Information Synthesis: A New Approach to Explore Secondary Information in Scientific Literature

Catherine Blake
School of Information and Library Science
214A Manning Hall
University of North Carolina
(919)843 5636
cablake@email.unc.edu

ABSTRACT
Advances in both technology and publishing practices continue to increase the quantity of scientific literature that is available electronically. In this paper, we introduce the Information Synthesis process, a new approach that enables scientists to visualize, explore, and resolve contradictory findings that are inevitable when multiple empirical studies explore the same natural phenomena. Central to the Information Synthesis approach is a cyber-infrastructure that provides a scientist with both primary and secondary information from an article and structured information resources. To demonstrate this approach, we have developed the Multi-User, Information Extraction for Information Synthesis (METIS) System. METIS is an interactive system that automates critical tasks within the Information Synthesis process. We provide two case-studies that demonstrate the utility of the Information Synthesis approach.

Categories and Subject Descriptors
H.3.1 [**Information Systems**]: Linguistic processing; H.3.3 [**Information Search and Retrieval**]: Information Filtering; H3.3.7 **[Digital Libraries]** Dissemination, J.3. [**Life and Medical Sciences**]: Medical information systems.

General Terms
Documentation, Experimentation

Keywords
Text Mining, Information Extraction, Information Retrieval, Document Summarization, Information Synthesis

1. INTRODUCTION
It is rare that multiple empirical studies will report exactly the same findings, even when they explore the same natural phenomena. For example in medicine, empirical studies of patients that explore similar treatment regimes often produce different results. To manage and resolve contradictory evidence medical scientists have formulated a variety of techniques, such as the systematic review process [1,8,13]. The **Systematic Review** process has become a pseudo-standard to integrate evidence from biomedical literature.

The systematic review process requires that a user identify a comprehensive collection of articles, extract information from those articles, verify the accuracy of those extracted facts, and analyze the extracted facts using either qualitative or quantitative techniques. Although a systematic review accurately captures evidence, the process is time-consuming, taking 28 months from conception to publication [20] and 1139 hours [2]. With more than 12 million citations in MEDLINE and an additional 100,000 references added each year, the manual techniques currently used to retrieve, extract, verify, and analyze information for the systematic review process are becoming increasingly difficult to apply. Consider a breast cancer expert. It would be difficult, but necessary for her to consider the 12,600 articles typically published on breast cancer during the 28 months required to conduct a systematic review. Faced with the daunting task of sifting through currently available and recently added articles, our breast cancer expert may turn to other strategies to reduce the number of articles, such as constraining her hypothesis or her selection criterion. However, both of these constraints could introduce undesirable biases, and thus reduce the validity of her review.

The systematic review process is just one example of a more general class of information synthesis activities that users conduct to manage information in an information intensive environment. It is my position that the cyber-infrastructure of a digital library should support a user during their information synthesis activities. In this paper, I provide a specific example of how to achieve that goal.

In addition to supporting the existing systematic review process, the cyber-infrastructure provided by a digital library can provide a scientist with a new way to detect publication bias. **Publication bias** occurs when factors other than the quality of the study influence the acceptance of an article for publication [24]. One study that explored publication bias submitted several fake study results to a

JCDL'05, June 7–11, 2005, Denver, Colorado, USA.

variety of journals. The fake studies differed only with respect to the statistical significance of the result. Editors were three times more likely to publish the fake article with statistically significant results than the article that did not report significance [3]. One consequence of publication bias is that the published literature reflects an over-sampling of significant positive findings. Scientists in epidemiology go to great lengths to detect and avoid publication bias because such an over-sampling of significant findings would increase the perceived impact of a risk factor [4] [11]. In this paper, we describe how information synthesis can enable scientists to identify publication bias.

This paper is organized as follows. In Section 2, I introduce the Information Synthesis process using a detailed scenario. In Section 3, I provide details of the Multi-User Information Extraction for Information Synthesis (METIS) system that automates critical tasks within the information synthesis process. In Section 4, I provide two case studies that demonstrate the utility of the Information Synthesis approach.

2. INFORMATION SYNTHESIS

The Information Synthesis process uses the cyber-infrastructure provided by a digital library to extend the existing systematic review process in two ways. First, the information synthesis process uses a more relaxed inclusion criterion, than in a traditional analysis. Although combining studies with different study designs has been described as 'highly controversial' [22], and 'not generally appropriate' [16], we posit that such constraints are applied to manage information and that differences between findings based on study design should be explored empirically.

Our second extension requires a new characterization of scientific articles. We characterize information in a scientific article as either primary or secondary. **Primary information** reflects the main finding of an empirical study. For example, that Tamoxifen can be used to treat breast cancer. Our pragmatic definition of primary information is information that appears in the title, abstract, or keywords of an article. In contrast, **secondary information** provides the context of an empirical study. For example, the demographic characteristics of the breast cancer study, or the specific treatment regime followed. We define secondary information as information that appears only in the full-text of an article. Our second extension enables a scientist to incorporate secondary information into their analysis.

