Bennett Jones

AI: Is Patenting Really the Right Move?

A CORPORATE DECISION TOOL

A Business-Driven Decision Framework for Innovators, CTOs and In-House Counsel Making the Call on Patent or Trade Secret Protection for AI Innovation

Table of Contents

Decision Framework 1

FACTOR 1 Identifying the True AI Innovation 3

FACTOR 2 The Enforceable Patent Scope 5

FACTOR 3 Reproducibility of the AI Innovation 7

FACTOR 4 The Business Delivery Model 8

FACTOR 5 The Commercial Longevity 10

FACTOR 6 Competitor Defensive Positioning 11

FACTOR 7 Patentability Potential & Layered Strategies 12

Key Contacts

Ahmed Elmallah

Counsel, Patent Agent,
Trademark Agent

780.917.4265
elmallaha@bennettjones.com

Edward (Ted) Yoo

Partner, Patent Agent,
Trademark Agent

780.917.5231
yoot@bennettjones.com

Lorelei Graham

Partner, Head of Agribusiness
Industry Team

416.777.6547
grahaml@bennettjones.com

Stephen D. Burns

Partner, Trademark Agent,
Co-Head of Innovation,
Technology & Branding
Practice

403.298.3050
burnss@bennettjones.com

Benjamin K. Reingold

Partner

416.777.4662
reingoldb@bennettjones.com

J. Sébastien A. Gittens

Partner, Trademark Agent

403.298.3409
gittenss@bennettjones.com

Kees de Ridder

Associate, Patent Agent,
Trademark Agent

403.298.3122
deridderk@bennettjones.com

Stay head of the curve in intellectual property and artificial intelligence. Update your subscription preferences to receive timely insights and strategic analysis delivered straight to your inbox.

Meet Our Artificial Intelligence Team

Meet Our Intellectual Property Law Team

Bennett Jones

Decision Framework: Patents or Trade Secrets

As artificial intelligence innovation accelerates, inventors, chief technology officers (CTOs) and in-house teams are increasingly confronted with a fundamental protection question: *does it make sense to file a patent, or is the better strategy to rely on trade secret protection?*

While the patent lawyer may be biased towards "always patent", companies need to address more substantive business realities, including the true value add of a patent versus its cost.

Why a Decision Framework?

The question of patenting AI (or keeping it a trade secret and confidential) is, in many ways, not fundamentally different from the longstanding challenges surrounding **software patenting**, since AI innovation is primarily software-driven.

At the same time, AI introduces **distinct considerations** from other classes of software innovation. Most notably, the effectiveness of an AI platform often turns on the quality of its training data. That means that even if a patent discloses all the details of an applied AI or machine-learning model, that disclosure can – in certain cases – be of subdued consequence, especially if competitors lack access to that same quantitative and qualitative data repository.

As with many complex decisions, whether to pursue patent or trade secret protection for AI innovation is best approached through the lens of a **practical, objective framework**. That framework should focus on aligning patent and IP strategy with underlying business realities and moving beyond purely "legal" considerations.

From a **governance perspective**, such frameworks also enable decisions that are "**defensible**" and "**explainable**", both to internal stakeholders (e.g., executives and management) and external stakeholders (e.g., investors and shareholders).

Decision Tool: A 7-Point Framework

It's suggested that decision-makers weigh **seven factors** in determining whether patent or trade secret protection is the more viable approach for protecting AI innovation.

