

# USPTO Should Improve Governance to Promote Effective Oversight of Its Artificial Intelligence Tools

FINAL REPORT NO. OIG-25-018-A

APRIL 22, 2025



U.S. Department of Commerce  
Office of Inspector General  
Office of Audit and Evaluation



April 22, 2025

**MEMORANDUM FOR:** Coke Morgan Stewart  
Acting Under Secretary of Commerce for Intellectual Property  
and Acting Director of the United States Patent and Trademark  
Office

**FROM:** Kevin D. Ryan  
Acting Assistant Inspector General for Audit and Evaluation

**SUBJECT:** *USPTO Should Improve Governance to Promote Effective Oversight of  
Its Artificial Intelligence Tools*  
Report No. OIG-25-018-A

Attached is the final report on our audit of USPTO's governance of its artificial intelligence tools. We will post the report on [our website](#) per the Inspector General Act of 1978, as amended (5 U.S.C. §§ 404, 420).

Within 60 calendar days, please provide an action plan addressing the report's recommendations, as required by Department Administrative Order 213-5.

Any nongovernmental organization or business entity specifically identified in this report can submit a written response to clarify or provide additional context on any specific reference (Pub. L. No. 117-263, § 5274). The response must be submitted to Amni Samson, Director for Audit and Evaluation, Intellectual Property Program, at [asamson@oig.doc.gov](mailto:asamson@oig.doc.gov) and to [OAE\\_Projecttracking@oig.doc.gov](mailto:OAE_Projecttracking@oig.doc.gov) within 30 days of the report's publication date. We will post the response on our website as well. If the response contains any classified or otherwise nonpublic information, the organization should identify the information and provide a legal basis for redacting it.

We appreciate your staff's cooperation and professionalism during this audit. If you have any questions or concerns about the report, please contact me at 202-750-5190 or Amni Samson at 202-793-3324.

Attachment

cc: Will Covey, Acting Deputy Under Secretary of Commerce for Intellectual Property and  
Acting Deputy Director, USPTO  
Valencia Martin Wallace, Acting Commissioner for Patents, USPTO  
Jamie Holcombe, Chief Information Officer, USPTO  
Jay Hoffman, Chief Financial Officer, USPTO  
Sean Mildrew, Deputy Chief Financial Officer and Audit Resolution Officer, USPTO  
Nicolas Oettinger, Senior Counsel for Regulatory and Legislative Affairs, USPTO  
Stacy Long, Senior Counsel for Employment Litigation, USPTO





# Report in Brief

April 22, 2025

## Background

To deliver reliable intellectual property rights, the United States Patent and Trademark Office (USPTO) is leveraging emerging technologies like artificial intelligence (AI) that may enhance the quality of patent and trademark examinations and improve efficiency in its operations. For example, patent examiners are using AI tools to help assess whether the inventions in patent applications are patentable as claimed. USPTO also has public-facing AI tools intended to improve the usability of patent and trademark systems or to enhance customer service.

## Why We Did This Review

With AI tools, USPTO has an opportunity to improve efficiency and reliability. However, these tools come with inherent risks that USPTO must manage to ensure accountability and responsible AI use. An effective governance structure for AI use ensures accountability, transparency, and safety, which in turn fosters public trust.

Our audit objective was to determine whether USPTO had an effective governance structure and processes in place to manage its AI tools. To meet our objective, we tested two of the six AI tools USPTO had in use when we began our audit.

## UNITED STATES PATENT AND TRADEMARK OFFICE

### USPTO Should Improve Governance to Promote Effective Oversight of Its Artificial Intelligence Tools

OIG-25-018-A

## WHAT WE FOUND

Overall, we found that USPTO has begun developing its AI workforce but should strengthen key organizational and system-level governance practices needed to effectively manage and oversee its AI tools. Specifically, USPTO:

- has a governance structure that defines roles and responsibilities for key personnel, but it should improve internal stakeholder involvement;
- should promote transparency to external stakeholders on its AI tools;
- does not have the specific, measurable objectives needed to define system success;
- did not trace requirements or technical specifications to system objectives; and
- does not have an AI-specific risk management plan.

Together, these weaknesses increase the risk that USPTO will develop unreliable, untrustworthy AI systems.

Strong governance can drive and enhance internal practices and norms to facilitate the safe, reliable, and trustworthy development and use of AI. Adopting key governance practices will promote accountability, transparency, and safety in AI use now and in the future. It will also further USPTO's mission to drive U.S. innovation and global competitiveness through the efficient delivery of reliable intellectual property rights.

