

# The Learning Intelligent Distribution Agent (LIDA) and Medical Agent X (MAX):

## Computational Intelligence for Medical Diagnosis

Steve Strain, Sean Kugele, and Stan Franklin

Cognitive Computing Research Group (CCRG)

University of Memphis

Memphis, Tennessee USA

Corresponding author: [sfstrain@memphis.edu](mailto:sfstrain@memphis.edu)

**Abstract**—The complexity of medical problem solving presents a formidable challenge to current theories of cognition. Building on earlier work, we claim that the systems-level cognitive model LIDA (for “Learning Intelligent Distribution Agent”) offers a number of specific advantages for modeling diagnostic thinking. The LIDA Model employs a consciousness mechanism in an iterative cognitive cycle of understanding, attention, and action, endowing it with the ability to integrate multiple sensory modalities into flexible, dynamic, multimodal representations according to strategies that support specific task demands. These representations enable diverse, asynchronous cognitive processes to be dynamically activated according to rapidly changing contexts, much like in biological cognition. The recent completion of the LIDA Framework, a software API supporting the domain-independent LIDA Model, allows the construction of domain-specific agents that test the Model and/or enhance traditional machine learning algorithms with human-style problem solving. Medical Agent X (MAX) is a medical diagnosis agent under development using the LIDA Model and Framework. We review LIDA’s approach to exploring cognition, assert its appropriateness for problem solving in complex domains such as diagnosis, and outline the design of an initial implementation for MAX.

**Keywords**—*cognitive modeling; LIDA; medical diagnosis; topic modeling*

### I. INTRODUCTION

Historical attempts to apply computational solutions to medical diagnosis have met with notable successes (e.g. [1-2]). In spite of this, information technology has failed to transform healthcare in the same way it has so many other industries [3]. The reasons for this are legion; in a word, healthcare has thus far defied all attempts to reduce its problem space—which spans physiological, pharmacological, technological, social, economic, administrative, political, legal, and ethical domains—into closed, computationally tractable subspaces.

A government report identified lack of cognitive support for physicians as a key area behind the US healthcare system’s underperformance [3]. Physicians must spend an inordinate amount of time sifting through raw medical data and engaging electronic documentation and transaction

processing systems, leaving little time for clinical reasoning. However, data filtering, documentation, and transaction entering for physicians have proven difficult to automate, since clinical problems rarely present in a stereotypical or well-defined manner, and require significant medical knowledge and human-like intelligence to solve. Automated cognitive support for such tasks seems to require a new method of knowledge synthesis—or perhaps, as we suggest, it merits deeper study of a very old one.

Biological cognition exhibits massive parallelism under modest power requirements, while learning to utilize novel stimulus patterns in dynamic environments. It reliably addresses spatial and temporal complexity in a variety of reference frames and in the midst of multiple modes of interference and signal degradation, elegantly tailoring perceptual and attentional filtering to distinct contextual demands. It facilitates flexible, adaptive modes of action shaped to varying motivational states and goals. It enables flight, attack, transport, and maintenance modes of locomotion, spatial navigation, fine motor coordination, flexible communication, and multiple forms of learning. Art, civilization, science, and medicine have emanated from it. Evolution has produced a cognitive organ that elegantly satisfies a vast number of environmental, physiological, and behavioral constraints. Rather than relying solely on approaches that convert naturally occurring information sources into data structures native to the von Neumann computing paradigm, it might be wise to avail ourselves of evolution’s successes in order to inspire the design of software and hardware systems that organize processes in a fundamentally different way.

The Learning Intelligent Distribution Agent<sup>1</sup> (LIDA) Model represents one such biologically inspired approach to a deeper understanding of cognition. LIDA integrates findings from a broad range of cognition-related fields into an architecture for designing and implementing software agents that utilize biological cognition, including human-like intelligence. The discussion of the LIDA approach found

---

<sup>1</sup> LIDA is a domain-independent architecture generalized from IDA (Intelligent Distribution Agent), an implemented software agent that uses human-like intelligence to perform the job of a Naval detailer. For more, see Footnote 4 in Section II.C.

below presents an argument in favor of LIDA’s suitability for modeling human-like diagnostic reasoning. A discussion of the design and plans for implementation of Medical Agent X (MAX), a LIDA-based medical diagnosis agent, follows. In particular, MAX’s modalities for “sensing” meaning in medical record text utilize a novel design, and preliminary experimental results in support of this design are presented. Finally, we outline directions for future work.