Consider a collection of scientific articles that report the results of studies with breast cancer patients. If a scientist were to conduct a traditional meta-analysis of breast cancer and alcohol consumption, they would start the process by collecting all of the articles that report both breast cancer and alcohol consumption. Using that traditional approach it is unlikely that our scientist would consider an article that reports a study on residential magnetic fields because the primary information contains neither alcohol nor a synonym of alcohol [9]. Yet that magnetic field article reports information that is relevant to her study of alcohol and breast cancer. It reports that 163 subjects with breast cancer never consume alcohol, that 636 subjects consume more than one drink per day, and that 14 subjects have unknown alcohol consumption behaviors.

On the remote chance that our scientist did find the article on magnetic fields during their traditional analysis, she would be unable to include the alcohol consumption information because the magnetic field article provides only the rate of alcohol consumption for breast cancer cases (the case rate). The article does not report the rate of alcohol consumption for the subjects who are cancer free (the control rate). Thus, in a traditional analysis, our scientist would be unable to incorporate the alcohol consumption details reported in the magnetic fields study.

Although the magnetic field study does not report a control rate, our scientist could combine information in the study with other information in the digital library to estimate the missing information. A digital library contains a variety of survey results from the Center for Disease Control, such as the **Behavioral Risk Factors Surveillance system (BRFSS)**. The BRFSS is a state-based telephone questionnaire that captures risk behaviors and preventive health practices in the United States [6]. Even though our scientist could not identify the exact person from the magnetic field study who participated in the BRFSS, she could use context information reported in the magnetic article to identify a similar population from the BRFSS respondents. **Context information** are factoids provided by an author that describe the study participants, the study design, and factors that the author knows could influence the study outcome. In addition to the alcohol consumption rate of breast cancer subjects, the magnetic field article reports the following context information:

(1) the number subjects with breast cancer;
(2) the number of subjects with breast cancer at each level of exposure to the candidate risk factor;
(3) the minimum and maximum age of subjects;
(4) the gender of the breast cancer group;
(5) the location of the study; and
(6) the time frame of the study.

The magnetic field study provides enough context information for our scientist to estimate a control group. She searches for BRFSS respondents who are between 20 and 74 women, living in Washington State between 1992 and 1995. One-hundred and eighteen BRFSS respondents satisfy these criteria. Of that sample, 50.1% consumed more than one drink per day. Thus, the scientist can estimate the missing information from other resources in the digital library. The estimate may not be as good as a deliberately collected control rate. For example, some of the BRFSS respondents would have breast cancer; however, this methodology would enable our scientist to explore risk factors that were not the main point of an article.

Enabling a scientist to include information from an article that was not the main finding may also help them to detect publication bias. The underlying premise is that primary information is subject to publication bias, whereas secondary information is not.

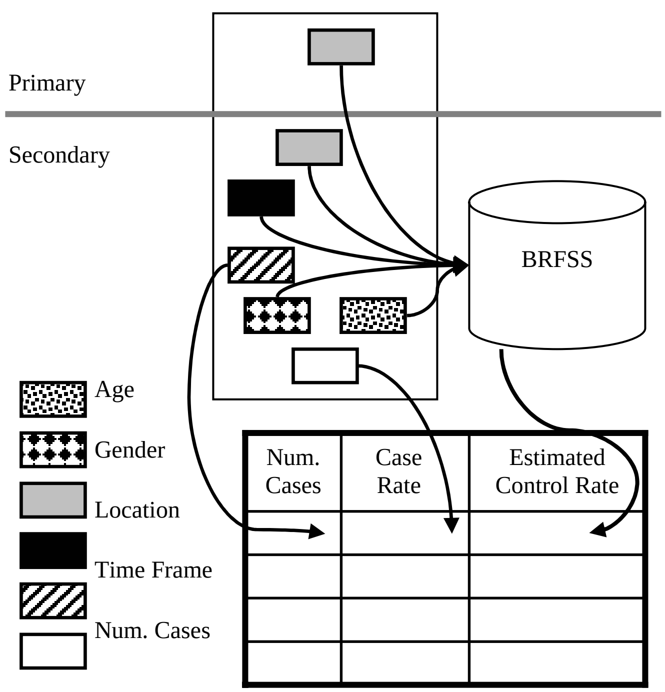

Figure 1 – Context information from each article to used to estimate missing information

The methodology described above (and shown in Figure 1) was used to quantify the association between smoking and impotence [23]. It took approximately six weeks to conduct the manual analysis on the 1008 impotence articles [Personal communication T.Tengs, 2002]. A similar analysis of the more than 100,000 breast cancer articles would take more than thirteen years to complete. Further, our scientist would need to consider the thousands of articles that were published as she conducted the review. Clearly, an alternative is required. In the following section, we show how the infrastructure of a digital library could enable such analyses.

3. METIS

To demonstrate the information synthesis approach we have developed the Multi-User Extraction for Information Synthesis (METIS) system. METIS automates critical components of the information synthesis process. In this section, we describe each of those tasks. The METIS Information Retriever identifies articles, the METIS Information Extractor identifies context information, the METIS Verifier enables a user to verify each extracted fact and the METIS Analyzer estimates the missing information and generates the visual summary.

