

The Future of Knowing
Andrew Abbott

"Brunch with Books"
sponsored by the University of Chicago Alumni Association
and the University of Chicago Library

6 June 2009

Thank you, Judi, for your kind introduction.

Good morning and welcome to Alumni House. I would like to thank the GLS Alumni Association and the library for giving me the chance to talk to you this morning. I would like also to thank them for sending not one but two reminders of it to me in my capacity not as speaker but as a university alumnus. Apparently they were a little worried that I might forget the speaking obligation, but thought that I would remember to come as an alumnus. Then when I showed up, they could remind me that I was supposed to give the talk. It's quite possible, I suppose, that I would give a better talk under those conditions. But as it is you'll have to live with the prepared one.

It has three basic parts, which is appropriate enough since I'm a catholic and today is Trinity Sunday. There's one section of logical games then two sections of data and a tiny conclusion.

My title is "The Future of Knowing." And my talk will answer one simple question, which has probably already occurred to you: Why is my title not "The Future of Knowledge?" After all, there is a famous ditty from a nineteenth century satire, which asserts that knowing and knowledge are the same, by identifying the two in the person of the then Master of Oxford's Balliol College, Benjamin Jowett. It goes like this:

First come I, my name is Jowett,
if it's knowledge, then I know it.
I am the master of this college,
If I don't know it, it's not knowledge.

Now this doggerel has something useful to teach us, but learning its lesson will require a walk down memory lane. Now you're all alumni of the University of Chicago, and so your memory lane probably has a small logic shop on it somewhere. And if you step into that logic shop for a moment, you will note that this little ditty's two conditional statements do not in fact really set knowing and knowledge equal, because the two conditionals are really only one conditional. That's because "if it's knowledge, then I know it" is logically equivalent to its contrapositive "If I don't know it, it's not knowledge," and that is precisely what the second line of the second couplet says. In either conditional, therefore, knowing is the necessary condition and knowledge the sufficient.

By contrast, consider the other ordering, that is, the converse and inverse of the original conditional. The converse of "If it's knowledge, then I know it" is "If I know it, then it's knowledge" and the inverse of "if it's knowledge then I know it" is "If it's not knowledge, I don't know it." These statements are each other's contrapositives, and hence are also really only one assertion, to wit - "If I know it, then it's knowledge." Here knowledge is the necessary condition and knowing is sufficient.

To make this elaborate analysis clear, let us move to Venn diagrams, which are in the adjacent case in the logic shop, next to the plain doughnuts.) What the doggerel says is that knowledge is a set of things entirely contained within the set of things that Benjamin Jowett knows, as jelly is within a jelly doughnut. Either knowledge is exactly equivalent to what Jowett knows (a doughnut that's all jelly - a bit hard to eat), or it is less, for it must be entirely within what Jowett knows. There is no knowledge outside Jowett's knowing, just as there is no jelly on the outside of a jelly doughnut. There is nothing that is simultaneously knowledge/jelly and at the same time outside the doughnut of things that Jowett knows.

By contrast, the other version, the inverse/converse of the doggerel,

means that the set of everything Jowett knows is contained within the set of knowledge, which implies that there might be things that are knowledge but that Jowett doesn't know. That is, in this second case, knowledge is the doughnut and Jowett's knowledge is the jelly. Obviously, it was preferable for the undergraduates who wrote this particular ditty to insult the Master for arrogance rather than ignorance. So they very properly went with "If it's knowledge, then I know it," using the helpful rhyme of "know it" and "Jowett" as a mnemonic.

But I might point out for the real logic aficionados in the audience that this rhyme would also work with the inverse of the original statement (that is, it works with "If it's not knowledge I don't know it"), and therefore that in addition to remembering the rhyme, it is also necessary to recall that the original line constitutes the major premise of a syllogism in the medieval syllogistic mode called Barbara. This is in turn easily remembered by remembering the phrase "Major Barbara," itself easily recalled because of the play of that name by George Bernard Shaw. So once you have left here and find that you have difficulty remembering the ditty that captures the relation of Jowett and knowledge, you can easily regenerate it by remembering the word "rhyme" and the phrase "Major Barbara," or maybe even just "rhyme" and "Shaw."

Now if that way of remembering seems laborious and somewhat medieval, we could try the more modern approach. Today's student would remember this by noting that the words "knowing," "knowledge," "Major," "Barbara," "Shaw," "Jowett," and "mnemonic" occur together on only 46 pages on the entire internet, and that the first such page is a guide to, of all things, the Charles Sanders Peirce papers at Harvard. So perhaps a contemporary student would remember the ditty simply by remembering the URL of the guide to the Peirce papers. That the guide also includes thousands of irrelevant words the student well knows, but figures that enlightenment will come, just as the

medievalists among us figure that enlightenment will come if we think long enough about rhyme and major and Barbara.

