

3-1-2019

Separating the Privileged Wheat from the Chaff - Using Text Analytics and Machine Learning to Protect Attorney-Client Privilege

Robert Keeling
Sidley Austin, LLP

Nathaniel Hubert-Fliflet
Ankura Consulting Group

Jianping Zhang
Ankura Consulting roup

Rishi P. Chhatwal
AT&T Services, Inc.

Follow this and additional works at: <https://scholarship.richmond.edu/jolt>

Recommended Citation

Robert Keeling, Nathaniel Hubert-Fliflet, Jianping Zhang & Rishi P. Chhatwal, *Separating the Privileged Wheat from the Chaff - Using Text Analytics and Machine Learning to Protect Attorney-Client Privilege*, 25 Rich. J.L. & Tech 1 ().

Available at: <https://scholarship.richmond.edu/jolt/vol25/iss3/2>

This Article is brought to you for free and open access by the Law School Journals at UR Scholarship Repository. It has been accepted for inclusion in Richmond Journal of Law & Technology by an authorized editor of UR Scholarship Repository. For more information, please contact scholarshiprepository@richmond.edu.

**SEPARATING THE PRIVILEGED WHEAT FROM THE CHAFF –
USING TEXT ANALYTICS AND MACHINE LEARNING TO
PROTECT ATTORNEY-CLIENT PRIVILEGE⁺**

Robert Keeling*
Nathaniel Huber-Fliflet**
Dr. Jianping Zhang***
Rishi P. Chhatwal****

Cite as: Robert Keeling, Nathaniel Huber-Fliflet, Jianping Zhang & Rishi P. Chhatwal, *Separating the Privileged Wheat from the Chaff – Using Text Analytics and Machine Learning to Protect Attorney-Client Privilege*, 25 RICH. J.L. & TECH., no. 3, 2019.

⁺ This article has been prepared for informational purposes only and does not constitute legal advice. This information is not intended to create, and the receipt of it does not constitute, a lawyer-client relationship. Readers should not act upon this without seeking advice from professional advisers. The content herein does not reflect the views of the organizations associated with any of the authors.

* Robert Keeling is a partner at Sidley Austin, LLP. He is an experienced litigator whose practice includes a special focus on electronic discovery matters. Robert is co-chair of Sidley’s eDiscovery Task Force, and he represents both plaintiffs and defendants in civil litigation and conducts internal investigations in the U.S. and throughout the world.

** Nathaniel Huber-Fliflet is a Senior Managing Director at Ankura Consulting Group, based in Washington, D.C. He has 15 years of experience consulting with law firms and corporations on advanced data analytics solutions and legal technology services.

*** Jianping Zhang is a Senior Managing Director at Ankura Consulting Group, based in Washington, D.C. He received his Ph.D in Computer Science from University of Illinois. He has about 30 years of experience in working on applications of A.I. and Machine Learning to solve real world problems.

**** Rishi P. Chhatwal is Assistant Vice President and Senior Legal Counsel for Enterprise eDiscovery at AT&T Services, Inc. He received his J.D. from the University of Georgia School of Law. This article was prepared in his personal capacity, and does not represent the views of AT&T.

ABSTRACT⁺

The digital age has created unique challenges for parties that engage in large-scale litigation. Safeguarding the attorney-client privilege is a critical task for litigators during discovery—one that becomes more difficult and expensive every year. Document review is now responsible for the vast majority of costs in the average legal matter, and costs are only rising. The volume of digitally-stored data doubles roughly every two years, driving up discovery costs and increasing the risk of inadvertent disclosure of privileged information. As the digital world evolves, the legal community has sought to evolve with it, particularly in the document review process.

Keyword searching has been the dominant method of identifying digitally-stored, privileged documents for the last several decades, but attorneys have conducted little research about the most efficient ways to use this method. Most legal teams rely on a combination of intuition and conventional wisdom. To subject those intuitions to the rigor of scientific experiments, we used three data sets and search term lists from real legal matters to determine which search terms were effective in identifying privileged communications. The results from our study revealed that thoughtfully crafted keyword term lists do identify a significant portion of the privileged document population. What may be surprising to experienced practitioners is that many commonly used terms that are believed to be imprecise proved quite effective at identifying privileged documents, while limiting the volume for review. Other popular terms proved to be ineffective. The study also compared the effectiveness of identifying privileged communications using predictive modeling and machine learning. The insights provided in this article can, if implemented by practitioners, add additional client protections against the disclosure of privilege documents and make privilege review more defensible and less costly.

⁺ The authors wish to thank Matthew Letten of Sidley Austin LLP for his assistance in writing this Article.

Table of Contents

I. INTRODUCTION..... 4

**II. THE PROBLEM: SOARING DISCOVERY COSTS AND
INEFFICIENT SEARCH TECHNIQUES..... 7**

**A. What is the Attorney-Client Privilege,
and Why is Protecting Privilege So Important?..... 7**

**B. The Challenge and Cost of Identifying
Privileged Documents in the Digital Age..... 15**

1. Rising Costs 15

**2. Current Methods: Keyword Searching
and Predictive Modeling..... 17**

a. Keyword Searching 17

b. Predictive Modeling 23

**III. THE WAY FORWARD: EMPIRICAL STUDY ON
KEYWORD SEARCHING AND PREDICTIVE MODELING
AT IDENTIFYING PRIVILEGED DOCUMENTS..... 29**

A. Our Experiment..... 30

B. The Results..... 32

1. Effectiveness of Keyword Searching..... 32

a. Outside Counsel Keywords..... 34

2. Effectiveness of Predictive Modeling..... 40

3. Keyword Searching vs. Predictive Modeling..... 42

C. A Way Forward..... 43

IV. CONCLUSION..... 44

I. INTRODUCTION

[1] The complexities of the digital age present a unique challenge for litigators: they must identify privileged documents within a universe of data that nearly doubles every two years, while keeping quality legal services affordable.¹ Far too often, practitioners are forced to sacrifice one of these two goals—providing affordable legal services and protecting attorney-client privilege—for the sake of preserving the other. This article provides insights that, if implemented, will reduce the pressures of this dilemma.

[2] The attorney-client privilege is one of the oldest principles in the Anglo-American legal system. The privilege protects clients from being forced to disclose confidential communications to or from their attorneys for the purpose of seeking or obtaining legal advice.² Closely related to the attorney-client privilege is a similar protection called the work product doctrine. The work product doctrine protects against the disclosure of any document prepared by an attorney in anticipation of litigation.³ The goal of both doctrines is to encourage clients to be candid in their communications with their attorneys and to incentivize full disclosure. Exposing clients to legal liability based on the contents of their

¹ See Steve Lohr, *The Age of Big Data*, N.Y. TIMES (Feb. 11, 2012), <http://www.nytimes.com/2012/02/12/sunday-review/big-datas-impact-in-the-world.html> [<https://perma.cc/6A9R-LNL8>]. In 2018, an estimated 281 billion emails are sent and received per day, world-wide. That number is expected to reach 333 billion by the end of 2022. The Radicati Group, Inc., *Email Statistics Report, 2018-2022*, THE RADICATI GROUP, INC. (Mar. 2018), https://www.radicati.com/wp/wp-content/uploads/2018/01/Email_Statistics_Report,_2018-2022_Executive_Summary.pdf [<https://perma.cc/AA7S-RL5L>].

² See FED. R. EVID. 502(g)(1) (stating that attorney-client privilege is “the protection that applicable law provides for confidential attorney-client communications . . .”).

³ See FED. R. EVID. 502(g)(2) (stating that work-product protection is “the protection that applicable law provides for tangible material (or its intangible equivalent) prepared in anticipation of litigation or for trial.”).

communications with their legal counsel would chill candor and make adequate representation more difficult.⁴

[3] Since protecting privileged information is so critical, legal teams place a premium on identifying privileged documents during discovery. The traditional method of filtering privileged information from document productions was straightforward—teams of attorneys pored over rooms full of documents, manually inspecting each one for privileged content.⁵ The digital era, and the vast data growth that accompanied it, quickly made the traditional model cost-prohibitive: a typical commercial litigation matter between two large corporations often includes millions of documents.⁶ To combat this information glut, legal teams started using keyword searches to target documents that were most likely to contain privileged material.⁷ As the art of keyword searching has developed over the last few decades, many practitioners have developed internal lists of search terms that they believe to be effective.⁸ These search term lists, however, are typically the product of trial-and-error learning.⁹

[4] Until now, no researchers have examined individual search terms' effectiveness at identifying privileged material. Furthermore, our study

⁴ See *Upjohn Co. v. United States*, 449 U.S. 383, 389 (1981) (asserting that the purpose of the Federal Rules of Evidence regarding attorney-client privilege is to “encourage full and frank communication between attorneys and their clients . . .”).

⁵ See THE SEDONA CONFERENCE, *The Sedona Conference Best Practices Commentary on the Use of Search & Information Retrieval Methods in E-Discovery*, 15 THE SEDONA CONF. J. 217, 222 (2014).

⁶ See *id.* at 243.

⁷ See *id.* at 227. See also Mia Mazza et al., *In Pursuit of FRCP 1: Creative Approaches to Cutting and Shifting the Costs of Discovery of Electronically Stored Information*, 13 RICH. J.L. & TECH., no. 3, at ¶ 20 (2007).

⁸ See THE SEDONA CONFERENCE, *supra* note 5, at 234.

⁹ See *id.* at 233–34 (explaining different approaches to building keyword searches).

reveals that many long-held assumptions within the legal community are antiquated. The word “counsel,” for instance, is widely considered too broad of a search term and therefore not an efficient indicator of privileged information.¹⁰ Our research shows otherwise. Similarly, the conventional wisdom among practitioners is that predictive modeling is inferior to keyword searching and too unreliable to be useful in privilege reviews¹¹—our results contradict this belief.

