The Coming Revolution in Professional Patent Searching

A White Paper Prepared by
the Leaders of Ensemble IP LLC

ensembleip

The Coming Revolution in Professional Patent Searching

A White Paper Prepared by the Leaders of Ensemble IP LLC

TABLE OF CONTENTS

The Problem with Patent Searching 1

Brief History of Patent Searching 3

The Promise of Artificial Intelligence 8

Machine-Assisted Patent Searching (AI Patent Searching) 11

Comparison of AI Patent Search Providers 15

The Coming Revolution in Prior Art Searching 17

ensembleip

The Problem with Patent Searching

Professional patent searching has evolved slowly over decades and has only marginally improved in that time. This is because the patent search field is human-dependent in a field that requires the ability to find and sift through numerous records from a corpus of millions and assess technological relevance under tight time and budget constraints. Humans can think deeply and make informed judgments but are unable to process thousands of records efficiently.

The job itself requires many skills. The ideal patent analyst is an expert in technology, understands patent law, has access to multiple data sources, and is theoretically able to search patents and published literature in every language since the beginning of time. On the surface it sounds impossible, and in practice, it is.

Moreover, patent searches for legal matters are custom requests that result in custom deliverables. Often, the search cannot be shared, resold, or even used again. This fact means that a completed search will not improve with time. It is rarely revisited or revised. It may be used to draft a legal opinion, advise a client, assert an intellectual property right, or to invent around a protected invention, but nothing more. Were it to be shared and reworked by others, it could improve over time. Instead, it must be done correctly the first time, every time, to be of use to patent practitioners.

In Short, Humans Are the Problem.

When you invent a new technology, you can receive a patent if the invention is proven not to have existed, for which you should conduct a search of other patents and written words in every language. Because this is humanly impossible, most people hire an engineer or scientist to search several patent and technical literature databases (in their spoken language) which

have millions of records. They search text using Boolean search methods, review drawings and diagrams, and arrive at their best judgment of the relevance of whatever “prior art” that they have gathered.

Unfortunately, the human-driven method is incomplete for five reasons:

1. A human is incapable of finding every instance of relevant information worldwide given inevitable time and budget constraints. This may be referred to as information overload.
2. When found, there is simply too much information for a human to review and the volume of accessible data is accelerating. This is a constraint on human intellectual capacity.
3. The relevance of what has been found depends on the purpose of the patent search (different for patentability, infringement, and validity). This requires specific subject matter expertise on each search.
4. Some relevant data is in non-native languages (not understood by the human analyst). This is the constraint that few patent analysts are literate in more than one language, let alone the five major languages in which patents are granted.
5. Each human expert will arrive at different opinions as to the relevance of the patent or non-patent literature that is analyzed. This is known as subjective ground truth. There is no objective ground truth such as distinguishing a cat from a dog.

The Real Solution Rests with Innovation.

This paper proposes an alternative strategy for achieving better results. Every search requires precise technological expertise, access to relevant and complete data, the retrieval of documents in multiple languages, and judgment that reflects an understanding of patent law. The best approach combines these skills with both Boolean search tools and powerful AI tools and machine learning, applied to the full corpus of patent documents and scientific and technical literature.

The discussion that follows make the case for this more complete approach to patent searching.

Brief History of Patent Searching

Before we offer solutions that will result in a revolution in searching, it is useful to understand the history of professional patent searching, the periodic innovations in the field, the rise of alternative business models that address these shortcomings, and the current state of the industry.

Physical Presence at the USPTO Search Room

For decades, patent attorneys relied on skilled engineers and scientists to conduct quality patent searches. Law firms and inventive companies would hire them as scientific advisors or contractors who would then commute to the patent office (e.g., USPTO) and work from the search room flipping through patent copies by hand.

Unfortunately, those hand searches were incomplete as patent offices only had copies of their own patents, PCT filings, and an incomplete collection of foreign patent documents. Patent searchers would visit the scientific library at the PTO to search non-patent literature or, commonly, contact examiners in that subclass for search advice or for a copy of a relevant scientific article kept among their personal files. If foreign patents were searched, they were usually English language abstracts. For searches requiring an in-depth review of non-digitized information, analysts might use world class physical libraries like the Library of Congress, the British Library in London, or the Tsukuba library in Tokyo. A more complete search could be done using federated databases such as Dialog and STN or by having a subscription to Derwent. However, these resources were expensive and required special training to use given their command line interfaces.