## II. THE LIDA APPROACH TO COGNITIVE MODELING

### A. The LIDA Model and Framework

The LIDA Model [4] integrates findings from experimental and theoretical branches of neuroscience, psychology, biology, and artificial intelligence into a domain-independent architecture for investigating cognitive processes. Inspired also by Sloman’s work on virtual machines and mind-brain supervenience [5], LIDA views mind as a space of processes, together with the principles that govern them, implemented on a physical system such as the brain.

The larger LIDA Model comprises two sub-models. The conceptual model offers a functional landscape for mapping out the causal and temporal organization of mental events including perception, memory, understanding, attention, consciousness, deliberation, action selection, and learning. The computational model serves as a virtual machine coupling the mind to its physical substrate, and provides a high-level specification of the structure and dynamics of cognitive content. Together these two seek to give a complete substrate-independent theoretical account of cognition, known as the LIDA Model.

LIDA facilitates the creation of simulations in restricted domains that test cognitive hypotheses along one of two arcs: 1) the science fork, which compares simulation behavior to data from experimental psychology or neuroscience; or 2) the engineering fork, for cognitive applications such as robotics, natural language communication, or complex problem solving. The LIDA Framework [6] consists of a Java API that supports the design of cognitive agents according to the Model.

The LIDA Model is consistent with biological cognition, but it does not require implementations to be faithful to biological models. On the contrary, LIDA formulates abstract cognitive principles independent of the specific physical substrate on which cognition may supervene [7]. Agents implemented along the science fork must adhere to a biological model in order to compare simulation behavior with experimental results. On the other hand, agents implemented along the engineering fork do not seek to investigate or explain fundamental cognitive principles, and thus may hybridize human-style reasoning with approaches from machine learning, robotics, natural language processing, computational neuroscience, symbolic AI, and so forth.

### B. LIDA’s Modules and Cognitive Cycle

LIDA’s conceptual model comprises a number of modules that serve as a functional taxonomy for mental processes (Fig. 1). The modules provide *descriptions* of functionally related cognitive processes, and imply neither anatomic modularity in the brain, nor a strong commitment to a particular taxonomy of mental functions [8]. Rather, our taxonomy summarizes a current understanding of cognition, and provides a structure for conceptualizing cognition with a view toward implementation. We present a brief summary of the modules here; for a more complete discussion, see [9].

The LIDA Model contends that human cognition can be viewed as continuing sequences of potentially overlapping<sup>2</sup> “cognitive cycles” within which *understanding*, *attending*, and *action selection/learning* occur. Higher-level cognitive processes, such as deliberation, planning, and problem solving, particularly in complex, dynamic environments, require multiple cognitive cycles.<sup>3</sup> Due to LIDA’s asynchronous, parallel nature, the cycle does not have a true beginning or end; however, it is useful to consider a single cognitive cycle beginning with sensory input and concluding with the selection and execution of an action. In principle, the LIDA Model allows for an unlimited number of sensors operating over multiple sensory modalities. During the understanding phase, sensors receive stimuli from the external and internal environments, which are processed and eventually integrated into the Current Situational Model (CSM). The CSM conceptually resides within the Workspace (see below), and contains the agent’s preconscious representation of its current situation.

Before integrating with the CSM, incoming stimuli are processed by feature detectors in Sensory Memory and Perceptual Associative Memory (PAM). Low-level feature detectors (for example, those that detect primitive shapes and colors) reside in Sensory Memory, and high-level feature detectors (for example, those that detect categorical associations, complex objects, and relations) reside in PAM. Features recognized by the low-level feature detectors are used by the high-level feature detectors to activate internal representations in PAM that correspond to those features. The representations in PAM that receive sufficient activation will be added to the Workspace as *percepts*.