## WHAT WE RECOMMEND

We recommend that the Under Secretary of Commerce for Intellectual Property and USPTO Director direct the Commissioner for Patents and the Chief Information Officer to:

1. Revise policies and practices to ensure that future AI systems develop measurable objectives and system success criteria during project planning and ensure that such objectives trace to system specifications.
2. Develop clear, specific, and measurable objectives and system success criteria for the two AI tools we reviewed.
3. Revise project management policies and practices to encourage increased engagement, as appropriate, with end users and others affected by the system, such as the public, earlier in the project lifecycle.
4. Create an AI-specific risk management plan (or revise existing policies and procedures to reflect the same) that requires a system-specific risk management approach, including identifying, documenting, analyzing, and managing risks.
5. Implement AI-specific risk management practices for the two tools in our review.

# Contents

<b>Introduction.....</b>	<b>1</b>
<b>Objective, Finding, and Recommendations .....</b>	<b>5</b>
I. USPTO Should Strengthen Key Governance Practices That Are Needed to Effectively Manage and Oversee Its AI Tools .....	6
A. <i>USPTO has a governance structure that defines roles and responsibilities for key personnel, but it should improve stakeholder involvement.....</i>	<i>6</i>
B. <i>USPTO should promote transparency to external stakeholders on its AI tools .....</i>	<i>8</i>
C. <i>USPTO does not have the specific, measurable objectives needed to define system success .....</i>	<i>8</i>
D. <i>Requirements or technical specifications could not be traced to system goals and objectives .....</i>	<i>10</i>
E. <i>USPTO does not have an AI-specific risk management plan .....</i>	<i>11</i>
F. <i>Rapidly changing AI technology and inadequate planning led to the need for stronger key governance practices .....</i>	<i>12</i>
Conclusion.....	13
Recommendations .....	13
<b>Summary of Agency Response and OIG Comments .....</b>	<b>14</b>
<b>Appendix A: Objective, Scope, and Methodology .....</b>	<b>15</b>
<b>Appendix B: Agency Response .....</b>	<b>17</b>

Cover: Herbert C. Hoover Building main entrance at 14th Street Northwest in Washington, DC. Completed in 1932, the building is named after the former Secretary of Commerce and 31st President of the United States.

# Introduction

The United States Patent and Trademark Office (USPTO) is responsible for granting patents and registering trademarks. Since 2000, patent and trademark application filings have increased exponentially. To deliver reliable intellectual property (IP) rights, USPTO is leveraging emerging technologies like artificial intelligence (AI) that may enhance the quality of patent and trademark examinations and improve efficiency in its operations.

## AI's Role at USPTO

While descriptions vary, AI can be defined as “a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.”<sup>1</sup> At USPTO, examiners are using AI tools to find prior art<sup>2</sup> related to patent applications to assess whether the inventions in the applications are patentable as claimed. USPTO also has public-facing AI tools intended to improve the usability of patent and trademark systems or to enhance customer service. In addition, USPTO is working toward creating AI tools for image searching and identification of goods, which are intended to assist examiners reviewing trademark applications.

With AI tools, USPTO has an opportunity to improve the efficiency of its operations and strengthen the reliability of IP rights. However, these tools come with inherent risks that USPTO must manage to ensure accountability and responsible AI use.

In 2020, an executive order established principles to ensure that agencies design, develop, acquire, and use AI in a manner that fosters public trust and builds confidence in AI.<sup>3</sup> In 2021, in consideration of that order, the U.S. Government Accountability Office (GAO) identified a framework of principles and key practices for federal agencies to follow when developing and implementing AI systems.<sup>4</sup> Of those, the first principle that needs to be in place is governance structures and processes for managing, operating, and overseeing these systems. An effective governance structure for AI use ensures accountability, transparency, and safety, which in turn fosters public trust in AI. Our audit focused on USPTO's governance to ensure responsible development and use of AI technology.

## AI Governance and Oversight in the Department and at USPTO

Agencies across the federal government are using AI to better serve the public, making AI a top priority for government leaders. The December 2020 executive order encouraged innovation and use of AI to improve government operations and services. The order highlighted the

---

<sup>1</sup> 15 U.S.C. § 9401(3), “Artificial intelligence.”

<sup>2</sup> Prior art consists of publicly available information such as patents and published patent applications, journal and magazine articles, books, catalogs, and websites.

<sup>3</sup> Executive Office of the President. December 2020. Executive Order 13960: *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government*, 85 Fed. Reg. 78939.