3.1 METIS Information Retriever

To ensure a comprehensive collection, a two-phased retrieval strategy is required. In the first phase, the METIS Information Retriever identifies all articles that report the medical condition as primary information (in the preceding example the medical condition breast cancer). In the second phase, the METIS information retriever identifies all the articles that report the risk factor as secondary information (in the preceding example alcohol consumption). Searching primary information for both the medical condition and the risk factor would be insufficient because risk factor information could be reported either as primary or secondary information.

The digital library that enabled a user to search the full text would be invaluable during the Information Synthesis process. Unfortunately, such a digital library was not available, so the METIS Information Retriever automatically downloads full-text articles with a keyword *Breast Neoplasm's* using tools provided by the National Library of Medicine (nih.nlm.gov). To maximize the number of empirical studies the system excludes articles with a publication type of *Comment, Editorial, Letter, Review, Meta-Analysis, News, or Evaluation Studies*. We considered constraining publication types to include a collection of study designs, such as the *Randomized Clinical Trial*, and *Clinical Trial;* however, our preliminary analyses revealed that the publication type field was often empty. Thus, the system employs an exclusion rather than inclusion criteria.

A domain expert helped to focus our retrieval strategy by providing her pre-existing subscription to 29 breast cancer journals. Of the journals on her list, she classified ten as critical (nine were available electronically at our institution) and nineteen journals as non-critical (twelve were available electronically). Of the 3354 articles in MEDLINE that satisfied both the search and journal criterion, the system downloaded 1345 (40.1%) automatically. We then manually downloaded 596 (17.8%) that failed to download automatically. The full-text of the remaining 1413 (42.1%) articles was not electronically available at our institution. Thus, our breast cancer corpus comprises 1968 full text articles.

During the second retrieval phase, articles that contain the required context information are identified. The most selective context information for our breast cancer scenarios were the risk factors. Thus the METIS information retriever identifies the breast cancer articles that contain one or more of the following alcohol terms: *wine*, *beer*, *liquor (including hard liquor)*, *ethanol*, *drinking*, and *alcohol*. The second analysis used the terms *tobacco*, *cigarette*, *smoke*, *smoker* or *smoking/* Of the 1968 articles in our corpus, 568 articles contained one of these alcohol terms. We then manually identified the articles that reported alcohol consumption and tobacco consumption, and loaded them into a mysql database.

Other required context information are: (1) the number subjects with breast cancer; (2) the number of subjects with breast cancer at each level of exposure to the candidate risk factor (for our example alcohol consumption); (3) the minimum and maximum age or the mean age of subjects; (4) the gender of the breast cancer group; (5) the location; and (6) the time frame of the study. Articles that do not contain this context information are excluded from the analysis.

3.2 METIS Information Extractor

Our analysis of meta-analytic methodologies [15,18] and of users as they conducted meta-analyses [5] revealed that researchers required similar factoids, despite conducting analyses in very different research areas. Thus, our approach has been to identify context information that generalizes, rather than information extraction techniques that generalize.

The METIS Information Extractor uses shallow natural language processing to identify context information from each full-text article. We developed the semantic grammar using a small collection of 30 breast cancer articles. For example the semantic grammar to identify 1359 as the number of subjects from the sentence fragment *'In total, 1,359 breast cancer cases were included…'* is: pre-filler {}, slot filler {numeric value}, post-filler {semantic type=('*Neoplastic process*', or *'Disease or Syndrome')*, immediately preceding the term=(*'case'* or *'cases'*)}.

Our grammar employs semantic types from the **Unified Medical Language System (UMLS)** [19], a knowledge base developed and maintained by the National Library of Medicine. In the previous example, semantic types include '*Neoplastic process*' and *'Disease or Syndrome'*. We anticipate that these semantic types will enable the grammar to generalize to articles with different medical conditions.

