

January 2017

ROSS Intelligence

Artificial Intelligence in Legal Research



BLUE HILL
— RESEARCH —

ROSS Intelligence and Artificial Intelligence in Legal Research

Published: January 2017

Report Number: A0280

Share This Report



What You Need to Know

The growing availability and practicality of artificial intelligence (AI) technologies such as machine learning and Natural Language processing within the legal sector has created a new class of tools that assist legal analysis within activities like legal research, discovery and document review, and contract review. Often, the promised value of these tools is significant, while lingering cultural reluctance and skepticism within the legal profession can lead to hyperbolic reactions to so-called “robot lawyers,” both positive and negative. What is often lacking is evidence-based assessments of the impact of the growing market of AI-enabled legal tools on both the successful practice and business operations of legal organizations.

In order to assist organizations with their assessments of this new class of legal solutions, Blue Hill is committed to developing research programs that provide grounded assessments of the measurable impact of AI solutions within real-world legal use cases. As a product of that effort, this Benchmark Report summarizes the observations made in a study conducted by Blue Hill with a research panel of 16 legal researchers. This study compares the impact of traditional legal research tools, such as Boolean search and Natural Language search, with use cases of the [ROSS Intelligence](#) AI-supported legal research platform to supplement these traditional tools. For the purposes of this comparison, the analysis is in the context of United States Bankruptcy law research, although ROSS Intelligence uses the same underlying technology for all research products.

Key areas of comparison contained in this report include: (1) the quality of information retrieval in the search results produced by the observed use of Boolean search, Natural Language search, and the ROSS tool; (2) user feedback with respect to ease of use and confidence in the results retrieved among the use cases studied; (3) the impact on the time required for users to complete research activities; and (4) the ultimate business value and return on investment (ROI) derived from these efficiency gains.

AT A GLANCE

Technology Assessed

[ROSS Intelligence](#) AI legal research platform

Research Objective

To assess the impact of ROSS-assisted use cases in bankruptcy law research with respect to:

- Information Retrieval Quality
- Usability and User Confidence
- Research Efficiency

Research Methodology

Blue Hill used a panel of 16 legal researchers to benchmark primary ROSS use cases with those involving Boolean and Natural Language search capabilities of research platforms.

ROSS: Impact Identified

Reduction in Research Time from Incorporating Use of ROSS

- **30.3%** over Boolean alone
- **22.3%** over Natural Language alone

Increase in Information Retrieval Quality Compared to Boolean and Natural Language Search

- **42.9%** more relevant authorities retrieved
- **30.3%** more results constituted relevant authorities
- **86.9%** higher Normalized Discounted Cumulative Gain

Estimated Business Impact & ROI

- **\$8,466 - \$13,067** annual revenue increase per attorney based on a 25% conversion of unbillable time to billable time
- **176.4% to 544.5%** resulting return on investment

ROSS Intelligence: Overview and Business Case

The [ROSS Intelligence](#) tool is an artificial intelligence (AI) platform supporting legal research activities. Built on ROSS Intelligence's proprietary legal AI framework, Legal Cortex, combined with technologies such as [IBM Watson's](#) cognitive computing technology, ROSS uses Natural Language processing and machine learning capabilities to identify legal authorities relevant to particular questions. Users conduct searches by entering questions in plain language, rather than by complex search strings. ROSS's cognitive computing and semantic analysis capabilities permit the tool to understand the intent of the question asked and identify answers "in context" within the searched authorities.

ROSS Intelligence positions its platform as a case law research supplement to traditional Boolean search and Natural Language parsing approaches used by electronic legal research tools. In this context, ROSS promises to provide increased research output quality (by collecting the most relevant authorities among its initial returned results) as well as a resulting improvement in the efficient execution of legal research activities when compared to the use of traditional tools alone.

Research Objectives and Methodology

In order to assess the potential impact of the ROSS AI-assisted research platform, Blue Hill Research conducted a benchmark study intended to record and compare the utility of the solution with traditional electronic legal research tools with respect to three primary categories:

- Quality of information retrieval
- User satisfaction and confidence
- Impact on research efficiency

Table 1 (below) provides definitions for these categories as they were measured in the course of this research.