<sup>4</sup> GAO. June 2021. *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities*, [GAO-21-519SP](#). Accessed February 2024.



importance of trustworthy AI and described principles for agencies to follow, including that AI use is accurate, reliable, and effective; that AI tools are understandable and traceable; and that agencies are transparent in disclosing information to stakeholders.<sup>5</sup>

Consistent with the federal government's initiatives to promote trustworthy AI, the U.S. Department of Commerce established multiple AI governance bodies:

- The Commerce AI Governance Board, chaired by the Deputy Secretary of Commerce, serves as the highest agency-wide governance body that coordinates and governs issues tied to the use of AI. This board includes USPTO representatives.
- The Commerce AI Council, chaired by the Chief AI Officer, serves as the primary agency-wide advisory body for AI strategy, operations, and reporting at the Department.
- The AI and Open Government Data Assets Working Group, chaired by the U.S. Census Bureau's chief scientist, was established by the Commerce Data Governance Board and is tasked with developing guidelines for publishing the Department's AI-ready open data. The group includes USPTO representatives.

In January 2024, USPTO established the AI Innovation and Governance Council,<sup>6</sup> tasked to advance the effective and responsible use of AI. The council's planned activities include facilitating initiatives that improve AI governance and risk management across USPTO.

Apart from the bodies discussed above, USPTO's AI tools follow established IT policies, processes, and procedures. USPTO's Office of the Chief Information Officer develops and implements IT plans and budgets for patent and trademark systems as well as operating the bureau's computer facilities, equipment, and network. This office also has the responsibility to develop, maintain, and oversee the USPTO IT security program, which includes authorizing USPTO systems for use through an authority to operate (ATO).<sup>7</sup> The Office of Information Technology for Patents manages the development and implementation of IT systems, tools, and resources that are specific to patents.

## USPTO AI Tools We Reviewed

When we initiated this audit in March 2024, USPTO had six AI tools in use. To evaluate USPTO's system-level governance practices, we selected two of these tools as case studies for detailed testing: the Cooperative Patent Classification (CPC) AI tool and the patent search AI tool. A description of each tool follows.

**CPC AI tool.** CPC is a classification structure for grouping all patent documents by their subject matter. The CPC structure is hierarchical, starting with a broad category (called a section) and adding increasingly specific subcategories denoted by alphanumeric codes. At the

---

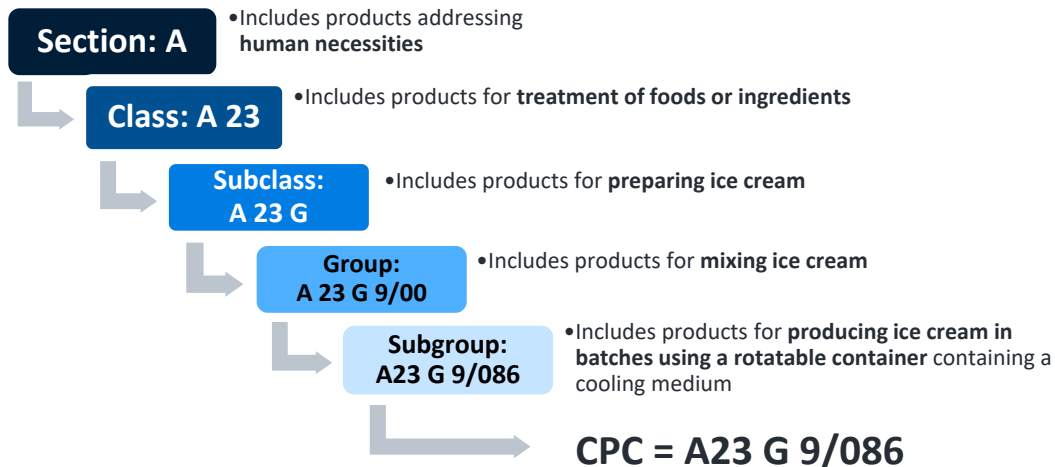
<sup>5</sup> Executive Order 13960: *Trustworthy Artificial Intelligence in the Federal Government*, 85 Fed. Reg. 78939–78942.

<sup>6</sup> USPTO. n.d. AI Innovation and Governance Council charter.

<sup>7</sup> ATO is the official management decision by a senior official to authorize operation of an information system. Before issuing an initial ATO for a system, all applicable security controls must be assessed.

lowest subgroup level, the final code is known as a CPC symbol. As an example, figure I shows the CPC structure and symbol for a type of ice cream maker.

