

“What If There Was No Oxygen?”: Responding to Hypothetical Questions in an Intelligent Tutoring Agent

(Extended Abstract)

Neil Yorke-Smith
American University of Beirut
and University of Cambridge
nysmith@aub.edu.lb

Stijn Heymans
Artificial Intelligence Center
SRI International
stijn.heyman@sri.com

Vinay Chaudhri
Artificial Intelligence Center
SRI International
vinay.chaudhri@sri.com

ABSTRACT

Our aim is for intelligent tutoring agents to replace traditional and even online textbooks with personalized, adaptive, one-to-one instruction. We focus on science subjects, and describe an approach to answering hypothetical questions from the student, such as “Would cellular respiration continue in the absence of oxygen?”

Categories and Subject Descriptors

I.2.11 [ARTIFICIAL INTELLIGENCE]: Distributed Artificial Intelligence—*Intelligent agents*.

Keywords

tutoring systems; question answering; natural language

1. INTELLIGENT TUTORING BY ASKING AND ANSWERING QUESTIONS

We focus on assisting high-school students with their learning, homework preparation, and exploration of science subjects. The challenges in developing such an intelligent tutoring agent are many: from knowledge representation and information processing, to user modelling and adaptation, to natural language understanding, to pedagogy and adoption.

This extended abstract describes an approach to answering hypothetical questions from the student, such as “Would cellular respiration continue in the absence of oxygen?” To respond to such questions, our agent *AURA* must understand the question, formulate a set of answers by reasoning over its scientific knowledge, formulate and present a response for the individual student, and ready for follow-up.

Our work is part of Project Halo [6], a multi-institution effort to develop a ‘Digital Aristotle’—an artificial agent containing large volumes of scientific knowledge and capable of applying sophisticated problem-solving methods to answer novel questions. This long-term vision holds two applications: an intelligent tutor (instructing and assessing students), and a research assistant (helping students and even scientists with research projects). *AURA* is a knowledge-centric agent comprising of components for knowledge entry

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May 6–10, 2013, Saint Paul, Minnesota, USA. Copyright © 2013, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

and representation [3], a user interface rendered on a desktop or a tablet device (iPad), and a suite of reasoning capabilities that leverage the scientific knowledge base (KB).

Distinctives of our problem and the knowledge-centric approach, in contrast to efforts such as *Watson* [5], are that we focus on a single domain rather than general knowledge, we explicitly codify knowledge from a single authorized source in a formal representation, and we aim for a tutoring agent rather than simply question answering.

Among intelligent tutoring systems (ITS), Graesser’s work includes dialogue agents for scientific domains, notably biology [9]. *AutoTutor* and descendants feature e.g. conversational dialogue and animated avatars. The technology component in *AutoTutor* closest to our work is a question-answering module based on Latent Semantic Analysis.

Our focus domain is pre-university biology, as found in a standard American high school biology textbook [10]. A team of Subject Matter Experts (SMEs), supported by knowledge engineers, have entered the knowledge from this textbook into *AURA*’s KB. Earlier work showed that *AURA*’s frame-based knowledge representation and knowledge entry interface are expressive enough to accommodate physics and chemistry domains also, while being usable by SMEs [3].

2. HYPOTHETICAL QUESTIONS

Hypothetical questions are valuable to advance the user’s learning, because they concern what *could be*. Through either the student asking the agent these questions, or the agent asking the student, higher levels of the Bloom taxonomy [8] can be reached—compared to more straightforward definitional, structural, functional, or even relationship questions. Consider three more examples:

- In the absence of oxygen, can yeast cells obtain energy? If so, by which process?
- What happens if rubisco production is blocked?
- The rate of reaction of the electron transport chain that functions in oxidative phosphorylation can be reduced by removing what substance?

As well as stretching the student, hypothetical questions stretch all three aspects of *AURA*’s architecture: question interpretation, reasoning, and answer presentation. These questions have in common a supposed change from the normal state of affairs: i.e., a hypothetical variant of the KB is posited. We can generalize many, if not almost all, examples of hypothetical questions to the abstract pattern: *What are the (possible) effects of a change in X upon Y, in context Z?*

For question interpretation, *AURA* accepts English in a restricted form of natural language (NL). This dialect is a ‘sweet spot’ between formal logic (which *AURA* uses internally) and full NL, aiming to be both human-usable and machine-understandable [4]. The output of question interpretation is a precisely-specified logical query, containing ground concepts, relations, constraints and roles from the KB, and possibly containing variables (unknowns).

As we have for other forms of questions, we asked SMEs to author a suite of representative hypothetical questions from selected chapters of the textbook. Of these nearly 200 questions, some are relatively straightforward and can be answered from searching the fulltext of the book or searching the web, or from human-generated question-answering sites (e.g., Yahoo! Answers). Other questions are significantly more challenging, such as the earlier examples. (It does not suffice to only obtain an answer: it must agree with the content of the textbook.) We analyzed the question suite to identify common question forms and wrote generalized templates to direct question interpretation.