Benchmark Assessment Trial

To benchmark the performance of the examined research tools, Blue Hill employed a research panel of legal researchers to complete an assessment trial of sample questions. The research panel consisted of sixteen experienced legal research professionals, randomly separated into four assessment groups of four members each.

Each research participant received a standard set of seven questions modeling real-world questions posed to legal practitioners. To enable its analysis, Blue Hill collected: (1) recorded research time, (2) participant answers, (3)

Technology Classes Discussed

Boolean Search

Method of search using keywords to identify documents containing particular words and Boolean connectors or operators that narrow results based on the relationships between the terms.

Natural Language Search

Method of search where a query is entered in plain language and is parsed by the search algorithm to identify content addressing the same topic.

Natural Language Processing

Artificial intelligence technology that derives computer-processable semantic and contextual meaning from natural, human language.

Machine Learning

Artificial intelligence capabilities that permit a computer or application to alter operations without explicit programming as it is exposed to new data.

search histories, and (4) user surveys from each participant. To provide a consistent basis of comparison, Blue Hill limited the scope of subject matter for the questions to United States federal bankruptcy law.

Table 1: Key Assessment Factors Benchmarked

Category	Factor	Definition	Measurement
Information Retrieval Quality	Thoroughness	Portion of the total pool of existing relevant authorities that were retrieved	Percentage of the total set of relevant authorities that were retrieved
	Accuracy	Portion of total results retrieved that included relevant authorities	Percentage of the total results retrieved that represent relevant authorities
	Ranking Effectiveness	Relative placement of relevant authorities within the list of top results retrieved	Normalized Discounted Cumulative Gain
User Satisfaction	Ease of Use	Participant's satisfaction with the ease of use of the research approach.	Self-reported Likert Scale responses to standardized satisfaction questions
	Confidence	Participant's perceived confidence that he or she obtained a complete answer	Self-reported Likert Scale responses to standardized satisfaction questions
Research Efficiency	Time to Complete	Amount of time required for a participant to obtain a satisfactory answer the question.	Total time the researcher spent using the research approach to obtain an answer

Source: Blue Hill Research, January 2017

Assessment Group Composition and Constraints

Blue Hill assigned each participant to research and provide a written answer to the legal question within a pre-determined answer period of two hours. Each participant was at liberty to complete research and prepare an answer subject to constraints in available research tools imposed by the participant's group assignment.

In order to model the ideal use case identified by ROSS Intelligence, Blue Hill constrained the research tools available for use by each group as follows:

- **Boolean Search** – participating researchers were only permitted to use Boolean keyword search capabilities of assigned major legal research platforms.
- **Natural Language Search** – participating researchers were only permitted to use the Natural Language search capabilities of assigned major legal research platforms.
- **ROSS and Boolean Search:** participating researchers were directed to use the ROSS platform and Boolean keyword search capabilities of assigned major legal research platforms as they saw fit.
- **ROSS and Natural Language Search:** participating researchers were directed to use the ROSS platform and Natural Language search capabilities of assigned major legal research platforms as they saw fit.

Blue Hill assigned each group to a mixture of qualifying tools, using [Westlaw](#) and [LexisNexis](#) research platforms to ensure that results were representative of the most common legal research experiences. In addition, members of each research group were limited to individuals with no prior experience with the assigned tool(s) and relatively minimal experience with bankruptcy law. This provided Blue Hill with a common level of experience to establish a baseline user profile for comparison and analysis. Similarly, in order to provide a common field of comparison in the output of the tools, Blue Hill limited its analysis used to the top 20 results retrieved in response to each query performed with the tool.

Research Outcomes: Performance Benchmarked

Blue Hill Research aggregated the experiences of participants in each assessment group in order to identify the average performance of the sample group related to each of the four legal research use cases identified.

Study results demonstrate significant improvements in research time as well as researcher experience resulting from the use of ROSS when used with both Boolean and Natural Language capabilities. Analysis of the search results produced by each tool largely echoes this as well, with ROSS generating better results in terms of:

- Identification of relevant authorities
- Minimization of non-relevant results
- Prioritized placement of relevant authorities in search results.

These findings indicate clear advantages resulting from the addition of the ROSS tool to electronic legal research involving traditional tools. The sections that follow describe the impact Blue Hill observed with respect to the above assessment factors.