**Figure I. CPC Hierarchical Structure Example**



Source: OIG analysis of [a CPC scheme](#) (accessed October 2024)

The desired outcome of the CPC AI tool is to improve the quality and efficiency of patent application processing by completing two processes, CPC autotclassification and auto C-star (or C\*):

- In the CPC autotclassification process, the AI tool uses natural language processing and machine learning algorithms to classify and automatically assign a CPC symbol or symbols to a patent application.<sup>8</sup> A single application could have dozens of CPC symbols.
- When a CPC symbol represents at least one concept specifically claimed for patent protection, the tool automatically adds a C\* label to the symbol. At least one CPC symbol in an application must be designated C\*. USPTO then uses the CPC symbols with the C\* designation for application routing and examination timing.<sup>9</sup>

USPTO must fully classify all nonprovisional utility patent applications when it first receives them.<sup>10</sup> Once an application is classified based on its subject matter, USPTO can route it to an examiner for a patentability determination.<sup>11</sup>

<sup>8</sup> Natural language processing uses machine learning to enable computers to understand and communicate with human language. Machine learning algorithms use data to infer rules or decision procedures that aim to predict specified outcomes.

<sup>9</sup> Examination timing refers to the time allotted to an examiner to review a specific application.

<sup>10</sup> The statutory authority and requirement to classify patents is explained in USPTO's July 2022 *Manual of Patent Examining Procedure* (MPEP), 9th ed. See MPEP §§ 903, "Classification in USPC," and 905, "Cooperative Patent Classification." CPC symbol assignments apply only to utility applications.

<sup>11</sup> For our other oversight work related to patent classification and routing processes, see our August 30, 2023, report, *USPTO Needs to Improve Oversight and Implementation of Patent Classification and Routing Processes* ([OIG-23-026-A](#)).

The CPC AI tool was initially designed and developed in-house by IT staff. However, due to performance, technical, and sustainment limitations with the in-house tool, USPTO decided to implement only the C\* process and halted rollout of the tool's CPC symbol assignment capability. In October 2022, USPTO awarded a contract to improve the CPC symbol assignment system's infrastructure and model performance. Throughout this report, we refer to this contracted effort as the "next-generation CPC AI tool."

**Patent search AI tool.** When examining a patent application, patent examiners must thoroughly search for prior art. To improve the efficiency and quality of these searches, USPTO released two new AI features as part of the Patents End-to-End search suite. The "More Like This Document" feature was deployed in October 2021, followed by "Similarity Search" in September 2022:

- More Like This Document allows the examiner to select a document generated during a prior art search as an "anchor" document. It then uses an AI algorithm to return a list of documents that are like the anchor document.
- Similarity Search allows the examiner to select specific text in an application, such as the application abstract. It then uses trained AI models to return a list of patent documents that are similar to the patent application information.<sup>12</sup>

Both the CPC and the patent search AI tools are critical to the patent examination process. With a strong governance structure and processes in place, these tools can improve the efficiency of USPTO operations and strengthen the reliability of IP rights.

---

<sup>12</sup> A trained AI model refers to a program that has been trained on one or more sets of data. Once trained, the model can be applied and used to make future predictions without further human intervention.



## Objective, Finding, and Recommendations

Our audit objective was to determine whether USPTO had an effective governance structure and processes in place to manage its AI tools. Specifically, we assessed six key practices, identified in GAO's AI accountability framework, that USPTO should incorporate into its governance structure and processes.<sup>13</sup> We focused on USPTO's organizational and system-level governance processes and practices for its AI tools, considering information from 2019, when USPTO implemented its first AI tool, through October 2024. Appendix A details our scope and methodology.

Overall, we found that USPTO has begun developing its AI workforce but should strengthen key organizational and system-level governance practices that are needed to effectively manage and oversee its AI tools. Specifically, USPTO:

- has a governance structure that defines roles and responsibilities for key personnel, but it should improve internal stakeholder involvement;
- should promote transparency to external stakeholders on its AI tools;
- does not have the specific, measurable objectives needed to define system success;
- did not trace requirements or technical specifications to system objectives; and
- does not have an AI-specific risk management plan.

These issues exist, in part, because AI is a relatively new and rapidly changing technology with limited government-wide direction for agencies to rely on. In addition, while USPTO's AI tools, as IT products, must adhere to USPTO's IT policies and procedures, USPTO did not conduct the necessary project planning early in the development of the AI tools we reviewed as our case studies. Finally, USPTO relies on its ATO process, which focuses on cybersecurity and privacy, to assess AI-related risks.

Together, these weaknesses increase the risk that USPTO will develop and use AI systems that are unreliable and untrustworthy. Using AI for its benefits also requires mitigating substantial risks. To protect the public's trust, it is crucial to have governance that proactively addresses the complexities, risks, and potential consequences that come with AI.

