Perennial, Permuted, and Pervasive Search in Ambient Intelligence*

Javed Mostafa[1], Xiaopeng Lu[2] and Sandeep Avula[3]

Abstract— **Searching in ambient intelligence is complex because it requires fine-grained tracking of human intentions, adapting to varying interests and intentions, and maintaining search accuracy in the context of fast content changes and increasing diversity of information. The paper describes details on some of the computational challenges by framing them at three levels: 1) a human interacting with a machine, 2) a small group of humans and machines collaborating, and 3) decision-support based on diverse and fast evolving information in the context of a large community of users and machines. Along with the challenges, some early glimpses of potential solutions are also discussed.**

I. INTRODUCTION

A. Foundation

We live in an age when 90% of the information the world has ever produced was generated in the last two years[1]. Information volume is growing, but so is the rate of information production (paced by faster compute cycles). To calibrate the pace, consider the fact that by the time you finished reading the last sentence, Google had processed about 80,000 search queries[1].

Information has become a critical component not just for conducting complex tasks. Information has become essential for supporting our livelihood in a day-to-day or perhaps even minute-to-minute manner. However, we are only at the early stages of conceptualizing and planning services that will ensure that humans can receive, interpret, use, and interact with information in a natural and timely way.

Many challenges lie ahead. A critical one has to do with supporting search services. It is clear that we are moving toward broader deployments of ambient intelligence: environments that are all around us and rely on intelligent technologies such as IoTs, video trackers, and cloud-based information services (large knowledge bases, data analytics, and search tools) to support essential day-to-day services. It is also clear that without conceptualizing how search would function in such an environment, both at the machine and at the user level, little progress will be made for humans to function in ambient intelligence in a productive and safe way.

*This work was partly supported by a grant from the United Health Foundation, Project ENABLE A18-0207-002 and the NIH Grant 5-T15-LM012500-03

[1]J. Mostafa is with the iSchool and the School of Medicine, University of North Carolina at Chapel Hill (UNC-Ch), 216 Lenoir Drive, NC 27599, USA, jm@unc.edu

[2]X. Lu is with the Language Technologies Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA, xiaopen2@andrew.cmu.edu

[3]S. Avula is with the iSchool, University of North Carolina at Chapel Hill (UNC-Ch), 216 Lenoir Drive, NC 27599, USA, asandeep@live.unc.edu

B. Aims

In this vision paper, the aims include providing research directions in three specific ambient intelligence contexts: 1) human-machine communication in one-to-one settings, 2) humans and machines collaborating in small group settings, and 3) many humans and machines collaborating in complex environments asynchronously or synchronously (i.e., near real-time).

The subsequent sections of the paper is organized as following. Critical search challenges will be discussed categorized under the three ambient contexts. Upon discussions of potential limitations and gaps, some likely solutions that can address the challenges will be presented as well.

II. SEARCH CHALLENGES

The next three sections are: 1) Deeper Human Search Patterns, 2) Talking and Searching, and 3) Searching Diverse and Changing Knowledge.

A. Deeper Human Search Patterns

Imagine a human-machine interaction scenario where the human is engaged in a complex activity such as a medical diagnosis task. The user, in this case an expert clinical neurologist, is reviewing a set of magnetic resonance imaging data recently produced from a brain scan. The goal is to pinpoint the regions-of-interest (RoI) that are likely to be indicators of an early stage ailment such as Alzheimer's disease. Although, the demand for information may vary depending on the specialty and the style of analysis, the latter task is intensely search dependent. Some of the searches are explicit and user-initiated, for example, looking up relevant cognitive scores[1] or the genomics analysis profile. But many searches are context-dependent and they are triggered by user-behavior such as the cursor moving to a particular brain image location which draws the user's attention. There is a need here for interest-profile driven personalized search (i.e., a good profile incorporating information from the patient would generate clinical guidelines and other relevant evidence). There is, however, a need for continuous search based on deep knowledge of the domain and near real-time searches that track the user's behavior. We refer to this type of search as MAPS (micro-accuracy perennial search).

In a complex task, MAPS need to execute almost simultaneously as users interact with the information displayed. For example, as the user's cursor moves proximal to the CGI score, the system should already be in a search mode

[1]One of type of standard scores is known as Cognitive Global Impression or CGI, which is a score assigned by clinicians as an assessment of severity of mental illness relative to past patients with similar diagnosis.

