

What We Talk About When We Talk About Information Retrieval

Justin Zobel

School of Computing & Information Systems
University of Melbourne, Parkville, Australia

Abstract

The field of information retrieval (IR) is typically defined, in a variety of different wordings, as concerned with retrieval of documents that satisfy an information need. In this essay, I argue that these definitions are inaccurate, fail to capture major threads of activity in IR research, and in particular are flawed because they omit the element of human participation in the retrieval process. After outlining some perspectives to consider in formulating better definitions, I offer an option, as an illustration of how the field might be presented; this option is centred on the purpose of IR, namely, support of cognition. There is an obvious need for a clear statement of the purpose of the discipline: information access is recognized as a human right and IR is the basis of a critical technology for providing that access – one that is deeply intertwined with daily life and is changing human psychology. Well-grounded descriptions can encourage IR researchers to embrace a view of the field that enables richer connection with other disciplines, and should embody a vision of what IR research can accomplish.

Definitions of ‘Information Retrieval’

The discipline of information retrieval (IR) has a long history. Its computational roots can be traced to the end of the 1940s, with practical retrieval systems being demonstrated by 1960. It has been described (perhaps contentiously) as a subfield of information science, and, within computer science, has overlap with or close relationship to numerous areas including database systems, human-computer interaction, language technology, image recognition, subfields of compression and algorithmics, and subfields of artificial intelligence.

The arrival of the Web created a historical disjunct in IR research and technologies. Pre-Web, true IR was largely restricted to laboratories, with deployed systems having only limited capabilities; both research and commercial systems were standalone, expensive, restricted, and applied to defined, curated repositories. Early retrieval systems could be caricatured as ‘what a library would be like if it was properly computerised’; while researchers in the field might not have shared that perspective, it is certainly the case that the only substantial document collections were those managed by librarians or created by editors.

Post-Web,¹ large-scale IR systems rapidly emerged in the form of search engines, which were applied to ad hoc, uncurated data that was collected by crawling. These have been so successful that IR and Web search are often equated with each other, but this false equivalence misrepresents

¹Strictly speaking, post origin-of-Web.

both activities. The field of IR embraces collections that do not have the properties of the Web, and indeed, if the goal of early IR was to provide computation-based libraries,² then that goal is far from achieved; search facilities in even major libraries continue to be ineffective and frustrating, compared to the ease of search on the Web.

The first textbooks in IR appeared in the 1960s, and offered definitions such as that given by Salton [1968] (or with minor wording changes by Salton and McGill [1983]):

Information retrieval is a field concerned with the structure, analysis, organization, storage, searching, and retrieval of information.³

Over forty years later, Croft et al. [2010] quoted and re-used this definition, describing it as ‘still appropriate and accurate’. Other textbooks offer related definitions:

Information retrieval is concerned with representing, searching, and manipulating large collections of electronic text and other human-language data. [Buettcher et al., 2010]

Information retrieval deals with the representation, storage, organisation of, and access to information items such as documents, Web pages, online catalogs, structured and semi-structured records, [and] multimedia objects. [Baeza-Yates and Ribeiro-Neto, 2011]

Two further textbook definitions are similar, but with a critical distinction:

Information retrieval is ... a means by which users of an information system or service can find the documents, records, graphic images, or sound recordings that meet their needs or interests. [Meadow et al., 2000]

Information retrieval (IR) is finding material (usually documents) of an unstructured nature (usually text) that satisfies an information need from within large collections (usually stored on computers). [Manning et al., 2008]⁴

These versions differ from the others in that they introduce the concept of ‘information need’, which is widely used in descriptions of IR, and is a signpost of the purpose of retrieval: to assist humans.

Other definitions can be found in a dictionary and a crowd-contributed encyclopaedia:

The techniques of storing and recovering and often disseminating recorded data especially through the use of a computerized system. (Merriam-Webster Dictionary online, accessed 22 August 2017.)

Information retrieval is the science of searching for information in a document, searching for documents themselves, and also searching for metadata that describes data, and for databases of texts, images or sounds. (Wikipedia, accessed 22 August 2017.)

These definitions are all essentially similar, and in my view all of them are flawed.⁵ Notably, they

²I’ve already admitted that this is a caricature.