Strong governance can in turn drive and enhance internal practices and norms to facilitate the safe, reliable, and trustworthy development and use of AI. The need for effective governance practices extends far beyond the two tools we reviewed. AI is still a developing technology, with many potential uses that have yet to be explored. Adopting key governance practices will promote accountability, transparency, and safety in AI use now and in the future. It will also

---

<sup>13</sup> Throughout this report we cite GAO's *AI Accountability Framework*. However, the practices described in the framework are similar to requirements and best practices found in sources that precede GAO's framework, such as the Office of Management and Budget's July 2016 Circular No. A-130, *Managing Information as a Strategic Resource*, and the Software Engineering Institute's November 2010 *CMMI® [Capability Maturity Model Integration] for Acquisition*, ver. 1.3.

further USPTO's mission to drive U.S. innovation and global competitiveness through the efficient delivery of reliable IP rights.

## I. USPTO Should Strengthen Key Governance Practices That Are Needed to Effectively Manage and Oversee Its AI Tools

Federal AI guidance has focused on ensuring that agencies are accountable and responsible in their design, development, and continuous monitoring of AI systems. To promote accountability, USPTO needs governance structures and processes to provide oversight, manage risks, and ensure that AI systems meet performance requirements.

We found that USPTO has begun developing its AI workforce but should strengthen key organizational and system-level governance practices in the areas of stakeholder involvement, transparency, system objectives, specifications, and risk management. Limited AI guidance and USPTO's own lack of early planning contributed to the issues we found. Without effective governance, USPTO cannot be certain that the CPC and patent search AI tools operate successfully. Our assessments of these areas are described in the following sections.

### A. *USPTO has a governance structure that defines roles and responsibilities for key personnel, but it should improve stakeholder involvement*

According to GAO, federal agencies should implement governance structures that assign clear roles and responsibilities and involve diverse perspectives from a community of stakeholders throughout the AI lifecycle. Management should define roles and responsibilities and delegate authority for the various stages of the AI system's lifecycle, including design, development, deployment, and monitoring. The roles and responsibilities of personnel should be appropriate and clearly understood, and stakeholder involvement should include engagement with individuals who may use or be affected by the AI system.<sup>14</sup>

To evaluate USPTO's governance structure, we reviewed organizational charts, IT guidance, and other relevant documents. We found that USPTO had an established oversight structure for managing its IT products, which include its AI tools. For example, USPTO has a quarterly review board that provides oversight of the progress of IT products.

Specific to AI, as previously stated, USPTO established the AI Innovation and Governance Council in January 2024. According to its charter, the council was created to facilitate collaboration on initiatives that will improve AI governance and risk management across USPTO. The council organized its activities under four workstreams:

- Fostering agency AI innovation
- Use case assessment and documentation

---

<sup>14</sup> GAO, *AI Accountability Framework*, 27–28.

- Controls and mitigations
- Workforce policies and outreach

We found that USPTO defined roles and responsibilities for its IT investments through various policies and procedures that were applicable to AI system development. Through these policies, USPTO defined key personnel for system-level governance, such as the product line lead, lead product owner, and product owner. In addition, the personnel assigned to these roles for the CPC AI and patent search AI tools understood their responsibilities as they related to the various stages of the AI system's lifecycle.

However, USPTO can improve its stakeholder involvement prior to AI system development. While we determined that some stakeholders were involved in all stages of the AI lifecycle, we also found that examiners should have been engaged earlier. Feedback during the design/planning phase from users or others outside the development team who are affected by a tool (examiners, for example) can provide valuable insights on matters related to the tool's context, such as whether to pursue a project or how to effectively design a project.<sup>15</sup>

**CPC AI tool.** USPTO did not engage with examiners in the early lifecycle phases of the CPC AI tool, though it later did so after a prototype was developed for the legacy tool. According to a USPTO official, the Patents group made a business decision to keep the tool's audience small. We surveyed patent examiners and found that only 56 percent of respondents were even aware that USPTO used AI to apply the C\* label to patent applications.<sup>16</sup> Additionally, according to a USPTO official, the bureau has yet to obtain examiner feedback on the next-generation CPC AI tool.

As part of the patent examination process, it is the duty of each primary examiner to personally review the original classification.<sup>17</sup> Therefore, including examiners as early as possible on AI tools that could affect their work would help to ensure that their concerns are considered before the tools' development.

**Patent search AI tool.** We found that examiners' perspectives were also not included before the tool's development. As the tool's primary users, examiners have the most firsthand experience and therefore are best positioned to provide valuable insight and suggest potential improvement. Of the survey respondents who had used the patent search features, only 40 percent of users were satisfied with More Like This Document and just 36 percent were satisfied with Similarity Search.