> ***The Result:*** *With the help of a skilled searcher, patent practitioners might be assured of a quality U.S. Patent only search at the USPTO only. The foreign patent search and non-patent literature search would be incomplete and likely unreliable. Moreover, the attorney would need to wait until the searcher had the capacity to conduct the search as the patent analyst could do only one at a time.*

Electronic Search Tools

The move to electronic records improved the quality of patent search results because it increased the number of experts who had access to tools, and the tools themselves became more comprehensive in their coverage. Several patent offices offered electronic search tools, including the USPTO (EAST) and the European Patent Office (ESPACENET). Private companies also developed powerful Boolean-based tools with expanded data. Other companies digitized articles and publications and made them available by subscription. Specialized databases allowed improved searching.

Patent Publications

By 2008, anyone could subscribe to a large number of search tools and a variety of patent and non-patent literature resources. Patent analysts had digital access to patents and publications across 95 patenting authorities that included a mixture of full-text, machine-translated full-text, and bibliographic patent document collections.

Scientific and Technical Publications

The vast availability of non-patent literature sources included resources from Elsevier, Dialog, Thomson, ProQuest, IEEE, ACM, and many others. Non-patent literature searching was conducted similarly to patent searching by beginning with the broadest coverage databases. Then analysts would strategically investigate different types of scientific literature, including index and abstract files, industry specific journals, product information, and news and press releases, among others. With the rise of the internet and access to subscription-based databases, analysts could simultaneously search thousands of individual journals at once. Specific journals identified during

this step were then individually investigated further electronically or manually. Many resources also became available online in the German, Japanese, Chinese, Korean, and other languages.

The resulting coverage was breathtaking and comprehensive, encompassing many different types of documents and virtually every technology area.

> ***The Result:*** *The availability of online records improved the breadth of patent search results as it allowed for more comprehensive coverage. It facilitated access to full-text native language searching in major languages, led to a proliferation of specialized tools by technology interest, and brought enhancements in Boolean search functions. It also led to quicker turnaround as more patent experts at inventive companies, law firms, and universities had access to these tools.*
>
> *However, professional patent searching still relied on human effort to know which databases to search, how to design and implement a search strategy, how to use specific software tools, and the time to do it. In many cases it added to the time needed to do the search.*

It did not address the fundamental problem that faces any labor-intensive process, which is the lack of economies of scale. More data results in more noise unless the data can be sifted through faster while maintaining its relevance. The patent analyst will need to review more references from more sources under similar time constraints. Searches are more comprehensive, but not necessarily more accurate.

The Rise of the Comprehensive Patent Search Firm

Despite the limitations, the combination of searching-by-hand and the use of online tools led to the rise of the comprehensive patent search firm. Such companies offered more complete patent search results by searching patent and non-patent literature from multiple sources and providing quicker turnaround.

The innovations of the comprehensive patent search firm were in business management, personnel development, and the strategic use of database tools. They did not fundamentally change the way patent searches were conducted. Searches still relied on human experts to find, read, and determine the relevancy of each reference. However, customer service improved as the best search firms competed on quality and reliability and implemented business practices that added value for clients.

Semantic Search Engines

From 2005 until 2011 some firms attempted to innovate by creating semantic search engines to replace Boolean searching altogether. These tools intended to simplify the patent search process by having the patent analyst enter natural language terms, upon which computer algorithms would predict the meaning of the words and would cite relevant search results. This latent semantic approach was promising but failed because they: (1) were positioned as replacements to Boolean search systems which were already useful, (2) relied on extensive computing power before graphic processing units (GPUs) were reconfigured to expedite machine learning[i], (3) required immense storage to successfully train algorithms on tens of millions of patent records, and (4) depended on one algorithmic approach to achieve results as opposed to using the modern ensemble approach to achieving better results.

Crowd Sourcing

At the same time, another innovator offered a crowd-sourced approach to patent searching. The company assembled thousands of independent engineers and scientists who searched without compensation but cited art in return for a reward for finding the most relevant prior art references. This was promising but flawed for several reasons, which are: (1) invention disclosures are confidential so the approach was useful for validity studies but not all search types, and (2) the art cited by the crowd requires human analysts to review its relevancy, which could take longer than conducting the search itself. Crowdsourcing has not gained a foothold beyond a few loyal customers.