The Workspace contains ongoing representations of post-perceptual, pre-conscious cognitive processes. In the Workspace, percepts may cue the retrieval of associated content from additional long-term memory modules, as well as from Transient Episodic Memory. For example, associated facts and beliefs, or autobiographical memories of events may be sent to the Workspace from Declarative Memory. Specialized processors called *structure building codelets* continually probe the Workspace, looking for the specific types of content that they are designed to process. Once the sought content is found, the structure building

---

<sup>2</sup> Cycles overlap when a subsequent one begins before a previous one has finished.

<sup>3</sup> LIDA’s cognitive cycles are weakly periodic in that they do not necessarily occur at constant frequency. Empirical evidence suggests a “quasiperiodicity” in the range of 4-10 Hz in humans [10].




Fig. 1. The LIDA Model diagram. Sensory stimuli enter from the environment on the upper left, and the cognitive cycle (see text) proceeds in a clockwise circle to action execution, bottom left. Not all elements of the diagram are discussed here. See [9] for details.

codelets elaborate on it, supplying additional associations that give the content meaning. In turn, these associations may cue additional content from PAM, Declarative Memory, Transient Episodic Memory, etc. The recent percepts, the cued associations, and the content constructed by the structure building codelets taken together form the agent's understanding of its current situation; i.e., the Current Situational Model.

The CSM may contain an abundance of information. *Attention codelets* are processors that specialize in determining the portions of the CSM that deserve the agent's attention. Similar to structure building codelets, attention codelets continually watch the CSM, selecting portions of the CSM relevant to the specific concerns of that codelet. The selected structures, referred to as *coalitions*, are sent to the Global Workspace module to compete for the agent's limited focus, consistent with Baars' Global Workspace Theory [11]. The winner of this competition will be the structure to which the agent attends during the current cognitive cycle, and will have its contents globally broadcast to all of the various LIDA modules. This global broadcast, also known as the conscious broadcast, initiates the learning mechanisms in the LIDA Model, and starts the action selection phase when it is received by the Procedural Memory module.

Procedural Memory contains data structures called *schemes*. A scheme comprises an action template, a context that suggests when the action should be taken, and the expected outcome. Schemes with contexts resembling the contents of the global broadcast instantiate; that is, they bind variables within the templates to the specific details of the current situation. The scheme instances then propagate to the Action Selection module, which selects exactly one high-level behavior for execution. Finally, the Sensory Motor Memory module executes the algorithmic steps that implement the selected action using the agent's actuators.

### C. Virtues of the LIDA Approach

LIDA comprises a broad theoretical account of cognition that supports the design and implementation of agent software in relatively narrow domains. Such implementations will attempt either to replicate psychological studies of mental tasks, or to apply cognitive approaches to technological ends, termed above as the science and engineering forks of LIDA, respectively.

The LIDA Model was inspired by the implementational success of the Intelligent Distribution Agent (IDA), an autonomous software agent that fully automates the role of a




Fig. 2. MAX's cognitive cycle.

human detailer for the US Navy [12, 13].<sup>4</sup> LIDA (for “Learning Intelligent Distribution Agent”) constitutes a domain-independent generalization of IDA’s cognitive cycle, augmented by the Conscious Learning Hypothesis of Baars’ Global Workspace Theory of consciousness [8, 9].

Due to its breadth, LIDA requires much additional work to empirically justify its theoretical claims (eg [8]). However, applications designed using the LIDA approach offer unique value to the study of cognition in that its dual modes of simulation enhance one another in a manner paralleling the reciprocal interplay of science and technology. LIDA’s modes or forks resemble Polanyi’s dual modes of control in a biological system, which he terms the “test tube” and “machine” types of control [14]. Agents built along the science fork can serve as a metaphorical test tube for investigating cognitive principles, whereas those built along the engineering fork—machine-type agents—harness these principles for another purpose. This endows the LIDA approach with an evolutionary flavor in which scientific and technological criteria select for soundness among a variety of cognitive hypotheses.

