INTELLIGENT NATURAL-LANGUAGE UNDERSTANDING OF COMPUTERIZED PATIENT MEDICAL RECORDS: Proposal for the Multidisciplinary Pilot Project Program

January 28, 2000

1. We propose to adapt an existing knowledge-representation, reasoning, and acting (KRRA) system that has natural-language (NL) competence so that it can read, understand, and make recommendations about hospital discharge summaries (DS) focusing on patients diagnosed with dehydration. This is a pilot project designed both as a "proof of concept" and to explore the problems involved. Upon successful completion of the project, we will extend the system to handle more complex diagnoses.

2. One of the major goals in the application of computers to medicine (medical informatics) is to produce a computerized patient record (CPR). Ideally, the CPR would contain all of a patient's medical information in a secure, reliable fashion that would allow clinicians ready access to the data needed to make informed medical decisions. In addition, information in the CPR could be used as input to an array of support programs. For example, based on the information in the CPR, reminders to perform routine screening, warnings of drug interactions, or even suggested diagnoses could be provided to clinicians by a computer.

A critical aspect of the CPR is the encoding of the medical data. The advantage of the CPR over paper records lies in the ability to easily extract information from it, and incorporate the information into a deductive knowledge base (KB). Some medical data (especially strictly numerical data) lend themselves naturally to encoding in, and processing by, a computer. There is, however, a large body of medical information that is not easily encoded, namely the clinician's notes.

Clinical notes are usually free-text, often with little structure and irregular grammar. They contain many terms from a specialized vocabulary and grammar (e.g., the use of 'present' as a non-transitive verb, as in "The child presented to the hospital with the heart murmur"), often with multiple ways of specifying the same concept. There are often misspellings, errors in wording, and mistranscribed words in dictated notes (e.g., the name 'Kay Seal' for the chemical 'KCl' (potassium chloride)). This makes the process of encoding clinical notes difficult. Although the free-text data could simply be captured in an unprocessed form, this fails to provide any of the benefits that CPRs have over paper records and may even be less efficient. In such a capture, no attempt is made to understand the text or to allow anything but simple word searches on the information it contains. Another method of encoding is to control the vocabulary that clinicians can use. Instead of allowing free-text entry, a clinician is forced to use terms selected from a defined list. Although this makes the task of encoding data trivial, it severely limits the ability of the clinician to express the nuances of the medical findings (Baud et al. 1992b).

An ability to interpret the natural language of the clinical note is, although more difficult than either method described, the "preferred medium for human beings" (Rassinoux et al. 1994: S80) and the only method that promises to deliver the benefits from computerization of the medical record (Hripcsak et al. 1995: 681, 687). The computer must in some sense "understand" the note so that it can then answer questions about the medical data. This, in fact, will be our methodology: to get a KRRA system with NL competence (the ability to understand and generate NL) to read and "understand" a medical discharge summary in the way a human would, so that the system can then be queried, give answers, and make recommendations in the way a human would.

The major goal of previous work in medical NL processing has been "extracting clinical findings from [free-text] medical narratives" (Spackman & Hersh 1996: 155) and to automate the following tasks: (1) to facilitate data collection, perhaps by integrating information from different sources, about single or multiple patients, into a standardized format (Evans et al. 1996: 388) for storage in an electronic medical record "central repository" (Johnson & Friedman 1996: 537), or even a software agent (Tuttle et al. 1996: 150); such information could be used, e.g., to determine whether treatments are successful (Sager, Lyman, Bucknall, et al. 1994: 151; Spackman & Hersh 1996: 155). (2) To learn new vocabulary, perhaps in the service of updating online medical lexicons (Spackman & Hersh 1996: 155) and integrating lexicons from different sources for standardization (e.g., as in the Unified Medical Language System; National Library of Medicine 1997), or even just to be able to cope with unfamiliar words (Baud et al. 1992a: 119). (3) To determine patient needs on the basis of observations ("findings"), so as to be able to pose clinical queries (Nangle & Keane 1994: 208) and to monitor and assess patient care, e.g., for treatment of asthma (Sager, Lyman, Bucknall, et al. 1994; Sager, Lyman, Tick, et al. 1994; Ertle et al. 1996).