Quality of Information Retrieval

To provide a basis for comparing the effectiveness of the information retrieval of the tools, Blue Hill analyzed the first 20 results produced in response to each query performed in the course of the assessment. This analysis was used to compare the quality of information retrieval of the three tool classes used in this research: (1) Boolean Search, (2) Natural Language Search, and (3) the AI-supported search of the ROSS tool. Blue Hill compared the performance of the tools with respect to three categories: Thoroughness, or the number of relevant authorities identified; Accuracy, or the amount of retrieved results that represented relevant authorities; and [Normalized Discounted Cumulative Gain \(NDCG\)](#), a standardized measure of the ranking of search results compared to an idealized ranking according to the relative value each result has to a user.

Determining Relevance

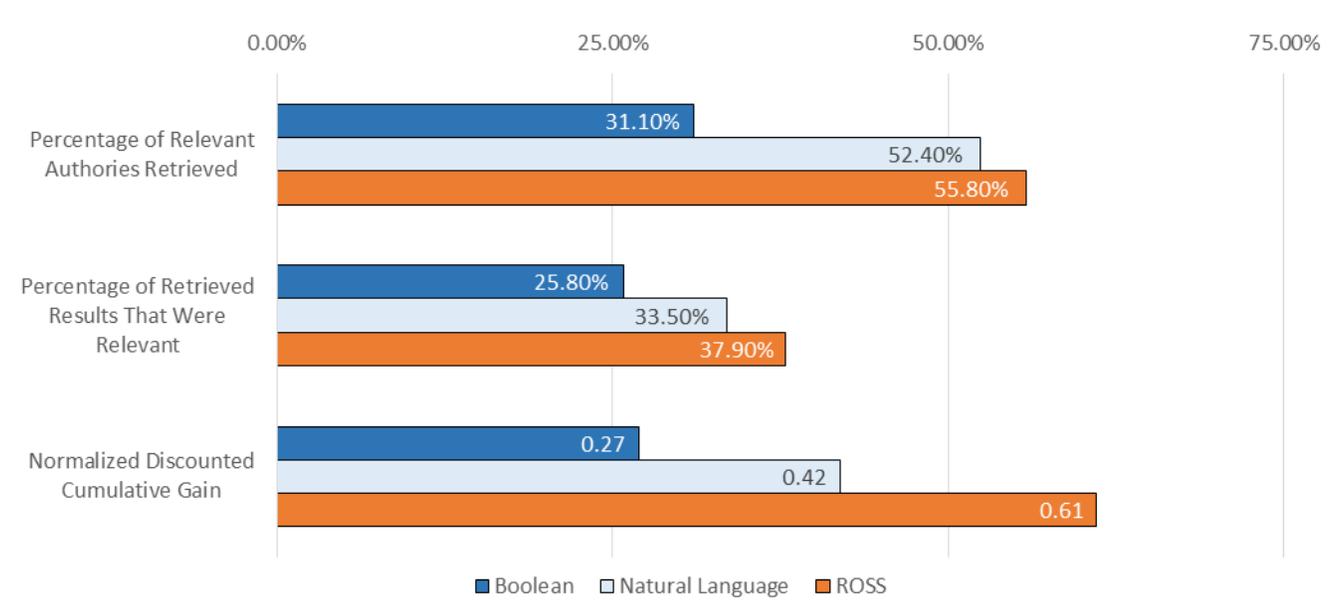
A determination of the relevance of a case to a particular legal question can be subjective and often turns on the circumstances involved. Accordingly, Blue Hill ranked the relevance of authorities retrieved based on their utility to the study participants, rather than using an objective determination of relevance.

Blue Hill scored cases as relevant if they had been cited by a participant in order to answer the question researched. One citation by a participant was sufficient to identify a case as relevant. Each additional citation by a participant increased the weighted "relevance score" of a case. Blue Hill used these scores to inform the analysis used to determine Thoroughness, Accuracy, and NDCG.

Because Blue Hill constrained the number of search results analyzed to the first 20 results produced in response to each query, not all cases cited by participants necessarily appeared in the search results reviewed. Similarly, the use of the subjective basis for relevance potentially means that not all cases that are actually relevant to the questions presented were identified. These factors result in a conservative analysis of the effectiveness of the tools considered.