978-1-7281-6737-4/19/$31.00 ©2019 IEEE
DOI 10.1109/CogMI48466.2019.00027

to access and retrieve the score guideline, i.e., the definition and the clinical validation anchoring the guideline, and the recent cohort data on similar patients.

Another example can be drawn from the annotation stage of the scan. Imagine that as the neuro-radiologist is conducting the annotation, the system retrieves similar regions of interest (RoI) from previously annotated scans in near real-time. Simultaneously, the system volumizes the RoI to improve clarity and provides key quantitative measurements of the RoI such as contextual distribution measures to indicate the relative aberrations associated with the RoI. The latter annotation scenario is not unlike what happens in a modern Computer-Aided Design (CAD) tool, whereby the designer receives context-sensitive assistance on the size and the possible shape of an object being designed, near real-time, based on intent and constraints such as proxi- mal mechanisms/components, potential functions, and other fully-designed components that are similar with respect to mechanisms and functions.

A critical challenge here is to deploy search functions that can operate based on ultra-fast intent detection – calibrated with the brain's processing speed (sub 20 milliseconds). A broader challenge is to study behavior and human responses based on signals that are not post facto[2]. Fundamentally, to be able to develop theory-anchored, generalizable, and powerful modalities for supporting user's actions it is critical that we create new methods for collecting and interpreting physiological and neurological data associated with user-machine interactions (called neurophysiological or NP data). Unfortunately, current methods to capture and track NP data such as eye gaze, focus, pupil dilation, skin temperature, electrodermal response, electroencephalogram (EEG), electrocardiogram (ECG), and functional magnetic resonance imaging (fMRI) have only been sparely utilized for studying search. Experimental methods that use NP for elucidating search are at an early stage of development and not well understood[2].

Few years ago, the famous AI researcher Tom Mitchell and his colleagues demonstrated that it is possible to train a neural network to predict the neuro-activation patterns of the brain based on nouns that the brain has not been exposed to before[3]. A corollary of this approach, if it were to be developed, could potentially reverse the process. In other words, at near real-time, as the user is formulating a thought or reflecting on a particular concept, then the corresponding noun term can be predicted. If such an inversion is possible, and it is a big caveat[3], then it can become possible to calibrate and refine intent in a much more granular way and achieve significant improvement in perennial search.

In another line of research it has been demonstrated that there is a strong correlation between user's relevance judgement and the items they gaze upon[2]. As can be seen in the figure below, there is a distinct difference in user's reading patterns when he/she responds to items that are considered irrelevant vs. relevant. As the user transitions into interpreting and/or reading content it may be possible based on the gaze patterns to determine if an item is of interest to the user and begin to anchor intent-specific cues on such a pattern.

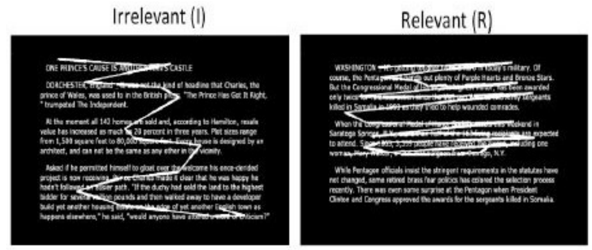

Fig. 1. It is possible to segment the irrelevant content from the relevant content displayed based on users' gaze-patterns[2]

The challenge is that detecting deep and engaged patterns of use as shown in figure 1 demands significant time. However, on the positive side, the latency may provide affordances for MAPS to achieve better results. Another related approach that may boost MAPS results is anchoring intent on a different source of physiological signal. It has been shown that subtle changes in pupil dilation and pupil contraction are behavioral markers for relevance and irrelevance respectively[2]. An area which deserves closer examination for tracking physiological signals is the rate and quality of signal production based on reasonably priced and currently available devices (e.g., eye trackers). The trade-offs among cost vs. quality vs. usability and the influences on search results require further investigations.

In general, ML-based classifier accuracy, in the context of streaming data, is very much dependent on stream-complexity (volume, format, and noise) and the rate of data ingest. However, if the computation can be conducted using high-density parallel processor platforms it may be possible to achieve both efficient and effective response[5].