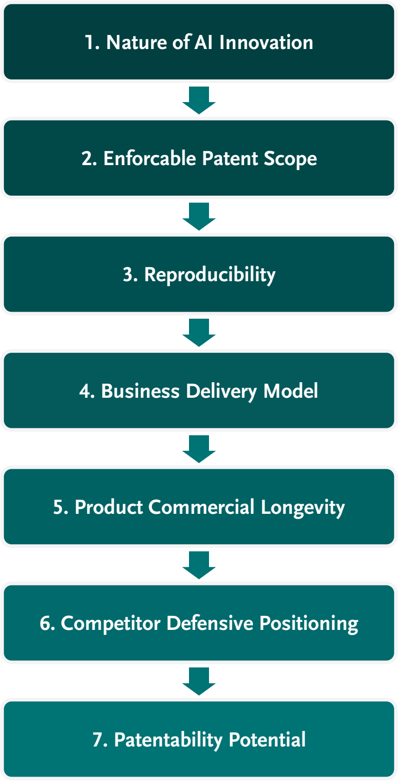

1. Nature of AI Innovation

2. Enforcable Patent Scope

3. Reproducibility

4. Business Delivery Model

5. Product Commercial Longevity

6. Competitor Defensive Positioning

7. Patentability Potential

Taking a step back, AI innovation is broadly categorized into "**core AI**" and "**applied AI**".

Core AI refers to advances in the underlying engines themselves, such as new mathematical model architectures, learning paradigms or training methods. **Applied AI**, by contrast, involves adapting existing AI engines and models to solve specific, real-world problems, often through domain-specific training data and deployment choices (see for example AI in Oil and Gas).

For ease of discussion, this framework focuses only on **applied AI**, as it represents the most common and commercially relevant form of AI innovation in practice.

Please note as well that, while the proposed framework is relevant to patenting generally, it has been tailored to address the specific and unique aspects of AI technologies, and their development and deployment. The framework may also share overlap with broader software patenting considerations.

For a further discussion on layered IP strategies for AI innovation, please also see our on-demand video Intellectual Property: Key Considerations at Every Stage of the AI Value Chain.

If your organization needs assistance evaluating which aspects of its AI innovation are better suited to patent protection versus trade secret protection, our team can help. Our team can also support patent filing and the development of a broader IP strategy.

FACTOR 1

Identifying the True AI Innovation

The first question to ask is *what exactly is the AI innovation, or more specifically, what incremental advance over existing technology is proposed?*

Innovation in applied AI often sits somewhere along the spectrum between the big picture level (or application-level innovation) and the small picture level (or implementation-level innovation). As discussed below, this distinction informs much of the remaining framework.

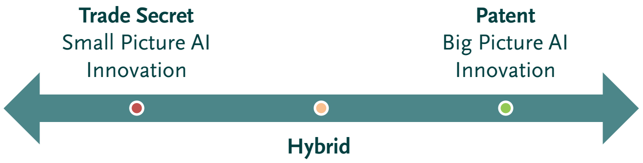

Patents: Big Picture AI Innovation

Big picture AI innovation encompasses new applications or system-level approaches for using AI. This type of innovation is often better suited to patent protection because it represents more visible, high-level advances. Such innovations are easier for competitors to observe and replicate, which increases the value of securing patent rights.

Example: Crop Stress AI Detection

In an agricultural application, an innovation is developed to analyze images of crops to identify crop stress using broad AI-based image analysis techniques.

The innovation here lies at the application level, in recognizing crop stress detection as a suitable application for image-based AI analysis. The innovation may also involve identifying which image features are most informative for training an accurate crop stress model, as well as the nature of the resulting model outputs.

The novelty of the innovation therefore is not the specific type of AI architecture used (i.e., other than image analysis AI models, broadly), but its general application.

Trade Secrets: Small Picture AI Innovation

Small-picture AI innovation involves improving performance or efficiency through implementation-level refinements within an AI system. It often lies in model tuning or architectural details that are difficult to detect externally and are therefore better protected as trade secrets. In some cases, small picture AI may still warrant patent protection where other factors counterbalance the narrow technical scope (e.g., commercial value, etc.)

Example: Geological AI Image Analysis

The use of computer vision (broadly) to identify mineralization patterns in geological imagery is well known (i.e., in this hypothetical case). However, improved accuracy can be achieved by deploying a specific and more complex computer vision model architecture tailored to the visual characteristics of core samples or rock surfaces.