Figure 1 summarizes Blue Hill’s evaluation of each electronic research tool used in the study with respect to the three information retrieval quality assessment factors. On every measure, ROSS outperformed the traditional tools evaluated. These figures are described and explained in detail below. Readers should note that “Thoroughness” and “Accuracy” are represented as percentages, while “NDCG” is represented as a number between 0 and 1.0, in keeping with convention.

Figure 1: Information Retrieval Effectiveness of Legal Search Tools Based on Observed Queries



Source: Blue Hill Research, January 2017

A review of Figure 1 reveals clear limitations with Boolean search. Participants using Boolean tools retrieved less than a third of the identified relevant authorities within their first 20 results. At the same time, only 25.8% of the results these searches produced were relevant authorities, meaning that 74.2% of the retrieved results with Boolean search represented noise that did not assist the participants to answer the presented legal questions. These outcomes were exceeded by Natural Language search, which produced a significantly higher percentage of the relevant authorities (52.4%), while approximately 30% more of the results retrieved were relevant authorities. The ROSS tool outperformed both evaluated search technologies on these measures, producing 55.8% of the relevant authorities within its top 20 results, with 37.9% of all retrieved results representing relevant authorities.

ROSS’s most significant observed improvement over both Boolean and Natural Language search fell within the ranking of results by their ultimate value to the participants. NDCG is a measurement of how closely the prioritization of results generated by a search platform matches an ideal ranking of results. An NDCG assessment assumes that the earlier a search result appears in a list of results, the more valuable it is to a user. Accordingly, NDCG measures the utility of search based on how well the value of a result to a user is correlated to its position in the list of results. NDCG scores fall between 0 and 1.0 (a perfect match to the ideal ranking). The ROSS tool achieved an NDCG of .61, 46.1% higher than Natural Language search and 127.8% higher than Boolean search. For the user, this means that not only is the ROSS tool generating more relevant results with less noise than

Boolean or Natural Language search, these results are highly concentrated within the first results that the user sees.

Below, we observe the impact these differences make on the overall time the user spends conducting research.

User Satisfaction and Confidence

In order to understand user perceptions regarding the effectiveness of their respective tools, Blue Hill required participants to complete a user satisfaction survey at the completion of the assessment. Blue Hill provided participants with a list of statements regarding the tools used and asked to indicate their relative agreement or disagreement using a five-point Likert scale, with 1 indicating strong disagreement, 5 indicating strong agreement, and 3 being neutral. Table 2 summarizes the average responses of each assessment group with respect to key statements regarding both the usability of and confidence in the toolset that they were assigned.

Table 2: Average Agreement with Statements Describing User Experience with the Toolsets Employed

		Boolean	Natural Language	ROSS & Boolean	ROSS & Natural Language
Usability	I found the tool's user interface to be intuitive and easy to use.	3.3	4.0	5.0	5.0
	The search results returned by the tool were concise and primarily contained cases that were relevant to my legal questions.	4.0	3.5	5.0	5.0
	The search results returned by the tool did not include a large number of results that were NOT RELEVANT to my legal questions.	2.0	2.7	4.7	4.7
Confidence	It was easy to find all of the cases required to give a complete answer to my legal questions using the tool.	3.5	3.3	4.8	5.0
	I am confident that the tool returned all of the cases required to give a complete answer to my legal questions.	3.3	3.5	4.8	5.0

Scale - 1: Strongly Disagree, 2: Disagree, 3: Neutral, 4: Agree, 5: Strongly Agree-

Source: Blue Hill Research, January 2017

Taken as a group, the survey results summarized in Table 2 reveal strong indicators of user satisfaction of the participants using a ROSS-supported toolset with respect to the usability, presentation of search results, and inclusion of relevant authorities within the search results. Participants within these groups also indicated high levels of confidence in the ability of the tool to identify all authorities relevant to the matter. In nearly all cases, the responses indicated by participants using ROSS and another tool often exceeded those of organizations using only Boolean search or only Natural Language search by at least a full point. Blue Hill found no significant variation in the difference between responses when comparing Boolean & ROSS and Boolean groups or Natural Language & ROSS & Natural Language. In both cases, groups using the ROSS tool in addition to the other legal research tool reported higher levels of satisfaction and confidence.