Offshore Providers

Because of its inability to achieve scale economies and low barriers to entry, the patent search field remains fragmented.

This situation was compounded in the 2000s, when many low-cost offshore providers entered the market and conducted patent searches at a fraction of the price of search providers in North America, Europe, and Japan. Offshore providers are not more innovative than traditional patent search firms and compete merely on price.

> ***The Result:*** *Although access to online records and the rise of comprehensive patent search firm has improved search results, none of these firms can assure complete patent and non-patent literature searches in multiple languages, or even one language. The singular reliance on Boolean search methods in the native language of the analyst (usually English) provides a disservice to patent practitioners who have the critical role of delivering informed legal advice and to innovative companies who spend immense capital to create and protect intellectual property worldwide.*
>
> *The following table summarizes deficiencies in searching, by type of patent search innovator. Until the industry addresses these shortcomings, patent practitioners will switch from firm to firm, hoping for the complete set of prior art references that allows them to provide clear and confident legal advice to their clients. Without more innovation, they will come up short.*

Coverage and Deficiencies by the Type of Patent Search Innovator

Search Objectives	**Search Room** Pre-2002	**Online Tools** 2002-Present	**Comprehensive Search Frm** 2004-Present	**Offshore Provider** 2006-Present	**Latent Semantic Tool** 2005-2011	**Crowdsourced** 2011-2017
Reliable U.S. patent search	●	●	●	●		●
Reliable foreign patent search		●	●			●
Reliable non-patent literature search		●	●			●
Native language patent searching			●			●
More breadth of data		●	●	●		●
Complete patent search (all patent authorities)						
Complete non-patent literature search						
All search types	●		●	●	●	
Quick turnaround / adequate capacity			●	●	●	●

blank = deficient

The Promise of Artificial Intelligence

The patent community has begun to learn about the promise of AI as a supplement to traditional search methods. Approximately a dozen innovators have designed and are developing AI-assisted patent search tools; they claim the ability to reshape the search for prior art. Ensemble IP believes that AI is the innovation that will revolutionize patent searching.

Ensemble IP has tested and assessed the performance of some of these AI-assisted patent search tools; our findings, and our perspective on the appropriate role for such tools in patent searching going forward, are shared in Parts 4 and 5 of this paper.

However, prior to discussing our assessment, it helps to review the current state of AI and machine learning and which methods are likely to revolutionize professional patent searching. This introduction will help to understand why some methods have promise and others should be discounted.

Expert Systems

Expert systems are rule-based algorithms that approximate the decisions that humans make. Similar to "if-then" arguments, expert systems cannot learn but instead attempt to replicate the decision making of experts. For this reason, they are rigid and require rework as human experts improve their own knowledge. Their rigidity rests in the fact that they need re-coding. They cannot recognize new terms or their definitions and cannot learn on their own. We believe expert systems are an ineffective approach to finding and proposing relevant prior art and will not have a meaningful impact on future innovations in professional patent searching.

Machine Learning

Machine learning enables computers to learn through a set of software code in specialized programs.[ii] This approach will have a major impact on the development of AI-assisted patent search tools as it allows the human analyst to receive a curated list of prior art references from highly trained algorithms without having to find and curate the list themselves. There are a few forms of machine learning that are presented here:

SUPERVISED LEARNING

Supervised learning is a form of machine learning whereby a computer is presented with input and output pairs, and the algorithm is trained to predict an output given an input.[iii] Outputs are referred to as labeled data (e.g., a photo of a cat with a label that identifies the photo as that of a cat).

With patent searches, the input is either all the text of a patent, some of the text of a patent, or an invention disclosure. The output is the specific data used to train and test the mathematical algorithm. The output supervises the algorithm.

With this form of learning, there may be several ways to classify patents and scientific articles as relevant or not.

UNSUPERVISED LEARNING

Unsupervised learning is a form of machine learning whereby a computer is given unlabeled data and is programmed to find patterns in the data. Unlike supervised learning, the algorithm is not trained and corrected to match inputs with outputs but to find the similarities and differences. Unsupervised learning offers opportunities to improve patent search results.