For reasoning, *AURA* has a suite of general-purpose and dedicated query engines. When *AURA* is running on limited processing devices, such as an iPad, it can communicate with a server for heavy computation.

The core of our technical approach is a default reasoning technique operating over process flowcharts derived from the KB. We take a question identified as hypothetical in nature, and reformulate it as a form of disjunctive query using a new *affects* relation and other relations. The *affects* relation is ontology-driven, using dynamic chaining over our graph-based concept representation [6]. We employ a finite materialization of the KB derived from skolem instances for each class in order to quickly obtain solutions, and invoke a general-purpose first-order logic solver when needed.

We explored several simulation approaches [2, 7] over the process flowcharts, and decided against them because of issues with model alignment, scaling, and expressivity.

For answer presentation, *AURA* is endowed with a sophisticated set of heuristics developed in conjunction with SMEs. Factors in the heuristics include relation and concept importances, the question form, and a biological ‘interest-iness’ measure. *AURA* selects the best answer or answers, and presents them contextually using natural language [1], diagrams, and summary tables, together with hyperlinks to sections of the textbook, suggestions for follow-on questions, and access to dynamically-generated glossary pages.

3. CONCLUSION AND DIRECTIONS

This extended abstract outlined how a tutoring agent can answer hypothetical questions about high school biology. *AURA* functions as an evolving research prototype; it has been used by some 100 students and teachers. A demonstration of *AURA* in operation, including comments from SMEs, may be viewed at www.youtube.com/watch?v=fTiW31MBtFA.

AURA does not yet attain our vision of an intelligent tutor: future work is to add the user model and other necessary ITS features. The capability reported here could be embedded in an existing ITS framework, such as AutoTutor.

Hypothetical questions are valuable in engaging student learning at a higher level, and in assisting with more sophisticated homework and research topics. Progress continues to improve answer quality, advised by SME feedback, and to address representational issues that arise, such as qualitative relations and state changes.

Q. *Would respiration continue in the absence of oxygen?*

A. No, oxygen is a raw material of cellular respiration.

Media

Figure 9.6 An overview of cellular respiration. [...]

Further reading

9.1 Catabolic pathways yield energy by oxidizing organic fuels
9.5 Fermentation and anaerobic respiration enable cells to produce ATP without the use of oxygen

Related questions

How is cellular respiration achieved?
What is the relationship between cellular respiration and oxygen?
In the absence of an oxygen, how does a cell obtain ATP?
What does cellular respiration result in?
What are the steps of cellular respiration?

Figure 1: Response to a hypothetical question about oxygen. Key terms, and related reading and questions, are hyperlinked. Figure © [10] omitted here.

There are several metrics for assessing *AURA*’s success. One is end-user satisfaction of the agent’s question-answering dialogue: its usability and usefulness to students. Another metric is the learning outcomes achieved by students with *AURA* versus without it. A third metric is the ability to answer questions from standardized US university entry (SAT), as judged by SMEs. Such formal evaluation of the work is ongoing on an annual basis.

4. ACKNOWLEDGMENTS

This work was supported by Vulcan, Inc. We thank the members of the *AURA* team, particularly Michael Wessel.

5. REFERENCES

- [1] E. Banik, C. Gardent, D. Scott, N. Dinesh, and F. Liang. KBGen: Text generation from knowledge bases as a new shared task. In *Proc. INLG’12*, 2012.
- [2] B. Bredeweg et al. DynaLearn: Architecture and approach for investigating conceptual system knowledge acquisition. In *Proc. ITS’10*, 2010.
- [3] V. K. Chaudhri et al. Enabling experts to build knowledge bases from science textbooks. In *Proc. K-CAP’07*, 2007.
- [4] P. E. Clark and P. Harrison. Exploiting paraphrases and deferred sense commitment to interpret questions more reliably. In *Proc. COLING’10*, 2010.
- [5] D. Ferrucci et al. Building Watson: An overview of the DeepQA Project. *AI Magazine*, 31(3), 2010.
- [6] D. Gunning et al. Project Halo update – progress toward Digital Aristotle. *AI Magazine*, 31(3), 2010.
- [7] D. Inclezan and M. Gelfond. Representing biological processes in modular action language ALM. In *Proc. AAAI SSS on Common Sense Reasoning*, 2011.
- [8] D. R. Krathwohl. A revision of Bloom’s taxonomy: An overview. *Theory into Practice*, 41(4), 2002.
- [9] A. Olney et al. Guru: A computer tutor that models expert human tutors. In *Proc. ITS’12*, 2012.
- [10] J. B. Reece, L. A. Urry, M. L. Cain, S. A. Wasserman, P. V. Minorsky, and R. B. Jackson. *Campbell Biology*. Benjamin Cummings, 9th ed edition, 2011.