The innovation here lies at the implementation level of the AI itself. It is localized within the model architecture and focused on system optimization. It is not focused on the big-picture application of image-based geologic mineral analysis.

As noted above, patent protection may still be appropriate in these cases where other factors, such as commercial relevance, outweigh the narrow technical scope.

Hybrid Approach (Patents/Trade Secrets)

Technological innovation will most often lie somewhere between the two extremes. In many cases, in fact, AI innovation may include elements of both big picture and small picture innovation. This can require a hybrid patent/trade secret approach.

Example: AI-Driven Analysis of Oilfield Imagery

A novel application is proposed for AI-based image analysis to identify subsurface features or anomalies in oilfield downhole well imaging data (e.g., downhole camera images). Therefore, patent protection may be pursued over the application-level concept.

Further, to satisfy patent disclosure and novelty requirements, the patent can describe the overall system flow, data inputs, expected outputs and relevant feature categories.

At the same time, specific implementation details used to optimize performance of the image analysis (e.g., proprietary model architectures, tuning parameters, feature weighting strategies and training heuristics), can be retained as trade secrets. These details are not required for enablement of the patented invention or to establish novelty and can be kept confidential to preserve a competitive advantage beyond the patent term.

FACTOR 2

The Enforceable Patent Scope

An AI patent with meaningful scope can block competitors from implementing commercially viable alternatives. Assessing enforceable scope therefore requires focusing on some of the following factors: **(a) competitor blocking potential; (b) detectability of infringement;** and **(c) divided infringement.**

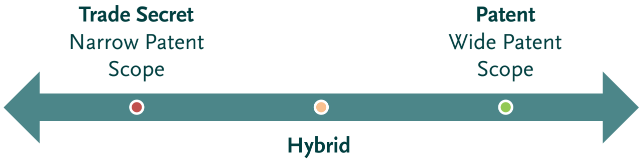

a. Competitor Blocking Potential

If the AI patent scope is too limited, competitors may be able to design around the patent with little effort while still benefiting from the disclosed ideas. That said, narrow scope patents can still be valuable where the inventor has identified a highly valuable and specific combination or configuration that competitors are likely to adopt.

Examples: Competitor Blocking Potential

- **Patents (High Competitor Blocking)**
  As noted in the previous factor, patents that emphasize the overall functional outcome and system-level use of AI tend to offer stronger competitor blocking because the scope of protection is wide and covers a broad application (e.g., big picture AI innovation.
- **Trade Secrets (Low Competitor Blocking)**
  Patents that are narrowly tied to specific technical AI implementations provide weaker blocking potential (e.g., small picture AI innovation). This is because competitors can avoid infringement by making modest technical adjustments to the implementation, while delivering similar functionality.

  Having said this, patents directed to specific technical implementations may still be valuable where those implementations are particularly effective and therefore likely to be sought out or emulated by competitors.
- **Hybrid Approach (Patents/Trade Secrets)**
  A patent may be used to protect overall function and system-level use. Further, narrow technical AI implementation details (to the extent not required for patent enablement requirements), may be protected with a trade secret.

b. Detectability of Infringement

AI patents have limited enforceable value if infringement cannot be practically identified or enforced. Where AI innovation is externally observable/detectable in a competitor's product or published material, patent protection is often viable. Where detectability is low, trade secret protection may be the more effective option. The deployment environment, such as cloud-based, local or edge deployment, plays an important role in determining how easily infringement can be detected.

Examples: Detectability of Infringement

- **Patents (High Detectability)**
  Infringement is more readily detectable where the AI innovation is of a type that, if adopted by a competitor, would be deployed in a way that allows its use to be observed or evaluated. This includes innovations that would normally be implemented in edge-deployed systems or other user-facing products and applications, and that can be assessed through direct product analysis or input–output testing. In some cases, the nature of the innovation is

such that its adoption would also be reflected in a competitor's public disclosures, such as promotional materials or user documentation.