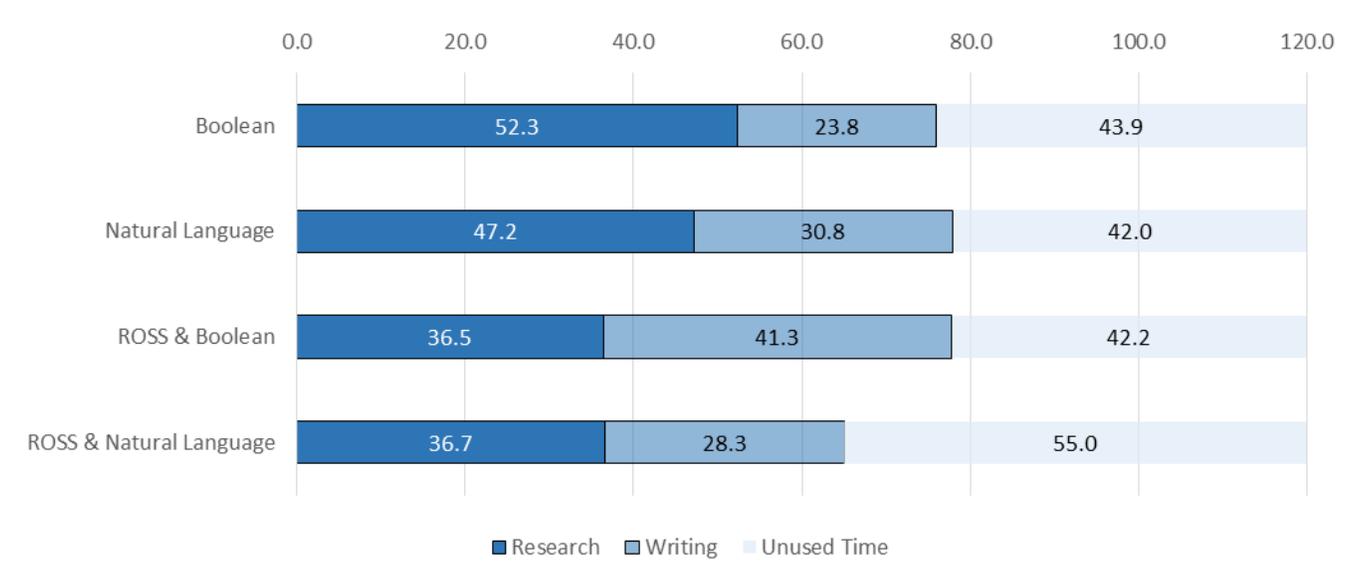
These results relate purely to the perceived experiences of the research participants. Blue Hill observes a clear connection between the increased information retrieval effectiveness identified in the ROSS tool above and the user perceptions identified here. Anecdotal responses in supplemental portions of the user satisfaction survey suggest that the ROSS tool's higher concentration of relevant authorities among initial search result positions

played a role with the higher satisfaction with the ease of use and confidence. This is reinforced by participant behavior over the course of the study, as participants in the ROSS-assisted groups tended to increase use of and reliance on the ROSS tool as they progressed through the question sets.

Research Efficiency

As a group, the participants completed research efforts and provided finished answers well within the two-hour window given for all questions. To this end, the average time to research a question across all groups is 43.1 minutes. Blue Hill Research observed clear differences in the research times achieved across assessment groups. Groups using combined ROSS and traditional tools completed research in an average of 36.5 minutes (ROSS and Boolean) and 36.7 minutes (ROSS and Natural Language). This represents a 30.3% and a 22.3% reduction in research time compared to the groups relying solely on Boolean search or Natural Language search. Figure 2 identifies the average research times of the respective groups as well as the average allocation of time for the remaining answer period.

Figure 2: Average Allocation of Time to Activities by Assessment Group (in Minutes)



Source: Blue Hill Research, January 2017

Blue Hill did not identify any notable correlation between the tools used and the time used to write an answer to the questions presented. The time dedicated to answer varied considerably across the assessment groups as well as among individual participants within each group without discernable pattern. Accordingly, no observation may be extracted from this study regarding potential impact on legal analysis and written application. Nonetheless, the impact on research time alone demonstrated by participants using the ROSS tool to support their research efforts is significant enough to net overall reductions in the labor and time required to obtain answers to legal questions.