---

<sup>15</sup> National Institute of Standards and Technology (NIST). *Artificial Intelligence Risk Management Framework (AI RMF) Playbook*, "[Govern](#)," Accessed September 25, 2024.

<sup>16</sup> We surveyed patent examiners, primary patent examiners, and supervisory patent examiners. The results are based on completed surveys from 1,822 out of 8,805 examiners. The results cannot be generalized to all examiners due to the survey's low response rate. Appendix A details our survey methodology.

<sup>17</sup> MPEP § 905.03, "Classifying in CPC."

*B. USPTO should promote transparency to external stakeholders on its AI tools*

A key system-level governance practice is to promote transparency for anyone who interacts with an AI system, knowingly or not.<sup>18</sup> Transparency means that the government makes information about its AI systems available to external stakeholders, including the public, and provides clear and sufficient information to promote awareness of the tool's use. Such information gives stakeholders a greater understanding of the AI system, including what the system is, how it may affect them, and how its output will be used.

When effectively implemented, transparency into AI systems can increase stakeholders' confidence in the systems because it enhances the systems' explainability, data, outputs, and intended outcomes.<sup>19</sup> We found that USPTO should improve transparency by communicating with external stakeholders on its use of AI.

**CPC AI tool.** USPTO decided not to seek public feedback on the tool and shared limited notifications, meaning these stakeholders were not informed about how the tool will be used or not used (e.g., limitations). While members of the public are not direct users of the tool, it is used to classify and categorize patent applications; therefore, the applicants and patent practitioners who submit the applications are indirectly affected and should be aware of the tool's use and purpose. Similarly, since the CPC is the result of a partnership between the European Patent Office and USPTO, sharing more information on this AI resource will increase international awareness of the tool and its potential to help the global IP community enhance classification and improve patent searching.

**Patent search AI tool.** The tool's patent search features have not been made available to the public, despite a notification from USPTO that they would be and a legal requirement to make search systems available for public use.<sup>20</sup> One of the initial goals of this tool was to deliver relevant prior art to applicants as well as examiners. If USPTO released the tool's publicly shareable features, it would help promote the tool's transparency and explainability.

*C. USPTO does not have the specific, measurable objectives needed to define system success*

A key first step in the lifecycle of an AI system is to clearly define the system's goals and objectives.<sup>21</sup> To ensure that intended outcomes can be achieved, the goals and objectives should be specific, measurable, and clear. All levels of the organization should understand what is to be achieved, how it is to be achieved, and how the AI system will provide functions more effectively and efficiently than previously used methods.

<sup>18</sup> GAO, *AI Accountability Framework*, 5, 30–31, 36.

<sup>19</sup> GAO, *AI Accountability Framework*, 5, 30–31; and NIST, January 2023, *Artificial Intelligence Risk Management Framework* (AI RMF 1.0), [NIST AI 100-1](#), 15. Accessed September 2024.

<sup>20</sup> 35 U.S.C. § 41(i)(2), *Availability of Automated Search Systems*.

<sup>21</sup> GAO, *AI Accountability Framework*, 5, 26–27.

While USPTO identified goals for both tools we reviewed, the goals were not accompanied by specific, measurable, and clear objectives and metrics. Measurable objectives define system success, provide a basis to evaluate whether the AI tools are operating effectively and efficiently, and help ensure that all stakeholders understand the tools' success criteria. Table I provides initial goals for each system, as described by USPTO.

**Table I. Initial USPTO Goals by AI Tool**

AI Tool	Goal Description
<b>CPC AI tool</b>	<ul style="list-style-type: none"> <li>Accurately route applications, apply the appropriate time for examiner credit, and support quality dissemination of IP</li> <li>Use AI technologies to automate the addition of CPC symbols to the application and to determine C* claim indicators</li> </ul>
<b>Patent search AI tool</b>	<ul style="list-style-type: none"> <li>Use AI technologies to provide additional suggestions of prior art during the examiner searching function</li> <li>Put the agency on a path towards pre-search that can be provided to applicants</li> </ul>

Source: Unfunded request to support Office of the Under Secretary and Patent AI efforts, dated September 2019

**CPC AI tool.** During the tool's initial design and development, USPTO did not have specific, measurable, and clear objectives so that performance could be assessed. We identified an initial business case specifying that USPTO should measure success by requiring that "less than X% of the documents require a C\* challenge,"<sup>22</sup> but the percentage was never defined. Moreover, this success measure does not address the system goal of automating the assignment of CPC symbols to applications.