We distinguish between (medical) NL *processing* and (medical) NL *understanding*: The former uses whatever computational and statistical techniques will efficiently convert NL text into a form usable in a CPR. The latter is more concerned with cognitive issues: the ability of a software agent to *understand* the medical narrative. There is a spectrum leading up to and connecting these two extremes, measured by levels of access and representation. At the pure computer-processing end, data are stored in a computer record or database-management system and retrieved via queries, keyword searches, or statistical analysis. More sophistication can be added by inferring new information from stored information, as in a deductive database or—if commonsense, background (e.g., specialized medical knowledge), or world knowledge is added—in a KB. The use of computational-linguistic techniques to capture "meaning" (Nangle & Keane 1994: 209) adds even more value, with the modeling of a cognitive agent adding more human-like understanding.

3. We propose to accomplish the latter via Cassie, a computational, cognitive, software agent implemented in SNePS (Shapiro 1979, 1991; Shapiro & Rapaport 1987, 1992). SNePS, the Semantic Network Processing System, is a KRRA system with an English lexicon, morphological analyzer/synthesizer, and a generalized augmented-transition-network parser-generator for NL understanding and generation that, rather than building an intermediate parse tree, translates English input directly into a semantic interpretation represented as a propositional semantic network and can generate English sentences expressing its "knowledge" (Rapaport 1988, 1991; Shapiro 1982, 1989; Shapiro & Rapaport 1995). Nodes in a SNePS network represent concepts, linked by labeled arcs. All information, including propositions, is represented by nodes; propositions about propositions can be represented without limit. Arcs form the underlying syntactic structure of SNePS. Arc-paths can be defined for path-based inference, including property inheritance within generalization hierarchies. Nodes and represented concepts are in 1-1 correspondence. This uniqueness principle guarantees that nodes are shared whenever possible and that nodes represent intensional objects (e.g., concepts, propositions, properties, and objects of thought). In particular, this facilitates dealing with multiple specifications for a single concept. SNePS's inference package accepts rules for deductive and default reasoning, allowing the system to infer "probable" conclusions in the absence of contrary information. When combinations of asserted propositions lead to a contradiction, the SNeBR belief-revision package allows the system to remove from the inconsistent context one or more of those propositions (Martins & Shapiro 1988). Once a premise is no longer asserted, the conclusions that depended on it are no longer asserted in that context. **STU: WE NEED A SENTENCE OR 2 ABOUT THE ACTING COMPONENT, WITH A REFERENCE.**

The extant literature suggests that no one has yet attempted to use a robust KRRA system on medical data. Most researchers seem to be using straightforward NLP software to process the records (Baud et al. 1992a; Friedman et al. 1994; Nangle & Keane 1994; Sager, Lyman, Bucknall, et al. 1994; Schröder 1994). However, no one seems to be using more advanced AI techniques that would allow the computerized record to be used as a deductive KB. Such a KB could not only be queried in natural language by users, but could be expected to perform inferences on its "knowledge", thus being able to reason about the information in the medical record and serve as an "intelligent agent" in assisting physicians in their diagnoses. In addition, a KRRA system that has some background knowledge of medical information and jargon ought to be able to handle the frequent "switching" between "medical English" and standard English that we have found in the DSs, and overcome many of the parsing problems discussed in the literature (Evans et al. 1996, Johnson & Friedman 1996, Spackman & Hersh 1996).

We will begin by viewing the DSs as written in a "sublanguage" (a specialized technical language embedded in NL; Harris 1968, §5.9: 152–155; Kittredge 1982). The most suitable grammar for dealing with sublanguages is a "semantic grammar". A typical *syntactic* grammar parses a sentence into a phrase-structure tree (a sentence is analyzed as a noun phrase followed by a verb phrase; a noun phrase might be analyzed as an article + adjective string + noun combination; a verb phrase might be analyzed as a verb + object combination; etc.). A *semantic* grammar parses a sentence into "semantic" categories determined by the subject matter (e.g., the noun phrase 'epigastric pain' might be analyzed as a body-part + symptom combination, as in Sager, Lyman, Bucknall, et al. 1994: 143, rather than as a more general adjective + noun combination). Our semantic grammar, however, would directly output a SNePS network representing the "meaning" of the sentence.

Cassie would also be supplied with appropriate background information that would help her to process the input. Such information would be the kind of information that a typical informed reader of an DS (e.g., a physician) would normally bring to the task. Of course, we do not expect to be able to provide Cassie with the equivalent of a full medical education; only the information *necessary* for understanding the DS would be supplied (as in a related project in learning new vocabulary from context; Ehrlich & Rapaport 1997, Rapaport & Ehrlich 1999).

For our project, the background information would include knowledge about the dehydration protocol in an emergency-department setting for pediatric patients with viral gastroenteritis. This consists of knowledge about signs and symptoms of dehydration on initial history and physical exam, and algorithms for dealing with mild, moderate, and severe dehydration (see Appendix).