- **Trade Secrets (Low Detectability)**
  By contrast, detectability is lower where AI functionality is embedded within opaque or distributed systems. Where models are deployed exclusively in backend or cloud environments, where inputs and outputs are heavily abstracted, or where public disclosures are limited or vague, it becomes more difficult to evaluate how the AI operates. In these situations, infringement may be difficult to detect or prove, even where similar functionality is suspected.
- **Hybrid Approach (Patents/Trade Secrets)**
  In some cases, an AI innovation includes both observable and non-observable elements. The externally visible aspects of the system, such as user-facing behavior or high-level functional outcomes, may be detectable if adopted by a competitor and therefore suitable for patent protection. At the same time, internal implementation details that operate within backend or cloud environments may remain difficult to observe and are less readily detectable. This mixed detectability supports a hybrid approach.

c. Divided Infringement

AI systems are often distributed across multiple parties. For example, one party may train the model (e.g., developer), another may host it (e.g., cloud service provider) and a third may deploy it within a product or service (e.g., customer). This creates a risk of divided infringement, where no single party performs all steps of a patented method, making enforcement difficult. Although workarounds exist to catch divided infringement, they are complex and uncertain. Patent claims are therefore most effective when they are performed by a single party, and where that party is ideally a competitor rather than a customer.

Examples: Divided Infringement

- **Patents (Low Probability of Divided Infringement)**
  The patent scope is focused on a single locus of activity, such as model training alone, or model deployment alone. Alternatively, it's focused on both, however each provides standalone innovative value.
- **Trade Secrets (High Probability of Divided Infringement)**
  The innovation and patent scope spans both model training and model deployment. In practice, these steps are often performed by different parties, such as a technology provider that trains the model and a customer that deploys it.
- **Hybrid Approach (Patents/Trade Secrets)**
  The patent scope is directed to multiple system-level AI capabilities that can each be implemented and controlled by a single party (e.g., platform provider or service operator). For example, training-related and deployment-related functionality are each separately novel.

  More granular interactions between training and deployment, which in practice may be split across multiple actors, are not relied upon for patent enforcement and are instead maintained as trade secrets. For example, detailed processes for updating a model based on deployment feedback, such as how user data is selected and incorporated into retraining. In certain cases, this may require coordination between a service provider and customers and are therefore more likely to be split across multiple actors.

FACTOR 3

Reproducibility of the AI Innovation

Reproducibility considers whether a competitor could realistically replicate the AI innovation by reviewing your public materials or otherwise by obtaining limited access to the system itself.

If the core functionality can be inferred or reverse-engineered from these sources, trade secrets are of little value since the innovation is readily discoverable. In these cases, patent protection may be more useful to prevent straightforward copying.

To that end, reproducibility is often driven by the deployment environment in which an AI innovation operates. Cloud-based systems, local installations and edge deployments present different levels of visibility, which in turn affect how easily a competitor can access or study the system.

The factors indicating reproducibility are closely aligned with those discussed in the previous factor under “detectability of infringement,” as they represent two sides of the same coin.

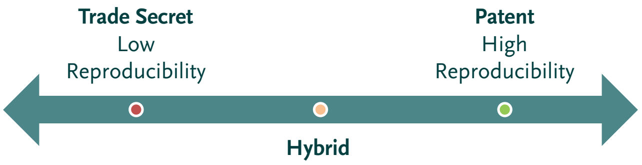

Examples: Reproducibility

- **Patents (High Reproducibility): Local or Edge-Deployed, User-Facing Product**
  An AI feature embedded in an edge-deployed consumer product that performs real-time image or signal analysis and produces observable, repeatable outputs. Because the functionality runs locally and its behavior can be tested by varying inputs and measuring outputs, competitors can broadly infer the overall processing logic, making the capability easy to reproduce.
- **Trade Secrets (Low Reproducibility): Cloud-Deployed, Bundled Product**
  An AI capability implemented as a subcomponent of a bundled, cloud-based platform where the functionality is distributed across multiple backend services, data pipelines and orchestration layers. Inputs may be abstracted, outputs may be aggregated or post-processed and internal workflows are hidden from users.