Let's step back from this elaborate joke, however, and think about what we have just done. I want to talk about content first and form second. First about content.

My logical analysis showed that although it sounds as if the ditty sets Jowett's knowing equal to knowledge, in fact it states that knowledge is at best coextensive with Jowett's knowing, for Jowett may know many things which might not be regarded as knowledge in the sense of being important things that all would agree should be known by everyone or even by most educated people. This collection of Jowettiana which was not such knowledge might include his parlor-maid's mother's state of health, and things like that. But to avoid that triviality, let's transform the ditty a bit, to make it a little more general. Suppose the ditty said:

First come we, humans inchoate;
If it's knowledge then some of us know it

This is quite different. For it asserts - quite properly in my view - that the set of human knowing contains all of knowledge, and indeed that knowledge in the narrower sense of symbolic content with human importance beyond the local practices of individual lives must always be the object of knowing by someone. Put in more extreme form, knowledge is not like apostolic ordination in that once a human has known something it becomes knowledge forever. Rather it is like the parochial call in parts of Protestantism; something is knowledge while it is the object of some corporate knowing, but not otherwise. So my first lesson is that even this little ditty suggests that knowing is somehow more important than knowledge.

Now second, I want to draw a lesson from the FORM of my little analysis. It doesn't matter how you remember the ditty. It doesn't matter whether you think about Major Barbara, or Bernard Shaw, or Peirce. All the little memory

aids I spoke of are pieces of knowledge in some trivial sense: that Shaw wrote a play called Major Barbara, that the syllogism with universal premises and conclusion is the form mnemonized in the Middle Ages as Barbara, that Charles Sanders Peirce as philosopher inevitably had papers involving the words "mnemonic" and "syllogism" and inevitably commented on Jowett's new translation of Plato and inevitably knew someone named Shaw (as it happens it was not Bernard Shaw, but the great librarian John Shaw Billings). And I hope that the randomness - indeed the silliness - of this knowledge will persuade you that neither medieval way of remembering things - which gives us the ditty "Barbara, Celarent, Ferio, Darii" to remember the valid forms of syllogism - nor the modern way of remembering things - which is to cruise the internet with search engines that turn up finds us references to John Shaw Billings when we are looking for George Bernard Shaw - is of any real help in remembering the actual IDEA that I am talking about, which is that there is something fundamentally different about knowing and knowledge, at least as we usually use those words, and that figuring out what knowing is may be more important than figuring out where knowledge is going.

Put another way, my second insight is that there are words and there are ideas, and we should not mistake the one for the other. Medieval mnemonics and modern concordance indexes are just "word junk" in the way of the idea. It is the IDEA of knowing that we want to encounter and to remember. Yet the word "knowledge" looms so large in our thinking about this whole area of affairs that it seems completely natural to say "the future of knowledge" while it seems a little odd to say "the future of knowing." We can't see the idea that I'm trying to point to with "knowing" without getting caught up in the simple word "knowledge." Indeed, we turn out to have trouble imagining the IDEA that is behind the word "knowledge" itself. Certainly our students do.

So you see the problem. We have to talk about what is happening in the world of ideas - of books, libraries, databases, kindles, search engines, and

so on - without allowing our vocabulary to predirect our thinking. Once we free ourselves from the trammels of customary language, we can see that the whole debate about knowledge and information, about knowledge automation, about bringing all knowledge to everyone, is a sham. Knowledge is not a thing; it is the result of an activity. That activity is knowing. The internet has changed nothing whatever about that activity. It has changed the conditions under which that activity is done - the setting, if you will, of knowing. But it has not changed knowing itself, either as a personal activity nor, I think, as a corporate one. If knowing is changing, it is because those of us who know, and who teach the practice of knowing, have lost our bearings, and we are failing to teach knowing to the next generation. In this talk, I want to remind us what knowing is.

Now you have all heard and experienced the new world of information or knowledge or whatever you want to call it. And I could go ahead and comment on the large-scale changes in that world: both the macro-scale events of digitization and knowledge overload and the meso-scale changes in knowledge production, in publication, and in knowledge use or consumption. But I think what is probably of much more interest to you are the changes that I see at the micro-level, in the practice of knowing as it occurs among undergraduates, and graduate students, and - we can talk about this in the Q and A - among my colleagues on the faculty.

In the past year, I spent much time researching practices of knowing at the undergraduate and graduate levels. I don't mean that I gave surveys. You can't learn about these things from surveys. The students are highly selected as master test-takers, and a survey is for them just another test. No, you have to do ethnography. You have to be in the classroom, and make interventions, and try crazy expedients, and solicit continual commentary and response, both anonymous and identified. Then you find out what they are

actually doing, when they are faking, and when they do not understand what you are talking about.