Before awarding the contract for the next-generation CPC AI tool, USPTO developed objectives such as providing an automated system to comprehensively classify patent documents, including a CPC full classification picture with C\* designations, and the following system success measures to track the project's outcome:

- Generating classification information in a way that meets USPTO's routing and search needs
- An overall efficient and cost-effective system
- A system that is scalable to be used by stakeholders throughout the global IP community
- A system with a quality management system sufficient to allow for continuous improvement<sup>23</sup>

<sup>22</sup> A challenge occurs when an examiner determines that a new application has one or more improper C\* classifications. The examiner may submit a classification challenge to have a C\* designation added or removed.

<sup>23</sup> USPTO. March 30, 2022. *Patents Project Charter*.

However, the objectives and system success measures are not specific, clear, or measurable, so USPTO has an insufficient basis to determine whether the tool—on which USPTO has spent, according to its own estimates, approximately \$6.7 million for the legacy and next-generation versions as of July 2024—is meeting its goals.

**Patent search AI tool.** USPTO developed objectives, such as to develop technology to make examiners more efficient by expanding their search tools. However, the objectives are not specific and measurable to ensure that the tool would meet its intended goal. Additionally, USPTO had metrics focused on examiners' use of the tool that were not related to stated goals or objectives.<sup>24</sup> For example, for the goal of expanding examiners' prior art searches, USPTO tracks how many examiners use the AI tool but does not track whether the examiners' searches were improved or expanded.

According to a USPTO official, the rollout of the tool was intended to meet strategic initiatives, and success was broadly defined by implementation alone. In other words, success was measured by whether the tool was released, and USPTO did not establish metrics to determine whether the tool improved examiners' ability to search for prior art.

As a result, approximately \$71.5 million, including contract costs, has been spent on the patent search tool as of July 2024, but USPTO cannot say whether the tool provides its intended benefit. A USPTO official stated that the bureau has started collecting information to assess whether the tool adds business value, including assessing whether it provides more-comprehensive prior art searches that reduce rework, but that work was still in process.

*D. Requirements or technical specifications could not be traced to system goals and objectives*

A key system-level governance practice is to establish and document specifications to ensure that the AI system will meet its intended purpose.<sup>25</sup> This includes defining system technical specifications and requirements that align with system goals, and identifying any system boundaries or limitations.

We reviewed each tool's system design documents, which include architecture diagrams and process workflows, as well as contract documents and other relevant material, and we found that USPTO did not document how the tool's technical specifications align with system goals and objectives. The findings for each tool are described below.

**CPC AI tool.** USPTO incorporated the legacy CPC AI tool's technical specifications into the system design document for USPTO's data repository.<sup>26</sup> However, we found

<sup>24</sup> USPTO had data dashboards for the AI tool; however, there were no stated targets for the data collected.

<sup>25</sup> GAO, *AI Accountability Framework*, 5, 29–30; and Software Engineering Institute, *CMMI*.

<sup>26</sup> A data repository is a centralized location in which an organization can integrate, analyze, and create reports on large amounts and different sources of data.



that USPTO did not clearly document or describe how the specifications in the system design document align with system goals and objectives.

USPTO did not fully implement the legacy tool due to concerns about maintaining it. Instead, as previously noted, USPTO awarded a contract to develop the next-generation version of the tool.<sup>27</sup> USPTO officials stated that the next-generation tool will be cloud based, use updated models, and have a different approach to model training, all of which aim to increase the tool's speed and reliability. However, we found that USPTO still did not describe how the system specifications align with system goals and objectives.

**Patent search AI tool.** Although USPTO documented the patent search AI tool's specifications in a system design document, we could not identify how or from what source (e.g., system purpose, goals, objectives, requirements, or business needs) USPTO derived the tool's specifications. For example, USPTO stated that the tool uses similarity models because those models "address issues of breadth and depth" when conducting similarity searches. However, USPTO did not explain how the chosen similarity models trace to measurable objectives related to "breadth and depth" to ensure that the tool, as designed, will meet its intended purpose.

Without system specifications that derive from system objectives, USPTO cannot be sure that the specifications provide the tool's intended benefit. Until USPTO corrects this practice, the bureau could deploy AI tools that underperform or otherwise do not meet user needs or organizational targets. If USPTO must redesign an AI system or architecture to more fully meet needs, address performance gaps, or respond to undocumented or unclearly linked objectives, it may also encounter cost increases, uncontrolled growth in project scope, and schedule delays.

#### *E. USPTO does not have an AI-specific risk management plan*

Effective governance promotes accountability by establishing processes to identify and mitigate risks throughout the AI system lifecycle. Risk management should distinguish between risks that are inherent to the business or subject matter (e.g., patent classification) and those directly associated with the AI systems.<sup>28</sup> An AI-specific risk management plan is a key component in the responsible development and use of AI systems.