  The lack of visibility into the integrated platform architecture makes the capability difficult to decipher or replicate, favoring trade secret protection.
- **Hybrid Approach (Patents/Trade Secrets): Multi-Component AI Product**
  An AI capability deployed as a module within a larger software platform that provides defined inputs and outputs with consistent, testable behavior, while relying on backend processing that is not fully visible.

  Core functional logic and system-level interactions are observable and can be disclosed and protected through patents (big picture AI). In contrast, certain internal data transformations and optimization routines (small picture AI) are not directly visible and remain difficult to reverse engineer.

  This intermediate level of reproducibility supports a hybrid approach, combining patent protection for the externally discernible aspects with trade secret protection for internal implementation details.

FACTOR 4

The Business Delivery Model

How an AI system is commercialized and deployed affects whether patent protection is appropriate. In many cases, the degree of control retained by the provider over the deployed system acts as a counterweight to other considerations that might otherwise favor trade secret protection.

Where the deployment model is **customer-controlled**, such as when the innovation is licensed for use to the customer, sold as a stand-alone product or deployed within a customer's own IT environment, patent protection can play an important role. These delivery models necessarily expose the technology to customers or integration partners, increasing the risk that key aspects may be accessed, replicated or reused. In such cases, patents provide enforceable rights that extend beyond contractual use restrictions and confidentiality obligations.

By contrast, where the deployment model is **provider-controlled**, such as when the AI is offered as a hosted SaaS service or used internally within the company, access to the underlying implementation is more tightly controlled. In these scenarios, trade secret protection may be more suitable for certain aspects of the technology, particularly where customers interact only with outputs rather than the system itself.

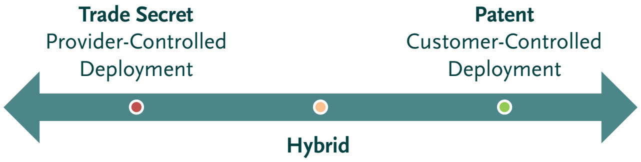

Examples: Business Delivery Model

- **Patents (Customer-Controlled Deployment): Financial Risk-Scoring AI Platform**
  An AI risk-scoring platform that analyzes transaction data to generate fraud or compliance scores and is deployed within a financial institution's own IT environment.

  Under this customer-controlled deployment model, the platform is licensed to third-party financial institutions for local installation and operation on their internal training data repositories.

  Because the developer does not retain operational control over the deployed system, patent protection plays an important role in protecting the core technology once it is transferred to customers for independent use within their IT infrastructure.

- **Trade Secrets (Provider-Controlled Deployment): Logistics AI Optimization for Manufacturing Workflows**
  An AI-based supply chain visibility platform offered to third-party customers as a centrally hosted SaaS service. The platform ingests logistics and operational data from multiple customers to provide customer-specific real-time insights, forecasting and alerts. The AI models and core system logic remain fully controlled and operated by the provider.

  In this example, customers may interact with the service through dashboards and APIs but do not receive access to the underlying models or deployment environment.

  Because the provider retains control over the AI and its execution, exposure to competitors is limited, making this deployment model more conducive to protecting key aspects of the innovation as trade secrets.

- **Hybrid Approach (Patents/Trade Secrets): Retail E-Commerce AI Recommendation Engine**
  A company develops an AI-based recommendation engine for the retail e-commerce industry using a hybrid delivery model that combines customer-controlled and provider-controlled elements.

  The core recommendation engine is licensed to merchants and integrated into their online storefronts, where it operates within the merchant's environment to generate personalized product recommendations based on local user behavior and product data. This customer-controlled deployment exposes the system architecture and integration interfaces to third parties, making this component well suited to patent protection and external licensing.