This last problem is actually the reason surveys don't work. It turns out, for example, that most undergraduates don't actually know what the word reading means to person of my age. They think it means parsing the sentences of a text and rendering them into internal representations. They don't even imagine MY concept of reading, according to which reading means using a text as a stimulus for a complex reflection; parsing may or may not be part of that.

Or take the word index. To students both undergraduate and graduate, this word means search by text words - what is properly called concordance indexing. Most current students at both levels are quite astounded by the fact that books had indexes before the era of searchable text, and they believe - this is no joke - that back-of-the-book indexes before searchable text were produced by people who laboriously searched the text by hand for the appearance of certain words. That indexers searched the texts not for words but for ideas is quite literally not imaginable to them. They don't have any concept for subject indexing; if you use the phrase "subject indexing" in class they assume you mean search by keywords. This is not hyperbole. It is a simple fact. Our students - among the best in the land - literally do not have in their minds a general concept of indexing. The word "index" means nothing to them but computer-based word search.

Let me then talk a little bit about concepts of knowledge, knowing, and so on among our students. I should start by pointing out that our present first-year undergraduates probably started using the internet in elementary school; thus their entire educations have taken place in an electronic world. Their high school research assignments sent them to the web, not the high school library, and indeed the web probably had better resources than their

high school library. While doing those assignments they instant-messaged their friends constantly. Reading and research, that is, were done in a multi-tasking environment, in parallel with friendly conversations, ordering stuff from Amazon, and so on.

These things are only slightly less true for our first-year graduate students. Most of them spent their college years on the internet, and the vast majority of them have no more idea how to use a library or, in fact, how to use the internet in a sophisticated fashion, than they do of how to fly to the moon. They are unaware of basic issues of provenance, quality, and organization, and they have a touching faith - a more appropriate phrase would be an astonishing naivete - about the electronic research universe, whose quality, accuracy, and good intentions they in fact take for granted, although as we shall see they think themselves very sophisticated users indeed.

So let's spend some time discussing knowing in the wild. I report first from my undergraduate core course in the fall. This was a class of 16 students, to whom I was to teach Adam Smith, Karl Marx, and Emile Durkheim. I had concluded from previous classroom experience that I should forget the content and focus on teaching them the more basic skills of how to read and how to think. But I had little idea how difficult this would turn out to be in practice. So I gave them questionnaires and a wide variety of personal assignments, which produce the data I will give you over the next few minutes.

A first task I set them was the task of describing a book solely in the terminology of the internet. I was assuming that they understood the internet and wanted them to reflect about how a book is different from - but not completely unintelligible as - a web structure. I got some surprising results.

After "sifting through the meaningless and occasionally incorrect information available to me in a peer-edited article," one student was

"intrigued to find that the movie "A Clock Work Orange" was based on a book by an English philosopher." So first of all, for him the book had already disappeared behind the movie. Second, although my student thinks he can tell reliable information from garbage he nonetheless believes that A Clockwork Orange was written by a philosopher, when Anthony Burgess is a writer who had previously been a composer and a civil servant.

Another student, writing a brilliant parody of internet lingo, which sadly does not reproduce well aloud, tells me "How can the cumulative knowledge of the ppl, updated by the second, ever be wrong?" Another tells me "because of the inconveniences of page-turning and lack of a "search" function, the paper book can truly be called a time-wasting version of the internet webpage."

One particularly wonderful response gives an ethnographic account of reading The Great Gatsby online. I would love to read you the whole hilarious thing, but will read only that point where he describes what I learned this past fall is now the standard way for young people to "read" on line:

I find a section of the story that I feel is important so I decide to highlight it and press ctrl-C to copy it and then I open microsoft word and press ctrl-v to paste it. I continue to do this until I feel I have a good set of notes to study for the night as a summary of the story. I finish a page and there is a link at the end of the page to connect me to the next chapter. I double click it but before I can go on to the next webpage I am shown a Google ad with an opportunity to win a getaway cruise ship online. Reading a hardcopy of the novel would have saved me from this absurdity.

I want to emphasize that what you have just heard is not a parody. This is how most first year students "read" texts: they try to extract a set of words from the text itself that serve as summary, which can in turn be memorized. The text is reduced to bits, the bits are copied somewhere, and the result becomes a possession of the reader, his "reading" of the text. Many of them have been taught to write papers this way in high school - my own son was so taught at the Lab School, by teachers obsessed with "getting enough evidence."

It turns out that many students write papers by first selecting

appropriate quotes, then interspersing one sentence of interpretation between each quote. Since the students haven't spent any time trying to regenerate the book's actual argument in their own heads, such papers are gibberish. Half of my first paper assignments were of this type.