Blue Hill Analysis: Quantifying the Business Impact of Research Efficiency

Blue Hill's research indicates clear advantages resulting from the use of ROSS to supplement traditional electronic legal research practices. Placing these results in meaningful context for the operations of a legal organization requires an understanding of how these efficiency gains impact the underlying business needs of the organization. For example, non-profit legal aid organizations and in-house legal departments will primarily evaluate this impact in terms of the operational efficiency and labor capacity of attorney resources. In other words: the ability to increase the number of matters and activities that the organization's attorneys are able to perform. For private law firms and legal services providers, the primary question will be: to what extent will the gain in research efficiency impact revenue and profitability.

There are a number of ways in which research effectiveness can relate to firm profitability and revenue generation, depending on the business model of the firm and how the organization makes use of the time saved. Reduced research time can affect the number of billed hours that go unpaid by clients. It can also impact the firm's ability to effectively take on additional clients and matters or to ensure that contingency or other flat fee matters are accomplished at optimal costs to the firm. Organizations thus need to consider their own circumstances and business objectives when assessing the potential impact of an investment in a tool similar to the ROSS AI platform.

However, in all cases, assessing the potential value of an investment requires a comparison of the net gain provided by the tool with the costs associated with acquiring the technology. In order to assist organizations with their own assessments, the following section presents a model of the financial impact of the research efficiency gain based on the conversion of "written off" hours to billable time.

Investment Impact Estimate: Conversion Model

The conversion model that appears below is adapted from Blue Hill Research's [Calculating the Value of Legal Research](#) Analyst Insight. This model is based on the notion that a reduction in the amount of unbilled time or billed time that is ultimately unpaid by the client creates the opportunity to generate new billable time. The degree to which this opportunity is realized represents a conversion of the previously wasted and written-off time to productive, revenue-generating time. For law firms, the revenue loss associated with this wasted time is significant. Based on industry calculations, Blue Hill estimates that the average associate works [an average of 743.6 hours a year completing legal research](#). [Twenty-six percent \(26%\) of legal research time](#) is written off as unbilled or unpaid by clients.

Reducing the amount of work performed on a per-client or per-matter basis helps reduce the need and likelihood that work performed needs to be written off, thus creating the opportunity for conversion and a net increase in revenue. A firm is far from guaranteed to be able to make a total conversion of this unbilled research time. However, even partial conversion can generate significant benefits for the firm's business. To this end, Table 4 models the potential value and ultimate return on investment (ROI) of this research efficiency based on this potential for increased revenue, based on conversion rates ranging from 10% to 100%. This model also used an estimated associate hourly billing rate of \$320, as a conservative skew of the most recent [Laffey Matrix](#) data for attorneys with 1 – 3 years of experience and hourly billable ranges identified by [the National Consumer Law](#)

Center. To develop this model, Blue Hill uses the average research time and write-off rate identified above as caps for the net revenue gain that is obtainable by conversion. In all circumstances, Blue Hill recommends that firms use this model as a framework to build their own projections employing their actual rates and expected costs.

Table 3: Impact of Research Efficiency by Conversion of Write-Offs

		22.3% Reduction	30.3% Reduction
Time Savings	Number of Hours Saved (Year)	165.8	225.3
	Number of Hours Saved With Potential Positive Revenue Gain (Year)	165.8	193.3
Annual Revenue by Conversion of Time Saved to Billable Hours			
10% Conversion		22.3% Reduction	30.3% Reduction
	Net Increase in Number of Billable Hours Paid (Year)	16.6	19.3
	Increase in Revenue (\$320 / Hour)	\$5,306	\$6,187
25% Conversion		22.3% Reduction	30.3% Reduction
	Net Increase in Number of Billable Hours Paid (Year)	41.5	48.3
	Increase in Revenue (\$320 / Hour)	\$13,266	\$15,467
33.3% Conversion		22.3% Reduction	30.3% Reduction
	Net Increase in Number of Billable Hours Paid (Year)	55.2	64.4
	Increase in Revenue (\$320 / Hour)	\$17,670	\$20,602
50% Conversion		22.3% Reduction	30.3% Reduction
	Net Increase in Number of Billable Hours Paid (Year)	82.9	96.7
	Increase in Revenue (\$320 / Hour)	\$26,532	\$30,934
75% Conversion		22.3% Reduction	30.3% Reduction
	Net Increase in Number of Billable Hours Paid (Year)	124.4	145.0
	Increase in Revenue (\$320 / Hour)	\$39,797	\$46,401
100% Conversion		22.3% Reduction	30.3% Reduction
	Net Increase in Number of Billable Hours Paid (Year)	165.8	193.3
	Increase in Revenue (\$320 / Hour)	\$53,063	\$61,868

