
BENCHMARKING THE DOCUMENT REVIEW PROCESS

H5 WHITE PAPER

COPYRIGHT © 2007 – H5. All Rights Reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without the express written permission of H5.

The H5 logo and banner are trademarks of H5. The absence of a trademark or service mark from this list does not constitute a waiver of H5's trademark or other intellectual property rights concerning that trademark or service mark.

Other company names and product names mentioned in this document may be trademarks or service marks of their respective owners.

WHAT IS BENCHMARKING?

Benchmarking is a decision-making tool intended to improve performance through comparative quantitative evaluations of different systems and processes.

Benchmarking is particularly useful when an organization, or even an entire industry, has preconceived notions about how particular tasks must be performed. As circumstances change, the system that was once ideal for handling a particular task may no longer be effective. Incumbent systems are often maintained due to a bias for stability, uncertainty about the effectiveness of alternatives, or simple inertia. Benchmarking allows organizations to quantify trade-offs between alternative approaches and identify optimal strategies to best meet their current needs.

WHY BENCHMARK DOCUMENT REVIEW?

Document review is a complex, expensive, and labor-intensive process. In addition, it is usually undertaken in response to pressing and sometimes unanticipated legal needs. This means that every document review event is unique in scope and execution, adding further to the complexity (and cost) of performing it.

In the past, fully manual review was the only process available for evaluating a population of documents, and the nature of document creation and retention enabled this method to produce adequate results. Now, however, circumstances have changed: the explosion of electronically stored information (ESI) has resulted in document collections subject to review that are orders of magnitude larger than they were five or ten years ago. Yet traditional manual review by attorneys is still used as a primary method of assessment.

In recent years, the market has responded to the emerging challenges of ESI: numerous vendors offer state-of-the-art search tools and technologies to assist attorneys in the review process; consultancies have built competencies around this field; and other service providers have developed alternative methods of assessment, such as technology-enabled automated review systems.

The objective of a benchmarking exercise in the context of document review is to compare alternative review processes in order to determine which offering will provide the highest utility given the specific needs of a case. Absent such an evaluation, a vendor's claims to meet an organization's particular document review needs cannot be validated. Furthermore, a benchmarking exercise employed by any vendor or service provider should rigorously compare its process to the client's best alternative, so as to avoid artificial inflation of differences.

DOCUMENT REVIEW MODELING FRAMEWORK

Accuracy is an often-stated goal for document review tasks, and, when used in a general non-technical sense, accuracy is indeed an essential characteristic of an effective document review process. When, however, we want to quantify precisely the performance of a document review process, we need to select our metrics carefully. What we want are metrics that are sensitive to the risks that matter most to the companies and attorneys that undertake document review projects: (1) the risk of missing documents that should have been retrieved ("false negatives") and (2) the risk of retrieving documents that should not have been

retrieved (“false positives”). Recall and precision are the metrics that best gauge these risks: *recall* measures how many of the relevant documents in a document population have been retrieved, and *precision* measures how many of the retrieved documents are actually relevant. By measuring both the recall and the precision achieved by a process, we obtain the most meaningful measure of the process’s overall effectiveness.

It is fairly easy to achieve high precision alone – that is, retrieve far more relevant documents than non-relevant documents – through carefully crafted searches. Most approaches to document retrieval can achieve an acceptable level of precision of about 80 percent. The hard part is maintaining precision, while managing time and cost and achieving high recall – that is, retrieving most of the relevant documents in the population. Missing potentially relevant documents is a real risk that should not be underestimated.

A benchmarking model that allows for an appropriate evaluation of document review alternatives should therefore address the client’s unique objectives in terms of *time*, *cost*, and *risk*. This often forces certain trade-offs. Attempting to achieve high precision and high recall simultaneously can lengthen duration of a review and raise costs considerably. But ignoring either of these aspects of accuracy during a review will lead to unacceptable levels of risk. Following are some possible objectives at the outset of a review and their implications:

1) REDUCE TIME

- *Definition:* *Time* is defined here as time-to-completion – the amount of time required to complete the document review project, starting with the first document collected and ending with the last document reviewed.
- *Example:* “My objective is to complete the full review in three months because I have tight deadlines that cannot be changed. What is the resulting cost and anticipated level of risk?”
- *Implications:* In order to conduct an effective document review in a limited time, it will be necessary to employ more reviewers or purchase more sophisticated software thus raising costs. If fewer reviewers are used, increased review rates to meet short timelines will lead to more mistakes and a higher risk of overlooking relevant documents.

2) REDUCE COST

- *Definition:* *Cost* is defined here as all-in cost. Traditional review is often divided into discrete segments with varying costs that can increase over the lifecycle of a review project. The benchmarking exercise is intended to account for total cost of the entire project, which includes the cost of the core service, all related professional fees, all travel and lodging expenses, etc.
- *Example:* “My objective is to not exceed \$6 million in review expenditures because I do not have a larger budget. What review rate and expected risk result from this budget?”
- *Implications:* A restricted budget means fewer reviewers and possibly an extended timeline. Having fewer resources may also severely limit the number of documents that can be reviewed. Reviewing fewer documents greatly increases the risk of missing relevant material.