<sup>4</sup> A detailer is responsible for determining new assignments when current assignments end. The determination must take into account sailor preference as well as numerous naval policy requirements, and the needs of particular jobs. IDA communicates with sailors via emails in natural language, reads and understands relevant language in emails from sailors, and uses a cognitive cycle to consult appropriate human resource databases and deliberate on which options to offer [12,13].

It should be stressed that the LIDA approach does not require the current hypotheses of the Model to be correct; rather, they are tentative commitments forming the basis for the construction of science or engineering agent simulations. The Model evolves as new evidence accumulates, be it the evidence of LIDA simulations, or the evidence of empirical domains related to cognition. As an example of the former, the success of a test tube agent model of the attentional blink has greatly influenced LIDA’s attentional hypotheses [15]. An example of the latter is the effect on the LIDA Model produced by recent progress in the study of brain rhythms and of non-linear dynamics in cognition (see [7]), which has produced several Model commitments detailed elsewhere [4,7].

Regarding the domain of medical diagnosis, highly trained humans are the only agents currently capable of a cognitive performance adequate to the needs of real-time healthcare. The LIDA Model provides a path for incrementally constructing an artificial agent with the necessary cognitive skills. The performance of this agent, Medical Agent X (MAX), offers the promise of computational evidence to help refine the Model.

### III. MEDICAL AGENT X (MAX)

#### A. Medical Reasoning

Medical diagnosis comprises a complex, high level set of cognitive processes that includes data- and hypothesis-driven reasoning as well as perceptual, episodic, attentional, and procedural memories [16,17]. MAX's initial implementations will focus on hypothesis-driven reasoning, and, as an "engineering fork" implementation (see II.A), will integrate strategies that do not mimic biological models. Nonetheless, its central aim is to endow a computational agent with human-like intelligence for diagnostic thinking in the context of a clinical encounter with a patient.

Clinical hypothesis-driven reasoning is known as *differential diagnosis*. The clinician generates a list of possible diagnoses according to likelihood and medical urgency, and then proceeds to investigate along lines designed to eliminate (rule out) or confirm (rule in) potential causes. Differential diagnosis can also be used to refer to the list of clinical hypotheses under current consideration.<sup>5</sup> These hypotheses need not be mutually exclusive. A more in-depth description of differential diagnosis may be found in our earlier work [17].

#### B. MAX's Cognitive Cycle

MAX is equipped with sensors and actuators that allow it to read electronic medical records, and interact with medical professionals to requests tests, suggest diagnoses and treatments, and receive instructions and feedback in natural language. MAX has a "sensory modality" that specializes in processing electronic medical records, in particular, the History of Present Illness (HPI).<sup>6</sup> This modality uses a topic model,<sup>7</sup> trained on clinical data from the MIMIC-II clinical database, a large corpus of de-identified electronic medical records [20]. MAX's topic model will infer a pattern of diagnostic categories from an HPI and use this pattern as a starting point for MAX's diagnostic reasoning. Further details regarding clinical topic modeling can be found in Section IV.

MAX will then refine its diagnostic impression using standard lines of clinical investigation over multiple cognitive cycles. The following steps describe MAX's cognitive cycle (Fig. 2).

(1-3) A collection of topic model feature detectors infers likely diagnostic categories from a partial History of Present Illness (HPI). Concurrently, NLP feature detectors extract lexical and syntactic information. Both pass activation to corresponding representations in PAM. Representations with activations above a user-defined threshold propagate to the Workspace as percepts.

<sup>5</sup> When used without an article ("a" or "the"), differential diagnosis refers to the process; when used with an article or other specifier, it refers to a list of diagnostic possibilities under consideration.

<sup>6</sup> The HPI details a patient's progression of symptoms and is the starting point for medical reasoning. See [17] for details.

