

Description of the PatientGenesys Dialogue System

Leonardo Campillos Dhouha Bouamor Éric Bilinski Anne-Laure Ligozat

Pierre Zweigenbaum Sophie Rosset

LIMSI - CNRS

Campus universitaire Bâtiment 508

Rue John von Neumann

91405 ORSAY cedex

firstname.lastname@limsi.fr

Abstract

This paper describes the work-in-progress prototype of a dialog system that simulates a virtual patient (VP) consultation. We report some challenges and difficulties that are found during its development, especially in managing the interaction and the vocabulary from the medical domain.

1 Introduction

Virtual Patients (VPs hereafter) are used in health care education. PatientGenesys is an interdisciplinary project that aims at developing a computer tool to provide continuing education to medical doctors. Trainer doctors will create new tailored clinical cases for the medical students to train consultation skills, in a 3D environment. At present, three prototype clinical cases have been created (anesthesiology, cardiology, and pneumopathy). This paper presents an overview of previous VP systems in Section 2 and describes the architecture of our system in Section 3; then, in Section 4, we put forward some conclusions. The system only supports French, but all examples are in English for the sake of understandability.

2 Previous work and motivation

VPs have been applied to several medical education tasks, ranging from history taking communication skills (Deladisma et al., 2007), dealing with mentally ill patients (Hubal et al., 2003; Kenny and Parsons, 2011) or training Pharmacy students (Park and Summons, 2013). We refer to (Cook et al., 2010) and (Kenny and Parsons, 2011) for a recent overview of VP systems. Although most systems are available for English, there are VPs for other languages (López Salazar et al., 2012). The PatientGenesys system is one of the few for the French-speaking community. However, some of the challenges we have found

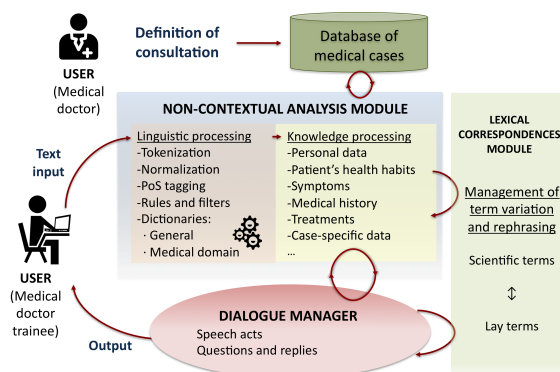


Figure 1: Overview of the PatientGenesys system

in designing a VP dialogue system may be raised regardless of users' language. The first difficulty concerns the lack of available corpora to train the system, which hinders using machine learning approaches. The second challenge concerns the variability of medical discourse. A concept may be referred to with different acronyms and jargon terms (e.g. *tonsillectomy* and *surgical removal of tonsils*) and lay terms from other registers (e.g. *tonsils operation*). The third one concerns the design of a core dialogue system that will be able to address dialogues of new clinical cases robustly. Future challenges will be raised when the system is to be adapted to other languages, mainly due to the ambiguities of medical terms in each language.

3 Architecture of the system

The system uses a user-initiative strategy (i.e. it will not ask questions). This is due to its pedagogical goals, which focus on training doctors in consultation skills. Input is text data, whereas output is spoken (text-to-speech, TTS). Four modules make up the system as shown in Figure 1: non-contextual analysis, lexical matching module, database of medical cases, and dialogue manager.

3.1 Database of medical cases

Knowledge on a medical case, provided by the instructor, defines patient state and knowledge. Frame-based structures organize the information in schemata. Cognitive frameworks already exist to model patient data and discourse (Patel et al., 1989). We use the YAML formalism (Ben-Kiki et al., 2005) to code information. General sections of patient data correspond to those proposed for VP data standards (Triola et al., 2007).

- Personal data: patient's name, family status, profession, height and weight.
- Lifestyle data: activities, diet habits, social behavior and addictions.
- Patient history data: family history, past diseases and treatments, allergies and surgeries.
- Symptoms data: type of symptom, anatomic place, onset time or duration, observations.
- Current treatments: International Nonproprietary Name, dose and method of administration, frequency and observations.

3.2 Non-contextual analysis module

Two main processes are involved in this stage: linguistic and knowledge processing. Linguistic processing consists of the following steps:

- Tokenization, normalization, downcasing, and Part-of-Speech (PoS) tagging with the French TreeTagger (Schmid, 1995).
- Spelling correction, to fix misspellings that may hinder text recognition.
- Linguistic annotation identifies verb tense, inflectional and derivative variants of terms referring to the same concept (e.g. *to operate* and *operation*), and other information, based on syntactic and semantic grammars written using wmatch, a word-based regular expression engine (Galibert, 2009). An example is the rule ANATOMY + *operation*, which tags the entity *tonsils operation* as a surgery.

Knowledge processing involves these steps:

- Entities are recognized using wmatch semantic rules and lists of medical terms. Vocabulary lists were drawn from the French component of the Unified Medical Language Sys-

tem (UMLS) (Bodenreider, 2004) and the VIDAL drug database.¹ Affixes are also applied: e.g. the suffix *-tomy* is used to detect surgical procedure entities (e.g. *appendectomy*). There are three broad types of named entities: general entities (e.g. date, frequency or age), domain-specific entities (e.g. drugs, symptoms), and discourse entities to classify speech acts (e.g. telling hello).

- Domain knowledge processing is used to enhance the understanding of input questions about patient illness. Medical knowledge comes from hierarchical relations extracted from the UMLS (e.g. *hypertension IS_A cardiovascular disease*).