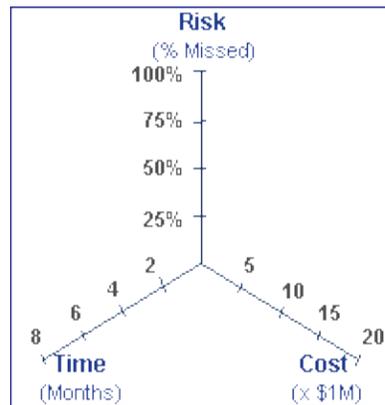
3) REDUCE RISK

- *Definition:* *Risk* is defined as the likelihood of missing relevant documents during your review. The less accurate a review is, the more risk is associated with the results.
- *Example:* “My objective is to retrieve at least 70 percent of all relevant documents (i.e., 70 percent recall) because the stakes are high and my case depends

heavily on discovery. How long will it take to achieve this, and how much will it cost?"

- *Implications:* It is relatively easy to avoid missing relevant documents by performing broad searches and retrieving more documents than are necessary. However, the resulting loss of precision will require more time and resources to sort out the relevant documents from the numerous retrieved non-relevant documents. On the other hand, opting for increased precision through narrow searches, will save time and money and ensure that most of the documents retrieved are relevant. But the consequent loss of recall will mean a greater risk of missing relevant material.

The relationships between these objectives can be visualized using a three-axis diagram:



The time axis gives time-to-completion in months increasing from 0 at the origin (the point where the three axes intersect). The cost axis gives the all-in cost of the document review project in millions of dollars increasing from 0 at the origin. The risk axis gives the percentage of relevant documents missed, increasing from 0 at the origin. An "ideal" review would be completed instantaneously, have no cost associated with it, and miss no relevant documents. On this diagram, such a review would be represented by a point at the origin.

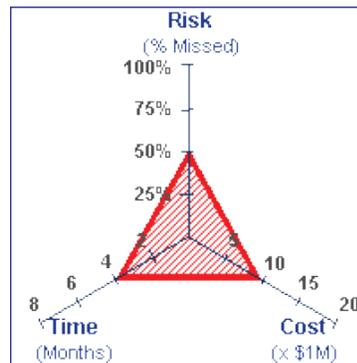
A real-world document review project will necessarily depart from this ideal: it will take some time to complete, a financial investment will be required, and some relevant documents will unavoidably be missed. Connecting the points for each of these factors on this diagram results in a triangle, which shows the relationship between the three variables. The smaller the area covered by this triangle, the closer the proposal is to the ideal: infinitely fast, infinitely cheap, and completely risk free.

Use of this modeling framework will:

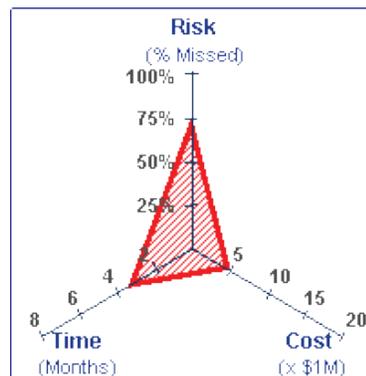
- 1) Capture the three factors under consideration in a single model;
- 2) Conduct sensitivity analyses through manipulation of the factors under consideration (time-to-completion, cost, and risk), and quantification of the resulting trade-offs,
- 3) Minimize uncertainty and risk in decision-making, thus ensuring that the actual review project will achieve the maximum possible utility (i.e., it will optimally meet the review objectives along each dimension); and
- 4) Avoid cost creep during the actual review project by ensuring that all costs have been thoroughly documented and understood up front.

ILLUSTRATION: DOCUMENT REVIEW MODELING FRAMEWORK

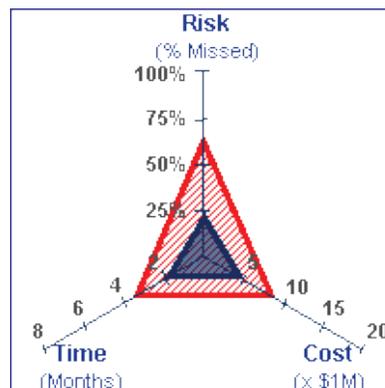
A single option for a document review project is visualized as a triangle on the three-axis model. With the variables shown along the time, cost, and risk axes, the model will appear as follows:



Change in any one of the three factors under consideration will result in corresponding changes in the other variables. For instance, decreasing the cost will likely force increases in risk, as illustrated below:



The primary purpose of the benchmarking exercise is, however, to compare two alternative offerings or methods for completing a document review project. In this case, the diagram will show two triangles, each representing a different offering:



The smaller the triangle, the closer it approaches the “ideal” review. Overlaying the variables that describe each alternative offering in this manner allows for a simple visual representation of the differences between the offerings. The option with the smaller triangle is in general “better,” but this evaluation may vary depending on the importance assigned to each factor under consideration. An offering that takes significantly less time and has less risk, but that costs significantly more than the slower, riskier alternative may not be acceptable if cost is of paramount importance. If, however, the smaller triangle is completely contained within the larger triangle, it would be considered undoubtedly better than the alternative in that it would take less time, be less expensive, and involve less risk than the alternative under consideration.

BENCHMARKING EXPERTISE

Several expert competencies are involved in conducting a benchmark for document review. These include financial modeling expertise, an understanding of accuracy from an information retrieval standpoint (i.e., precision and recall) and its relationship to time-to-completion, cost, and risk, as well as knowledge of the performance of various processes and technologies under consideration. Such competencies are necessary in order to conduct an accurate benchmark.

CONCLUSION

The specific data provided by a benchmarking exercise will equip an organization to make informed decisions by helping determine the optimal equilibrium between competing priorities of time-to-completion, cost, and risk. An organization’s investment of time and effort required to complete a benchmark is small, but the return on this investment is considerable. The results may or may not validate the originally anticipated review process; but in any event, the benchmark will provide an organization with an assessment of how to best meet its objectives while maximizing the value derived from its review expenditures.