  At the same time, the company retains a provider-controlled AI system that operates centrally across the platform. This internal system analyzes aggregated interaction data across multiple merchants to identify platform-level performance patterns and guide ongoing product development. Because this functionality is never deployed to customers and derives its value from cross-merchant aggregation under the provider's control, it is not offered for licensing and is better protected as a trade secret.

FACTOR 5

The Commercial Longevity

Commercial longevity considers whether the AI innovation is likely to remain commercially relevant long enough to justify patent protection.

Innovations with short life cycles may not justify the time and cost of patenting, whereas durable capabilities that persist across product generations and address stable problem domains are stronger candidates. It is important to note that a patent can take several years to grant, and enforceable rights arise only after issuance.

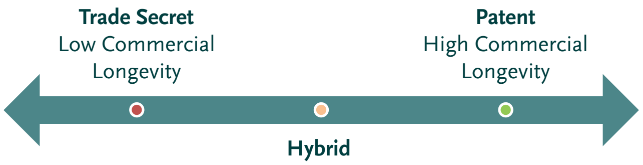

Examples: Commercial Longevity

- **Patents (High Commercial Longevity): Medical AI Imaging Analysis**
  An AI-based medical image analysis system for detecting common pathologies in diagnostic imaging, where the clinical need, imaging modalities and core input features evolve slowly over time. As a result, the underlying model logic and outputs remain relevant across multiple product generations, allowing the innovation to retain long-term commercial value.

  In this example, product iterations/evolutions are not necessarily based on revising the big picture AI innovation (e.g., using AI to detect medical pathologies), but on the small picture AI innovation (e.g., tweaking the implementation for greater accuracy).

- **Trade Secrets (Low Commercial Longevity): AI-Enabled Inspection Module**
  An AI-enabled inspection module designed to be deployed with handheld or mounted scanning hardware to support a one-time infrastructure upgrade program, such as the rollout of a new generation of smart utility meters.

  During the upgrade period, the AI analyzes sensor readings and device identifiers to verify installation correctness and compatibility with legacy systems. The module is sold or licensed to utilities and contractors specifically for the duration of the rollout.

  Once the upgrade program is completed and legacy meters are retired, the need for the module largely disappears, giving it limited commercial longevity. This case may therefore favor trade secret protection.

- **Hybrid Approach (Patents/Trade Secrets): Smart Fashion AI Platform**
  A smart fashion platform may use AI to analyze images captured by a wearable camera or smartphone to identify clothing items and basic visual attributes (e.g., color, fit and layering). This core capability addresses a persistent and reusable problem across many fashion applications and does not depend on short-term trends, making it well suited to patent protection.

  A second, interrelated component may support a specific brand collaboration or limited-run campaign. In this case, the AI is re-adapted or retrained to recognize collection-specific garments or styling rules that apply only to that collaboration. Because this behavior is tied to a time-limited commercial initiative and may not be reused, its commercial longevity is uncertain, making it better suited to trade secret protection.

FACTOR 6

Competitor Defensive Positioning

Defensive blocking considers whether a patent can prevent others from patenting or controlling adjacent technical or commercial space.

Defensive blocking matters for two main reasons: (i) to stop competitors from obtaining AI patents that could later restrict your ability to operate, and (ii) to provide leverage for counter-assertion if a competitor brings a patent claim against you.

This consideration also extends to the risk that competitors may gain access to the technology through your former employees with detailed knowledge of your systems (i.e., despite the internal function of the system not being reproducible).