This view of the text - that the key to itself is some subset of the sentences written in it - is what drives the earlier quote saying that a book is useless because it doesn't have a reader driven search function. The key to the book is written somewhere in it, it consists of a sentence using certain words, and it can be found by finding those words. Many students feel this way about books, for the obvious reason that this is how one finds pages on the internet, and, beyond that, for the less obvious reason that finding things, and in particular finding things on the internet, is their principal model for cognition itself.

This pattern continued in the questionnaire responses. I asked them "In your current work, how do you physically analyze a text - underline? highlight? take notes? prepare an outline?" 18 of 21 told me they underline or highlight. 17 take notes separately or write notes in the margin. 12 do one of each of those (underline/highlight AND take notes/marginalize). Only six underline/highlight AND create a separate outline of the text, and only one person did all three.

Once again the pattern is to passively parse the text into important bits and less important bits. That is essentially what underlining does. And I can tell you all about their marginal notes. I looked at them. They're pathetic. It is only the outliners who are really getting the structure of the text.

This became woefully clear on the second day we spent on Adam Smith, Thursday of second week. I was wondering whether they retained anything from Tuesday's readings. After all, Adam Smith has an argument, and if you don't remember the first sixty pages of The Wealth of Nations you aren't really

going to make much sense of the next sixty pages. So I stopped class ten minutes before the end - we were having a discussion that wasn't going anywhere anyway - and I asked them all to write a six-sentence summary of the readings for the preceding class meeting - the first sixty pages of Wealth of Nations. I said I wanted no names, just a measure of how much they were retaining. After they recovered from their shock, they dutifully wrote out summaries. (Aside) I do have to say these kids were a terrific class, who learned much and put up with all my experimentation.

The TA typed them all up and we handed them out in the next class. Of course the students split their sides laughing, as the TA and I had before them. The summaries were full of vague recollections, lunatic errors, and minor asides. (My favorite was "Smith does not like haircuts, because there is no continuation after the expense.") Of course the ideas were all related to Adam Smith, most of them were substantially correct, and, indeed, the core arguments of Smith could be discovered here and there. But nobody came close to being able to give a step-by-step summary of Smith's argument, even though these were readings done for a class only two days before, in a book that is quite evidently a progressive argument.

This is not because these are stupid students. They aren't. Nor is it that they are unmotivated or that I'm a bad teacher. They simply don't understand that books have arguments and that arguments have logic and direction. The internet has taught them that you can enter a text anywhere for any reason. So they take that lesson on to their college courses. And they approach reading as an extraction task, to be done once for each set of readings. And as I noted earlier, they believe that the crucial statements are always in the text itself. With Adam Smith, they were lucky - a lot of the crucial arguments are indeed summarized somewhere in the text. With Marx, they would not be so lucky. But by then they had learned something about reading.

I decided that I would force them to read slowly by demanding that for

each reading they select two or three sentences, memorize them, and then meditate silently on them for fifteen minutes outside of class. We modeled this in class with one page of Smith, and it seemed to go well. Furthermore, I demanded that at the beginning of each succeeding class, they would write down their meditation phrases from memory, write the page number, sign the paper, and then turn it in. Every student thus came to every class knowing by heart at least a couple of sentences from the reading.

For well over half the students, this experience was transformative, as I found out when I surveyed them at the end of the quarter. Several had switched to outlining. Many had REDUCED their underlining. "I have begun reading texts more slowly and carefully," one student said. "I have had to tone down my underlining because it was compulsive and I was more focused on what to underline than on the implications of what was underlined." Another said "I use post-its with more frequency now and have abandoned writing in the book." These are students who have begun to realize that reading is not a process whereby they simply select text, but rather a process of knowing wherein they must engage with, reflect about, and respond to the text.

That this experience was transformative doesn't tell you so much about me as a teacher as it tells about what these kids are like when we first see them in the class. They are smart, multi-talented, and highly motivated. But they have no real reading skills at all. Moreover, they don't know that they have no reading skills, but think quite the contrary that they are pretty good with texts. But their model of reading and indeed of knowing comes from the internet and is worthless with complex texts.

One sees this the more clearly - and becomes more able to teach against it - by reading about website design. Just go to the web and look up sites on good site design. These sites will tell you that users "don't read, they scan." They "don't make optimal choices." They "want to have control" and

"follow their intuition." Web design sites tell you to keep the page length short, to divide pages into "logical chunks of information" so that each page is "information on a single topic." Site designers are told ALWAYS to allow cross movement and NEVER to force a reader to follow a prearranged, hierarchical structure. (In short, that is, DON'T have an argument.) Design sites will tell you that "users use a site like they shop at a store." Therefore, don't make them think, don't squander their patience, focus their attention, don't have long blocks of uninterrupted text, strive for simplicity, organize the text so that it is simple, clear, and distinctive. Can you imagine such a recipe for a core text? This might be a description of Wealth of Nations, although even it makes readers think and has long blocks of uninterrupted text. But it sure isn't a description of the first volume of Marx's Capital.