SEMI-SUPERVISED LEARNING

Semi-supervised learning is a form of machine learning that attempts to make sense of noisy data or when it is unclear what label should be assigned to a data element. In the patent field, semi-supervised learning might be applied when the computer queries the software user to help discover the correct output. This is known as Active Learning. This type of

learning has been used effectively to optimize recommender systems such as which products to purchase on Amazon.

REINFORCEMENT LEARNING

Reinforcement learning is when an algorithm learns by repeatedly being measured on the accuracy of its actions. Then it improves its actions iteratively, in a sense becoming its own teacher. This is similar to how you might train a dog by giving them a treat for taking a correct action or withholding a treat for an incorrect action.

DEEP LEARNING

Deep learning is a subset of other machine learning models. These are algorithms in multi-layered artificial neural networks, which are mathematical structures loosely inspired by how biological neurons are structured. Deep learning requires the iterative design and testing of these layered neural networks.

The retrieval and assessment of relevant prior art is essentially a topic matching problem that compares either an invention disclosure or a granted patent to both publicly available technical literature and other published patent documents. As a matching problem, the logical approach to modeling the algorithm is to use some form of Natural Language Processing (NLP). Deep learning has been shown to work well on NLP problems and holds promise in the development of an AI-assisted patent search tool.

Whatever methods are used, it is an iterative process of training and discovery, and simple models often outperform complex deep learning approaches. In fact, ensemble methods combining several machine learning models have performed better than deep learning in several industry applications. After extensive study, Ensemble IP believes the “ensemble” method of combining several machine learning models will yield the best AI-driven search results. Our name is based on this belief.

Machine-Assisted Patent Searching (AI Patent Searching)

With respect to patent searching, AI has great promise. This section of the paper describes the machine learning process as applied specifically to patent and technical literature searching and describes the key performance measures for evaluating AI-assisted patent searching tools.

Designing the Machine Learning System

GROUND TRUTH

In machine learning, the ground truth is an objective target that an AI algorithm uses to make a prediction. It is used in supervised learning as a target by which to predict future outcomes. In the field of medical diagnostics, for example, ground truth may be defined as whether or not a patient has cancer. A radiologist or a mathematical algorithm may misdiagnose the presence of cancer, but the target never changes.

THE CONFUSION MATRIX

This concept can be understood with a table known as a Confusion Matrix[iv], which is a framework for visualizing the performances of either a mathematical algorithm or a human prediction.

The rows are the predictions and the columns are the actual conditions. If the radiologist or an algorithm predicts that you have cancer and you do not, the diagnosis results in False Positive. If they predict that you do not have cancer and you do, the diagnosis results in False Negative. Nonetheless, the ground truth is objective and verifiable.

		Actual Condition	
		Cancer (You have it)	**No Cancer** (Or you don't)
Predicted Condition	**Cancer**	**True Positives** *Predicted cancer and it is cancer*	**False Positives** *Predicted cancer, but it is not cancer (Type 1 error)*
	No Cancer	**False Negatives** *Predicted it is not cancer, but it is cancer (Type 2 error)*	**True Negatives** *Predicted it is not cancer and it is not cancer*

In the field of patent searching the ground truth target is defined as whether or not a cited reference is relevant. However, unlike cancer detection, ground truth in patent practice is approximated because disagreements among patent experts about the relevance of cited art are common. Equally capable patent analysts, attorneys, examiners and agents will disagree on whether a cited reference is relevant, and this fact is fundamental to the practice of patent law. It is why a patent application requires examination and, once granted, may be opposed or litigated in the courts.

The following confusion matrix demonstrates how a cited reference may be deemed relevant and how that decision may be viewed by others in the patent field. This situation requires that AI software companies experiment with various approaches to assessing ground truth and determine which ones arrive at what most patent experts would deem to be relevant references.