Once Cassie has read and "understood" the DS (i.e., incorporated it into her KB), she would critique the DS with respect to whether the dehydration protocol should be followed, or was followed appropriately. It is this ability to critique and make recommendations that distinguishes our project from others that merely aim to extract useful information from CPRs for manipulation by other systems. Cassie would be, in a sense, an "expert system" that would both read/"understand" and act on the CPR, much as a human health-care practitioner would.

We would evaluate the system by testing it against medical students or residents at Children's Hospital of Buffalo. **PETER IS CHECKING ON WHETHER WE NEED HUMAN-SUBJECTS APPROVAL.** The students or residents would be given the same DSs and asked to evaluate them in the same way that Cassie does. This comparison can cut both ways. Cassie could also critique DSs written by the medical students or residents, not only for whether they handled cases of dehydration "correctly", in particular, but also for whether they wrote a comprehensible DS, in general. Humans reading DSs often think to themselves "I don't understand this" or "Is this information in the right section of the DS?" or "Only one number is given for blood pressure; shouldn't there be two?". Cassie could pose the same questions to the students or residents. Our system could thus find a use in medical education as well as patient care.

4. We request funds to support *n* graduate students to assist in the development of the semantic grammar, encoding of the background knowledge, and testing of the system. Anonymous DSs will be provided by one of the PIs. Medical students or residents participating in the evaluation component will come from that PI's classes.

Appendix: Dehydration Protocol in an Emergency-Department Setting for Pediatric Patients with Viral Gastroenteritis

PETER: THIS IS BASED ON YOUR EMAIL REPLIES TO MY QUERIES; PLEASE CHECK IT OVER CAREFULLY; THANKS.

1. During initial history and physical exam (normally in the emergency room), determine degree (%) of dehydration, on basis of the following signs and symptoms; in general, the more s/s, the more dehydrated the patient is; some signs occur later, however, and indicate more severe dehydration.

   Decreased urine output
   Increased heart rate
   Decreased moisture in mouth
   Decreased tear production
   Sunken anterior fontanel
   Lethargy
   Prolonged capillary refill
   Tenting
   Sunken eyes

2. IF mild (5%) dehydration, THEN:
   (a) Trial oral rehydration
   (b) IF retains fluids, THEN may discharge with follow-up
   ELSE GOTO moderate protocol.

3. IF moderate (10%) dehydration, THEN
   (a) IV fluid push of 20cc/kg NS or LR over 1 hour; check electrolytes
   (b) Trial oral rehydration after IVFP complete.
   (c) IF retains fluids and electrolytes normal, THEN may discharge with follow-up
   ELSE GOTO algorithm for IV rehydration

4. IF severe (15%) dehydration, THEN
   (a) IV fluid push of 20cc/kg NS or LR over 1 hour
   (b) Check electrolytes
   (c) REPEAT up to 2 times if necessary:
      i. IV fluid push of 20cc/kg NS or LR over 1 hour
   (d) IF IVFP fails, THEN GOTO Algorithm for IV rehydration

5. Algorithm for IV rehydration:

(a) Admit to hospital.

(b) IF sodium is normal (130–160), THEN rehydrate with D5 1/2NS, add 20mEq/L KCl after first void. Run fluids to give deficit and maintenance in 24 hours.

(c) IF sodium is high (>160), THEN rehydrate with D5 1/4NS, add 40mEq/L KCl after first void. Run fluids to give deficit and maintenance in 48hours. Monitor sodium frequently if necessary.

References

1. Baud, Robert; Rassinoux, Anne-Marie; & Scherrer, J.-R. (1992a), "Natural Language Processing and Semantical Representation of Medical Texts", *Methods of Information in Medicine* 31: 117–125.

2. Baud, Robert; Rassinoux, Anne-Marie; & Scherrer, J.-R. (1992b), "Natural Language Processing and Medical Records", in K. C. Lun et al. (eds.), *Proceedings of MEDINFO92* (Amsterdam: North-Holland): 1362–1367; <http://mbi.dkfz-heidelberg.de/helios/doc/nlp/Baud92b.html>.

3. Ehrlich, Karen, & Rapaport, William J. (1997), "A Computational Theory of Vocabulary Expansion", *Proceedings of the 19th Annual Conference of the Cognitive Science Society (Stanford University)* (Mahwah, NJ: Lawrence Erlbaum Associates): 205–210.