My point is that our students have been brought up spending much of their time - the time that we spent reading magazines, second-rate novels, and the occasional piece of fine literature - surfing an internet that has been optimized in terms of these retail-oriented principles of web design. That's where their model of cognition is formed. Ours was based on rubbish texts, to be sure. But at least they were texts. The current generation of students has been raised on a cognitive form that is deliberately designed to be as indulgent, as "user friendly," as preorganized as is humanly possible, all in order to hold the reader's attention long enough to sell him something. This is not a model of knowing, but of selling, and what is sold is - by definition - commodities. And indeed, it is very helpful to realize that our students tend to understand knowledge as a commodity. That's yet another reason for focusing them on knowing rather than knowledge.

The students very much WANT to escape from this commodity approach to knowledge. I saw this clearly when I set up my assignments for them. Because I found the meditation system to be working well, and because I didn't look

forward to reading sixteen slightly wrong versions of "what was the concept of human nature in Adam Smith?", I told them the first paper assignment was simply to take their five meditation passages and turn them into an argument about Adam Smith. Of course the result was a train wreck, but it was a good train wreck. Here was a paper that demanded that they simply think, not that they cruise Wikipedia and the web for ideas. This was their own bed and they had to lie in it. (Indeed, it was around this time that I surveyed them for their use of the internet with respect to my course and was pleased to find that few of them had used it for anything. I had designed the course to guarantee this, and was pleased to see that it worked.)

When I got the papers, what they showed was that the students had no idea whatever how to perform such a task, to look behind their five passages and see common themes and ideas. What they can do is parse texts, find crucial sentences, and replicate them. If they have trouble, they can find someone else who has done it for them. Their concept of knowing, once again, is finding something, knowing the URL where the answer is. This is not some derogation on their part; it is the natural result of the cognitive environment in which they are growing up. We have to recognize it and adjust our pedagogy to it.

The fact is that the students ended up LOVING my kind of paper assignment. It felt personal. It was their paper not mine. It gave them confidence. It taught them skills they didn't know themselves capable of. It's for that reason that I say that they WANT to escape from a commodity approach to knowledge.

Another piece of evidence about their ways of knowing came from notebooks. I took my own idea notebooks to class on the first day and showed them how I develop my own ideas. My notebooks are books of short, one-page, hand-written essays on various topics. At the top of each page are written the subject

topics that I feel capture the ideas on the page, and I enter the page numbers immediately under the subject terms in an ongoing subject index at the back of the notebook. As a result, I always have a continuous, subject-indexed record of my thinking about various topics. I asked my students to do the same, and three times during the quarter I collected these notebooks and read them.

The result was disappointing but revealing. They are clearly not used to doing this. They want to read a text then put it away, not reflect about it. It is difficult for them to elaborate, to let themselves go. They do not seem able to just think, to speculate, to dream on paper. Nor can they see that this kind of reflection is necessary to produce real argument. My own sense - I don't have evidence here - is that this is because they are in a rush. They envision knowing - or perhaps their processes of knowing - as being efficient. They want to accomplish the most in the shortest time. They want to calculate the minimal necessary understanding of a text, then find out the minimum time necessary to achieve that. Serious reflective thought is something they simply have not yet envisioned. It is not part of their experience, except perhaps about personal identity. These notebooks, then, show how far they have to go.

Finally, I should note the evidence of what students wanted from me. At the end of the class, I asked about their preferences for formats of class meetings and found that the formats they most liked were the few early in the quarter - before I went on my experimental rampage - in which I forced them to produce, under my leadership, lists of key terms and building blocks for Adam Smith's argument. I quickly realized that they needed to learn to do this for themselves. I was just serving them as a personal version of wikipedia, even though I was only emphasizing what WERE the important terms, NOT necessarily how they related to each other. But even that, I felt, was too much help for them. So I stopped. But although they wanted the lists of key terms and building blocks, they didn't realize that they had to produce them whether I

helped or not. Once I stopped leading them through it, they stopped doing it, as the notebooks showed. Another year I shall have to focus on that. As we shall see with the graduate students, this is a lesson they too need to learn.

So we see in summary that typical undergraduates have a process of knowing that is very internet based. It is intuitive, disorganized, non-hierarchical. It presumes simplified texts preadapted for use by simple minds. It believes the world of thinking is organized in discrete and almost arbitrarily-related bits of knowledge. It takes "finding things" as its basic model for knowing. It is not particularly reflective, and can only be driven to reflection by the coercion of meditation.

Such is knowing for today's undergraduates. Their strangeness makes us all the more aware of how unusual and specific is our own scholarly mode of knowing. In years past, students grew up in modes of knowing far closer to those of scholars than do these our students today. It is a curse, in that we must teach them the very basics, but also rewarding, because it makes us much clearer about what good knowing actually is.