[5] This article gives practitioners evidenced-based practices for conducting keyword searches, compares keyword search terms to predictive modeling, and provides legal teams with techniques to use both technologies in a complementary manner to achieve maximum effectiveness. We begin by emphasizing the importance of privilege protections and explaining how inadvertent disclosure of privileged information can influence the outcome of a case. We then examine the current methods of identifying privileged material in large-scale litigation and conclude the article by discussing the results of our study and outlining our recommendations about the best ways to identify privileged information in the future.

¹⁰ See, e.g., Kaushal Jha, *How to Kick Privileged Information Out of Your Production Set*, RELATIVITY BLOG (Jan. 22, 2016), <https://www.relativity.com/blog/how-to-kick-privileged-information-out-of-your-production-set/> [<https://perma.cc/RU3W-72JZ>] (noting that terms such as “legal,” “attorney,” “lawyer,” “privilege,” and “counsel” may “return a great number of false hits” and “capture documents containing confidential disclaimers in the footer of an email”).

¹¹ See generally Tonia Hap Murphy, *Mandating Use of Predictive Coding in Electronic Discovery: An Ill-Advised Judicial Intrusion*, 50 AM. BUS. L.J. 609, 611, 621 (2013) (discussing advantages of keyword searches); Charles Vaccaro, Note, *Look Before You Leap into Predictive Coding: An Argument for a Cautious Approach to Utilizing Predictive Coding*, 41 RUTGERS COMPUTER & TECH. L.J. 298, 318 n.136 (2015) (discussing the unreliability of predictive modeling for identifying privilege).

II. THE PROBLEM: SOARING DISCOVERY COSTS AND INEFFICIENT SEARCH TECHNIQUES

[6] Inadvertently disclosing privileged information can undermine a client's position and jeopardize her legal claim altogether. Therefore, it is critical for attorneys to take practical measures to safeguard such information. Diligent privilege review in high-stakes, large-scale litigation is an expensive endeavor, however. Each additional hour spent scouring documents adds to the client's bill, and every year the number of documents attorneys must review substantially increases. The result? Law firms are forced to strike a balance between their clients' privilege and cost-effective review. Yet, as the volume of digital data increases each year, that balance becomes more difficult to maintain.

A. What Is the Attorney-Client Privilege, and Why Is Protecting Privilege So Important?

[7] Robust protections of privileged information are a critical part of the U.S. legal system and essential to sound legal advice and strategy. Legal scholars have long maintained—as early as the 16th Century, if not earlier—that legal privilege constructs allow attorneys to be as informed as possible when rendering legal advice.¹² In his highly influential and often-cited 18th Century treatise on the common law of England, Blackstone wrote that the right to protect communications with one's attorney was a logical and essential corollary to the common-law right against self-incrimination.¹³ The Louisiana Supreme Court observed that

¹² See 1 EDNA EPSTEIN, *THE ATTORNEY-CLIENT PRIVILEGE AND THE WORK-PRODUCT DOCTRINE* 4–5 (5th ed. 2007). Some scholars believe the English common law may have recognized attorney-client privilege as early as 1280. See *In re Selser*, 105 A.2d 395, 403 (N.J. 1954).

¹³ See 3 W. BLACKSTONE, *COMMENTARIES ON THE LAWS OF ENGLAND* *370 (1768) (“[N]o man is to be examined to provide his own infamy. And no counsel, attorney, or other person, intrusted with the secrets of the cause by the party himself, shall be compelled, or perhaps allowed, to give evidence of such conversation or matters of privacy, as came to his knowledge by virtue of such trust and confidence.”); see also

the vital importance of the privilege was already an “unquestioned” principle in English common law by the beginning of the reign of Queen Elizabeth I in 1533.¹⁴ Some legal scholars even cite the laws of Ancient Rome, which prohibited “advocates” from testifying against their clients, as the origin of the attorney-client privilege.¹⁵

[8] The Supreme Court has likewise recognized the attorney-client privilege as “the oldest of the privileges for confidential communications known to the common law.”¹⁶ The purpose of the privilege, in the Court’s words, “is to encourage full and frank communication between attorneys and their clients and thereby promote broader public interests in the observance of law and administration of justice.”¹⁷

[9] Though the precise formulation may vary between jurisdictions, as a general matter, the attorney-client privilege protects communications: (1) between a client and his or her attorney, (2) that intend to be, and are in fact, kept confidential, and (3) for the purpose of obtaining or providing legal advice.¹⁸ The privilege extends to electronic transmissions, as well as

Daniel Northrop, *The Attorney-Client Privilege and Information Disclosed to an Attorney with the Intention That the Attorney Draft a Document to be Released to Third Parties: Public Policy Calls for at Least the Strictest Application of the Attorney-Client Privilege*, 78 *FORDHAM L. REV.* 1481, 1491 n.62 (2009).

¹⁴ See *State v. Green*, 493 So. 2d 1178, 1180 (La. 1986) (“The inception of the attorney-client privilege can be traced to the reign of Elizabeth I where the privilege already appears unquestioned.”).

¹⁵ See James A. Gardner, *A Re-Evaluation of the Attorney-Client Privilege Pt. I*, 8 *VILL. L. REV.* 279, 289–90 (1963) (noting the privilege may have its origins in Roman law).

¹⁶ *Upjohn Co. v. United States*, 449 U.S. 383, 389 (1981).

¹⁷ *Id.* at 389.

¹⁸ See *United States v. Mejia*, 655 F.3d 126, 132 (2d Cir. 2011).

to more traditional forms of communication like letters or verbal conversations.¹⁹ It also applies to non-verbal communication.²⁰

[10] There are additional considerations when considering attorney-client privilege in the corporate context. Under federal common law, the privilege will generally protect communications between the corporation's in-house and outside counsel, and the corporation's employees, made for the purpose of obtaining or providing legal advice.²¹ But in some state jurisdictions, the privilege will only extend to communications between the corporation's counsel and the corporate "control group"—the final decision makers and top advisers of the corporation.²²

[11] The attorney-client privilege is not absolute and can be waived, with the costs of waiver being potentially significant.²³ When the privilege is waived, the communication is no longer confidential, and it cannot be shielded from disclosure to third parties. Disclosing otherwise privileged communications to one's opponent could alter the course and outcome of the litigation. For example, in a litigation context, there may be privileged communications discussing the merits of the litigation, evaluating the key evidence that could be used against a party, or summarizing other potential

¹⁹ See RESTATEMENT (THIRD) OF THE LAW GOVERNING LAWYERS § 69 cmt. b (AM. LAW INST. 2000).

²⁰ See *id.* at § 69 cmt. e.

²¹ See *Upjohn Co.*, 449 U.S. at 394–95.

²² See, e.g., *Consolidation Coal Co. v. Bucyrus-Erie Co.*, 432 N.E.2d 250, 252 (Ill. 1982) (“The appellate court considered the attorney-client privilege inapplicable because there was no allegation by B-E that the disputed documents were received from members of B-E’s ‘control group.’”).

²³ According to one scholar, there are many ways that attorney-client privilege can be nullified. See EPSTEIN, *supra* note 12, at 1–2, 4–5, 407–576.

legal claims that have not been brought.²⁴ Courts have allowed such potentially damaging communications to remain in the case.²⁵ Thus, it is not surprising that one federal court described the disclosure and waiver of privileged information as “the misstep feared by all litigators.”²⁶

[12] Inadvertent disclosure is by far the most common method of waiver.²⁷ Historically, under the subject matter waiver doctrine, a party that inadvertently disclosed privileged material in a negligent or reckless fashion could be found to have waived the privilege, and the court could order the party to disclose all other documents related to the same topic.²⁸ The rigid application of the subject matter doctrine in the digital age

²⁴ See, e.g., *Starr Int’l Co. v. United States*, 121 Fed. Cl. 428, 434–35 (2015), *aff’d in part, vacated in part* by 856 F.3d 953 (Fed. Cir. 2017) (government produced an email exchange with in-house counsel noting that the government’s legal basis for taking over AIG was on “thin ice”).

²⁵ See, e.g., *irth Sols., LLC v. Windstream Commc’ns., LLC*, 2017 WL 3276021 (S.D. Ohio 2017); *In re Google Inc.*, 462 Fed. Appx. 975 (Fed. Cir. 2012); *In re Columbia/HCA Healthcare Corp. Billing Practices Litig.*, 293 F.3d 289, 304 (6th Cir. 2002); *Westinghouse Elec. Corp. v. Republic of the Philippines*, 951 F.2d 1414, 1418 (3d Cir. 1991); *Gruss v. Zwinn*, No. 09 Civ. 6441 (PGG)(MHD), 2013 WL 3481350, at *1-2 (S.D.N.Y. 2013); *Lenz v. Universal Music Corp.*, No. 5:07-cv-03783 JF (PVT), 2010 WL 4789099 (N.D. Cal. 2010); *Rambus, Inc. v. Infineon Techs. AG*, 222 F.R.D. 280, 299 (E.D. Va. 2004) (displaying circumstances where privileged communication remained in litigation).

²⁶ See *Amgen Inc. v. Hoechst Marion Roussel, Inc.*, 190 F.R.D. 287, 288 (D. Mass. 2000); see also Paul W. Grimm et al., *Federal Rule of Evidence 502: Has It Lived Up to Its Potential?*, 17 RICH. J.L. & TECH., no. 3, ¶ 1 (2011) (“Nothing causes litigators greater anxiety than the possibility of doing, or failing to do, something . . . that waives attorney-client privilege or work-product protection.”).

²⁷ See Henry S. Noyes, *Federal Rule of Evidence 502: Stirring the State Law of Privilege and Professional Responsibility with a Federal Stick*, 66 WASH. & LEE L. REV. 673, 690 (2009).