<sup>7</sup> A topic model utilizes the statistical structure of word occurrences within a corpus of documents to infer "what a document is about" [18,19].

(4-6) The percepts in the Workspace cue associations in Declarative Memory such as diagnostic hypotheses, expert knowledge, and previously encountered medical cases. Structure building codelets create/refine the differential diagnosis, a structure which forms the core of MAX's Current Situational Model. Attention codelets determine salient feature sets in the CSM, which they send to the Global Workspace as coalitions.

(7-9) Coalitions compete and the winning coalition is broadcast globally. Procedural Memory receives the broadcast, activating schemes most relevant to what MAX is currently attending. These are instantiated into the action selection module, creating a pool of relevant actions that can compete for selection. Example actions include initiating a particular line of investigation, asking a particular question, requesting the results of a diagnostic study or treatment, and suggesting the likely diagnosis. The selected action generates natural language for a human user or creates appropriate Workspace codelets and data structures to track expectations and plans.

(10) A Natural Language Processing (NLP) sensor analyzes the medical practitioner's response and MAX repeats the above process with the additional information. The topic modeling feature detectors continue to monitor the HPI as the medical practitioner adds information to the document.

It should be noted that all the modules continue to operate asynchronously as the cycle advances; for example, while an action is being selected, the perceptual processes continue to update the CSM, cues to Declarative Memory continue, and codelets continue to operate in the Workspace.

### IV. CLINICAL TOPIC MODELING FOR MAX

Much of the information required to perform diagnostic reasoning resides in free-text portions of the medical record [21]. Free-text medical records are generated by humans for humans, and exhibit narrative structure with clinical facts and causal clues--both explicit and implicit--embedded together in a variety of temporal contexts [21]. Although such data is extremely useful to human providers, the automated extraction of complete and accurate information from this data has long presented a challenge that remains to be surmounted.

Typical approaches to this problem involve data mining, named-entity recognition (NER), negation detection, sentence detection, question answering, and other standard NLP techniques [21, 22]. Topic modeling is often included among these standard techniques. However, we are investigating the use of topic models as a knowledge base for the "intuitive" recognition of semantic patterns in clinical text. Here we present the preliminary results of our investigations.

#### A. Topic Modeling

Topic modeling refers to a hierarchical Bayesian generative model for a corpus of text documents, originally

A.				B.			
Genetics	Evolution	Disease	Computers	Coronary disease	Pulmonary disease	Head trauma	Acute abdomen
human	evolution	disease	computer	artery	pulmonary	head	pain
genome	evolutionary	host	models	coronary	respiratory	ct	abdominal
dna	species	bacteria	information	left	lower	trauma	diarrhea
genetic	Organisms	diseases	data	catheterization	oxygen	hemorrhage	nausea
genes	life	resistance	computers	disease	treated	fracture	vomiting
sequence	origin	bacterial	system	cardiac	failure	scan	days
gene	biology	new	network	fraction	chest	intubated	fever
molecular	groups	strains	systems	anterior	pneumonia	injury	back
sequencing	phylogenetic	control	model	descending	ray	dilantin	began
map	living	infectious	parallel	bypass	distress	negative	bowel
information	diversity	malaria	methods	ejection	extremity	extremities	chills
genetics	group	parasite	networks	graft	lobe	arrival	change

Fig. 3. A. Sample topics from a 100-topic model of *Science* articles [19]. B. Sample topics from a 30-topic model trained on HPI text from MIMIC-II discharge summaries.

developed by David Blei and colleagues [18]; see [19] for an excellent review. We confine our discussion here to the simplest form of topic modeling, Latent Dirichlet Allocation (LDA) [18]. In brief, topic modeling begins with the idea of a topic distribution, which assigns a probability to each word in a vocabulary of size  $V$ . Therefore, a single topic distribution can be represented by a real-valued vector of length  $V$ , with the constraint that its components are non-negative and sum to one. Thus, a random sample from a topic will produce a single word. Words that correspond to what that topic is about will have the highest probability of being sampled. A document is modeled as a “bag of words” without regard to word order.