4. Ertle, Alan R.; Campbell, Emily N.; & Hersh, William R. (1996), "Automated Application of Clinical Practice Guidelines for Asthma Management", in James J. Cimino (ed.), *Proceedings of the 1996 AMIA Annual Fall Symposium: Beyond the Superhighway: Exploiting the Internet with Medical Informatics, Journal of the American Medical Informatics Association, Symposium Supplement* (Philadelphia: Hanley & Belfus): 552–556.

5. Evans, David A.; Brownlow, Nicholas D.; Hersh, William R.; & Campbell, Emily M. (1996), "Automating Concept Identification in the Electronic Medical Record: An Experiment in Extracting Dosage Information", in James J. Cimino (ed.), *Proceedings of the 1996 AMIA Annual Fall Symposium: Beyond the Superhighway: Exploiting the Internet with Medical Informatics, Journal of the American Medical Informatics Association, Symposium Supplement* (Philadelphia: Hanley & Belfus): 388–392.

6. Friedman, Carol; Alderson, Philip O.; Austin, John H. M.; Cimino, James J.; & Johnson, Stephen B. (1994), "A General Natural-Language Text Processor for Clinical Radiology", *Journal of the American Medical Informatics Association* 1: 161–174.

7. Harris, Zellig (1968), *Mathematical Structures of Language* (New York: Interscience Publishers).

8. Hripcsak, George; Friedman, Carol; Alderson, Philip O.; DuMouchel, William; Johnson, Stephen B.; & Clayton, Paul D. (1995), "Unlocking Clinical Data from Narrative Reports: A Study of Natural Language Processing", *Annals of Internal Medicine* 122: 681–688.

9. Johnson, Stephen B., & Friedman, Carol (1996), "Integrating Data from Natural Language Processing into a Clinical Information System", in James J. Cimino (ed.), *Proceedings of the 1996 AMIA Annual Fall Symposium: Beyond the Superhighway: Exploiting the Internet with Medical Informatics, Journal of the American Medical Informatics Association, Symposium Supplement* (Philadelphia: Hanley & Belfus): 537–541.

10. Kittredge, Richard (1982), "Variation and Homogeneity of Sublanguages", in Richard Kittredge & John Lehrberger (eds.), *Sublanguage: Studies of Language in Restricted Semantic Domains* (Berlin: Walter de Gruyter): 107–137.

11. Martins, João P., & Shapiro, Stuart C. (1988), "A Model for Belief Revision", *Artificial Intelligence* 35: 25–79.

12. Nangle, Brenda, & Keane, Mark T. (1994), "Effective Retrieval in Hospital Information Systems: The Use of Context in Answering Queries to Patient Discharge Summaries", *Artificial Intelligence in Medicine* 6: 207–227.

13. National Library of Medicine (1997), *UMLS Knowledge Sources, 8th edition.*

14. Rapaport, William J. (1988), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding", in James H. Fetzer (ed.), *Aspects of Artificial Intelligence* (Dordrecht, The Netherlands: Kluwer Academic Publishers): 81–131; reprinted in Eric Dietrich (ed.), *Thinking Computers and Virtual Persons: Essays on the Intentionality of Machines* (San Diego: Academic Press, 1994): 225–273.

15. Rapaport, William J. (1991), "Predication, Fiction, and Artificial Intelligence", *Topoi* 10: 79–111.

16. Rapaport, William J., & Ehrlich, Karen (1999, forthcoming), "A Computational Theory of Vocabulary Acquisition", in Łucja Iwańska & Stuart C. Shapiro (eds.), *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language* (Menlo Park, CA/Cambridge, MA: AAAI Press/MIT Press); pre-printed as *Technical Report* 98-05 (Buffalo: SUNY Buffalo Department of Computer Science) and *Technical Report* 98-1 (Buffalo: SUNY Buffalo Center for Cognitive Science), <ftp://ftp.cse.buffalo.edu/pub/tech-reports/98-05.ps.Z> or <http://www.cse.buffalo.edu/∼rapaport/Papers/vocab.cogsci.tr.ps>

17. Rassinoux, Anne-Marie; Michel, Pierre-André; Juge, Christian; Baud, Robert; & Scherrer, Jean-Raoul (1994), "Natural Language Processing of Medical Texts within the HELIOS Environment", *Computer Methods and Programs in Biomedicine* 45, Suppl.: S79–S96.

18. Sager, Naomi; Lyman, Margaret; Bucknall, Christine; Nhan, Ngo; & Tick, Leo J. (1994), "Natural Language Processing and the Representation of Clinical Data", *Journal of the American Medical Informatics Association* 1: 142–160.