²⁸ See *Amgen*, 190 F.R.D. at 288.

caused discovery costs to soar, spurring Congress to amend the Federal Rules of Evidence in 2008 to reduce the scope of subject matter waivers.²⁹

[13] The Supreme Court has not established a definitive test for the inadvertent disclosure doctrine, and three different approaches have developed in the lower courts over the past few decades.³⁰ The first approach, adopted by a small minority of courts, is the “strict waiver rule,” which treats even unintentional disclosures of privileged information as waivers.³¹ The second approach, followed by even fewer courts, only considers privileged information to be waived by an express and voluntary act by the privilege holder.³² The third approach uses a balancing test very

²⁹ See FED. R. EVID. 502 committee note 2 (“This new rule . . . responds to the widespread complaint that litigation costs necessary to protect against waiver of attorney-client privilege or work product have become prohibitive . . .”); *Olaoye v. Wells Fargo Bank NA*, No. 3:12-CV-4872-M-BH, 2013 U.S. Dist. LEXIS 181358, at *3 (N.D. Tex. Dec. 30, 2013) (noting that congress added 502(b) in response “to the widespread complaint that litigation costs necessary to protect against waiver of privilege have become prohibitive due to the concern that any disclosure will operate as a subject matter waiver of all protected communication.”); see also Symposium, *The Sedona Conference Best Practices Commentary on the Use of Search and Information Retrieval Methods in E-Discovery*, 8 THE SEDONA CONF. J. 189, 192 (2007) [hereinafter Sedona Conference Symposium].

³⁰ See Anthony Francis Bruno, Note, *Preserving Attorney-Client Privilege in the Age of Electronic Discovery*, 54 N.Y. L. SCH. L. REV. 541, 547 (2009); see also *Gray v. Bicknell*, 86 F.3d 1472, 1483 (8th Cir. 1996) (noting the court’s use of three distinct approaches to attorney-client privilege waivers).

³¹ See, e.g., *In re Sealed Case*, 877 F.2d 976, 980 (D.C. Cir. 1989); see Bruno, *supra* note 30, at 547; see also *SEC v. Lavin*, 111 F.3d 921, 929–30 (D.C. Cir. 1997).

³² See, e.g., *Helman v. Murry’s Steaks, Inc.*, 728 F. Supp. 1099, 1104 (D. Del. 1990) (“It would fly in the face of the essential purpose of the attorney/client privilege to allow a truly inadvertent disclosure of a privileged communication by counsel to waive the client’s privilege.”); see also Bruno, *supra* note 30, at 547.

similar to the standard now adopted by Federal Rule of Evidence 502.³³ Under this test, inadvertently disclosing documents does not act as a waiver if the court determines that the disclosing party took reasonable precautions to safeguard the information and promptly took reasonable action to rectify the disclosure.³⁴ Rule 502 explicitly rejects the strict waiver approach.³⁵

[14] Courts have found that the failure to use keyword searches to find privileged documents will be grounds for finding a waiver.³⁶ Similarly, courts have found inadequate keyword search techniques—caused, in turn, by inadequate preparation by attorneys during document review to be sufficient grounds for a waiver of privilege.³⁷ A federal court in Maryland, for example, held that an attorney’s failure to sample the documents that his keyword searches flagged as “non-privileged” was sufficiently

³³ See *Baranski v. United States*, No. 4:11-CV-123, 2015 WL 3505517, at *4 (E.D. Mo. 2015) (noting the similarities between the “middle ground” approach and Rule 502’s factors).

³⁴ See Bruno, *supra* note 30, at 547–48.

³⁵ See FED. R. EVID. 502 Advisory Committee Note Subdivision (a) (“It follows that an inadvertent disclosure of protected information can never result in a subject matter waiver.”); Bruno, *supra* note 30, at 548.

³⁶ See *Talismanic Props., LLC v. Tipp City*, 309 F. Supp. 3d 488, 494–95 (S.D. Ohio 2017) (“The City presents no evidence that any person performs any further review of those search results for potential privileged information before producing them -- such as a review for emails sent or received by the City’s in-house or outside attorneys. Absent evidence showing any steps taken to review the product of email search results for privileged material, the undersigned is unable to conclude that the City’s precautions in this regard are ‘adequate.’”).

³⁷ See, e.g., *Mt. Hawley Ins. Co. v. Felman Prod. Inc.*, 271 F.R.D. 125, 136, 138–39 (S.D. W. Va. 2010) (finding that steps taken to prevent inadvertent disclosure of privileged content within a large document set were unreasonable because attorneys failed to test the reliability of their keyword searches).

negligent to waive the attorney-client privilege.³⁸ In doing so, the court observed that “all keyword searches are not created equal.”³⁹ As noted above, the quality of a legal team’s keyword searches can have case-determinative effects, because disclosing privileged documents—and therefore waiving the privilege—can easily influence the outcome of a dispute.

[15] The amendment to Federal Rule 502 was a necessary step to curb discovery excesses in the digital age, but the amendment’s changes do not protect clients from the dangers of inadvertent disclosure entirely. Rule 502(d) contains a “claw back provision” that allows parties to reassert the privileged nature of inadvertently disclosed documents and request a court order protecting the information contained within them.⁴⁰ The protections against inadvertent waiver under 502(d) can be expansive, and can protect against waiver in the specific litigation at issue, as well as “in any other federal or state proceeding.”⁴¹ However, not all proceedings are controlled by a 502(d) court order. Without such an order, the protections against inadvertent waiver are guided by the following requirements laid out in Rule 502(b): (1) the disclosure was truly inadvertent, (2) “the holder of the

³⁸ See *Victor Stanley, Inc. v. Creative Pipe, Inc.*, 250 F.R.D. 251, 256–57 (D. Md. 2008).

³⁹ *Id.*; see also *United States v. Brewington*, No. 15-CR-00073-PAB, 2018 U.S. Dist. WL 1046804, at *3 (D. Colo. Feb. 26, 2018) (finding that party did not take reasonable precautions to prevent inadvertent disclosure when they did not search the text of emails for the names of attorneys). Commentary to recent revisions to Rule 502 suggest that courts will also review a party’s use of other computer-assisted tools in a privilege review with an eye towards whether the party took reasonable steps to prevent the inadvertent disclosure of privileged materials. See Dennis R. Kiker, *Defensible Data Deletion: A Practical Approach to Reducing Cost and Managing Risk Associated with Expanding Enterprise Data*, 20 RICH. J. L. & TECH., no. 2, ¶ 11 (2014) (citing FED. R. EVID. 502(b) Advisory Committee Notes, Subdivision (b)).

⁴⁰ FED. R. EVID. 502(d).

⁴¹ See U.S. Magistrate Judge Andrew J. Peck, Rule 502(d) Order, https://www.txs.uscourts.gov/sites/txs/files/q27FzLwUTjurtSVVYE02_ILTA001_Rule_502d_Order.pdf [<https://perma.cc/9DGL-LEU3>].

privilege took reasonable steps to prevent disclosure” of the information, and (3) the disclosing party “took reasonable steps to rectify the error.”⁴²

[16] The application of 502(b)’s elements, however, is more complex than may appear at first glance. As one federal court explained, “Rule 502(b) sets out three elements that must be met to invoke the protections of the inadvertent disclosure rule, [but] waiver of privilege under the Rule is a flexible analysis.”⁴³ Indeed, as Professor Ann Murphy notes, “[t]he new inadvertent privilege evidence rule has been interpreted in many different ways by courts, creating uncertainty, and it is not a panacea for the attorney who inadvertently discloses privileged material.”⁴⁴

[17] Furthermore, no matter how courts apply Rule 502(b) elements, the mere existence of the claw back provision does not eliminate the potential headaches caused by disclosing privileged information in the first instance, for several reasons. At the most basic level, information that has been disclosed cannot be “unlearned” by the opposing party, even if the opposing party cannot explicitly use it. Inadvertent disclosure also creates the risk of additional, tangential litigation—the opposing party may choose to contest the privileged nature of the inadvertently-disclosed documents.⁴⁵ As Michael Correll of Morgan, Lewis & Bockius LLP observed, “a disclosing party will be faced with the very high likelihood that the receiving party will work vigorously to admit these particularly adverse privileged documents.”⁴⁶ If the receiving party successfully

⁴² FED. R. EVID. 502(b).

⁴³ U.S. *ex rel.* Schaengold v. Mem’l Health, Inc., No. 4:11-cv-58, 2014 U.S. Dist. LEXIS 156595, at *14 (S.D. Ga. Nov. 5, 2014).

⁴⁴ Ann M. Murphy, *Federal Rule of Evidence 502: The Get Out of Jail Free Provision – or Is It*, 41 N.M. L. Rev. 193, 194 (2011).

⁴⁵ *See, e.g.*, Rhoads Indus. Inc. v. Bldg. Materials Corp. of Am., 254 F.R.D. 216, 216 (E.D. Pa. 2008).

⁴⁶ Michael Correll, *The Troubling Ambition of Federal Rule of Evidence 502(d)*, 77 Mo. L. REV. 1031, 1070 (2012).

preserves the inadvertently disclosed material, it may further press the issue to force the disclosure of all other documents relating to the same topic. Parties that use this tactic successfully will effectively recreate the subject-matter waiver—one of the very outcomes the Rule 502 amendments were designed to prevent.

B. The Challenge and Costs of Identifying Privileged Documents in the Digital Age

[18] Since protecting privileged information is so vital, it absorbs a significant amount of resources and attention during the discovery phase of litigation. Safeguarding that information has become more difficult in the digital age, and attorneys have sought innovative solutions to keep up. The legal community has had some success in these efforts, but much work remains to be done.