A central intuition of topic modeling is the idea that documents can exhibit multiple topics [19]. In a topic model with  $T$  topics, the topic proportions for a document can be expressed as a real-valued vector of length  $T$ , also with non-negative elements summing to one. Thus, when a set of topic proportions is randomly sampled, a single topic is returned. Finally, a topic model must have a “distribution over distributions” that returns a random set of topic proportions. In LDA, the Dirichlet distribution plays this role.

Thus the words are observed, and the topics and topic proportions represent hidden variables that reflect the underlying statistical structure of the corpus.

Approximate Bayesian inference can produce the parameters for an LDA topic model from an actual set of documents. In this way,  $D$  documents of length  $N^8$  (the corpus) are modeled as having been generated by the topic model in the following manner:

- 1) Initialize  $d = 1$ , the index into documents, and  $w = 1$ , the index into document words;
- 2) From the Dirichlet distribution, randomly generate a vector of topic proportions for this document;
- 3) For each word  $w$  in the document, sample the document’s topic proportions for a topic  $t$ , and from  $t$  randomly generate a single word;

- 4) Increment  $w$  to  $w+1$  and repeat 3) until document  $d$  contains  $N_d$  words;
- 5) Increment  $d$  to  $d+1$ , reset  $w = 1$ , and return to 2) to generate the next document, until the corpus contains  $D$  documents.

Given the number of topics  $T$  and an actual corpus, the inference algorithm returns the topics, topic proportions, and word assignments that maximize the probability of observing the actual corpus. For example, Blei and colleagues fit a 100-topic model to 17,000 articles from the journal *Science*. Among the inferred topics were those shown in Fig. 3A [19].

Since no information other than the text data itself is available to the algorithm, LDA can be thought of as an unsupervised machine learning algorithm. However, it must be emphasized that the topic modeling algorithm is not a classifier or clustering algorithm in the traditional sense: rather than being assigned to a single class or cluster, documents exhibit multiple topics in varying proportions. The labels in Fig. 3A (bold) were chosen by the modelers based on the topic’s most probable words. The algorithm has no semantic knowledge. Rather, it learns hidden statistical structure in the corpus that corresponds to semantic categories.

### B. Clinical Topic Modeling: Preliminary Results

Much human-generated clinical record text, especially narrative text such as the History of Present Illness (HPI), contains language specialized to facilitate diagnostic reasoning [17]. Therefore, we expected such text to display topical structure corresponding to diagnostic categories and/or other semantic categories relevant to clinical concerns. Our initial investigations support this intuition.

U.S. legislation that protects the privacy of medical records greatly complicates access to medical record data for analysis. The MIMIC-II clinical database [20] offers controlled access to a large repository of deidentified clinical data from an intensive care unit (ICU). Deidentification is an expensive and time consuming process that enables access to clinical data in a manner compliant with legislative privacy mandates.

We extracted a corpus of HPIs from 8,358 MIMIC-II discharge summaries, filtered the corpus vocabulary for stop

<sup>8</sup> If the document size is uniform,  $N$  will be a scalar; otherwise, it will be a  $D$ -dimensional vector.

words (common words with low predictive value), and created several topic models using jGibbLDA<sup>9</sup>. Since human-generated medical record text contains numerous irregularities such as misspellings, alternative spellings, inconsistent ordering of document sections, and varying section headings and separators, the HPI extraction was non-trivial. Thus, pre-processing of the MIMIC-II text files required extensive manual review of the raw files to develop regular expressions capable of detecting the desired narrative sections. Such challenges typify the case of free-text medical records; see [22] for a review. Nevertheless, the extracted corpus proved quite capable as a basis for inferring topic models using jGibbLDA.

The first author, a physician, reviewed the inferred topics of the various models and evaluated them subjectively for diagnostic relevance. All of the models showed a tendency to discover diagnostic categories; however, the 30-topic model performed best, with nineteen topics corresponding to diagnostic categories, and eleven classified as “noise.” A sample of topics from this model, shown in Fig. 3B, demonstrates the clinical relevance of the topical structure. Labels were chosen by the first author and confirmed by several physician colleagues.