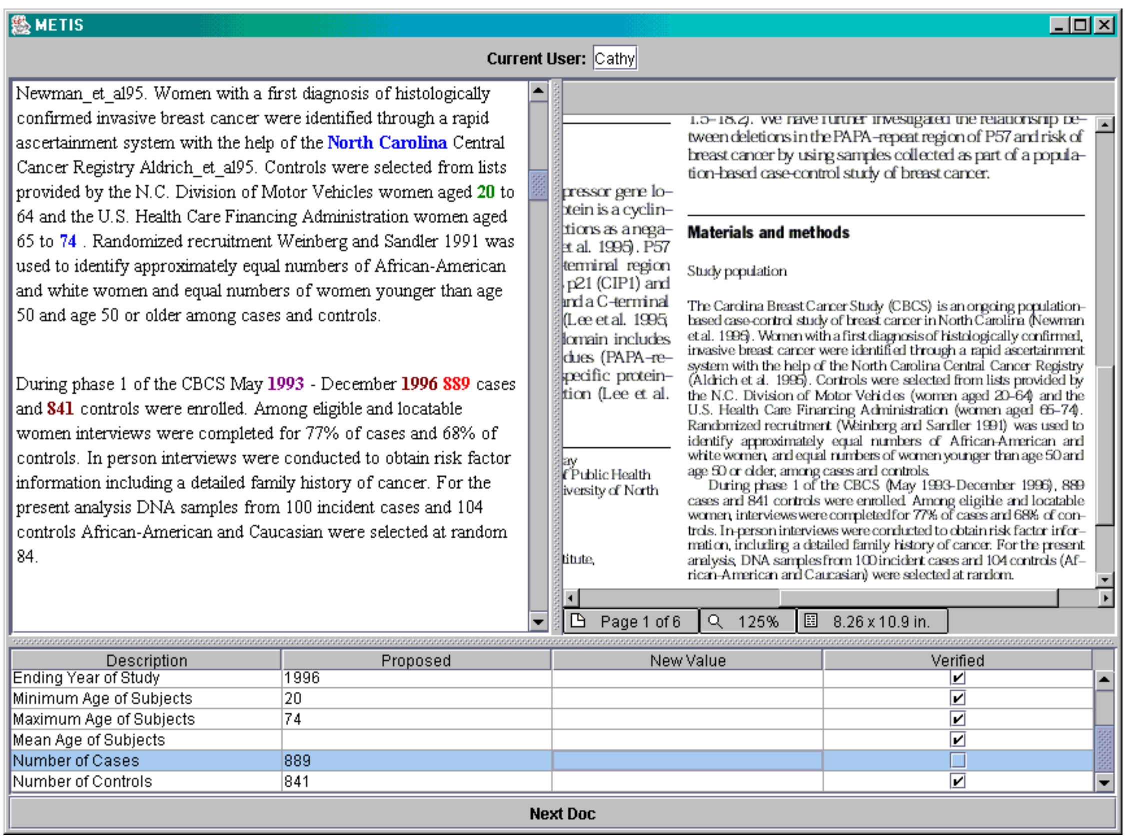

Figure 2 – The METIS Verifier Interface. The right pane provides the original document, the left pane shows generated text, and lower pane shows the verification progress of the user.

In addition to the pre, slot and post fillers, our extraction rules refer to the relative location of the factoids within the full-text document. The rule in the preceding section only identifies the number of subjects when the extraction rule conditions are satisfied, and the sentence appears on page 1 or 2. In addition to identifying information from prose, our extraction rules also operate on factoids that appear in tables using a combination of shallow natural language and two features that are specific to tables - row and column.

The current METIS information extractor comprises rules to identify 14 predefined factoids from each full-text article. Factoids related to the study include the number of subjects with breast cancer, the number of subjects without breast cancer, the start and end month, the start and end year of data collection, and the geographical location of the study. Factoids related to the population include the minimum, maximum, and mean age of participants, and their gender. Factoids related to the risk factor are the level exposure and the number of cases at each level. We also identify alcohol consumption, which we found was typically reported in a table.

3.3 METIS Verifier

The METIS Verifier enables a user to establish the accuracy of facts proposed by the METIS Information Extractor. The interface that supports verification is comprised of the panes (see Figure 2). The left pane comprises HTML that has been generated from METIS Information Extractor. The METIS Verifier shows each factoid in a different color. The right pane shows the document in the original PDF format. We show both because of conversion errors between the PDF and HTML versions. The bottom pane enables a user to accept or replace each proposed value. The last step of the verification process is to remove articles that refer to the same population group. In this prototype, we conducted this step manually; however, we are exploring the use of context information to automate the removal of duplicates.

A digital library could ease the Information Synthesis process by providing the full text of all breast cancer articles using a representation, such as ASCII, that avoids conversion errors.

3.4 METIS Analyzer

The METIS Analyzer is comprised of two components. The first uses verified context information from the METIS Verifier to estimate a control rate. The current implementation uses the Behavioral Risk Factors Surveillance system. The estimate considers the age and gender of the breast cancer subjects, and the time-frame and geographical location of the study.

The second component of the METIS Analyzer generates both an effect-size and the graphical representation of the available evidence **Meta-analysis** is “a statistical technique for assembling the results of several studies in a review into a single numerical estimate” [7]. Such techniques require that each study included in the analysis report a quantitative outcome. The result of a meta-analysis is an effect size, which is a unitless measure (see Figure 3). The METIS Analyzer uses a standardized mean difference as the effect size statistic because the facts from each study conform to this measure. Our variable names are consistent with the convention that a positive effect size indicates subjects who are receiving treatment (the case population) are more likely to have a positive outcome than subjects who are not receiving treatments (the control population). Thus, the effect size captures the difference between cases are controls.

```
Rate_Case^i ← rate of exposure in article i
Rate_Control^i ← estimated control rate for article i
Subjects_Case^i ← number of subjects in the article i
Subjects_Control^i ← number of database respondents for article i
Var_Case^i ← (Rate_Case^i * (1-Rate_Case^i))/ Subjects_Case^i
Var_Control^i ← (Rate_Control^i * (1-Rate_Control^i))/ Subjects_Control^i
for each article {  Var_Total += [1/(Var_Case^i+Var_Control^i)]}
for each article {
          newRate_Case^i = ((1.0/(Var_Case^i+Var_Control^i))/Var_Total)* Rate_Case^i
          newRate_Control^i = (1.0/(Var_Case^i+Var_Case^i))/Var_Total)* Rate_Control^i
}
for each article {
          TotalRate_Case += newRate_Case^i
          TotalRate_Control += newRate_Control^i
}
EffectSize = TotalRate_Case - TotalRate_Control
```

Figure 3 – Algorithm to Calculate Effect Size Based on a Random-Effect Assumption (Adapted from [10])

In contrast to a traditional meta-analysis, where a researcher would establish heterogeneity by using the Q-statistic, the METIS Verifier always assumes a random-effects model because we are deliberately using information from studies that report different primary information.