1. Rising Costs

[19] A Rand survey of parties in fifty-seven separate cases found that parties frequently spend millions of dollars simply preparing documents for production.⁴⁷ In one case, the total cost of document production alone was \$27 million.⁴⁸ An average of 73% of the costs incurred during document production occurred during the document review phase.⁴⁹ A 2013 study conducted by Microsoft revealed that the software giant stores an average of sixty million pages every time a party files a claim against

⁴⁷ See Nicholas M. Pace & Laura Zakaras, *Where the Money Goes: Understanding Litigant Expenditures for Producing Electronic Discovery*, RAND INST. FOR CIVIL JUSTICE 17 (2012), http://www.rand.org/content/dam/rand/pubs/monographs/2012/RAND_MG1208.pdf [<https://perma.cc/6CWQ-RGAC>].

⁴⁸ See *id.* at 17.

⁴⁹ See *id.* at xiv.

them.⁵⁰ Microsoft's legal team pares that number down to 350,000 documents by filtering by issue, source, and dates, but the remaining documents must be reviewed manually.⁵¹ Microsoft estimates that it has spent roughly \$600 million over the last decade on outside services—namely, counsel and e-discovery vendors—to assist with discovery.⁵² Yet, despite these massive expenditures on review, mistakes are still made and privileged documents are still produced.⁵³

[20] The burdens of discovery in the digital age—preserving, reviewing, and producing millions of documents—recently spurred a new amendment to the Federal Rules of Civil Procedure. The 2015 amendment to Rule 26(b)(1) requires discovery requests to be proportional to the needs of the case and strike a proper balance between the benefit of information and burden of producing it.⁵⁴ The amendment explicitly acknowledges the high discovery costs of the new era, and explains:

The burden or expense of proposed discovery should be determined in a realistic way. This includes the burden or expense of producing electronically stored information. Computer-based methods of searching such information continue to develop, particularly for cases involving large volumes of electronically stored information. Courts and

⁵⁰ See *Needles in Haystacks: The Secret Burden Holding Back our Economy*, MICROSOFT (Nov. 25, 2013), <https://blogs.microsoft.com/on-the-issues/2013/11/25/needles-in-haystacks-the-secret-burden-holding-back-our-economy/#pwAdZg3Fr7Clxwi1.99> [<https://perma.cc/K5H3-WVHR>].

⁵¹ See *id.*

⁵² See *id.*

⁵³ See Ralph C. Losey, *Predictive Coding and the Proportionality Doctrine: A Marriage Made in Big Data*, 26 REGENT U. L. REV. 7, 14 (2013) (summarizing a lengthy construction case from 2012 where both sides inadvertently produced thousands of privileged documents despite spending tens of millions of dollars on review).

⁵⁴ See FED. R. CIV. P. 26(b)(1).

parties should be willing to consider the opportunities for reducing the burden or expense of discovery as reliable means of searching electronically stored information become available.⁵⁵

[21] To summarize, the vital task of safeguarding communications made between clients and their attorneys has never been more expensive than it is today, and attorneys are eagerly seeking solutions to the problem.

2. Current Methods: Keyword Searching and Predictive Modeling

a. Keyword Searching

[22] Keyword searching has become the common practice for identifying privileged documents. The keyword-search method requires legal teams to develop lists of search terms, tailored to the details of the individual matter, that they believe are most likely to indicate the presence of privileged information.

[23] Recall that the privilege protects confidential communications between client and attorney made for the purpose of obtaining or providing legal advice. Particular keywords will target different elements of the privilege. Since the privilege protects communications involving counsel, a legal team might develop a list of law firm names and the names of in-house and outside counsel that have done legal work for the client.⁵⁶ The thinking behind this is that a communication involving one of

⁵⁵ FED. R. CIV. P. 26(b)(1) Advisory Committee's Comments on the 2015 Amendment.

⁵⁶ See *United States v. Brewington*, No. 15-cr-00073-PAB, 2018 U.S. Dist. LEXIS 30425, at *8 (D. Col. Feb. 26, 2018) (noting that searching for the names of individuals in the email address field was reasonable calculated to prevent disclosure of privileged emails); see also *Cole's Wexford Hotel, Inc. v. UPMC*, No. 10-1609, 2016 U.S. Dist. LEXIS 15035, at *6–7 (W.D. Pa. Feb. 8, 2016) (noting that counsel took reasonable steps to protect the privilege, including searching for in-house and outside counsel names in the full text of documents).

these lawyers or law firms is more likely to be a communication made for the purpose of obtaining or providing legal advice. In addition, legal teams might draft a list of terms that are potentially indicative of a request for legal advice or the provision of legal advice. These terms might range from narrow phrases—“attorney client communication” or “prepared at the request of counsel”—to very broad terms—“legal,” “counsel,” “confidential,” or “privileged.” Legal teams will then apply those lists of search terms to the documents eligible for production for that particular matter and often conduct a manual review of the resulting documents to confirm or deny the presence of privileged information in each one.

[24] At best, a list of keyword search terms will be an imperfect predictor of privilege, as there is no standardized set of terms that are used when requesting or providing legal advice. Nor is there any standardized format for a privileged document. At one end of the spectrum, there are formal memos from lawyer to client that clearly denote that the memo was created for purpose of advising a client about a legal matter. In today’s world, fewer and fewer privileged documents will follow this format. Legal advice is increasingly conveyed in email messages between lawyer and client that may not provide much context about the subject matter of the advice. That legal advice may find its way into PowerPoint presentations, Word documents, and other loose electronic files where the content of the legal advice may not be apparent from the face of the document.

[25] The effectiveness, or lack thereof, of the search terms selected by a legal team can have wide-ranging impacts on both the cost and outcome of litigation.⁵⁷ A keyword search term list’s performance depends upon the legal team’s understanding of both the document set and their client’s

⁵⁷ See Raymond Biederman & Sean Burke, *Biederman and Burke: Is Use of Keywords in E-Discovery a Game of ‘Go Fish?’*, THE INDIANA LAWYER (Nov. 16, 2016) (describing “poorly thought-out” searches as time-consuming and expensive), <https://www.theindianalawyer.com/articles/42021-use-of-keywords-in-e-discovery-a-game-of-go-fish> [<https://perma.cc/Q3T6-QSVM>].

business history.⁵⁸ Choosing words that are too broad will create a high number of “false positives” requiring costly and unnecessary manual review.⁵⁹ Choosing words that are too narrow will result in an incomplete review that inadvertently discloses privileged material. Legal teams with limited knowledge of the documents and business history related to a litigation matter may develop term lists that are over- or under-inclusive, leading to poor search results.⁶⁰ Judge Paul Grimm of the Southern District of New York highlighted this risk in *Victor Stanley, Inc. v. Creative Pipe*, noting “a growing body of literature that highlights the risks associated with conducting an unreliable or inadequate keyword search or relying exclusively on such searches for privilege review.”⁶¹ The U.S. District Court for the Western District of Virginia likewise remarked that it was “aware of [keyword searches’] limitations.”⁶² The Court observed that “simple keyword searching is inadequate . . . because simple keyword searches end up being both over and under-inclusive. . .”⁶³

[26] Even when search lists are created competently, however, locating privileged documents in discovery is still costly. Terms like “privilege”

⁵⁸ See Andrew Peck, *Search, Forward*, LAW.COM (Oct. 1, 2011), https://law.duke.edu/sites/default/files/centers/judicialstudies/TAR_conference/Panel_1-Background_Paper.pdf [<https://perma.cc/RMJ5-UVZ8>].

⁵⁹ See Biederman & Burke, *supra* note 57.

⁶⁰ See *Am. Capital Homes, Inc. v. Greenwich Ins. Co.*, No. C09-0622-JCC, 2010 WL 11561400, at *3 (W.D. Wa. Aug. 3, 2010) (internal citations omitted) (“[T]he only way to properly test the reliability of a keyword search is to sample the documents so as to determine whether the search was over or under-inclusive.”) *Id.*; see also Biederman & Burke, *supra* note 57.

⁶¹ *Victor Stanley, Inc. v. Creative Pipe, Inc.*, 250 F.R.D. 251, 257 (D. Md. May 29, 2008) (holding that a party waived a privilege after mistakenly disclosing 165 documents due to a failure to use adequate keyword search terms).

⁶² See *Adair v. EQT Production Co.*, No. 1:10cv00037, 2012 WL 1965880, at *5 (W.D. Va. May 31, 2012).

⁶³ *Id.* at 5 (quoting Sedona Conference Symposium, *supra* note 29, at 201).

and “confidential” are frequently included in privilege search term lists and they often include “wildcard” syntax to account for variations of word usage within documents. When combined with a wildcard, “privilege” and “confidential” become “priv*,” and “confid*,” increasing the possibility of false positives, but hopefully capturing more privileged documents.⁶⁴ In this example, the search will retrieve documents that contain words like “private” or “confident,” in addition to “privileged” or “confidential.” The names, email addresses, and web domains associated with attorneys will also increase the number of false positives in the search results—especially when the individuals involved have common names, such as Smith, Williams, Brown, or Adams. Confirming these search results requires teams of lawyers to spend hours conducting the costly and time-consuming manual review process—often at the cost of hundreds of dollars per hour, per lawyer.

[27] To create effective keyword search lists, legal teams often expend considerable effort identifying the legal parties that have interacted with the client and its employees and the nature of those interactions. Companies that retain multiple outside counsel, have a long history of litigation, or a history of investigations by enforcement agencies are especially challenging for legal teams developing privilege keyword term lists. Companies of this nature could require thousands of terms to account for every potentially privileged name, word, and legal domain. If legal teams do not obtain clear insight into all of the key legal players and events, their privilege keyword term lists will be incomplete by definition, creating the risk that privileged material could “survive” the keyword search and make it into the production to the opposing party.