Let me turn now to the graduate level. The story here is different, but nonetheless similarly alarming. Again, the state of their knowing makes us all the better aware of what it is that we should be teaching.

In my graduate course on Library Research this past winter, I had 22 potential students and kept 16. (Thanks to Judi Nadler, by the way, I was able to teach this class in the library. That's great.)

These students were strongly committed to library work. Prior experience had taught me to require that students commit ahead of time to a course project, since otherwise weeks are lost floundering around. So students took the course only if they were able to walk in with some idea of their research projects on day one. Half were MAPSS students who had decided on library-based MA papers; bear in mind that MAPSS students must take all their courses AND

produce an MA paper in nine months - so these students were expecting the course to jump-start their work. The other half were second to fourth year PhD students, most of them aiming to produce a qualifying paper or similar major output. So again they were intensely committed. In summary, we can safely assume that these students are representative of committed graduate students aiming at library- and net-based research in the humanities and the humanistic social sciences.

As in all my classes, I started off with a questionnaire. I asked first about general familiarity with physical and online research. A little over half of them thought themselves only "basic" at physical library research, while most thought themselves pretty good or highly skilled at on-line research. Their essays for me at the end of the course made it clear that they would now say that both of these were considerable overestimates.

This high-self-assessment was echoed in the facts that half of them had never asked a JRL librarian a question and that most of them had not heard of or had never used LENS - the fancy word-cloud, metacrawler research tool beloved of some of our younger librarians. It would also turn out - I didn't know to ask this at the time - that most of them had never heard of, much less used, the wonderful subject guides to research that are available one step off the library's top page. I wondered about this, but it seemed that they didn't feel they needed all these aids. I would later conclude that an odd combination of pride and diffidence had prevented their reaching out to these librarians' efforts.

On the other hand, there WERE some chinks in the armor. Although some of them had used the library's recommended REFWORKS bibliographical manager, nearly half used no reference manager at all. Most of them still organized their hard drives by course - the pattern characteristic of undergraduates - rather than by project, the pattern characteristic of scholars. There was also

a sense of desperation in their responses to an open-ended question on "what do you most want to learn." Although there were some predictable things like "Archives" and "Government Documents", there was a disturbing number (five or six) who mentioned "narrowing down bibliography," "learning what NOT to read," "identifying rubbish sources," or some other skill for handling overload. A couple asked for advice on the organization and management of projects.

So the typical student enters such a course thinking s/he needs to learn about the physical library, but is already pretty good with online materials, yet that student turns out not to know the library's own on-line overview tools, has worries about organization and overload, and wants to know how to ignore things. This is a distinctly mixed message.

That message gets bleaker when we ask on what detailed basis the supposed online skills stand. The answer is - not much. In the online world, they reported relying on a very narrow range of tools. Most of them have used JSTOR, a high quality source, as well as Google Scholar, which is a piece of junk. About half of them had used Lexis/Nexis and Worldcat. Only half said they had used the local catalogue, but that's probably because they didn't know its actual name. Standard, high-quality online sources like the various Proquest databases, the Wilson databases, and Web of Science were almost completely unused. There were, to be sure, some oddments here and there. One or two might have used Anthro Source or the MLA search website or World Biographical Information System. But the standard high quality tools were unused by the vast majority.

For a limited number of tools, I asked "have you heard of this and if so have you used it?" It turns out they've all heard of the purely electronic tools: Google Books, Google Scholar, Lexis/Nexis, and Worldcat. Among the tools that were once print but that are now electronic, their knowledge was considerably weaker. About half had heard of Historical Abstracts, Dissertation Abstracts, Sociological Abstracts, and Web of Science, but things like

WilsonWeb and World Biographical Information System were known only to three or four. As for the purely print sources - things like PAIS Bulletin (actually, it's now electronic), the CRL catalogue, the NUC pre-56 imprints, the Union List of Serials in the United States and Canada, and so on - they're simply unknown. That the Union List contains an exact print history of every single pre-1950 periodical held in American libraries was completely dazzling to them.

Indeed, this experience captures a general belief (among graduate students, and even stronger among undergraduates) that online sources contain all the information available in print sources and that therefore print reference sources can be safely ignored. They believe this with the quiet, secure ardor of religious converts.

As a result, throughout the course, the astounding richness of print reference sources was an eyeopener to them, culminating in a moment in the map room, where I had taken them to meet Christ Winters, our map maven. I made them each think up - on the spot - a way in which maps might be relevant to their projects. Dan Huebner, a splendid student who is studying the collation of various archivally held students' course notes into the book that we now know as *Mind, Self, and Society* by George Herbert Mead, said in a casual voice that "I don't have anything related to maps unless you can tell me what was located at 65xx Kimbark Avenue in 1929." Chris thought a minute, and said, "well, you might use city directories, but, but, wait a minute, you could use the Sanborn Insurance Maps, because they always tell what usage a given location has at any time," and he pulled an old volume off a nearby shelf with the answer to Dan's question. The class simply stared with amazement. Calm, unflappable Dan was like a young man who had found salvation. This kind of experience happened again and again. The students simply have no idea what riches prior researchers have produced, and so they don't even know to look

for them.