3.3 Lexical matching module

The aims of this component are, first, to rephrase the technical descriptions found in the provided medical case into natural, patient-level language; and, second, to map the elements found in the question to those found in the medical case. This module relies upon different lists of medical vocabulary and concepts:

- Lists of medical term variants and UMLS concept unique identifiers (CUIs) are used to index each concept and map it to variants or acronyms. For example, C0020538 is the index for *HT* or *hypertension*. We also used the UMLF (Zweigenbaum et al., 2005).
- A non-UMLS list of medical terms, similar to the previous list, collects items that were not found in the UMLS.
- Lists of medical and lay terms map acronyms or technical terms (e.g. *ligamentoplasty*) to lay terms (e.g. *ligament repair*).

3.4 Dialogue manager module

The system uses a frame-based design in order to allow flexible interactions. The type of speech act and data contents of each turn are stored (e.g. *type: tell_past_disease; content: hypertension*). Information from the previous utterance is used to both repeat the previous turn and process anaphora and ellipsis. The domain model is based on each clinical case. Two types of anaphoric expressions are handled: co-reference and non-co-reference binding (respectively, *that* and *other* in Example 3.1).

¹<http://www.vidal.fr/>

Example 3.1.

- Which symptoms do you have?
- I have pain in my abdomen.
- Have you ever had that before?
- No.
- And do you have any other symptoms?
- No.

Ellipsis is related to short questions—usually by using *wh-* words—immediately after the system has given a piece of information (Example 3.2).

Example 3.2.

- I had a tonsils operation.
- When?
- I had a tonsils operation in my childhood.

The following types of speech acts are covered:

- Greetings: e.g. telling hello/bye and related speech acts (*How are you?*, *It is a pleasure*).
- General conversational management acts: e.g. showing agreement, lack of understanding, asking for repetition, or giving thanks.
- General questions: e.g. quantity or frequency.
- Clinical interview questions: these can be divided into general clinical questions and case-specific questions, which are specific to the actual clinical case.

4 Conclusion

We presented the on-going development of the PatientGenesys dialogue system, which aims at creating VP simulations to train medical students. The project is raising challenges regarding the lack of training corpora, the design of a robust dialogue system, and the variability of the medical jargon.

5 Acknowledgments

This work was funded by the FUI Project PatientGenesys (F1310002-P).

References

- Oren Ben-Kiki, Clark Evans, and Brian Ingerson. 2005. Yaml ain't markup language (yamlTM) version 1.1. *yaml.org, Tech. Rep.*
- O. Bodenreider. 2004. The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32(suppl 1):D267–D270.
- D. A. Cook, P. J. Erwin, and M. M. Triola. 2010. Computerized virtual patients in health professions education: a systematic review and meta-analysis. *Academic Medicine*, 85(10):1589–1602.
- A. M. Deladisma, M. Cohen, A. Stevens, P. Wagner, B. Lok, Th. Bernard, Ch. Oxendine, L. Schumacher, K. Johnsen, R. Dickerson, et al. 2007. Do medical students respond empathetically to a virtual patient? *The American Journal of Surgery*, 193(6):756–760.
- D. A. Evans, M. R. Block, E. R. Steinberg, and A. M Penrose. 1986. Frames and heuristics in doctor-patient discourse. *Social science & medicine*, 22(10):1027–1034.
- O. Galibert. 2009. *Approches et méthodologies pour la réponse automatique à des questions adaptées à un cadre interactif en domaine ouvert*. Ph.D. thesis, Université Paris Sud.
- R. C. Hubal, G. A. Frank, and C. I. Guinn. 2003. Lessons learned in modeling schizophrenic and depressed responsive virtual humans for training. In *Proceedings of the 8th International Conference on Intelligent User Interfaces, IUI '03*, pages 85–92, New York, NY, USA. ACM.
- P Kenny and T Parsons. 2011. Embodied conversational virtual patients. *Conversational Agents and Natural Language Interaction: Techniques and Effective Practices. Information Science Reference*, pages 254–281.
- V. López Salazar, E. M. Eisman Cabeza, J. L. Castro Peña, and J. M. Zurita López. 2012. A case based reasoning model for multilingual language generation in dialogues. *Expert Syst. Appl.*, 39(8):7330–7337, June.
- M. Park and P. Summons. 2013. A computer-generated digital patient for oral interview training in pharmacy. *Advanced Science and Technology Letters*, pages 28:126–131.
- V. L. Patel, D. A. Evans, and D. R. Kaufman. 1989. Cognitive framework for doctor-patient interaction. *Cognitive science in medicine: Biomedical modeling*, pages 253–308.
- H. Schmid. 1995. Treetagger— a language independent part-of-speech tagger. *Institut für Maschinelle Sprachverarbeitung, Universität Stuttgart*, 43:28.
- M. M. Triola, N. Champion, J. B. McGee, S. Albright, P. Greene, V. Smothers, and R. Ellaway. 2007. An XML standard for virtual patients: exchanging case-based simulations in medical education. In *AMIA Annu Symp Proc*, pages 741–745.
- P. Zweigenbaum, R. H. Baud, A. Burgun, F. Namer, É. Jarrousse, N. Grabar, P. Ruch, F. Le Duff, J.-F. Forget, M. Douyère, and S. Darmoni. 2005. A unified medical lexicon for French. *International Journal of Medical Informatics*, 74(2–4):119–124, March.