[28] Judge Andrew Peck of the Southern District of New York described many practitioners’ keyword selection processes as “the

⁶⁴ See, e.g., *Dornoch Holdings Int’l, LLC v. ConAgra Foods Lamb Weston, Inc.*, No. 1:10-CV-00135 TJH, 2013 WL 2384235, at *4 (D. Idaho May 1, 2013) (observing that “more general search terms” were less effective at identifying privilege documents where, for example “privilege* NOT w/25 (intended or received or dissemination or addressee)” correctly identified a privileged document 13% of the time).

equivalent of the child's game of 'Go Fish.'"⁶⁵ Judge Peck is not alone in his beliefs about the shortcomings of many keyword searches. In *Am. Capital Homes v. Greenwich Ins. Co.*, the Western District of Washington held that "it [was] not judicial micromanagement to note" that the plaintiffs relied on untested, "simple keyword searches" despite previous courts' criticism of such practices.⁶⁶ Other federal courts have undertaken similarly critical assessments of poorly constructed keyword searches. In its review of an inadvertent disclosure claim, for example, the U.S. District Court for the District of Maryland noted, "[w]hile the Court is not aware of how complex the corporate structure of the [defendant] might be, it would seem that identifying the name of in-house counsel would have been the first step of a reasonable privilege review."⁶⁷ The defendant's motion to protect the inadvertently produced document was denied.⁶⁸

[29] The problem posed by incomplete searches is compounded by many attorneys' overconfidence in the effectiveness of their search methods. In a famous study conducted in 1985, researchers David Blair and M.E. Maron gave legal teams a set of 40,000 documents from a real legal matter and asked them to locate relevant documents within the set.⁶⁹ After conducting their searches and subsequent manual review, each of the teams estimated that they had located at least 75% of the relevant

⁶⁵ *Moore v. Publicis Groupe*, 287 F.R.D. 182, 190–91 (S.D.N.Y. 2012).

⁶⁶ *Am. Capital Homes, Inc. v. Greenwich Ins. Co.*, No. C09-0622-JCC, 2010 WL 11561400, at *3 (W.D. Wash. Aug. 3, 2010).

⁶⁷ *LES Engineering, Inc. v. Corus*, No. WMN-08-2115, 2009 WL 10682245, at *3 (D. Md. Aug. 14, 2009).

⁶⁸ *Id.*

⁶⁹ See David C. Blair & M.E. Maron, *An Evaluation of Retrieval Effectiveness for a Full-Text Document-Retrieval System*, 28 COMM. OF THE ACM 289, 290–91 (1985).

documents in the set.⁷⁰ Their actual success rate? Roughly twenty percent.⁷¹

[30] The Sedona Conference’s Best Practices Guide to keyword searching summarizes the problem:

The limitations of keyword approaches to search and retrieval first exposed in the Blair and Maron study, and validated in subsequent research, have not faulted the ability of computers to locate documents meeting the attorneys' search criteria – but rather the inability of the attorneys and paralegals to anticipate all of the possible ways that people might refer to the issues in the case. The richness and ambiguity of human language causes severe challenges in identifying relevant information.⁷²

[31] In the Blair and Maron study, which used a document set related to a San Francisco public transportation accident, city officials referred to the accident as “the unfortunate situation,” while the victim and related parties referred to it as a “disaster.”⁷³ In other places, terms like the “event,” “incident,” “situation,” “problem,” or “difficulty” were used, further complicating the search process.⁷⁴ Blair and Maron focused on searches for relevant documents, but their insights are equally true for privilege-

⁷⁰ *See id.* at 293.

⁷¹ *See* THE SEDONA CONFERENCE, *supra* note 5, at 239.

⁷² *Id.* at 240.

⁷³ *See* Sedona Conference Symposium, *supra* note 28, at 206.

⁷⁴ Gregory L. Fordham, *Using Keyword Search Terms in E-Discovery and How They Relate to Issues of Responsiveness, Privilege, Evidence Standards, and Rube Goldberg*, 15 RICH. J. L. & TECH., no. 3, ¶ 14 (2009) (demonstrating the difficulties that language variation presents within the document review process).

related searches. A federal court recently noted that Blair and Maron's results have been confirmed and "replicated . . . over the past few years."⁷⁵

b. Predictive Modeling

[32] As the amount of digitally-stored data increases and discovery becomes more complex, the legal industry has sought innovative solutions to keep up. One innovation in particular, predictive modeling, also known as "technology-assisted review" or "predictive coding," has proven especially effective. Maura Grossman and Gordon Cormack, summarized the process as follows:

*A technology-assisted review process involves the interplay of humans and computers to identify the documents in a collection that are responsive to a production request, or to identify those documents that should be withheld on the basis of privilege. A human examines and codes only . . . a tiny fraction of the entire collection. Using the results of this human review, the computer codes the remaining documents in the collection for responsiveness (or privilege). A technology-assisted review process may involve, in whole or in part, the use of one or more approaches including, but not limited to, keyword search, Boolean search, conceptual search, clustering, machine learning, relevance ranking, and sampling.*⁷⁶

⁷⁵ Moore v. Publicis Groupe, 287 F.R.D. 182, 191 (S.D.N.Y. 2012).

⁷⁶ Maura R. Grossman & Gordon V. Cormack, *Technology-Assisted Review in E-Discovery Can Be More Effective And More Efficient Than Exhaustive Manual Review*, 17 RICH. J. L. & TECH., no. 3, ¶ 2 (2011) [hereinafter Grossman & Cormack] (citations omitted). Grossman and Cormack have defined predictive coding as "[a]n industry-specific term generally used to describe a Technology-Assisted Review process involving the use of a Machine Learning Algorithm to distinguish Relevant from Non-Relevant Documents, based on Subject Matter Expert(s)' Coding of a Training Set of Documents." Maura R. Grossman & Gordon V. Cormack, *The Grossman-Cormack Glossary of*

[33] Put differently, predictive modeling is a process by which “computers are programmed to search large quantities of documents . . . to mimic the document selection process of a knowledgeable, human document review.”⁷⁷ Underlying predictive modeling is a process of building a model using a machine learning algorithm. At a high level, machine learning is an artificial intelligence field that studies methods to organize or classify data by analyzing the patterns, information, and features within that data.⁷⁸ Machine learning algorithms are frequently used to build predictive models from historical data for making predictions and, by analyzing data, the algorithms can continue to improve their models and produce more accurate results.⁷⁹ In the legal context, the use of machine learning typically is known as “predictive coding.” The predictive coding algorithms build a predictive model that automatically classifies documents of legal interest into predefined categories, such as whether a document is privileged or responsive to a particular Rule 34 document request.⁸⁰

[34] The term “predictive coding” refers to how human reviewers, typically attorneys, code sets of documents that the algorithm uses to create a predictive model and the model analyzes other documents and classifies them as relevant or non-relevant.⁸¹ This process of “learning”

Technology-Assisted Review, 7 FED. CTS. L. REV. no. 1, at 26 (2013) [hereinafter *Glossary*].

⁷⁷ Charles Yablon & Nick Landsman-Roos, *Predictive Coding: Emerging Questions and Concerns*, 64 S.C. L. REV. 633, 634 (2013).

⁷⁸ See Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 89, 94–95 (2014).

⁷⁹ See Emily Berman, *A Government of Laws and Not Machines*, 98 B.U. L. REV. 1277, 1286 (2018).

⁸⁰ See Nathaniel Huber-Fliflet et al., *Empirical Evaluations of Preprocessing Parameters' Impact on Predictive Coding's Effectiveness*, IEEE, at 1 (2016).

⁸¹ See *Glossary*, *supra* note 76, at 26 (2013).

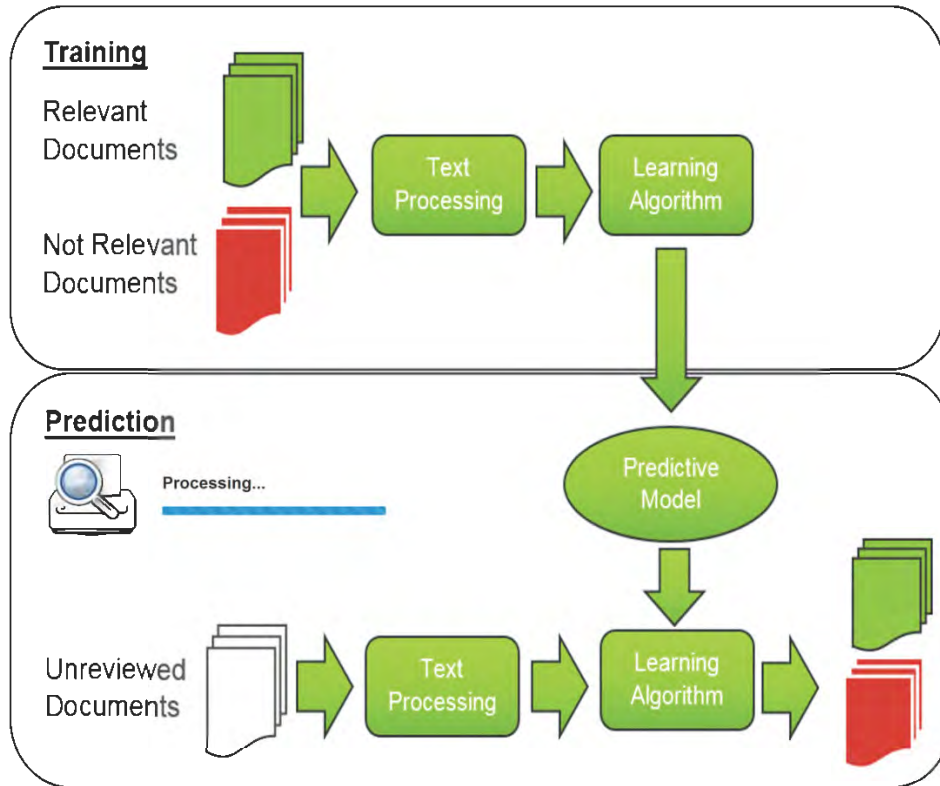
enables the predictive model to improve from experience and to increase its proficiency at identifying relevant material without being explicitly and repeatedly programmed. During this process, the predictive model assigns a probability score to every document in the set, indicating the likelihood that the document contains relevant material.⁸² The higher a document's score is, the greater the possibility that it contains relevant material.⁸³ The inverse is true of a lower score.⁸⁴ Depending on the specific predictive coding protocol being utilized, new human-reviewed documents might be fed into the machine learning algorithm to continue to improve the predictive model.⁸⁵ The predictive modeling process is illustrated by the following graphic:

⁸² See Yablon & Landsman-Roos, *supra* note 77, at 638.

