to a patient at all, then the entire guideline applies. For such a guideline, the considerations discussed here do not apply.

Summary

The purpose of this case study was to take a simple idea, the intuition that the richness of clinical data could affect the specificity of computer-generated recommendations, and make it concrete in the context of an example guideline. Also, we have outlined a general approach to assess in rough terms the degree to which this phenomenon applies to a particular guideline. Finally, we have discussed a number of the underlying issues that arise in the analysis. Other clinical guidelines are bound to raise further issues. This is a first step in helping to define the problem, and in exploring the various issues involved.

References ■

1. Field MJ, Lohr KN, eds. *Guidelines for Clinical Practice: From Development to Use*. Washington, DC: Institute of Medicine, National Academy Press, 1992.
2. Audet A-M, Greenfield S, Field M. Medical practice guidelines: current activities and future directions. *Ann Intern Med*. 1990;113:703-14.
3. Eddy DM. *A Manual for Assessing Health Practices and Designing Practice Policies*. Philadelphia: American College of Physicians, 1992.
4. Acute Pain Management Guideline Panel. *Acute Pain Management: Operative or Medical Procedures and Trauma*. AHCPR Publication no. 92-0032. Rockville, MD: Agency for Health Care Policy and Research, Public Health Service, U.S. Department of Health and Human Services, 1992.
5. Acute Pain Management Guideline Panel. *Acute Pain Management in Adults: Operative Procedures*. Quick Reference Guide for Clinicians. AHCPR Publication no. 92-0019. Rockville, MD: Agency for Health Care Policy and Research, Public Health Service, U.S. Department of Health and Human Services, 1992.