Source: Blue Hill Research, January 2017

Table 3 provides a range by which reader can evaluate the potential impact on revenue that can be derived from the conversion of unbillable or unpaid time to revenue-producing time. Understanding this range is particularly important when approaching the business impact of research efficiency on a conversion model, as a number of intervening factors will determine precisely how effectively an organization can redirect saved time into new billable time. Achieving a 100% conversion will be rare, but a 10% conversion requires an addition of less than 20 billed and paid hours over a year. Based on the efficiency gain observed in this research and the assumptions identified above, Blue Hill estimates the revenue yield at this level as falling between \$5,306 and \$6,187 per researcher. This impact is thus scalable across the organization.

Viewing the impact at subsequent tiers identified in Table 3 permits readers to understand how this impact scales based on an increasing conversion rate, up to a ceiling of \$61,868 per researcher. Again, readers should use their understanding of their own business and practices to estimate the conversion rate and resulting range of potential revenue gain that is obtainable.

Based on this analysis, we can also model the potential return on investment (ROI) available through the use of an AI solution consistent with the use case explored in this research. Notably, this use case does not involve a new business operation or replacement of an existing tool as much as the introduction of an additional tool to an existing work process. This provides a simple model for understanding ROI, based on the value of the net efficiency gain provided, balanced with the additional cost associated with the investment in the tool. Table 4 provides a model of this ROI calculation on a per user basis, based on the model presented in Table 3 and an estimated cost range for a seat of an AI-enabled legal research tool of \$2,400 to \$4,800.

Table 4: Estimated ROI Associated with AI-enabled Legal Research Platform

Cost of Solution	Estimated Cost Range		
	Cost of Artificial Intelligence Tool (Annual Cost Per Seat)		
	\$2,400 to \$4,800		
Annual Revenue by Conversion of Time Saved to Billable Hours			
10% Conversion		22.3% Reduction	30.3% Reduction
	Return (Increase in Revenue - Cost)	\$506 - \$2,906	\$1,387 - \$3,787
	Return on Investment	10.5% - 121.1%	28.9% - 157.8%
25% Conversion		22.3% Reduction	30.3% Reduction
	Return (Increase in Revenue - Cost)	\$8,466 - \$10,866	\$10,667 - \$13,067
	Return on Investment	176.4% - 452.7%	222.2% - 544.5%
33.3% Conversion		22.3% Reduction	30.3% Reduction
	Return (Increase in Revenue - Cost)	\$12,870 - \$15,270	\$15,802 - \$18,202
	Return on Investment	268.1% - 636.3%	329.2% - 758.4%
50% Conversion		22.3% Reduction	30.3% Reduction
	Return (Increase in Revenue - Cost)	\$21,732 - \$24,132	\$26,134 - \$28,534
	Return on Investment	452.7% - 1,005.5%	544.5% - 1,188.9%
75% Conversion		22.3% Reduction	30.3% Reduction
	Return (Increase in Revenue - Cost)	\$34,997 - \$37,397	\$41,601 - \$44,001
	Return on Investment	729.1% - 1,558.2%	866.7% - 1,833.4%
100% Conversion		22.3% Reduction	30.3% Reduction
	Return (Increase in Revenue - Cost)	\$48,263 - \$50,663	\$57,068 - \$59,468
	Return on Investment	1,005.5% - 2,111.0%	1,188.9% - 2,477.8%

Source: Blue Hill Research, January 2017

As with Table 3, Table 4 provides a scale against which organizations can evaluate their own circumstances to plot an expectation of the value available from the adoption of a tool such as ROSS in the use cases reviewed in

this research. Notably, using Blue Hill's estimated billable rate of \$320, a positive ROI is obtained at a 10% conversion rate, meaning that the investment drives a net gain to the firm with a minimal recovery of written-off hours. Similarly, a 176.4% to 544.5% ROI becomes possible with at least 25% conversion. Where exactly an organization will fall within these ranges will depend on a variety of factors related to its business, investment costs, and the scope of use adopted. In all cases, however, Blue Hill's model strongly suggests positive business gain is available from the investment in ROSS.