³This definition is curiously similar to the standard descriptions of information science, which however is generally seen as a much broader field. For example, in Wikipedia (accessed 24 September 2017) information science is described as ‘concerned with the analysis, collection, classification, manipulation, storage, retrieval, movement, dissemination, and protection of information’.

⁴These instances include all the major general textbooks in the field, at least, that I am aware of. There are many other books on IR, but most of them are intended for researchers or are concerned with subdisciplines. For example, Witten, Moffat, and Bell’s *Managing Gigabytes* and Frieder and Grossman’s *Information Retrieval: Algorithms and Heuristics* explicitly focus on retrieval mechanisms; moreover, neither offers a definition. Likewise, Voorhees and Harman’s *TREC: Experiment and Evaluation in Information Retrieval* has no definition, and is concerned with research methodology in the field rather than with the field itself.

⁵And they entirely fail to capture any sense of wonder or challenge. They are uninspiring.

are silent on the existence of a context in which retrieval is taking place and on how the retrieved material might be consumed, and struggle to explain what the material being searched consists of.⁶

An underlying issue with these definitions is one of terminology. The usage ‘information need’ isn’t particularly meaningful beyond the IR community,⁷ and (amongst other issues) is misleading in that it wrongly implies that the user always has a definite ‘information need’ that leads to a determinably satisfactory outcome. Even more problematic is ‘information’. Much of computer science is concerned with the ‘structure, analysis, organization, [and] storage’ of information; the term is far too sweeping to be (dare I say?) informative. ‘Recorded data’ is perhaps worse, because there is an ordinary interpretation of the term ‘data’ that entirely excludes textual material. ‘Electronic text’ seems to speak from a past era (today we usually say ‘digital’ rather than “electronic”) and focuses on the actual text to the exclusion of, for example, meta-data; while the clarification ‘other human-language data’ leaves one wondering what that data might be, if it isn’t text. ‘Of an unstructured nature’ suggests a norm in which data is ‘structured’, perhaps on the basis that it is created mechanically; the term isn’t quite meaningless, but it is very much open to interpretation. In my view these terminologies are all thoroughly unsatisfactory.

Three of the quoted definitions mention documents. But there is no agreement on what a document is. Does ‘text’ include Web links, annotations, or metadata? Is a piece of legislation one document, or many? Is a transcript a document? Traditionally, a ‘document’ was a physical artefact that embodied a written text: a newspaper article, a book, a piece of legislation. In pre-Web IR, such documents (transliterated into a digital form) were the primary target of retrieval. In that context it was entirely reasonable to describe IR as concerned with ‘retrieval of documents’. However, search systems and IR research are applied to many kinds of stored material that are not text.

Helpfully, though, the definition of ‘document’ has broadened over time. In legal contexts, broadly stated, ‘a document is an artefact (digital or physical) that embodies information and was created for human consumption’. (The wording is mine, encapsulating conversations with a lawyer.)⁸ Arguably this can reasonably be extended to include derivatives of documents, such as a snippet, as a fragment of a document is a document in its own right; and it unquestionably includes metadata, recordings, photographs, transcripts, and so on.⁹ Documents are typically but not always created by humans – security camera footage and machine-generated image captions being examples where people are not directly involved in creation. It follows that we can be comfortable identifying documents as the material to which retrieval is applied, and the term does not even need qualification.¹⁰

But these are side issues. The shortcomings of the definitions above are more profound than pedantic questions of terminology.

⁶A reasonable skeptical question might be: why have definitions at all? But we do need them. They communicate about the field to other disciplines, help to provide scope to expertise, help to identify experts, and identify where literature might be found. They help us set goals and define relationships between problems. They are also of value socially – we need to be in tribes in order to function in an organised way as a community. And, as in the previous footnote, they can motivate.

⁷A ‘fact’ that I have established by taking an informal, unrepresentative survey of a dozen or so colleagues in other fields, including a few from outside computing.

⁸For example, in the State of Victoria (Australia), under the *Evidence Act 2008*, a ‘document means any record of information, and includes (a) anything on which there is writing; or (b) anything on which there are marks, figures, symbols or perforations having a meaning for persons qualified to interpret them; or (c) anything from which sounds, images or writings can be reproduced with or without the aid of anything else; or (d) a map, plan, drawing or photograph’.