		Actual Relevance	
		Relevant (Deemed relevant)	**Not Relevant** (Deemed not relevant)
Predicted Relevance	**Relevant**	**True Positives** *Predicted relevant and it isrelevant according to most experts*	**False Positives** *Predicted relevant, but it is not relevant according to most patent experts (Type 1 error)*
	Not Relevant	**False Negatives** *Predicted not relevant, but it is relevant according to most experts (Type 2 error)*	**True Negatives** *Predicted not relevant and it is not relevant according to most experts*

As detailed later, the confusion matrix is the basis for the development of two key performance measures – precision and recall – that Ensemble IP used in our evaluation of current AI tools developed for patent searching. These measures provide a comparative basis for assessing the quality of the results provided by each tool.

SOURCES OF GROUND TRUTH

Since knowledgeable experts can disagree on the relevance of cited prior art for any given search, AI search system providers must leverage various sources to determine relevancy. This is an active area of research and differs with the type of search being executed.

The Machine Learning Process

ABOUT THE TRAINING METHOD

The training method is the approach that the data architect would propose to find relevant art. Multiple methods may be used in concert to predict more examples of relevant art (the ensemble method). The training method should be proposed based on a search expert's informed judgment.

ABOUT THE RESULTS

The results are how well the algorithm predicts which references should be cited as prior art. The results are measured using a confusion matrix, as previously described. Often many algorithms are compared to one another, using the

confusion matrix, for each training method. The goal is to cite only relevant and non-relevant references and to predict the level of relevance for other art.

ABOUT PERFORMANCE

In order to assess their performance, Ensemble IP measured the precision and recall achieved by each computer system using the same searches with each system.

Precision is the percentage of patents that were retrieved that are relevant. For example, if an AI search system retrieves 30 patents and 20 of them are relevant then the precision is 0.67 (2 of 3 patents cited are truly relevant). The 20 patents are "true positives" (cited as relevant and truly relevant). The other 10 patents are "false positives" (cited as relevant but not relevant).

Recall is the percentage of relevant patents that were actually retrieved. For example, if that same AI search system failed to cite 40 other patents that were deemed relevant it recalled only 20 of 60 relevant patents. The recall is 20/60 or 0.33. In this way, recall measures how complete the results are.

$$\text{Precision} = \frac{\text{Relevant}}{\text{Retrieved}} = \frac{\text{True Positives}}{\text{True Positives + False Positives}} = .67$$

$$\text{Recall} = \frac{\text{Relevant}}{\text{Relevant + Relevant Not Retrieved}} = \frac{\text{True Positives}}{\text{True Positives + False Positives}} = .33$$

A perfect patent search would cite only relevant art (perfect precision = 100%) and all relevant art (perfect recall = 100%). Of course, perfection is nearly impossible but should be the objective of patent search software tools and patent analysts who practice the profession. Without this objective, expensive innovations may not yield meaningful improvements.

> ***The Result:*** *AI and machine learning tools are being developed specifically for the patent searching field. Several of the approaches appear to be capable of adding value to the patent search process.*

Comparison of AI Patent Search Providers

Ensemble IP tested the performance of several AI-assisted patent search systems with more to follow on a regular basis. The tests included patentability, invalidity, and infringement searches across technology areas.

Finding Relevant Art

AI providers differ from one another in their approaches to finding prior art. Therefore, we compared the performances of these AI systems against one another and against human analysts.

With respect to assessing whether an AI system cites a relevant reference, you might ask how Ensemble IP could know the complete universe of relevant and non-relevant art for each search that was entered into the AI systems. The short answer is that we are not all-knowing but used searches that were conducted many times by several expert patent analysts and reviewed by other search experts. Then, the human results were compared to the AI results. When any of the AI systems cited art not found by the group of human analysts, but deemed relevant, that art was added to the list of relevant documents. In short, the tests were conducted to show relative performance of systems against one another, against professional searchers, and in combination with one another. This created a more complete list of relevant art by combining the results of human effort, machine effort, and limited crowdsourcing.

We will not achieve universal agreement of ground truth but believe that generally accepted approaches to assessing the level of relevance of prior art will develop. The most credible machine-learned predictions will be based on intuitive approaches that have been tested repeatedly in specific technology areas and across technologies.

Test Results

In order to protect confidentiality, this paper does not identify the actual AI providers whose systems were tested but provides a summary of the results. The better performers achieved recall (the percentage of relevant documents that were actually retrieved) between 0.455 and 0.625, and precision (the percentage of relevant documents among the set of AI-cited documents) between 0.125 and 0.238.