⁸³ See *id.* at 641.

⁸⁴ See *id.* at 642.

⁸⁵ See Grossman & Cormack, *supra* note 76, at 289–93 (describing a “Continuous Active Learning” protocol for predictive coding where the algorithm is continuously retrained as the human reviewer codes documents).



[35] Predictive coding’s effectiveness and cost-efficiency has led to it being “described as a fundamental change in the way discovery is conducted.”⁸⁶ For this reason, the 2017 Sedona Conference declared that predictive coding is “widely accepted [within the legal community] for limiting e-discovery to relevant documents and effecting discovery of [electronically stored information] without an undue burden.”⁸⁷ Indeed, courts frequently allow parties to use predictive coding to respond to

⁸⁶ See Yablon & Landsman-Roos, *supra* note 77, at 634.

⁸⁷ THE SEDONA CONFERENCE, *The Sedona Conference TAR Case Law Primer*, 18 SEDONA CONF. J. 1, 48 (2017) [hereinafter *Sedona Conference TAR*].

discovery requests.⁸⁸ A few courts have even suggested its use *sua sponte*.⁸⁹ For those courts that have supervised the use of predictive coding, several have set a 75% recall rate as the sufficient threshold for a predictive model used in discovery.⁹⁰ Multiple government agencies, including the Federal Trade Commission and the Department of Justice, have approved the use of predictive coding to identify and review documents for production during investigations as well.⁹¹

[36] Predictive modeling has proven to be very effective at identifying relevant documents, but there is a widely held belief in the legal community that it is incapable of mimicking the nuanced analysis required for privilege decisions. In a white paper on e-discovery, for example, one senior practitioner declared that predictive modeling has not proven particularly reliable for privilege calls.⁹² Another prominent e-discovery attorney concurred, saying “most predictive [modeling] engines have yet to demonstrate reliable results in identifying privileged, highly

⁸⁸ See, e.g., *Rio Tinto PLC v. Vale S.A.*, 306 F.R.D. 125, 127 (S.D.N.Y. 2015) (calling it “black letter law” that producing parties are allowed to use TAR for document review); see also *Nat’l Day Laborer Org. Network v. U.S. Immigr. & Customs Enf’t Agency*, 877 F. Supp. 2d 87, 109 (S.D.N.Y. 2012) (authorizing the use of predictive coding and noting that predictive coding “can significantly increase the effectiveness and efficiency of searches.”).

⁸⁹ See *Mot. for Partial Summ. J. Mot. to Dismiss Countercl. and Ruling of the Ct.*, at 66–67, *EOHRB, Inc. v. HOA Holdings, LLC*, (No. 7409-VCL), 2013 Del. Ch. LEXIS 336 (2012); see also *Civil Minutes of Status Conference at 1-2, Indep. Living Ctr. S. Cal. v. City Los Angeles*, (No. CV 12-551-FMO (PJWx)) (C.D. Cal. June 26, 2014).

⁹⁰ See *Sedona Conference TAR*, *supra* note 87, at 37–39 (describing the recall thresholds of courts that have addressed the issue of reasonability of TAR results).

⁹¹ See *id.* at 42–44.

⁹² See Wallis M. Hampton, *Predictive Coding: It’s Here to Stay*, PRAC. L.J., at 28, 30 (2014) <http://docplayer.net/amp/3638887-Traditionally-the-gold-standard-for-identifying-potentially.html> [<https://perma.cc/3RB3-JRCS>].

confidential or ‘hot’ documents.”⁹³ A 2013 study concluded that predictive modeling is used primarily to cull document productions for responsiveness, “because no predictive [modeling] solution has proven fully effective for privilege classification.”⁹⁴

[37] Predictive coding’s inefficiency at identifying privileged documents is largely due to the individualized nature of each privileged document.⁹⁵ For example, a non-privileged email might contain a nearly identical message as a privileged email—with the email’s intended recipient being the only, but nonetheless critical, difference between the two. The difference between privileged and non-privileged material can be determined by a single phrase, or even the context in which the phrase itself is used.⁹⁶ E-discovery teams can create algorithms that filter documents based on patterns and content, but it cannot teach them to make fact-based and context-specific judgment calls often involved in privilege determinations. Furthermore, the vast majority of documents within a document-review set are not privileged.⁹⁷

[38] Before we conducted our study, little research had been published about the use of predictive models to target privileged information, and no practice group had published side-by-side comparisons of predictive

⁹³ Jason Lichter & Michael Frankel, *Facts and Fictions Underlying the Predictive Coding Revolution*, PEPPER HAMILTON (Feb. 18, 2014), <http://www.pepperlaw.com/publications/facts-and-fictions-underlying-the-predictive-coding-revolution-2014-02-18/> [https://perma.cc/AEM6-BR63].

⁹⁴ Manfred Gabriel et al., *The Challenge and Promise of Predictive Coding for Privilege*, INT’L CONF. ARTIFICIAL INTELLIGENCE & L., at 3 (2014).

⁹⁵ *See id.* at 2 (noting various challenges to finding privileged documents using predictive coding).

⁹⁶ *See id.* (stating that the determination as to whether legal advice was sought or rendered may be nuanced and subtle).

⁹⁷ *See id.*

modeling versus keyword searching to identify privileged material.⁹⁸ Our results revealed that keyword searching sometimes identified privileged material with greater precision than predictive modeling, but also revealed scenarios in which predictive modeling was more accurate than keyword searching and ways that predictive modeling can be used to enhance the efficiency of keyword searches.

[39] For example, our results demonstrated that predictive modeling can enable document review teams in large-scale litigation to prioritize documents that are likely privileged by reviewing the highest-scoring documents first. In addition, practitioners can gain insight into the precision of a keyword term before document review even begins by using predictive modeling and keyword searching in a complementary manner—knowledge that legal teams could not otherwise obtain by using keyword searches alone, until after review has concluded. Lastly, predictive modeling can identify privileged documents that keyword searching misses.

[40] The results of this study confirmed that some keyword searches are an effective privilege-targeting method—and that other common keywords are not. We further found that predictive modeling using machine learning can provide innovative ways to locate privileged documents within an ever-expanding universe of digital data. Embracing these results, this research suggests the way forward points toward privilege review that uses a combination of targeted keyword searching and machine learning. Part III provides insights about how to do so.

III. THE WAY FORWARD: AN EMPIRICAL STUDY ON KEYWORD SEARCHING AND PREDICTIVE MODELING AT IDENTIFYING PRIVILEGED DOCUMENTS

[41] The purpose of this study was to evaluate the effectiveness of keyword searching in privilege review and to compare the performance of predictive modeling versus that of keyword searching. Our results will

⁹⁸ See generally *id.* (serving as a notable exception to this statement).

help practitioners identify privileged documents with greater accuracy and efficiency—enhancing protections over privilege and helping legal teams provide affordable legal services in the process.

A. Our Experiment

[42] To ensure the results were as realistic as possible, we performed a “look back” analysis using data sets from three confidential, non-public, real legal matters. All the documents—each data set included email, Microsoft Office documents, PDFs, and other text-based documents—were previously reviewed and received attorney coding during previous privileged reviews, providing us with objective data sets against which we could measure the results. Our search term lists included standard terms such as “privileged,” “legal,” and “attorney,” as well as terms that were unique to the case, like attorney names and email addresses. Table 1 summarizes the statistics of the three data sets used in our experiments and the number of keyword search terms used in each dataset.

Table 1: Summary Statistics of Data Sets and Keyword Search Terms

Project Name	Total Documents	Privileged Documents	Not Privileged Documents	Richness	Number of Keyword Search Terms
Project A	360,531	46,756	313,775	12.97%	845
Project B	397,289	14,326	382,963	3.61%	6,771
Project C	8,715,165	536,788	8,178,377	6.16%	7,140

[43] After applying each matter’s keyword search term list to its respective data set, we calculated the recall and precision rates of each list. Recall and precision rates, two commonly used performance measurement metrics,⁹⁹ were calculated to evaluate the effectiveness of the keywords. The recall rate quantifies the proportion of privileged documents in the

⁹⁹ See THE SEDONA CONFERENCE, *supra* note 5 at 237 (2014).

data set that are identified by the keyword search term list, or other privilege-targeting method—helping to establish the completeness of the privilege review. The precision rate quantifies the proportion of documents identified by the keyword search term list, or other privilege-targeting method, that are *actually* privileged—helping confirm the efficiency of the review for privileged documents. Recall and precision are usually inversely proportionate measures: as recall rates increase, precision rates usually decrease, and vice versa.¹⁰⁰

[44] We also conducted experiments to test the effectiveness of predictive modeling by creating predictive models using a typical supervised learning approach.¹⁰¹ Training sets were generated for each data set using 5,000 randomly sampled documents. Table 2 summarizes the statistics of the training document sets.

Table 2: Summary Statistics of Training Sets

Project Name	Total Training Documents	Privileged Documents	Not Privileged Documents	Richness
Project A	5,000	689	4,311	13.78%
Project B	5,000	170	4,830	3.40%
Project C	5,000	326	4,674	6.52%

[45] After the predictive models were created, all the remaining documents in each data set were scored using its respective model. The

¹⁰⁰ See *id.* at 238.