B. *Talking and Searching*

"Air Canada 301, IFR Clearance", states a pilot sitting on the tarmac of an airport. The response from the tower is "Air Canada 301, Los Angeles Clearance Delivery. Calgary International Airport via the Gorman Four Departure, Shafter Transition, then as filed. 5000, Expect Flight Level Three Four Zero. Departure Frequency 124.30, Squawk 7201." The latter dialog may seem somewhat exotic but it is an extremely efficient and accurate way pilots and the air traffic controllers (ATC) exchange messages[4]. One could argue that as both the pilot and the ATC staff are humans

[2]Human-computer interaction studies have traditionally relied on observational approaches that monitor visible actions carried out by humans and use performance on tasks as data.

[3]Miniaturization of brain scanning devices is advancing at a rather steady pace. For example, it is possible to create a scanning device based on magnetoencephalography which permits natural movements while capturing brain electrophysiological measurements at milliseconds resolution quite accurately[4]

[4]There are several other conventions and specialized vocabularies used by ATCs. For some examples, see: https://www.vatsim.net/pilot-resource-centre/ifr-specific-lessons/ifr-clearances.

they could easily communicate in English and why do they choose to "complicate" matters by "speaking" in this pidgin format. It turns out using such pidgin languages allows fast exchanges to occur without sacrificing details and accuracy. The communication described here is emblematic of many such compact sub-languages that humans use to accomplish highly specialized tasks. For example, consider the dialog which takes place among the members of a surgical team while in an operating theater. The conversation may be less structured, but yet full of specialized terms and well-known turn-taking and transitions.

As many years of research on natural language processing (NLP) has revealed, achieving smooth and natural communication between humans and machines remains elusive. In currently available NLP-based applications, for example in Amazon Echo, users discover brittleness rather easily and they also “adapt” to such brittleness after a few hiccups[5]. However, for certain mission-critical tasks such as completing a computer-aided design project or conducting an experiment with the aid of automated laboratory machines, it may be much more efficient and effective to rely on a highly compressed version of a human language. Such a language should have some of the advantages of human languages (e.g., English like words but with narrow and concrete semantic scopes) and contain few easy to remember syntactic rules (e.g., how nouns and verbs are combined).

Of course, many computer languages have human-like "English" words. Command sets of Unix and DOS come to mind. However, these command sets are large and their syntactic combinations are not easy to remember. With the wider acceptability of IoTs, devices with built-in intelligence will continue to proliferate. The form factor of the intelligent devices, for example their size, their design, and their location or proximity may make it difficult to interact with them using conventional user interface features such as icons, menus, and mouse.

The situation in an intelligent or ambient environments will become even more complex when multiple devices need to be used and collaborations among a small team of individuals and machines has to be orchestrated to accomplish critical tasks (e.g., conducting a team-based scientific experiment). The most efficient way of accessing and manipulating specialized functions of devices and to ensure that team members performing a task together can follow and understand the commands is to verbally communicate the commands. The latter situation demands that the commands are easily understandable and the commands are limited in number to facilitate ease of memorization. The latter characteristics of the command set thus rules out specialized search languages such as SQL. A compact, task-centric set of commands, which is understandable and memorizable by humans, and can be processed efficiently and accurately by machines could be the basis for a whole class of new languages to support team-based science and more generally complex team-based activities.

[5]Almost in the same way humans overcome communication barriers when they are required to interact with people from another culture or country. For known obstacles with current NLP appliances and applications and how large number of user-interactions are being harnessed to improve their accuracy, see: https://qz.com/1288743/.

A major recent effort by my research group involved development of a compact language, called the Minimum Dictionary Language (MDL), which we will deploy to support specialized tasks. One use case we are experimenting with is “hands-free” interactions with mobile phones. As a pilot project, to simplify the situation, we selected a single human to machine interaction use case, involving searching and reviewing email messages from a mobile phone. Below, we offer some early glimpses of the email-MDL we created. Our hope is that it will generate critical analysis of our approach and bring focus to the area of compact human-machine languages that are designed to support specialized tasks.

1) Language Definition: Our formal language definition is expressed using Extended Backus-Naur Form (EBNF). The intuition behind MDL will be first introduced and the formal definition will be presented at the end of this subsection.