Most of the precision scores were below 0.06, which indicates that the AI systems that were tested cite a huge amount of non-relevant art in order to capture some relevant art. Unless the algorithms are better tuned, this approach will overly burden professional patent analysts with review of unusable prior art references. It may then lead to more time needed on a search and possibly limit adoption of AI tools among patent practitioners.

Recall results were better than precision for all systems, but the highest performers only identified half of the relevant art at best. This means the AI systems alone (without human intervention and Boolean searching) do not provide complete results and are therefore currently inadequate as stand-alone search systems. In no situations did an AI system outperform a patent analyst using Boolean search tools. However, patent analysts using Boolean tools and the best performing AI systems produced the best overall results. The tests allowed Ensemble IP to find additional examples of highly relevant art that were not found by analysts alone. In some cases, Ensemble IP analysts identified as much as 20% more highly relevant art with the best AI tools. This means that patent analysts using traditional search methods and the best AI tools produced more complete results than they would have without those AI tools.

> ***An Important Insight:*** *Test results indicate that there is an important complementary role for AI-assisted patent search systems. There is no role for AI-only patent searching as the results are weak compared to traditional human-driven searching using Boolean patent search systems.*

Strengths of Current AI Systems

The AI systems that were tested are adept at creating deep connections, whereby they discover art in obscure classification areas that appear irrelevant to the search, but encompass useful art. They are also useful for determining the effort level to conduct search quotes (time estimates) and as quality control tools (cursory check of a patent analyst's results using Boolean and hand search methods).

Weaknesses of Current AI Systems

Most systems share the same weaknesses, which are:

- They lack the ability to determine inventive step (obviousness) combinations.
- They lack image search capabilities.
- They lack infringement or freedom-to-operate (FTO) specific search modules.
- They result in information overload, whereby they cite an endless amount of "relevant" references that require a human analyst to verify.
- They produce noise, meaning they tend to target areas the algorithm deems of interest, but not of interest to the patent expert.
- They do not accommodate the search of chemical structures and bio sequences.
- They usually have been trained in one language and not on non-native patent documents.
- They have not been trained on scientific and technical literature (a.k.a. non-patent literature).
- They do not measure the degrees of prior art relevance.
- They lack relative ranking of references, meaning how many records down a ranked list will the algorithm cite a highly relevant document in the mind of a patent expert

- They make it hard to determine if they are focused on the wrong technology area.
- None of the systems can cite more examples of relevant art than a patent search expert using traditional Boolean search tools.

The Added Value of Machine-Assisted Patent Searching

The companies that have experimented with algorithmic solutions to the patent search problem have tried everything from citation mapping to expert systems to machine learning. We believe that supervised or unsupervised learning that predicts the relevance of other patents or scientific literature to an invention may improve patent search results (if combined with human judgment).

This machine-assisted method might address the limitations of human searching in these ways:

1 Finding more (if not all) instances of relevant information.

2 Classifying information as to relevance (if extensively trained).

3 Creating different machine learning methods depending on the purpose of the patent search.

4 Facilitating non-native language searching with either robust machine translation or machine learning of document collections written in other languages.

5 Allowing experimentation to home in on useful approaches to selecting references and using those results to inform human judgment.

However, to truly add value, machine-assisted patent searching will have to be combined with human experts and Boolean search tools, as described in the conclusions that follow.

The Coming Revolution in Prior Art Searching

What We've Learned

This paper describes the persistent problem that no prior art search is complete and, at best, cites some examples of relevant art that assist the patent practitioner. Prior art searching has evolved slowly from searching country-specific patent documents by hand, to online databases with expanded coverage and faster retrieval. However, search methods remain Boolean-based and human-dependent while the volume of written documentation continues to increase rapidly in many languages.

Patent analysts have tools to cite relevant art from more sources, but not more relevant art from all sources. As a result, patent attorneys and agents practice their professions with diminished expectations and sometimes select the cheapest prior art search provider. This is because they do not expect patent search results to be fully complete - just "good-enough" for the matter at hand.

Innovations have proceeded slowly along three paths: (1) business models where operational consistency has led to reliable but incomplete results, (2) failed experiments in crowdsourcing and early latent semantic searching, and (3) more comprehensive Boolean search tools that encompass data from more patent and technical literature sources.