4. DISCUSSION

We evaluated the semantic grammar by comparing the METIS Information Extractor's proposed factoids with factoids identified by three independent reviewers. The highest ranked METIS proposed value achieved an average precision of 0.62, and an average recall of 0.59. To verify the accuracy of the estimated control group we compared the estimated value from the METIS Comparison Estimator with articles in the breast cancer corpus that reported control rates for alcohol and tobacco consumption. For each example, the METIS was similar to the control rates reported in each article. To verify the meta-analytic effect size, we compared the quantitative summary generated by the METIS Analyzer with textbook examples, and published meta-analyses. The effect size was the same. Details of those analyses are reported elsewhere[5]. In this paper, we emphasize the visual representation produced from the METIS Analyzer.

Figure 4 shows the visual representation generated by METIS to explore contradictions between studies that report alcohol consumption of breast cancer subjects. The representation extends existing meta-analytic results by providing both primary and secondary information sources.

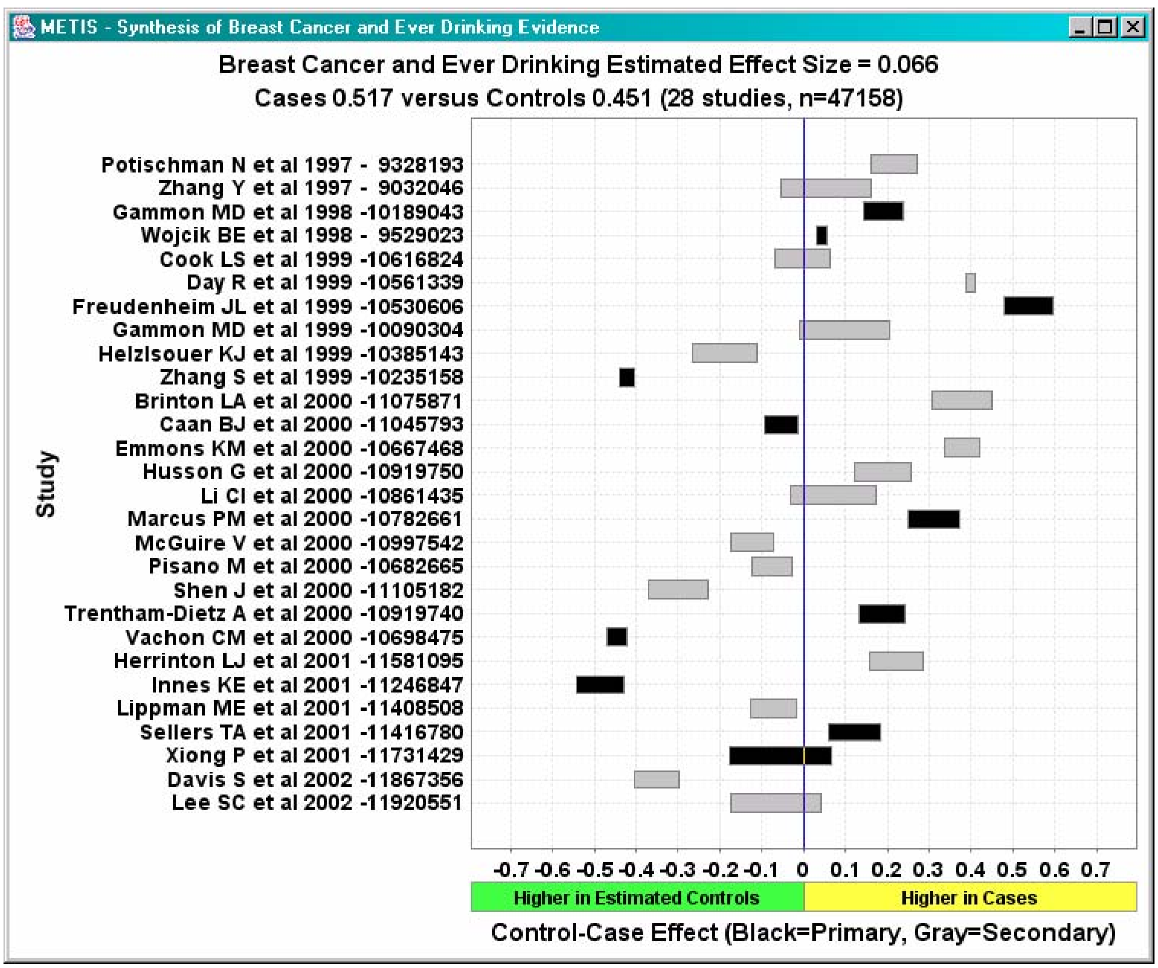

Figure 4 – The METIS Quantitative and Visual Summary for Articles in the Breast Cancer Corpus that Report Ever Alcohol Consumption as either Primary (black) or Secondary (grey) Information. The positive effect size suggests a correlation between ever alcohol consumption and breast cancer risk. Only the circled articles contain the information that is required to be included in a traditional meta-analysis.