⁹That is, the question of ‘what material is it that IR is applied to?’ is greatly simplified by explicit acknowledgement of the role of that material. As far as I know, this style of definition of ‘document’ hasn’t previously been proposed.

¹⁰And thus the distinction between documents and other material, as in the definition given by Baeza-Yates and Ribeiro-Neto [2011], is in some respects misleading.

Doubts and Questions

There are two broad respects in which the properties of retrieval systems, and trends in research in IR, raise questions about the standard definitions of the field. One respect is with regard to people and society; I explore these in the next section. Another respect concerns some of the capabilities of retrieval systems that have been developed post-Web, as I now discuss.

Both public research in IR and innovation in commercial search engines have led to remarkable improvements in and additions to the capabilities of retrieval systems; and likewise have greatly extended our understanding of the principles of IR. Some of these developments, however, confound the definitions above, although they may seem to IR researchers and practitioners to be natural elements of the field. I now consider some examples, while noting that this discussion is not intended to be exhaustive.

A results page is an entity of interest in its own right. A results page is more than the sum of its parts. It helps the reader interpret the listed documents, and may set them in a context. The introduction of query-biased snippets to results pages may seem in some respects to have been a small step, but in my view this change marked a transition to a dramatic new conceptualisation of what a page of search results is: an informative document in its own right, not just a guide to what documents are available.

A results page is just one possible form of presentation of the outcomes of retrieval (albeit an extremely common one), but the same argument also applies to other forms. There is an argument that, in response to a query, a search system does not retrieve documents so much as organise, summarise, and link to them – it provides on-the-fly information synthesis.¹¹ ◇

The documents in a collection may be richly interconnected. On the Web, in collections of journal articles, in medical reports, and in many other contexts the stored documents are explicitly interlinked (manually and automatically) in complex, informative ways. They may also be implicitly interlinked via indicators such as chronology, topic, or external event. These links and interconnections are yet to be fully exploited, but they are already being used in valuable ways, most prominently via ranking indicators such as PageRank or anchor text.

Moreover, interconnection challenges the notion of a document as a complete object: a Web page is not necessarily intended to be consumed independently, but explicitly forms part of a larger whole. That whole may not have defined boundaries (and is of unknown, unknowable size). Describing a network of documents as a mere collection or database, then, is to miss the point. A Web of documents is an entity, in which the relationships can be as significant as the individual nodes, and can be of great value in the retrieval process.

Connectedness is a further respect in which terminology fails us. Some authors have described collections of material as ‘information spaces’, which recognises the concept of locality but not necessarily the connectedness. We have ‘the Web’, but it is an instance of such a network, not the class. Without terminology, we will struggle to capture this characteristic in a definition; I don’t have a solution to offer.

Of course, we also search collections that are not explicitly interconnected, and perhaps for this purpose it is sufficient to recognise that, as was true for ‘document’, the term ‘collection’ may have a broadened meaning. ◇

¹¹That is, as a descriptor of the discipline, ‘retrieval’ may be a misnomer, just as is ‘information’.

The material being searched over can be uncurated, dynamic, and highly heterogeneous. In current collections (not just on the Web), the documents are not atomic. Each document can have implicit and explicit structure, be comprised of both static and dynamic components, and have content that changes depending on who is viewing it. Documents can also be transient, and viewing them can cause them to change. Such properties are a challenge to traditional definitions of ‘collection’ (which has already caused difficulties), and also to the concept of ‘retrieval’, as what the system sees may not be what the user sees.

Crawling introduces another characteristic that distinguishes many Web collections from the controlled collections of pre-Web IR. Links between pages are used to automatically determine the contents of the collection; the editing or curation is provided by the author of the links, not by the system provider. These links mean that the scope of a search is in a fundamental way unspecified, and correctness, completeness, and accuracy cannot be determined. ◊

Search systems employ sophisticated models of individual users. That is, knowledge of the context in which a query is presented is used to modify the search process. This is most obvious with regard to properties such as language and location, but user characteristics (potentially derived from other linked sources, such as authoring platforms or social media, as well as from prior searches) are also exploited.