Key Observations and Takeaways

We are currently at an early stage in both the development and adoption of AI-assisted legal research and other legal tools. As the sophistication of these tools grows and use expands from early adopters to the wider legal community, the cycles of hype and anxiety that have colored the initial discourse surrounding AI tools will begin to give way to demand for evaluations of the tangible impact and value available through these tools.

Based on its benchmark assessment of the ROSS platform used to supplement Boolean and Natural Language search in the context of bankruptcy law, Blue Hill finds that the ROSS tool provides significant, additive contributions to the effectiveness of legal researchers. These gains include between a 22.3% and 30.3% reduction in research time, stemming from substantial improvements in information retrieval, particularly in the ranking of research results identified by a .61 NDCG score. These results have the potential to unlock new gains in the efficient and profitable operation of legal organizations, as well as create opportunities for new revenue gain. It should be noted that none of these findings indicate that AI-assisted legal research constitutes a dramatic transformation in the use of technology by legal organizations. Rather, the use cases and impact reviewed indicate that tools like ROSS Intelligence more closely represent a significant iteration in the continuing evolution of legal research tools that began with the launch of digital databases of authorities and have continued through developments in search technologies. It is in this light that the potential of the tool are most accurately evaluated.

While often in danger of overstatement, the impact of this AI-led stage of evolution is substantial. [Blue Hill Research has previously observed](#) that the improved accessibility of legal data, authorities, and commentary has eroded the traditional value proposition of the online research database: the aggregation and interrogation of large sets of legal sources. The emergence of low-cost and no-cost alternatives puts pressure on the pricing models of traditional tools, while simultaneously driving the need for differentiation through new functionality and features that provide enhanced utility, either by increasing the speed and effectiveness of research, or by unlocking new value from legal data. With respect to the former, AI-assisted tools represent an improvement, by improving the interrogability of large sets of legal sources, while removing the labor-intensive manual indexing, classification, and passage identification traditionally used by legal publishers. Blue Hill's benchmark research

Blue Hill finds that the ROSS tool provides significant, additive contributions to the effectiveness of legal researchers. These gains include between a 22.3% and 30.3% reduction in research time, stemming from substantial improvements in information retrieval, particularly in the ranking of research results identified by a .61 NDCG score.

reveals a concurrent improvement in effectiveness and efficiency compared to established approaches to legal research. In this way, AI-assisted tools such as ROSS Intelligence represent a clear response to the present market needs, delivering value through both cost of ownership and contributed value vectors. It is this combination that permits these tools to demonstrate net business gains and ROI in use cases that enhance, rather than replace, traditional research strategies. These dynamics show a clear need for investigation as law firms and other legal organizations consider the role legal research plays in their own overhead and value provided to their clients.

David Houlihan, Esq.

Principal Analyst

David Houlihan researches enterprise risk management, compliance and policy management, and legal technology. He is an experienced advisor in legal and technology fields with a unique understanding of complex information environments and business legal needs.



CONNECT ON SOCIAL MEDIA



[@dwhoulihan](https://twitter.com/dwhoulihan)



[linkedin.com/in/houlihandavid](https://www.linkedin.com/in/houlihandavid)



bluehillresearch.com/author/david-houlihan

For further information or questions, please contact us:

Phone: +1 (617) 624-3600

Fax: +1 (617) 367-4210

Twitter: @BlueHillBoston

LinkedIn: [linkedin.com/company/blue-hill-research](https://www.linkedin.com/company/blue-hill-research)

Contact Research: research@bluehillresearch.com

Blue Hill Research offers independent research and advisory services for the enterprise technology market. Our domain expertise helps **end users** procure the right technologies to optimize business outcomes, **technology vendors** design product and marketing strategy to achieve greater client value, and **private investors** to conduct due diligence and make better informed investments.

Unless otherwise noted, the contents of this publication are copyrighted by Blue Hill Research and may not be hosted, archived, transmitted, or reproduced in any form or by any means without prior permission from Blue Hill Research.