The utility of these topics for the MAX project lies in the ability of a topic model to infer topic proportions for new documents. Equipped with the learned parameters of a good clinical topic model, MAX could get an “idea” of what a document is about by simply analyzing the document to obtain the topic proportions that best describe it under that topic model.

While these preliminary results show promise for the application of topic modeling to clinical record text, several challenges and limitations deserve mention. First, ICU records are not ideal for extracting a corpus representative of general medical knowledge, since they necessarily contain procedural data not normally present in narrative medical text, and represent the diagnoses of the critical care subspecialty rather than those of general medicine. Records from an internal medicine ward and clinic, or better still, a much larger corpus taken from a broad cross-section of medical specialties would provide more typical narrative text in a wider range of diagnostic contexts. However, due to privacy concerns and the high cost of deidentifying data, access to such a database remains problematic.

Second, no well established quantitative metric currently exists for evaluating a topic model as a clinical knowledge base<sup>10</sup>. Until such a metric exists, significant uncertainty will qualify any judgment as to the superiority of one model over another. However, since our purpose is to equip MAX with a useful “sensory” modality for extracting meaning from clinical text, the ultimate measure will be an improvement in MAX’s performance.

<sup>9</sup> jGibbLDA is an open-source implementation of LDA in the Java programming language, and is available at <http://jgibblda.sourceforge.net>.

<sup>10</sup> By “knowledge,” we mean the parameters of the topic model and the information they convey about the statistical structure in the corpus.

Third, our initial attempts at text pre-processing leave room for significant improvement, e.g. along the lines laid out by Meystre and Haug [22]. We expect a more thorough pre-processing would reduce the number of noisy topics as well as the number of noisy words (eg arrival, days, etc).

In spite of these challenges, we feel that topic modeling shows promise as a modality for quick recognition of topical content. MAX would use this modality to set a context and a focus for more lexically specific modalities, and could combine the different modes of information in its Workspace.

## V. FUTURE WORK

MAX will require a limited set of expert knowledge in order to reason over successive cognitive cycles, and thus provide a proof of concept. Such knowledge will allow it to identify symptom patterns as typical or atypical for particular diagnostic possibilities, and to instantiate learned lines of investigation for particular disease processes.

The LIDA Framework currently contains very simple implementations for structure building and attention codelets. MAX will require custom, domain-specific codelets, and their development may help specify the current claims of the LIDA Model regarding preconscious reasoning and attention.

Although our initial implementation of MAX will not include learning, the global broadcast provides a mechanism for this, in keeping with the LIDA Model. This would allow MAX to learn new clinical terms, identify new symptom patterns for disease recognition, and form new memories of clinical cases, procedures, and interventions. MAX’s abilities may be further enhanced by integrating additional sensors that utilize alternate sensory modalities.

Although much development will be needed, the LIDA Model provides an architecture for endowing MAX with the ability to reason causally, and to create narratives from clinical events in the context of its own diagnostic interpretations. In principle, the LIDA Model can accommodate cognition at any scale. It is hoped that this scalability will transfer to agent implementations.

## ACKNOWLEDGMENT

The authors gratefully acknowledge the assistance of Ryan McCall, Javier Snaider, Daqi Dong, Tamas Madl, Andrew Olney, David Blei, Xiangen Hu, Whitney Cade, Nopal Niraula, and Rajendra Banjade.