Users can be said to have ‘state’, and the aim of retrieval is and always has been to identify the documents that are most informative given that state. However, with the advent of user models the system too has state, which changes as the user model is informed and enriched. This dependence on the individual user was missing in earlier understandings of the field. ◊

Aggregate user behaviour provides feedback for future searches. Every user interaction with a current search system is an indicator of the value of the information encapsulated in a query list or results page. This implicit feedback is continually aggregated and employed to influence the output of future interactions, and thus system behaviour is a mechanical consequence of user behaviour. However, the perspective that a search engine is a mechanism for monitoring and exploiting human interactions is entirely absent from the definitions given earlier.

An emergent feature of current search systems is that, as they gather historical querying data, many interactions with search systems are used to explore past queries rather than indexed documents. That is, query completion technologies, developed to assist users to pose queries efficiently and accurately, are also used to search the query repository: the wording or spelling of a query may be the piece of knowledge that a user requires.

A collection of documents can, then, be regarded as a tool for collecting queries, and these queries can be used in a variety of ways in addition to direct support of future user interactions. For example, in this respect it could be argued that IR encompasses building of models of language use. ◊

It might be argued that these are largely technical inconsistencies, which can be addressed by clearer articulation of the terminology. However, this argument considers detail without examining the larger picture, namely, that the existing definitions focus on only a component of the field, while in other respects implying a scope that is much too wide. That is, it might be argued that the quoted definitions are adequate on the basis that they imply some of the elements that seem to be missing, such as user models, iteration, development of knowledge. However, other elements are simply not there – how returned results are consumed, for example¹² – while the definitions include activity (such as

¹²Ingwerson and Järvelin [2005] has a simple figure that characterises these shortcomings (p. 315). This rich, exploratory,

‘analysis of information’) that is far broader than IR. In any case, the argument is easily rebutted: surely there is no good reason to defend a definition in which the explanatory context is omitted and the terminology is either wrong or misleading.

Moreover, there is a deeper issue: the definitions do not reflect the fact that retrieval is a human-driven process that is intended to address complex human goals.

The Ghost in the Machine

Straightforwardly, the reason that people use search and retrieval is to acquire knowledge. This has been widely characterised as having an ‘information need’, that is, there is a knowledge gap that the user is seeking to fill. However, it is often the case that the user not only cannot articulate that need as a textual query, but may have an ill-formed understanding of what that need is [Belkin, 1980, Ingwerson and Järvelin, 2005].

A richer characterisation is that retrieval is a learning experience, in which each interaction with the system changes the user’s state of knowledge. This view of IR is the basis of interactive information retrieval, which is closely related to the concept of cognitive retrieval. Ingwerson and Järvelin [2005] presents an extensive overview of cognition, search, and retrieval, and also presents a variety of perspectives on how these concepts relate to each other.

A definition of cognition is “ ‘The mental action or process of acquiring knowledge and understanding through thought, experience, and the senses’ . . . Cognitive processes use existing knowledge and generate new knowledge.” (Wikipedia, accessed 22 September 2017.) Retrieval can play an obvious role in cognition, as a mechanical provider of knowledge. From the perspective of a user, what is arguably a very natural way to describe a retrieval system is that it is a kind of external knowledge bank, supplementing ordinary human memory. Without hyperbole or exaggeration, it is reasonable to describe search engines as knowledge tools for enhancing or supporting human cognition [Belkin, 1990]. The core of the approach, as Belkin explains, ‘is that it explicitly considers the states of knowledge, beliefs and so on of human beings’ and that with respect to IR systems the ‘goal is the eventual appropriate modification of [the] state of knowledge of the user’.

Ingwerson and Järvelin [2005] remark that ‘it seems prohibitive for the development of IR research if the IR community continues to consider the Laboratory Model in isolation of context’ (as they explain in detail in the introduction to Chapter 5). Belkin [2015] notes his frustration with:

the reluctance, or inability, of IR researchers to accept a broader and more realistic goal of their enterprise; that is, to go beyond identification of relevant, or even authoritative, information objects, to the goal which I think many would agree is what we should aspire to: *the support of people in achievement of the goal or task which led them to engage in information seeking*.¹³ [emphasis added]

Arguments for a cognitive approach have been made for several decades, not only in papers by Belkin and by Ingwerson and Järvelin [2005], but also by many other authors; I note for example Spink and Cole [2005]’s collection.