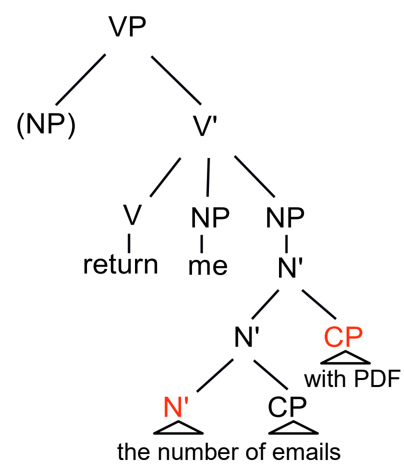

Fig. 2. Example of X-bar Structure of an Imperative Sentence

When people try to retrieve information through a visual assistant, a command is usually used (a non-command sentence can of course be used as well, however, it should be able to be transferred into a command. For example, "I want an apple" equals to "give me an apple"). As in English, commands can be expressed with imperative sentences. Figure 2 is an example of a part of the syntax structure of an imperative sentence illustrated by X-bar theory [6]. Since our study focuses on the content of the command, everything above the Verb Phrase(VP) was omitted. In a retrieval task, the intention of a user should be getting desired results through a query. Thus, for example, if the user wants to find how many emails there are with PDF as attachment an efficient approach would be to choose transitive verbs similar to "return" or "tell" under the Verb(V) which serves as the action; and "me" under the first Noun Phrase (NP) which serves as the first objective. To make our grammar succinct, we only keep the second Noun Phrase(NP) which contains the object we desire and the modifier of the object. The object can be "the number of emails" as shown in the example, or simply "emails". We call this part OPERATION, namely the object that we operate on. The second important part is the modifier which is composed of a Complementary Phrase(CP) "with PDF". We call this part SCOPE, namely the scope of the object that we care about. Therefore, we designed MDL as a tree structure with OPERATION and SCOPE as the first two branches.

For OPERATION, we structured it through four main aspects: *Who* (who is the sender), *When* (the time stamp associated with the origin or the creation of the message), *What* (the main topic of the message), and *How* (characteristics of the message, such as frequency, type of data contained, age of the message, language of the message, and importance of the message).

For SCOPE, as MDL serves as a search tool for a personal information bank for messages, we design the SCOPE to be based on two critical structures of messages, namely the *container* and the *content*.

2) Queries in MDL: Queries generated by users would typically concentrate on *Container* and *content*. For example, sometimes the focus of a query would be at the *container* level (i.e., not necessarily focus on the body of the message but on the labels used to describe the message container). A concrete example is *"?'Selim' LAST"*, meaning to retrieve the last message sent by Selim. Or, the user may request *"?'Selim' ALL"*, which means to retrieve all messages from Selim. Another similar query the user may issue is to retrieve all messages from an organization. For example: *"?'IBM' ALL"*. In this query the focus is not on the content of messages, rather, the focus is on the source of the messages. Therefore, it is a query with the focus on the *container* level.

A query associated with *container* could be more complex – the attachment format can also be included. Possible queries are *"?TOTAL EMAIL MSWORD"*: to retrieve all messages that contain MSWORD files; or *"?TOTAL EMAIL PDF"*: to retrieve all messages that contain PDF files. These types of queries are still focused on messages as *container*, instead of as *content*.

Content queries are another group of queries that focuses exclusively on what the messages actually are. A typical content query for retrieving all "Mike's" messages on the upcoming "picnic" would be *"?'Mike' ON 'picnic'"*. Also, retrieving the last message on the upcoming "picnic" would be *"?ON 'picnic' LAST"*.

Using the two primary dimensions, namely OPERATION and SCOPE, all MDL queries are created. Therefore, the grammar of the MDL query language can now be described in terms of the two dimensions and some additional details. Our goal is that by using MDL, users may maximize their need through succinct grammar and retrieve information as they desire. The structure of the design of MDL is shown in Figure 3.