## REFERENCES

- [1] B. G. Buchanan and E. Shortliffe, eds., *Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading, Mass.: Addison-Wesley Pub. Co., 1984.
- [2] S. G. Pauker, G. A. Gorry, J. P. Kassirer, and W. B. Schwartz, “Towards the simulation of clinical cognition: Taking a present illness by computer,” *The American Journal of Medicine*, vol. 60, pp. 981–995, June 1976.
- [3] W. W. Stead and H. S. Lin, Eds., “Computational technology for effective healthcare: Immediate steps and strategic directions

- (Technical Report),” Committee on Engaging the Computer Science Research Community in Health Care Informatics, National Research Council, 2009.
- [4] S. Franklin, S. Strain, R. McCall, J. Snaider, and U. Fagih, “Global Workspace Theory, its LIDA model, and the underlying neuroscience,” *Biologically-Inspired Cognitive Architectures*, vol. 1, pp. 32-43, July 2012.
  - [5] A. Sloman, “What cognitive scientists need to know about virtual machines,” in *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, N. Taatgen and H. van Rijn, Eds. Austin, TX: The Cognitive Science Society, 2009, pp. 1210–1215.
  - [6] J. Snaider, R. McCall, and S. Franklin, “The LIDA Framework as a general tool for AGI,” *The Fourth Conference on Artificial General Intelligence*, Mountain View, CA, 2011.
  - [7] S. Strain, S. Franklin, D. Heck, and B. Baars, “Brain rhythms, cognitive cycles, and mental moments: the matter-motion-mind connection,” unpublished
  - [8] S. Franklin, S. Strain, R. McCall, and B. Baars, “Conceptual commitments of the LIDA model of cognition,” *Journal of Artificial General Intelligence*, vol. 4(2), pp. 1-22, 2013.
  - [9] S. Franklin, T. Madl, S. D’Mello, and J. Snaider, “LIDA: A systems-level architecture for cognition, emotion, and learning,” *IEEE Transactions on Autonomous Mental Development*, vol. 6(1), pp. 19-41, 2014.
  - [10] T. Madl, B. J. Baars, and S. Franklin, “The timing of the cognitive cycle,” *PLoS ONE*, 6(4), e14803, 2011.
  - [11] B. Baars, *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press, 1988.
  - [12] S. Franklin, “An autonomous software agent for navy personnel work: a case study,” in *Human Interaction with Autonomous Systems in Complex Environments: Papers from 2003 AAAI Spring Symposium*, D. Kortenkamp and M. Freed, Eds. Palo Alto : AAAI, 2003.
  - [13] L. McCauley and S. Franklin, “A large-scale multi-agent system for navy personnel distribution,” *Connection Science*, vol. 14, pp. 371-385, 2002.
  - [14] M. Polanyi, “Life’s irreducible structure,” *Science*, vol. 160(3834), pp. 1308-1312, 1968.
  - [15] T. Madl and S. Franklin, “A LIDA-based model of the attentional blink,” *Proceedings of the 11<sup>th</sup> International Conference on Cognitive Modeling*. Berlin: Universitätsbibliothek, 2012, pp. 283-288.
  - [16] J. P. Kassirer, J. B. Wong, and R. I. Kopelman, *Learning Clinical Reasoning*, 2nd ed. Baltimore: Williams & Wilkins, 2010.
  - [17] S. Strain, and S. Franklin, “Modeling medical diagnosis using a comprehensive cognitive architecture,” *Journal of Healthcare Engineering*, vol. 2(2), pp. 241-258, 2011.
  - [18] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent dirichlet allocation,” *Journal of Machine Learning Research*, vol. 3, pp. 993-1002, 2003.
  - [19] D. M. Blei, “Probabilistic topic models,” *Communications of the ACM*, vol. 55(4), pp. 77-84, 2012.
  - [20] M. Saeed, M. Villarroel, A. T. Reisner, G. Clifford, L. W. Lehman, G. Moody, T. Heldt, T. H. Kyaw, B. Moody, and R. G. Mark, “Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): A public-access intensive care unit database,” *Critical Care Medicine*, vol. 39, pp. 952-960, 2011.
  - [21] D. Demner-Fushman, W. Chapman, and C. McDonald, “What can natural language processing do for clinical decision support?” *Journal of Biomedical Informatics*, vol. 24, pp. 760-772, 2009.
  - [22] S. Meystre and P. Haug, “Automation of a problem list using natural language processing,” *BMC Medical Informatics and Decision Making*, vol. 5(30), 2005.