A cognitive perspective removes the contentious terminology ‘information need’, and also introduces a way of describing the user models that are maintained by search engines. They are intended

wide-ranging text is not easily summarised, but its reflections on search, retrieval, models, learning and many other related concepts are an insightful critique of IR, and it provides a formal overarching framework in which individual IR activities can be informatively categorised.

¹³Obviously, as it is the topic of this essay, I share Belkin’s and Ingwerson and Järvelin’s concern with the narrowness of what can be seen as the rather mechanical views of the field.

to be a (highly imperfect) representation of the user's knowledge state and of what the user wishes to achieve.

This is not a new perspective, but it has implications that can be tested. A key implication is that a user who has access to an external knowledge bank might be expected to exhibit changed cognitive processes, and indeed that is the case: this phenomenon has been observed by psychological researchers in a range of studies. Small et al. [2008] found experience with Internet searching may affect how people undertake decision making and complex reasoning. Brown [2000] notes a range of respects in which Web access and search is changing behaviour, such as preparedness to use found material to learn tasks on-the-fly; he argues that web users are participating in a form of 'distributed intelligence'. Most dramatically, Sparrow et al. [2011] found that people's memories are adapting to accommodate the fact that some information is found online and are more likely to forget information if they know how to recover it – summarising, they found that the Web 'has become a primary form of external or transactive memory, where information is stored collectively outside ourselves'. Ward [2013] has similarly found that Web use is having significant impact on memory.

Some of these changes are profound. A skilled programmer, for example, may have only limited knowledge of the pertinent tools when the task is commenced. Experienced search users may place a relatively low value on experts and expertise,¹⁴ and indeed the nature of expertise itself is changing: an expert may be someone who has the ability to accurately integrate found knowledge, rather than someone who possesses that knowledge. In short, the existence of search systems, coupled with the extraordinary scope of Web resources, is altering how we think.

Other effects are also of relevance to IR research, in particular societal impacts. Access to information has long been recognized as a fundamental human right, in particular in Article 19 of the United Nation's Universal Declaration of Human Rights, which identifies the right 'to seek, receive and impart information and ideas through any media' (1948). Given that it is a long-standing right, and is now a critical resource in the daily lives of citizens, there are immediate implications with regard to issues such as accuracy, bias, and completeness: properties that we currently do not know how to audit or measure.

Another societal question is the legal status of commercial search engines, which cannot be neatly pigeonholed by previous categories of information provider. As discussed by Grimmelmann [2013],¹⁵ for example, such an engine is neither communication channel, editor, publisher, nor advisor, but has characteristics of all of those things, and also of curator and gatekeeper.¹⁶ Questions around the ethics, responsibilities, limits, and liabilities of search are also influenced by technical questions as to what can reasonably be achieved.

The use of editorial control and results-tweaking, introducing a human into the mechanical loop of retrieval (in ways that were probably not anticipated in pre-Web IR) presents an additional respect in which search sits uneasily in our existing social frameworks; again, there are questions for the IR research community as to how the integrity of an engine might be quantified or examined.

There is a further issue that is perhaps even broader: search is changing knowledge itself. Large-scale aggregation of knowledge or opinions wasn't possible prior to search, except at high cost, but now everyone does it routinely. We have seen in recent years how search has helped to create 'echo chambers' and environments in which opinion can triumph over fact, in a way that is qualitatively different to prior forms of experience, due to the wide variety of materials that can be drawn in.

¹⁴The behaviour captured in the catchphrase 'consulting Dr Google'.

¹⁵Grimmelmann has also published online *The Structure of Search Engine Law*, 2007, readily found with a Web search. This earlier document reviews a range of related issues.

¹⁶And, more insidiously, of surveillance, as these engines observe and monitor behaviour.

Knowledge, for both good and ill, is becoming less about facts, and more about frameworks.¹⁷

Nothing whatsoever of these human or social elements is captured in the current definitions of the field. A re-think is needed.

A Simple Proposition

Perhaps IR should always have been described primarily in terms of how it is used and what it is for, rather than attempting (inaccurately) to describe it in terms of how it works. If the scope of a discipline has been well stated, that scope would not be expected to greatly change over time; a quick review of major disciplines of computing and of science shows that they are typically explained now much as they were fifty years ago. It can be argued that the scope of IR hasn't greatly changed either, and thus, arguably at least, the definitions given at the start of this essay have always been inaccurate.