Figure 4 summarizes the 28 articles that report the number or rate of ever drinkers from the subjects who have breast cancer. Each bar depicts an individual article, and you can use the identifier shown on the left to identify the citation within MEDLINE. An *ever drinker* is someone who consumes one or more drinks per month (one glass of wine, one short of hard liquor, etc). The effect size (which we call the quantitative summary), is shown at the top of the image and reveals that subjects with breast cancer were more likely to be ever drinkers than their corresponding control population (0.517 versus 0.451).

A **black** bar indicates that the article reported tobacco consumption as **primary information** i.e. the title, abstract, or keywords contained of the article contain an alcohol consumption term. A **gray** bar indicates that the article reported alcohol as **secondary information**. Thus, Figure 4 shows that fewer articles in the breast cancer corpus reported ever alcohol as primary information (11 articles) than as secondary information (17 articles). Users are more likely to include the articles shown in black than those show in gray because titles, keywords and abstracts are more likely to be used to index articles than the text within the article.

The mid-point of each bar in Figure 4 indicates the difference between ever alcohol consumption reported for subjects with breast cancer, and the corresponding METIS estimate. The length of the bar indicates the confidence interval at the 95 percentile. For example, the tight confidence interval of article 10561339 indicates the result of this study could dominate the meta-analytic result. Of the 47,158 subjects included in this meta-analysis, 11,064 participated in the study reported in article 10561339.

Figure 4 shows that five articles found a lower rate of ever drinkers for subjects with breast cancer than subjects who do not have breast cancer. These are articles are shown as black bars with midpoints on the left of the zero line. These five articles refute the hypothesis that ever drinking increases breast cancer risk. Figure 4 also identifies studies that support the hypothesis that ever drinking increases breast cancer risk. These are depicted as six articles black bars that lie to the right of the zero line.

Figure 4 shows examples contradictory and redundant evidence. For example, articles 9328193 and 9032046 (the first two articles) **contradict** the evidence in article 11867358 and 11920551 (the last two articles) because their mid-points lie on opposite sides of the zero line. The first two articles are **redundant** because they both found that the rate of ever drinking in subjects with breast cancer was more than the METIS estimate. For some tasks, such as filtering, users consider redundancy undesirable. In the medical domain, redundancy can increase a user's confidence in the result.

Figure 4 suggests that ever alcohol consumption has a positive association with increased breast cancer risk. We generated additional METIS summaries that compared the rates of never, light, moderate or heavy drinking of subjects with breast cancer to our estimated population. Our results showed a positive effect size for light, moderate or heavy, and ever alcohol consumption and a negative effect size for never alcohol consumption. These statistically significant results ($P<0.05$) indicate that the rate of subjects who consume alcohol is greater than the rate of cancer free subjects who consume alcohol. Thus, the METIS summaries suggest a positive association between alcohol consumption and increased breast cancer risk.

If successful, the METIS generated summary should be consistent with analyses that consider the entire biomedical literature, rather than our small breast cancer corpus. To determine if the METIS summary was consistent with findings from the entire literature, we searched MEDLINE for published analysis that explored the relationship between breast cancer and alcohol consumption. We identified three articles that explored a possible association between alcohol consumption and breast cancer risk [12, 17, 21]. All three studies reported a positive association between alcohol consumption and increased breast cancer risk. Thus, the METIS summaries generated from our small collection of breast cancer are consistent with studies that consider the entire medical literature.

Figure 5 is the summary generated by METIS to explore a possible association between tobacco consumption and breast cancer risk. Figure 5 shows a positive effect-size between ever, former tobacco consumption, and a negative effect-size between current tobacco consumption and breast cancer. This result demonstrates that quantitative METIS summary captures both positive and negative associations. This property is useful to reduce false associations.

As with the alcohol analysis, we needed to see if the direction of the METIS summary was consistent with published analyses that considered the entire medical literature. Our MEDLINE search revealed only one meta-analysis that explored breast cancer and tobacco consumption [14]. The result of this published traditional meta-analysis was consistent with the METIS summary. We also conducted a traditional analysis using our breast cancer corpus. Similar to the alcohol analysis, only four articles satisfied the inclusion criterion of a traditional meta-analysis and the effect-size was not statistically significant.