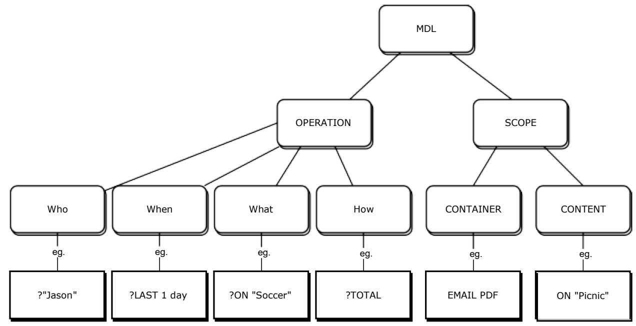

Fig. 3. Structure of email-MDL

The motivation behind developing the email-MDL was to build the corresponding parser for querying personal messages based on easy to learn commands and realistic use cases. Specifically, keyword commands are implemented so that it is relatively easy for inexperienced users to comprehend and learn. Also, as shown above, the email-MDL is clearly structured into SCOPE and OPERATION parts to match user-needs in a natural way. While the syntax of the email-MDL is limited to retrieval operations, a structure similar to the one shown in figure 3 in terms of organization of commands and the grammar can be easily created for other domains and operations.

As stated, the grammar of an MDL is based on EBNF (Extended Backus-Naur Form) and it is formally defined by a set of commands as well as fields that support granular but powerful operations. A *"?"* is used for the indication of the start of MDL query. Predefined commands at *op* (OPERATION) level contains *'TOTAL'*, *'LAST'*, and *'ON'*; predefined commands at *sc* (SCOPE) level includes *'EMAIL'*, *'EMAIL from'* and *'EMAIL last'*. User is supposed to give different inputs after these predefined commands for the completion of a MDL query. The formal EBNF definition for MDL is shown below.

The email-MDL grammar was developed through a bottom-up tree structure. Initially, a basic grammar was defined. To develop the basic grammar, we followed the initial design of SCOPE and OPERATION. Next, the top level grammar was elaborated and made more explicit by combining the lower level elements. The elaboration was achieved by extending the grammar to support four specific use cases: covering the cases "Who", "What", "When", and "How". The actual implementation of the email-MDL was conducted using the Python Parsimonious module.

With wide-spread availability of IoTs, development of smart devices with embedded intelligence, and creation of ambient places where many of us will work and live, a popular modality of interactions between humans and machines will be verbal communication. We can also be fairly certain that search functions, either conducted explicitly or implicitly, as part of human-machine dialogs in such smart environments will be frequent. Therefore, we must attend to how search can be supported seamlessly, efficiently, and effectively as part of human-machine dialogs.

C. Searching Diverse and Changing Knowledge

When Andrew Grove, the famous leader and the former CEO of Intel, was diagnosed with prostate cancer he initially felt he was in good hands with the highly qualified doctors

Algorithm 1: EBNF Definition for MDL

```
(all) ::= (space) (op) (space) (sc)
(sc) ::= (sc_EMAIL_from) | (sc_EMAIL_attach) |
    (sc_EMAIL_piece) | (space)
(sc_EMAIL_piece) ::= 'EMAIL last' (space)
    (op_LAST_piece)
(sc_EMAIL_from) ::= 'EMAIL from' (space)
    (op_lit_name)
(sc_EMAIL_attach) ::= 'EMAIL (space) (sc_attach)
(sc_attach) ::= 'MSWORD' | 'PDF' | 'GIF'
(op) ::= (op_trig) (space) (op_first) (space)
(op_first) ::= (op_TOTAL) | (op_LAST) | (op_lit_ON) |
    (op_lit_name)
(op_lit_ON) ::= (op_lit_name)* (op_ON) (space)
    (op_lit_topic) (space) (sc_attach)* (space)
    (op_LAST)*
(op_lit_topic) ::= ' (chars) '
(op_lit_name) ::= (' ((chars) (space))+ ' (space))+
(op_ON) ::= 'ON'
(op_LAST) ::= 'LAST' (space) ((op_LAST_time) |
    (op_LAST_piece))
(op_LAST_piece) ::= '[0-9]*'
(op_LAST_time) ::= '[0-9]*' (space) '[a-z]+'
(op_TOTAL) ::= 'TOTAL'
(op_trig) ::= '?'
(space) ::= ' '*
(chars) ::= '[A-z0-9]*'
```

he had access to. The journey he experienced to identify an appropriate treatment option, however, was frustrating, beguiling, and complex. Known as a person who preferred clarity and strict logic above anything, Andrew found himself shocked that what appeared on the surface to be a rather straight-forward situation turned out to be extremely circuitous and uncertain[7]. After all, among scientific fields, medicine prides itself as one strictly anchored in and driven by current and rigorous scientific evidence. As Andrew started receiving advice from various medical experts, his understanding of the prognosis and treatment became less clear and the path to recovery hard to chart.