19. Sager, Naomi; Lyman, Margaret; Tick, Leo J.; Nhàn, Ngô Thanh; & Bucknall, Christine E. (1994), "Natural Language Processing of Asthma Discharge Summaries for the Monitoring of Patient Care", in *Proceedings of the 17th Annual Symposium on Computer Applications in Medical Care: Patient-Centered Computing (SCAMC 1993; Washington, DC)* (American Medical Informatics Association): 265–268.

20. Schröder, Martin (1994), "Knowledge Based Analysis of Radiology Reports Using Conceptual Graphs", in Heather D. Pfeiffer & Timothy E. Nagle (eds.), *Proceedings of the 7th Annual Workshop, Conceptual Structures: Theory and Implementation (1992; Las Cruces, NM)* (Berlin: Springer-Verlag).

21. Shapiro, Stuart C. (1979), "The SNePS Semantic Network Processing System," in Nicholas Findler (ed.), *Associative Networks: Representation and Use of Knowledge by Computers* (New York: Academic Press): 179–203.

22. Shapiro, Stuart C. (1982), “Generalized Augmented Transition Network Grammars for Generation from Semantic Networks”, *American Journal of Computational Linguistics* 8: 12–25.

23. Shapiro, Stuart C. (1989), “The CASSIE Projects: An Approach to Natural Language Competence”, *Proceedings of the 4th Portuguese Conference on Artificial Intelligence (Lisbon)* (Berlin: Springer-Verlag): 362–380.

24. Shapiro, Stuart C. (1991), “Case Studies of SNePS”, Special Issue on Implemented Knowledge Representation and Reasoning Systems, *SIGART Bulletin* 2(3) (June): 128–134.

25. Shapiro, Stuart C., & Rapaport, William J. (1987), “SNePS Considered as a Fully Intensional Propositional Semantic Network”, in Nick Cercone & Gordon McCalla (eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge* (New York: Springer-Verlag): 262–315; earlier version preprinted as *Technical Report 85-15* (Buffalo: SUNY Buffalo Department of Computer Science, 1985); shorter version appeared in *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI-86, Philadelphia)* (Los Altos, CA: Morgan Kaufmann): 278–283; revised version of the shorter version appears as “A Fully Intensional Propositional Semantic Network”, in Leslie Burkholder (ed.), *Philosophy and the Computer* (Boulder, CO: Westview Press, 1992): 75–91.

26. Shapiro, Stuart C., & Rapaport, William J. (1991), “Models and Minds: Knowledge Representation for Natural-Language Competence”, in Robert Cummins & John Pollock (eds.), *Philosophy and AI: Essays at the Interface* (Cambridge, MA: MIT Press): 215–259.

27. Shapiro, Stuart C., & Rapaport, William J. (1995), “An Introduction to a Computational Reader of Narrative”, in Judith F. Duchan, Gail A. Bruder, & Lynne E. Hewitt (eds.), *Deixis in Narrative: A Cognitive Science Perspective* (Hillsdale, NJ: Lawrence Erlbaum Associates): 79–105.

28. Spackman, Kent A., & Hersh, William R. (1996), “Recognizing Noun Phrases in Medical Discharge Summaries: An Evaluation of Two Natural Language Parsers”, in James J. Cimino (ed.), *Proceedings of the 1996 AMIA Annual Fall Symposium: Beyond the Superhighway: Exploiting the Internet with Medical Informatics, Journal of the American Medical Informatics Association, Symposium Supplement* (Philadelphia: Hanley & Belfus): 155–158.

29. Tuttle, Mark S.; Sherertz, David D.; Olson, Nels E.; Nelson, Stuart J.; Erlbaum, Mark S.; Keck, Kevin D.; Davis, Alex N.; Suarez-Munist, Octavio N.; Lipow, Stephanie S.; Cole, William G.; Fagan, Lawrence M.; Acuff, Richard D.; Crangle, Colleen E.; Musen, Mark A.; Tu, Samson W.; Wiederhold, Gio C. M.; & Carlson, Robert W. (1996), “Toward Reusable Software Components at the Point of Care”, in James J. Cimino (ed.), *Proceedings of the 1996 AMIA Annual Fall Symposium: Beyond the Superhighway: Exploiting the Internet with Medical Informatics, Journal of the American Medical Informatics Association, Symposium Supplement* (Philadelphia: Hanley & Belfus): 150–154.