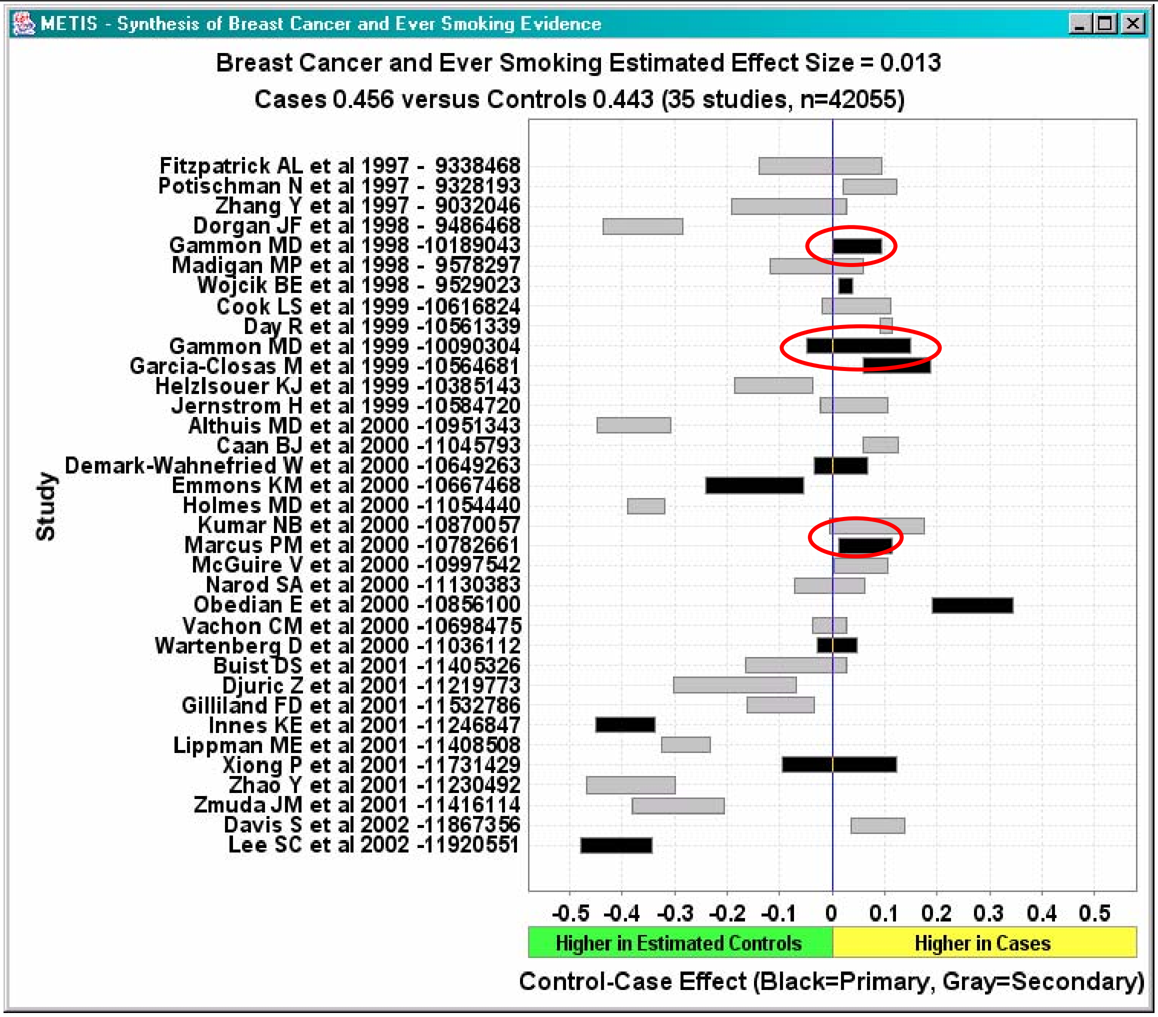

Figure 5 – The METIS Quantitative and Visual Summary for Articles in the Breast Cancer Corpus that Report Ever Smoking Consumption as either Primary (black) or Secondary (grey) Information. The positive effect size suggests a slight correlation between ever-tobacco consumption and breast cancer. Only the circled articles contain the information required to be included in a traditional meta-analysis.

The circled in Figures 4 and 5 identify articles that report exposure rates for both the subjects with breast cancer and a control population. The remaining 24 studies included in Figure 4, and the remaining 33 studies shown in figure 5 could not be included in a traditional analysis because they do not provide the rate of exposure for the control group.

In the ever-drinking example, three of the four studies that could be included in a traditional study show a positive association. In the tobacco consumption example, all four articles that could be included in a traditional analysis show a positive association. Although further analysis is required to rule out other factors that explain this finding, one possibility is publication bias. As mentioned in section 1, publication bias occurs when factors other than the quality of the study contribute to its publication. Our future work will explore the presence of systematic quality differences between studies that report risk factor details as primary and secondary information.

5. CONCLUSION

We have introduced the Information Synthesis process, which provides a user with a visual representation of findings reported in a collection of articles. In that visual representation, articles that contradict each other fall on opposite sides of the zero line. Thus, this representation enables a scientist to identify easily articles that have different findings.

In addition to the visual representation, the Information Synthesis process provides a scientist with a quantitative summary, called the effect size. The effect size captures the overall findings from a collection of articles; thus, enabling a user to resolve contradictory evidence.