¹⁰¹ We created the predictive model using the Logistic Regression algorithm, which our previous studies have proven to be highly effective. Our other modeling parameters were N-gram and normalized frequency, which our research has also shown to be advantageous. We used 20,000 tokens as features. See Huber-Fliflet et al., *supra* note 80, at 1–2.

probability scores and the previously applied attorney coding were used to calculate the recall and precision rates and evaluate the performance of predictive modeling.

[46] The recall and precision rates calculated in our experiments were used to measure the performance of each privilege-targeting method independently and to compare one against the other.

B. The Results

[47] Our experiments provided clarity about the effectiveness of keyword searches, compared precision rates of keyword searches and predictive coding in locating privileged information in a document set, and revealed that combining keyword searching and predictive modeling together to target privileged information maximizes discovery teams' performance. As we noted in Part II, roughly 73% of all production costs occur during the document review process.¹⁰² The insights provided by our study can reduce that figure if legal teams take advantage of them.

1. Effectiveness of Keyword Searching

[48] We evaluated the effectiveness of the unique keyword search terms intended to identify privileged content within the document population. Table 3 summarizes these results.

¹⁰² See *supra* Section II.B.1.

Table 3: Performance of Keyword Searching

Project Name	Search terms with at least one document hit	Total Privileged	Total Keyword Hits	Total Hit Privileged	Total Hit Not Privileged	Recall	Precision
Project A	812	46,756	193,017	43,847	149,170	93.78%	22.72%
Project B	4,270	14,326	368,506	13,571	354,935	94.73%	3.68%
Project C	5,547	536,788	2,493,846	508,549	1,985,297	94.74%	20.39%

[49] On all the three projects, keyword searching achieved an overall recall rate of close to 95%. The precision rates overall, however, were quite low: on Project A and Project C, only a little more than 20% of the documents identified by the search terms turned out to be privileged. On Project B the precision rate was almost same as the overall richness of privileged documents—3.65%—meaning keyword searching performed just about as well as randomly reviewing the documents.

[50] Using this data, we were able to observe the relative performance of different categories of search terms that are commonly used in privilege reviews, as well as the relative performance of particular terms. With respect to common categories of search terms, we observed differences in the performance of terms associated with different types of counsel, including outside counsel, junior in-house counsel, and senior in-house counsel. We found that terms associated with outside counsel and junior in-house counsel were extremely strong predictors for privilege, while terms associated with senior-in house counsel performed slightly worse in comparison. As for common generic search terms such as “counsel” or

“lawyer,” we found that these broad search terms performed better than many would have expected and outperformed the performance of the average search term for these projects.

a. Outside Counsel Keywords

[51] Beginning with common categories of search terms, a key takeaway from the results is that search terms associated with outside counsel will often be good predictors for privilege. Outside counsel keywords might include the names of law firms that have done work for the client, email domains associated with the firms, and the names and email addresses of particular outside counsel. Across all projects, precision rates for the names of outside counsel and outside law firms, especially email addresses and law firm domain names, greatly outperformed the average precision rate. For example, the average precision rate of keyword searches in Project C was 20.39%. Yet, keyword searches using outside counsel domain names returned rates of 80.88%, 92.10%, and 93.17% across the three projects. These high precision rates make sense, as it would be unusual for the client to involve (and pay) outside counsel if the client were not requesting or otherwise seeking legal advice. Courts recognize this fact and treat the presence of outside counsel on a communication as strong indicia of privilege.¹⁰³

[52] We also found that the email address and names of junior and mid-level in-house counsel were good predictors for privilege, although these terms were slightly less effective at identifying privileged documents than the outside counsel terms. For example, in Project C, precision rates for the email addresses of junior and mid-level in-house counsel ranged between 68.66% and 93.51%. These results are consistent with the different roles of in-house and outside counsel. While it is widely recognized that confidential communications between a client and in-house counsel are subject to the same protections as communications

¹⁰³ See U.S. *ex rel.* Baklid-Kunz v. Halifax Hosp. Med. Ctr., No. 09-cv-1002, 2012 WL 5415108, at *3 (M.D. Fla. Nov. 6, 2012) (“Communication between corporate client and outside litigation counsel are cloaked with a presumption of privilege.”).

between a client and outside counsel, it also may be more difficult to determine whether communications involving in-house counsel relate to the provision of legal advice.¹⁰⁴ This determination is difficult because in-house counsel may be involved with the business affairs of the company and may provide strategic business guidance in addition to legal advice.¹⁰⁵

[53] Finally, the precision rates for senior in-house counsel are lower on average than the precision rates for outside counsel and junior in-house counsel. In Project A, keyword searches for the general counsel (i.e. the most senior legal officer) had a precision rate of 46.75%, and the precision rate for similar searches in Project C was 44.34%. The differences in the performance of search terms related to senior and junior in-house counsel suggest that senior in-house counsel are often more frequently involved in the non-legal business affairs of a company than their junior colleagues.¹⁰⁶

¹⁰⁴ See *U.S. v. Mobil Corp.*, 149 F.R.D. 533, 537 (N.D. Tex. 1993) (“It is undisputed that communications between a corporation and its inside counsel are protected in the same manner and to the same degree as communications with outside counsel.”); see also *U.S. Postal Serv. v. Phelps Dodge Ref. Corp.*, 852 F. Supp. 156, 160 (E.D. N.Y. 1994) (The attorney-client privilege “protects communications with in-house counsel as well as outside attorneys.”); *Sec. & Exch. Comm’n v. Gulf & W. Industries, Inc.*, 518 F. Supp. 675, 681–82 (D.D.C. 1981) (finding that the burden is always on the proponent of the privilege to establish each element and that the burden is higher when the attorney in question is in-house counsel who also serves a business function.).

¹⁰⁵ See *Tex. Brine Co. v. Dow Chem. Co.*, No. 15-1102, 2017 WL 5625812, at *1 (E.D. La. Nov. 21, 2017) (“Determining whether the primary purpose of a communication with an attorney was to provide or receive legal advice can be complicated when the communication involves in-house counsel because these attorneys may serve in multiple roles (including non-legal).”); see also *AMP, Inc. v. Fujitsu Microelectronics, Inc.*, 853 F. Supp. 808, 830 (M.D. Pa. 1994) (“[I]n-house counsel may serve dual functions, acting as both legal counsel and business counsel. The privilege applies only to the former.”).

¹⁰⁶ In the projects used for this study, the lower precision was due in part to non-privileged marketing and news alerts sent to senior management, including the general counsel.

[54] In addition, we examined the individual precision rate of each search term used for each project.¹⁰⁷ Table 4 includes a short summary of the performance results of four commonly used terms that are essentially ineffective and imprecise because they are too broad.

Table 4: Performance of Commonly Used Privilege Terms

Term	Precision (Effectiveness)			Average
	Project A	Project B	Project C	
Legal	29.41%	6.55%	35.76%	23.91%
Privi*	31.17%	7.31%	37.55%	25.34%
Counsel*	38.18%	10.55%	51.89%	33.54%
Attorney*	44.50%	11.10%	49.80%	35.13%
Average of All Terms	22.72%	3.68%	20.39%	

[55] All the terms in Table 4 outperformed the precision rate of the search list in general, despite their breadth. Taken at face value, these results are surprising. These results suggest that attorneys searching on broad privilege search would find one privileged document for only every five reviewed. But digging deeper into the data reveals a slightly different conclusion. It turns out these terms return relatively high precision because they overlap with communications involving outside and in-house counsel. This intuitively makes sense because communications with counsel are more likely to include words indicative of privilege. Removing the communications with outside and in-house counsel returns the following results for the four commonly used terms:

¹⁰⁷ Since the data sets and the search term lists were from real legal matters, many of the terms are confidential and cannot be disclosed in this article.

Table 4(a)

Term	Precision (Effectiveness)			Average
	Project A	Project B	Project C	
Legal	4.38%	2.15%	15.91%	7.48%
Privi*	8.91%	2.87%	14.66%	8.81%
Counsel*	14.37%	5.03%	28.09%	15.83%
Attorney*	12.80%	4.82%	25.48%	14.36%
Average of All Terms	6.20%	1.54%	10.38%	6.04%

Once communications with counsel are removed from the data sets, the performance of common privilege terms decreases significantly. For example, in Project A, removing counsel communications causes the effectiveness of the term “Legal” to drop precipitously from 29.41% to only 4.38%. Removing counsel communications also highlights the differences between data sets. For example, the privileged key words performed poorly across the board for Project B. The privileged search terms “Privi*” was five times less likely to identify privileged documents in Project B as compared to Project C.

[56] That said, several of the terms continued to offer good results. The terms “Counsel*” and “Attorney*” identified privileged documents over 25% of time in Project C. Averaged across matters, these terms identified privileged information 15.83% and 14.36% of the time respectively. Even after removing counsel communications, the results in Table 4 suggest that the widely held intuitions about the effectiveness of these terms are incorrect, and legal teams can improve the precision of their keyword searches by including certain terms on search term lists. As identified above, the most striking result was the performance of “Counsel*” and “Attorney*” as search terms. These terms are widely considered to be too broad and therefore too inefficient for privileged keyword searching. But even after removing counsel communications, our study found them to be more precise than the overall performance of the lists across all the three data sets.