For many situations, more than one decision or conclusion may stand out as equally "optimum" or relevant. When information gets produced at a voluminous rate and it originates from different sources, the situation can become even more complex, as resolving contradictions and/or weighing differential consequences based on evidence becomes exceedingly difficult[6].

A potential way to reach a resolution when different facts appear to be equally correct or valid is to collect objective data from different sources. However, even when it comes to indexing content from a single source, today's advanced search systems perform rather poorly. First, the actual content of documents that are indexed is treated in a shallow way (i.e., indexing is not conducted in a domain- or knowledge-specific way). Second, when search engines do take into account the intellectual dimension of content, i.e., importance of evidence as determined by humans, it is also conducted in an indirect and rather perfunctory way: based on association of links to documents or association of keywords to documents clicked on.

Aggregation of content, data, and documents will need to be radically re-engineered for supporting subtle and complex decision-making in ambient environments. The infrastructure of search, particularly in large and complex environments, such as the type of ambient environments we envision in the future, need to accommodate incorporation of knowledge collected from diverse and many human experts. The human experts may hold different perspectives, even on the same evidence or same facts and such opinions may change depending on specific situations.

Additionally, with multitude of knowledge domains (or sources of evidence) and fast evolving knowledge, humans will need to rely on intelligent search systems that can keep up with them in a continuous and adaptive way. Two potential computational paradigms could provide some new directions for searching on diverse and evolving knowledge. The first is inspired by Distributed Artificial Intelligence (DAI), particularly multi-agent systems[9]. The intelligent search aids will have their own knowledge representations (i.e., means for new knowledge detection and knowledge classification) and can bring to the user's attention different evidence, along with confidence scores associated with the veracity of knowledge and the authority of their sources.

Our lab conducted projects on multi-agent systems to support critical search functions such as distributed identification, classification, and retrieval of information on behalf of users[10]. The environment, called Multi-Agent Collaboration for Classification of Information (MACCI) was designed for two broad purposes: 1) To understand agent dynamics during collaboration and 2) To establish trade-off between agent knowledge and the collaboration dynamics needed to achieve optimum performance. Agent dynamics refers to understanding collaboration patterns at a granular level, particularly how agents can identify collaborators, engage them appropriately, and utilize and update their knowledge to calibrate with global knowledge about domains and the agent community. The trade-offs are particularly important because agents cannot aspire to be "omnipontent"– possessing comprehensive knowledge about any domain or even the agent community. Such an approach would defeat

[6]Biomedicine is the most vigorous and the fastest growing scholarly area among all STEM fields[8].

the purpose of developing a distributed computing model and, more practically, it would not be viable to keep up with knowledge as typically produced, stored, disseminated, and updated. Hence, by design, an individual agent's knowledge is always incomplete and the goal is to find complementary knowledge in other agents and accomplish the selection and retrieval tasks through intelligent coordination schemes and collaborations.