Earlier I offered the definition:

A **document** is an artefact (digital or physical) that embodies information and was created for human consumption.

Considering this, and the discussion of collections and cognition, I suggest that:

Information retrieval is the study of techniques for supporting human cognition with documents, using material that is sourced from large document collections.

General community agreement on a specific wording is unlikely to arise, but an illustrative definition like this one provides a demonstration of how different, and compelling, a new explanation of the field can be.

In my view this concise definition has an immediate merit: it isn't wrong. Understanding it doesn't require knowledge of the discipline, and it has a focus on the human element. It provides a challenge in that it suggests lines of research that don't follow from the more mechanical definitions given earlier. These include issues such as assessment of the goodness, accuracy, or quality of information being retrieved; how to ensure neutrality, or to audit bias; closer integration of IR with tasks such as authoring; and consideration of how retrieved material might be annotated. These are examples, not in any sense an exhaustive list, but do demonstrate how the question of information consumption becomes significant when the purpose of retrieval is made explicit.

Of course, it has limitations. For example, there are elements noted in the discussion above that aren't explicit in this wording – such as connectedness, or the impact of search on behaviour.

The recent developments in IR haven't altered the discipline so much as highlighted shortcomings in past definitions that have always been present.¹⁸ Belkin and others argued for a cognitive viewpoint in the 1970s [Belkin, 1980]; the validity of these arguments has only been strengthened by developments in IR in the decades since. Given the importance of IR culturally, technologically, and commercially, and given its impact on people and on society, we need a shared view of the discipline that is accurate and that communicates clearly what is meant by 'information retrieval'.

¹⁷With documents being created by both automatic and human sources, automatic annotation of human-created material, and human annotation of automatically created material, there are also epistemological questions to address. A traditional document captures knowledge, but in the current context the distinction between information and knowledge is blurred.

¹⁸This is illustrated by Salton and McGill [1983], where IR is defined as given above but in the preface it is noted that 'any information system designed to augment the state of human knowledge and to aid human activities does utilize concepts and procedures from information storage and retrieval', thus indirectly acknowledging the primacy of the human focus.

References

- Universal Declaration of Human Rights, 1948. United Nations General Assembly Resolution 217A.
- R. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison Wesley, second edition, 2011.
- N. Belkin. Anomalous states of knowledge as a basis for information retrieval. *Canadian Journal of Information and Library Science*, 5, 1980.
- N. Belkin. The cognitive viewpoint in information science. *Journal of Information Science*, 16, 1990.
- N. Belkin. People, interacting with information. *SIGIR Forum*, 49(2), 2015. Salton Award lecture.
- J. S. Brown. Growing up digital: How the Web changes work, education, and the ways people learn. *Change*, 32(2), 2000.
- S. Buettcher, C. Clarke, and G. Cormack. *Information Retrieval: Implementing and Evaluating Search Engines*. MIT Press, 2010.
- B. Croft, D. Metzler, and T. Strohman. *Search Engines: Information Retrieval in Practice*. Pearson, 2010. Updated as a free online edition by the authors in 2015.
- J. Grimmelmann. What to do about Google? *Communications of the ACM*, 56(9), 2013.
- P. Ingwerson and K. Järvelin. *The Turn: Integration of information seeking and retrieval in context*. Springer, 2005.
- C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- C. T. Meadow, B. R. Boyce, and D. H. Kraft. *Text Information Retrieval Systems*. Academic Press, second edition, 2000.
- G. Salton. *Automatic Information Organization and Retrieval*. McGraw-Hill, 1968.
- G. Salton and M. J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, 1983.
- G. W. Small, T. D. Moody, P. Siddarth, and S. Y. Bookheimer. Your brain on Google: Patterns of cerebral activation during Internet searching. *American Journal of Geriatric Psychiatry*, 17(2), 2008.
- B. Sparrow, J. Liu, and D. M. Wegner. Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333, 2011.
- A. Spink and C. Cole, editors. *New Directions in Cognitive Information Retrieval*. Springer, 2005.
- A. F. Ward. Supernormal: How the Internet is changing our minds and our memories. *Psychological Inquiry*, 24, 2013.