[57] Our results indicate the effectiveness of other terms also decreases when counsel emails are excluded from the analysis. For example, in Project C, other terms that outperformed the average precision rate included “Lawyer*” (32%), “complainant” (32%), “statute” (31%), “legally” (30%), “atty*” (24%), “testimony” (22%), and “summons” (22%), and terms that proved to be less precise than average included “magistrate” (9%), “respondent” (14%), “testify” (17%), and “lawsuit” (17%). But when counsel emails were excluded, each of these terms showed less effectiveness in identifying privileged communications. For example, the below chart shows the effectiveness of three additional terms on the overall data set without removing counsel communications:

Table 4(b)

Term	Precision (Effectiveness)			Average
	Project A	Project B	Project C	
Lawyer*	30.01%	3.33%	32.42%	24.27%
legally	29.41%	6.55%	30.21%	22.06%
atty*	63.00%	4.45%	24.27%	30.57%

Compared to this second chart, which shows the decrease in effectiveness for these same terms after counsel communications are removed:

Table 4(c)

Term	Precision (Effectiveness)			Average
	Project A	Project B	Project C	
Lawyer*	5.28%	1.11%	14.17%	6.85%
legally	4.38%	2.15%	15.45%	7.33%
atty*	1.05%	0.93%	13.18%	5.05%

[58] As shown above, the effectiveness of terms such as “atty*” appear duplicative of counsel communications and perform poorly when used on communications that do not involve counsel. Finally, removing counsel communications highlights the differences in effectiveness of certain

terms depending on the matter. For example, comparing Projects B and C reveals significant differences in the performance of other common privilege search terms after removing counsel communications:

Table 4(d)

Term	Precision (Effectiveness)		
	Project B	Project C	Delta
complainant	3.13%	23.63%	20.50%
statute	4.43%	25.53%	21.11%
testimony	2.34%	12.09%	9.74%
summons	3.19%	14.21%	11.02%
magistrate	0.88%	5.64%	4.76%
respondent	0.81%	12.47%	11.66%
testify	1.36%	9.28%	7.92%
lawsuit	1.32%	9.17%	7.85%

[59] These results suggest that search terms that may be worthy of consideration in some matters (i.e. respondent at 12.47% for Project C) may result in ineffective and burdensome privilege reviews in other matters (i.e. 0.81% for respondent in Project B).

[60] In sum, our study revealed that the effectiveness of privilege keyword terms depends greatly on the category of information the term is meant to capture. Search terms associated with outside and in-house counsel—such as attorney names or email addresses—are good predictors of privilege. Outside counsel terms performed better than in-house counsel terms and search terms associated junior in-house counsel were slightly better predictors of privilege than the search terms associated with senior in-house counsel. As for general search terms, our results revealed that the conventional wisdom around keyword search terms—general terms like “counsel” or “attorney” are too broad, and therefore imprecise, and specific terms like “complainant” are better suited for identifying privileged information efficiently—is incorrect and wastes the resources of attorneys and clients alike. Not all privilege search terms are created

equal; practitioners should consider managing the effectiveness of their privilege term lists.

2. Effectiveness of Predictive Modeling

[61] The results of our experiments revealed that predictive modeling can effectively target privileged information within document review populations and provide a diverse set of implementation benefits. For example, the universe of documents in one project—referred to as Project A—contained 360,531 files. Instead of conducting an inefficient manual review of each document, reviewing attorneys applied a supervised machine learning algorithm—an algorithm built from a “human-reviewed subset of documents.”¹⁰⁸ First, they selected 5,000 documents to create the algorithm’s training set. They then reviewed each of these documents and coded them as either “privileged or “not privileged.” Finally, computers used the results of the training set to categorize the remaining 355,531 documents. This process noticeably increased the efficiency of the document review. Table 5 outlines the precision rate of each data set’s model at specific levels of precision indicating how efficiently the models can target privileged documents.

¹⁰⁸ See Huber-Fliflet et al., *supra* note 80, at 2.

Table 5: Precision and Recall Rates for the Predictive Models

Precision	Project A – Recall		Project B – Recall		Project C – Recall	
	Rate	Documents	Rate	Documents	Rate	Documents
50%	84.90%	39,696	5.73%	821	74.13%	397,921
75%	60.45%	28,264	2.00%	287	44.60%	239,407
80%	51.33%	24,000	2.00%	287	36.32%	194,961
90%	24.55%	11,479	2.00%	287	17.68%	94,904
95%	14.55%	6,803	2.00%	287	8.01%	42,997

[62] The precision and recall rates for the predictive models were inversely proportionate—as is typically observed when analyzing the results of predictive models.¹⁰⁹ The models achieved high precisions and identified large percentages of the privileged document populations. For example, for Projects A and C at 80% precision, the models identified 51.33% and 36.32% of the privileged documents in their data sets, respectively. In other words, 80 out of every 100 documents reviewed were privileged at this precision rate and provided review efficiency gains when compared to a random document review for privilege. The richness rates of Project A and C were 12.97% and 6.16%, respectively, and an 80% precision provided by the model resulted in a 600% increase in efficiency for Project A and nearly a 1,300% increase for Project C.

[63] Project B's precision was very low and just slightly better than random at 50% precision when compared to the data set's richness rate (Project B: 3.61%).

¹⁰⁹ See THE SEDONA CONFERENCE, *supra* note 5, at 238.

3. Keyword Searching vs. Predictive Modeling

[64] The experiments from this study enabled a comparison of the strengths and weaknesses of keyword searching and predictive modeling. Table 6 compares the precision rates of predictive modeling to the precision rates of keyword searching at similar recall rates for each of the three data sets.

Table 6: Precisions at Similar Recall Levels

Project Name	Keyword Searching		Predictive Modeling		Precision Comparison
	Recall	Precision	Recall	Precision	
Project A	93.78%	22.72%	93.78%	30.11%	-7.39%
Project B	94.73%	3.68%	94.74%	4.44%	-0.76%
Project C	94.74%	20.39%	94.74	17.43%	2.96%

[65] Project A's predictive model outperformed keyword searching by over 7% precision at a ~94% recall rate. For Project B, its predictive model was approximately .75% more precise than keyword searching at ~95% recall. However, on Project C, which had over eight million documents, keyword searching was roughly 3% more precise when compared to predictive modeling at approximately 95% recall. It is important to note that while these differences in precision may not immediately appear significant, single digit precision improvements can greatly impact the cost of review as data volumes rise. Precision improvements for Project A, using predictive modeling, resulted in reviewing nearly 22,000 fewer documents and for Project C, using keyword searching, resulted in reviewing nearly 245,000 fewer non-privileged documents.

[66] Our experiments demonstrated that predictive modeling can find privileged documents that keyword searching cannot. Table 7 reveals the

number of documents at 50% or greater precision identified by the predictive model and did not hit on a keyword search term. In all the three data sets, the predictive models identified privileged documents that did not hit on a keyword term, highlighting that predictive modeling can serve as a complementary privilege-targeting technology to keyword searching.

Table 7: Documents at 50% or Greater Precision and Not Keyword Hits

Project Name	Total Documents at 50% Precision or Greater and Did Not Hit on a Keyword Search Term	Coded Privileged by an Attorney	Coded Not Privileged by an Attorney
Project A	6,075	1,062	5,013
Project B	2	2	0
Project C	72,295	6,924	65,371

[67] Many practitioners' intuition that predictive modeling is an ineffective privilege-targeting technology and that keyword searching is more reliable does not always hold true. This study's side-by-side comparison proved that either keyword searching or predictive modeling could be more efficient at identifying privileged documents, depending on the specific project setting. This means neither of the two technologies should be used at the exclusion of the latter. Keyword searching and predictive modeling are entirely complementary and practice groups that aim to maximize their ability to protect their clients' attorney-client privilege should combine keyword searching and predictive modeling together to provide comprehensive protection.

C. A Way Forward

[68] The insights obtained from our study, if implemented, will have wide-ranging effects on future review practices. The results indicated that

practitioners who wish to maximize protecting privilege and efficiency should consider adopting the following practices:

- (1) use keyword searches as the primary method of identifying privileged documents and consider broad terms in the corresponding keyword search lists, such as “counsel” and “legal,” that were previously regarded as inefficient;
- (2) streamline the manual review process by using predictive modeling to prioritize documents returned by the keyword search terms based on their likelihood of containing privileged content;¹¹⁰ and
- (3) use predictive modeling as a complementary search method to identify documents that are highly likely to contain privileged material but do not contain any of the terms on the keyword search term list.

[69] Following this approach, legal teams can identify a greater number of privileged documents than by using conventional keyword searching wisdom alone, develop targeted review strategies by identifying the most sensitive documents in the data set and prioritizing them for review, and reduce the number of time-consuming and costly false positives among their search results.

IV. CONCLUSION

[70] The universe of electronically stored information is vast and expands each year, creating an ever-growing haystack of information in which legal teams must locate privileged documents. Whether a legal team preserves the attorney-client privilege or waives it through inadvertent

¹¹⁰ For more detailed information about best practices for predictive modeling, including which preprocessing parameters most effective for identifying privileged or relevant information, see our previous publication. Huber-Fliflet et al., *supra* note 80, at 1.

disclosure can alter the outcome of a case. Protecting privileged information is no less important now than it was centuries ago when our legal ancestors enshrined privilege protections in the common law system. In recent times, it has become overwhelming and expensive for legal teams to provide the same protections for their clients as they once did using traditional methods.

[71] The legal community has evolved in the digital age, using keyword searching and advanced text analytics to target sensitive privileged documents. These technological advances have been significant, but the legal community has conducted little research to confirm their strengths and weaknesses and identify best practices for implementing them. Instead, attorneys have chosen to rely on a combination of intuition and trial-and-error to evolve their review process. Wanting to take a scientific approach, we performed “look back” experiments on data sets from three real legal matters to test these longstanding intuitions and found many of them to be inaccurate.

[72] Our study demonstrated that both keyword searching and predictive modeling can identify privileged documents with varying degrees of precision, and that predictive modeling can manage the weaknesses inherent in keyword searching to identify privileged documents that keyword search terms miss. These results suggest that practitioners should consider conducting a privilege review using both technologies in a complementary manner.

[73] We used the insights generated by this study’s experiments to create practical considerations for legal teams seeking to maximize their ability to shield their clients’ privileged communications from disclosure in a cost-efficient manner. Our insights, if implemented, will help reduce the cost of discovery while simultaneously strengthening protections over privilege. They may not solve the legal community’s discovery cost dilemma altogether, but they are undoubtedly a step down that path.