On the questionnaire, I also asked some questions about basic library familiarity. Nobody knew what the A classification in LC stood for; those who thought they knew were wrong. None of them knew where the map collection was. None of them could name four important producers of comparative international data. Only one or two knew where the political science materials were. Most did not know where to find old copies of the Maroon. Most did not know what is the library of Congress heading for bibliographical materials. Only two knew who any of the social science bibliographers and reference workers were. Only five knew where the U of C dissertations were. In general, that is, they had NO idea of what is where in the library, no working sense of the materials there, nor any idea that the librarians or the physical materials could be of real use to them. The physical library and even the library of congress classification were terra incognita.

These responses suggested organizing the course around a familiarization with library materials and search strategies. That was fine with me, for I had taught the course five times before and had usually organized it that way; some early classes on general topics like setting up an overview and doing a bibliography, followed by classes on types of sources like archives, government documents, old quantitative data, biographical data, and so on.

Yet after only two or three weeks, I realized that the problems were much deeper than I had thought. The students were good at finding things, but bad at judging them. They were hoarders of references, misers of knowledge, unable to distinguish important sources from second-rate trash, unable to find their way amid the long lists of irrelevant mush turned up by Google and its friends. More important, they didn't have any idea of a BASIS on which to judge sources. They imagined research as an enormous stamp-collecting exercise, as if the end result would be the fullest binders of research "knowledge."

That is, they had no idea of how to conceive a research puzzle and to use that puzzle and the questions following from it, together with their own good judgment, as criteria for keeping or rejecting sources, for following or ignoring research trails, for working their way through a project to a final text. When I realized this, I completely restructured the course. I gave up the focus on "how to find it" and "places things can be found" and focused the course on how to organize and govern a project. I had always done some of this, but only in the background. I moved it to the foreground, and the course came suddenly alive.

Here I want to rephrase this insight in the terms of my theoretical arguments at the outset. I had been taking it for granted that the students understood that our quest for ideas is a matter of knowing, not knowledge. But even these graduate students thought - just as the undergraduates thought of the smaller task of reading a single classic text - that knowing was a matter of searching and finding. This is hardly surprising because their whole lives have trained them to think this. The internet is presented as the great repository, where everything can be found. Knowledge is stuff, and knowing consists of possessing the correct URL. The best source is out there somewhere and need only be found.

But of course, you can't recognize the best source unless you have a criterion by which to judge it, and you can't have a criterion unless you know "best for what?" Nothing, that is, is knowledge in and of itself. A given piece of information or interpretation is knowledge only with respect to a particular project of knowing. Until you have your project figured out, you can't say what knowledge is.

Now it turns out, of course, that your project shifts over time as you find things out, just as the things you find out shift in accord with the logic of the project at a moment. This non-linear, recursive quality of

library research (and of sophisticated web research) was very disturbing to most students at first, but came for most of them to be the central lesson they took from the course. As I kept telling them, the trick of library research is not finding things. It is knowing, when you happen on something, that you ought to have wanted to find it.

Indeed, we can gain further perspective on these students' concepts of knowing and knowledge by looking at the things they identified, at the end of the course, as the most important things they had learned. We can put that material together with my own conclusions from my day to day teaching experience to produce a sense of the image of knowing that these students start out from.

In the first place, they think the pre-electronic library world was prehistoric and ignorant. They believe this genuinely and without rancor, thinking themselves very fortunate to have been born in such wonderful times. For example, they think tagging is a wonderful novelty, even though what it really means is a downmarket version of what their elders meant by ongoing indexing. Tagging sounds hip to them, while indexing does not, precisely because they don't know that the former is actually a new name for the old version of the latter. Like all young people, they believe that something has to be young and hip to be good, and under its new name subject indexing is hot stuff, a fact that is hardly surprising.

Similarly, they think feeds are revolutionary, and are astounded to discover that there were things called clipping services a hundred years ago that delivered exactly the same functionality. Or they are astounded to discover that the circulation cards in the back of books made it obvious who had taken a book out before, providing a way to identify local peers who were working in one's area. They thought Amazon had invented such things with its "people who bought this book out also bought X" functionality.

Not only are they innocent about the power of prior library tools, they

have - perhaps this is a good thing - no idea whatever of the sheer mass of research, both pre and post the electronic revolution. For example, they have no idea that once you get out of the mainstream stuff our eight million books don't really overlap much with the eight million at Michigan. More generally, they have no idea whatsoever of the amount of research that has already been done, since they almost never check dissertations and other non-published sources and since they infer from the rhetoric about the digital revolution that no serious research can ever have been done before.