The industry is now witnessing the development of powerful machine learning algorithms (a.k.a. AI systems) that will lead to real improvements in patent searching. Some of the current tools, which are in their nascent stages, have the ability to learn what art is relevant and not relevant to patentability, freedom-to-operate, and validity studies.

These AI tools allow patent experts to conduct better patent searching,

though none outperforms human experts using the latest Boolean search software tools. The best of these AI tools help locate references not found with Boolean-assisted human searching alone. In some cases, patent analysts identified as much as 20% more highly relevant art with the best AI tools.

The Solution Rests with Human-Computer Interaction

Lasting innovations in patent searching would result in the accelerated retrieval of relevant prior art references, cite mostly relevant references, and cite a more complete set of relevant references.

The solution rests with blending the modern patent search firm with the use of the best machine-learned algorithms.

Because of the inability to predict absolute ground truth, AI systems need to be paired with humans (experienced patent analysts using traditional Boolean search methods). The AI system helps predict which art is relevant from the set that it found. The teaming of expert patent analysts with Boolean and AI tools appears to be the most appropriate solution to improve patent search results.

The authors led the innovation that applied operational improvements to the patent search field. These business model innovations occurred at the world's largest professional patent search firm, formerly owned by the current leader of Ensemble IP, and led to more reliable patent search results across technology areas and types of searches. In the Age of Artificial Intelligence, these innovations must be paired with more precise approaches to patent searching.

This is why Ensemble IP blends human expertise with proven technological tools (both Boolean and AI) to improve search results beyond what humans or machines can achieve by themselves. The company studies and uses the latest algorithms and software and puts them in the hands of the most capable patent analysts. This approach helps deliver the most reliable search results, and, as AI tools improve, will lead to complete search results.

Extensive tests have shown that AI tools perform well when paired with humans. This human-computer interaction (HCI) results in an ensemble approach to citing relevant prior art. Current AI systems complement human experts but do not replace them. Together, they deliver the most complete results to date.

> ***The Result:*** *Ensemble IP's testing shows that AI-only patent searching results are weak compared to traditional human-dependent searching using Boolean methods. We strongly recommend that decision makers pair AI tools with native language patent analysts who have extensive training in professional searching and patent law. The most complete results are achieved by blending human expertise, the best Boolean search software, non-patent literature, native language searching, AI algorithms, physical hand searching, and modern business operational practices that result in a high level of customer satisfaction. HCI promises dramatic quality improvements with a marginal increase in effort, to improve the quality of patent prosecution, legal advisory services, and results of litigation. All other options ensure an incomplete search. Because the search results are highly promising, this blended or "ensemble" approach followed by Ensemble IP, which combines multiple resources and methods, will lead to an inevitable revolution in professional patent searching.*

About Ensemble IP

Ensemble IP was founded by the team that created the largest and best-known patent search firm worldwide. Its company's leaders have led large teams of U.S.-, Europe-, and Japan-based technology experts, many of whom are patent agents and former patent examiners. The authors also designed and delivered professional search training at Patent Resources Group (PRG) where they educated nearly 1,000 people on patent searching best practices.

The company delivers more complete results by using multiple search tools, full-text native language searching, and hand searching at libraries. Their leaders have served the patent searching needs of many of the busiest litigators, the top patent law firms, organizations that have been granted patents, and the USPTO, as the first team to conduct PCT searches on their behalf.

Ensemble IP uses an "ensemble" approach to prior art searching, as described in this paper. This has resulted in marked improvements over traditional approaches to patent searching.

Ensemble IP's combination of expertise and experience makes them uniquely positioned to be the leader in revolutionizing patent searching.

To see the Ensemble approach in action, contact Ensemble IP today at mail@ensembleip.com and 202-869-0203.

---

[i] Led by Andrew Ng at Stanford University
[ii] Modern AI and Machine Learning Technologies - Applied Artificial Intelligence Handbook For Business Leaders
[iii] Modern AI and Machine Learning Technologies - Applied Artificial Intelligence Handbook For Business Leaders
[iv] From the confusion matrix, much analysis can be done to measure the accuracy of an algorithm. It is beyond the scope of this paper to discuss those types of measurements, but the confusion matrix is one of the most important diagrams in machine learning and predictive analytics.