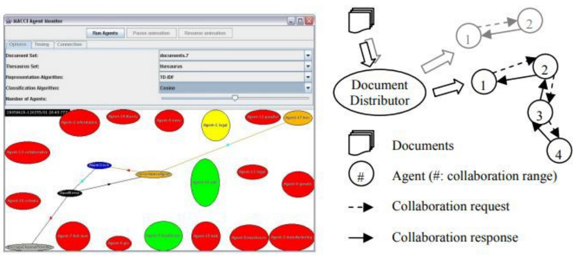

(a) MACCI interface (b) MACCI simulation

Fig. 4. MACCI: A Collaborative Agent Community

The task of searching for most relevant information in a vast, distributed, diverse, fast evolving, and sometimes contradictory[7] set of information cannot be solved by assuming information is aggregated and searched from a single service. The overall problem space, in terms of both demands and sources, is too large and too complex to rely on a single centralized search model. Alternative search solutions that decouple the aggregation from the retrieval of relevant information are likely to be far superior than current approaches. Again drawing inspiration from DAI or multi-agent systems, we are likely to be better served if an universal standard were to be agreed upon for agent services (e.g., similar in scope and wide acceptability as say TCP/IP). With the aid of such a standard, the current Internet would go beyond an environment which supports efficient information aggregation to one which supports *both intelligent aggregation and searching*. In such an environment every site with information available for access via the web, would also host information agents that offer two types of services: search on local content and communication with other agents it considers trustful. The latter functions, i.e., the communication and acquaintance list of trustful agents, would permit intelligent distributed search services to emerge. One viable approach for handling and resolving contradictions and ambiguities over diverse and changing information is the information market model[11], whereby agents can communicate with each other to find, validate, and generate potentially relevant information based on their past reputation in delivering reliable content. The information market approach has immense potential for supporting search in large and complex ambient intelligence. Information markets can draw upon mature theories from DAI and agent-based computational economics (e.g., see: https://www.predictit.org/support/what-is-predictit). In conclusion, it should be pointed out that permuted, perennial, and pervasive search in ambient intelligence is at an early stage of development and new and exciting prospects for improving search, perhaps based on some of the key contributions from the areas described in this paper demand serious attention.

[7]The problem of determining and retrieving accurate information in an environment which may in fact be deliberately set up for deception and influence peddling is related to identifying accurate information among contradictory set of evidence. However, there are other challenges such an environment poses such as dealing with metrics that anchor relevance on genuine but shallow sources such as links or click-throughs, or content fabrication that are near- or complete copies of the original. The area of avoiding the negative effects of adversarial or deceptive interactions in ambient search environments is a separate topic, which deserves its own treatment and I have not engaged in examining this area in depth here.

ACKNOWLEDGMENT

The paper benefited from vigorous discussions with colleagues in the Laboratory of Applied Informatics Research. A particular area covered, namely human-machine communication, was inspired by the original version of the movie "Blade Runner" – perhaps one of the greatest science fiction movies ever made.

REFERENCES

[1] B. Marr, "How Much Data Do We Create Every Day?", Forbes, May 21, 2018.
[2] J. Mostafa, and J. Gwizdka. Deepening the Role of the User: Neuro-Physiological Evidence as a Basis for Studying and Improving Search. In Proc. of the 2016 ACM on Conference on Human Information Interaction and Retrieval (CHIIR '16). ACM, New York, NY, USA, 2016, 63–70.
[3] M. A. Just, V. L. Cherkassky, S. Aryal, T. M. Mitchell. A Neurosemantic Theory of Concrete Noun Representation Based on the Underlying Brain Codes. PLoS ONE 5(1): e8622 DOI: 10.1371/journal.pone.0008622, January 13, 2010.
[4] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. Duque Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes. Moving magnetoencephalography towards real-world applications with a wearable system, Nature vol. 555, March 2018, pp.657-–661.
[5] Javed Mostafa, Snehasis Mukhopadhyay, Mathew Palakal, and Wai Lam. A multilevel approach to intelligent information filtering: model, system, and evaluation. ACM Transactions on Information Systems, vol. 15(4), 1997, 368–399.
[6] R. Jackendoff. X syntax: A study of phrase structure. Linguistic Inquiry Monographs Cambridge, Mass. vol. 2(1977), pp. 1—249, 1977.
[7] A. Grove. Taking on Prostrate Cancer. Fortune, May, 13th, 1996.
[8] L. J. Jensen, J. Saric, and P. Bork. Literature mining for the biologist: from information retrieval to biological discovery. Nature Reviews Genetics, vol. 7, February, 2006, pp.119–129
[9] N. R. Jennings, and M. J. Wooldridge (Eds.). Agent technology: foundations, applications, and markets. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 1998.
[10] K. Weimao, and J. Mostafa. Visualizing multi-agent collaboration for classification of information. Proc. Am. Soc. Info. Sci. Tech., 45, 2008, pp. 1–4.
[11] R. R. Raje, S. Mukhopadhyay, M. Boyles, A. Papiez, N. Patel, M. Palakal, and J. Mostafa. A bidding mechanism for Web-based agents involved in information classification. World Wide Web 1, 3 (March 1998), 155-165.