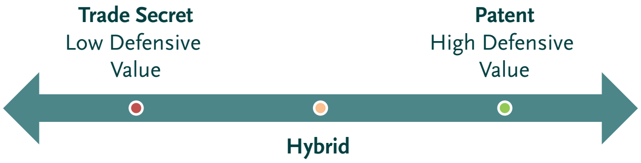

Examples: Competitor Defensive Positioning

- **Patents (High Defensive Value): AI-Based Documentation Classification and Compliance**
  An AI-based document classification or compliance analysis system deployed in a highly competitive market, where multiple vendors are developing similar solutions using overlapping techniques. Patents in this space can block competitors from patenting incremental variations and provide leverage in negotiations or disputes involving overlapping rights.

- **Trade Secrets (Low Defensive Value): Company-Specific Maintenance Operations**
  An AI system used internally to optimize maintenance scheduling for a company's proprietary equipment based on custom sensor configurations and operational constraints. The use case is highly specific to the company's internal processes and custom hardware, with few external competitors and little incentive for others to develop or patent similar solutions. This results in low defensive value for patent protection and making trade secret protection more appropriate.

- **Hybrid Approach (Patents/Trade Secrets): AI for Controlling Hydrogen Production with Client-Specific Adaptations**
  A company offers a licensable AI-based control platform for hydrogen production processes. The AI platform can manage electrolysis or reforming operations using standard process data including temperature, pressure and gas composition.

  The platform is offered broadly in a competitive market with multiple vendors providing similar AI-driven solutions. Patent protection is used to cover the system-level application of AI for monitoring and controlling hydrogen production. This provides defensive value against competing platforms.

  For individual customers, the platform is optionally further adapted to account for site-specific equipment configurations, or operating conditions, unique to a given hydrogen facility. These customer-specific AI adaptations are novel but are tightly coupled to the client's processes and provide little value outside that context. Because non-client competitors have little incentive to replicate them, these adaptations offer low defensive value and are better protected as trade secrets.

FACTOR 7

Patentability Potential & Layered Strategies

Patentable potential addresses whether the innovation is likely to satisfy the main requirements for patentability: novelty, non-obviousness and subject-matter requirements.

This assessment is highly fact-specific and depends on the prior art landscape and how the invention is framed. Early analysis helps determine whether patent protection is realistic and worth pursuing. For a more in-depth discussion on the topic, we encourage you to review our prior publication *Artificial Intelligence Patenting: Top Challenges and Key Considerations.*

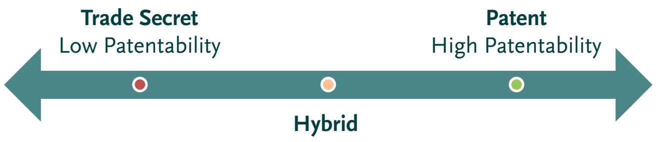

In a multi-layer analysis, you may choose to proceed with a patent application even if patentability is low, if other factors are in favor of patenting. For example, if there is a desire to position an application for defensive blocking, that may be reason to proceed even if patentability is in question.

Beyond Patents vs. Trade Secrets: A Layered Strategy

Patent protection decisions are rarely binary. Effective AI portfolios often rely on a layered strategy that aligns different forms of protection with different aspects of the technology.

- **Patents** are typically best suited for externally visible applications, system behavior and technical effects that can be observed, reverse-engineered or independently developed by competitors.
- **Trade secrets** are more appropriate for elements that derive value from remaining hidden, such as training data, data engineering workflows, model tuning strategies and internal performance optimizations.
- **Contractual controls and internal policies** then operate as a supporting layer, reinforcing confidentiality obligations, limiting misuse in partnerships or joint ventures and helping preserve trade secret status over time. As noted in our previous publications, AI companies are strongly advised to deploy internal Trade Secret, Confidentiality and IP Policies as added safeguards (see *X.AI Corp. v. OpenAI—Why Every Business (and Start-Up) Needs an Employee Governance Policy for Managing Confidential Business Information and IP Risk*).

Bennett Jones

The firm that businesses trust
with their most complex legal matters.

BennettJones.com

Calgary Edmonton Montréal Ottawa Toronto Vancouver New York