The Information Synthesis process is just the first step towards realizing the potential of the digital library cyber-infrastructure. We have shown that the combination of information extracted from a full text article and structured data in a digital library can extend the existing model of meta-analysis. Such an extension is not feasible without automated techniques.

As the quantity of information available to scientists continues to soar, their need for tools and methodologies that support information synthesis behaviors will continue to grow. The cyber-infrastructure provided by a digital library is critical to enabling the shift from retrieval to synthesis.

6. ACKNOWLEDGMENTS

This research was supported by a dissertation award from the California Breast Cancer Research Program of the University of California, Grant Number 8GB-0175. Thanks to Dr W. Pratt and Dr C. Carpenter for their guidance and earlier discussions.

7. REFERENCES

[1] Alderson, P., Green, S. and Higgins, J.P.T. (eds.). *Cochrane Reviewers' Handbook 4.2.2 [Updated March 2004]*. John Wiley & Sons, Ltd, Chichester, UK, 2004.

[2] Allen, E. and Olkin, I. Estimating Time to Conduct a Meta-analysis From Number of Citations Retrieved. *Journal of the American Medical Association*, *282* (7). 634-635, 1999.

[3] Atkinson, D.R., Furlong, M.J. and Wampold, B.E. Statistical significance, reviewer evaluations and the scientific process: Is there a statistically significant relationship? *Journal of Counseling Psychology*, *29*. 189-194, 1982.

[4] Begg, C. and Berlin, J. Publication bias: a problem in interpreting medical data. *Journal of the Royal Statistical Society* (151). 419-463, 1988.

[5] Blake, C. Information Synthesis: A Mixed-Initiative Meta-Analytic Approach to Facilitate Knowledge Discovery from Scientific Text *School of Information and Computer Science*, University of California, Irvine, 2003.

[6] Centers for Disease Control. Behavioral Risk Factor Surveillance System Survey Data, U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, Atlanta, Georgia, 2002.

[7] Davies, H.T.O. and Crombie, I.K. What is a Meta-Analysis? *Hayward Medical Communications*, *1* (8). 1-8, 1998.

[8] Davies, H.T.O. and Crombie, I.K. What is a Systematic Review? *Hayward Medical Communications*, *1* (5). 1-5, 1998.

[9] Davis, S., Mirick, D.K. and Stevens, R.G. Residential Magnetic Fields and the Risk of Breast Cancer. *American Journal of Epidemiology*, *155* (5). 446-454, 2002.

[10] DerSimonian, R. and Laird, N. Meta-Analysis in clinical trials. *Controlled Clinical Trials*, *7*. 177-188, 1986.

[11] Easterbrook, P.J., Berlin, J.A., Gopalan, R. and Matthews, D.R. Publication bias in research. *Lancet*, *337*. 867-872, 1991.

[12] Ellison, R.C., Zhang, Y., McLennan, C.E. and Rothman, K.J. Exploring the Relation of Alcohol Consumption to Risk of Breast Cancer. *American Journal of Epidemiology*, *154* (8). 740-747, 2001.

[13] Hunter, J.E. and Schmidt, F.L. *Methods of Correcting Error and Bias in Research Findings*. Sage, Newbury Park, CA, 1990.

[14] Khuder, S.A., Mutgi, A.B. and Nugent, S. Smoking and breast cancer: a meta-analysis. *Rev Environ Health*, *16* (4). 253-261, 2001.

[15] Light, K. *The Cochrane Library - Self Training Guide and Notes 2003 Issue 4*. Center for Reviews and Dissemination, University of York, 2003.

[16] Lipsey, M.W. and Wilson, D.B. *Practical Meta-Analysis*. Sage, 2000.

[17] Longnecker, M.P., Berlin, J.A., Orza, M.J. and Chalmers, T.C. A meta-analysis of alcohol consumption in relation to risk of breast cancer. *JAMA*, *260* (5). 652-656, 1988.

[18] National Coordinating Centre for Health Technology Assessment. The Health Technology Assessment Programme Information Sheet 1: An overview of the programme , Version 1.0, Southhampton, UK, 2003.

[19] National Library of Medicine. UMLS Metathesaurus Fact Sheet *Fact Sheet*, National Library of Medicine, 2004.

[20] Petrosino, A. Lead Authors of Cochrane Reviews: Survey Results. Report to the Campbell Collaboration, University of Pennsylvania, Cambridge, MA, 1999.

[21] Singletary, K.W. and Gapstur, S.M. Alcohol and breast cancer: review of epidemiologic and experimental evidence and potential mechanisms. *JAMA*, *286* (17). 2143-2151, 2001.

[22] Sutton, A.J., Abrams, K.R., Jones, D.R., Sheldon, T.A. and Song, F. (eds.). *Systematic reviews of trials and other studies*. York Publishing Services, York, UK, 1998.

[23] Tengs, T. and Osgood, N.D. The Link between Smoking and Impotence : Two Decades of Evidence. *Preventive Medicine*, *32* (6). 447-452, 2001.

[24] Thornton, A. and Lee, P. Publication bias in meta-analysis: its causes and consequences. *Journal of Clinical Epidemiology*, *53*. 207-216, 2000.