So their sheer ignorance of the past leads them to believe in the information revolution, when in fact most of the electronic tools are rediscoveries - in some cases second-rate rediscoveries - of the wheel. Their innocence about print materials and tools extends to those print-era tools designed to handle information overload, the one problem they most fear. For example, they do not have the concept of controlled vocabulary, and more broadly, of any set of fixed theoretical terms that will channel and focus their thinking. They even lack the equivalent in terms of empirical categorization; it is not the case that their first instinct with a long list of empirical facts is to organize that list into a structured system of categories. To me as a scholar, this implies the even scarier fact that they don't really have the concept of a rigorous analytic vocabulary.

Furthermore, they don't know that different sources have different controlled vocabularies, a fact that destroys the supposedly revolutionary character of mass electronic amalgamations like Nineteenth Century Masterfile. Nor do they realize that keyword searching returns unsteady results over time because of the drift of meanings of words.

As for the middle phase of library work, they in fact don't really believe in randomness, for all that they are supposedly the web-surfing, try-my-luck, networky generation. They actually believe devoutly that sources like

Google really do deliver what they claim to deliver, and use them as if they were truth machines. Paradoxically, it is precisely because no one has bothered to train them (and in part because they think they don't need to be trained) that they are petrified by the sheer volume of stuff on the internet and hence are not SKILLED at randomness, at effective browsing nor at following random chains quickly into whole new areas. Their notion of deepening their research is trying some new combinations in a keyword search query.

Luckily, we had an "aha" experience in this area as well. Brian Cody decided to write on the ways the US government gives out money, in particular about block grants. And he had found interesting material in Lexis/Nexis. He told us the idea of block grants went back far further than we thought. It wasn't just a Reagan invention, it went back to the 1930s. The only problem was that there was plenty of material in the 1980s, the 1970s, the 1960s, the 1940s and the 1930s, but nothing in the 1950s. Try as he would, Brian couldn't find anything.

But he's a New College grad - the University of Chicago of Florida - so he went after the problem. First he found a 1950s congressional speech sure to contain the phrase "block grant." Then he read the actual document and - sure enough - found the phrase. So it OUGHT to be found by the keyword query "block grant." And he ascertained that the keyword "grant" would locate this document, but the keyword "block" would not. So he tried various misspellings: black, blank, plonk, prank, blink, and so on. He hit paydirt with "blook." It turned out "blook grant" returned dozens of documents from the 1950s, all of which had in them the phrase "block grant." "Blook grant" of course, was an optical character recognition error. And it turned out, Brian discovered triumphantly, that the Federal government had changed its font around 1950. So the OCR algorithms, which are AI based, trained themselves on the old font,

and then couldn't read the new font when it showed up.

More than anything else that happened in class, this discovery finally persuaded the students that electronic sources are relatively unreliable. I already knew this, of course - Web of Science reports 1325 citations of my first book, and only 80% of them have the title correct. There are actually forty different erroneous titles for it, although you can be sure that in fact all of those who cited me got the title of the book right in their print versions. It's all OCR mistakes.

But it was news to the students that the electronic tools are not perfection. They had no idea that however comprehensive they may be, they are generally less accurate than the print sources that preceded them. Even less did they suspect that those inaccuracies could be systematic. Somewhere out there, they now realized, some idiot is writing a paper on how the concept of block grants disappeared from American political discourse in the 1950s. It is the fear of becoming that idiot that baptized them as serious library researchers.

Finally, and perhaps most important for my theme of knowing versus knowledge, the students didn't come to class with any sense of how to manage research. They don't know how to maintain files, and to keep to-do lists, and to constantly produce lists of open questions, and to endlessly rewrite their proposal documents, and so on. That is, they don't understand what a businessperson would call project management. Again, this is because they think that knowledge is a static thing, and that once they "have a source," usually by downloading it, they "have" what they need, because they possess it. They don't realize that the little facts they have in their files constantly change their meaning as the project evolves and defines this or that one of them as more or less important. Indeed, they knew so little about knowing as a dynamic activity that they didn't even know to keep a log of

everything they did so that they didn't end up doing things two and three times. Keeping logs seemed a revolutionary advance to most of them.

All of this takes us back to Jowett and the Balliol undergraduates. Those 19th century rhymesters saw that knowing was more important than knowledge, that it is knowing that makes knowledge, not the other way around. But I said I would talk about the future of knowing, but here I have spent all this time talking about its present. The reason is obvious. The kids are the future of knowing. We faculty can do anything we like. But if we don't figure out what knowing is and start to teach it, our students will have to spend decades figuring it out for themselves. We must get serious about understanding the present of knowing, so that it will have - a future.