We found that USPTO does not have an organizational AI-specific risk management plan to help leaders identify, analyze, and mitigate risks associated with its AI systems. Such a plan would help ensure effective oversight of USPTO's AI systems and promote trust in its AI use.

One reason USPTO lacks an AI risk management plan is that it relied on compliance with its ATO process to assess AI-related risks. However, this process focuses on ensuring that cybersecurity and privacy risks are at an acceptable level to operate within

---

<sup>27</sup> As of August 2024, the next-generation version had not been released.

<sup>28</sup> GAO, *AI Accountability Framework*, 5, 28–29.

the organization. Risk management procedures conducted as part of the ATO process do not include an assessment of many risks associated specifically with AI systems.

AI systems and the context in which they are used are complex, making it difficult to detect and respond to failures when they occur. AI risks can emerge from the interaction of technical aspects combined with factors related to how the system is used, who operates it, and the context in which it is deployed. These risks make AI a uniquely challenging technology to deploy and use. Therefore, an AI risk management plan is needed to ensure risk mitigation is continuous and timely.

*F. Rapidly changing AI technology and inadequate planning led to the need for stronger key governance practices*

The issues we describe above occurred in part because AI is a relatively new and rapidly changing technology with limited government-wide direction for agencies on AI use. As a result, USPTO did not implement key practices from GAO's accountability framework that are needed to effectively oversee its AI systems.

These issues also occurred in part because USPTO did not adequately plan early in the lifecycles of the CPC and patent search AI tools. For example, before funding a project, the IT need should be documented and analyzed in a business case or product roadmap.<sup>29</sup> A business case helps leaders understand the project's value.<sup>30</sup> However, USPTO's business case for both AI tools lacked information needed by decision-makers, such as an assessment of risks or system success measures. According to USPTO officials, the patent search AI tool was mandated by the then-USPTO Director, and both tools were expected to be implemented or put in operation within short timelines.

In addition, USPTO uses an agile approach to manage its IT projects.<sup>31</sup> With this approach, requirements may not be static, and it is expected that changes will arise over the life of the project. Still, some upfront planning must occur to ensure that risks and issues are mitigated. Adequate planning also helps to provide a baseline to measure progress and quality.

Poor project planning, such as undefined goals and objectives, lack of an AI-specific risk management plan, and lack of stakeholder involvement, increases the risk of project failure. This means time and resources could be wasted on a system that does not meet the needs of the user or the organization, as was the case with the development of the legacy CPC AI tool—a tool that was developed, but never fully put into use.

---

<sup>29</sup> USPTO Office of the Chief Information Officer. May 2019, *Capital Planning and Investment Control Guide*, version 4.4, 8–9; and April 2022, *Capital Planning and Investment Control Guide*, version 5.2, 3, 6.

<sup>30</sup> Project Management Institute. 2017. [A Guide to the Project Management Body of Knowledge](#) (PMBOK Guide), 6th ed., and *Agile Practice Guide*. ProQuest Ebook Central. Accessed July 2024.

<sup>31</sup> In its April 2022 *Capital Planning and Investment Control Guide* (version 5.2), USPTO describes *agile* as an iterative and adaptive approach to product development in which requirements and solutions evolve through continuous engagement between teams and the business.

## Conclusion

Strong governance would ensure that USPTO has processes in place to manage and oversee the implementation of its AI tools and that those tools meet performance requirements. It also ensures that USPTO has processes to identify and mitigate associated risks and reduces the likelihood of unintended consequences, which may include poor system performance and decreased system reliability.

Without effective governance, USPTO has no way to determine whether the tools we reviewed for this report—tools that represent a \$78.2 million investment to date—will achieve their intended goals. It is vital that USPTO implement governance practices to ensure that funds are used efficiently and that these tools will improve patent operations as intended.

The need for effective governance practices extends far beyond the two tools we reviewed. AI is still a developing technology, with many potential uses that have yet to be explored. Adopting key governance practices will promote accountability, transparency, and safety in AI use now and in the future, which will in turn foster public trust in its use. USPTO must improve its governance structure and processes to incorporate the organizational and system-level practices needed to manage, operate, and oversee its use of AI.

## Recommendations

We recommend that the Under Secretary of Commerce for Intellectual Property and USPTO Director direct the Commissioner for Patents and the Chief Information Officer to:

1. Revise policies and practices to ensure that future AI systems develop measurable objectives and system success criteria during project planning and ensure that such objectives trace to system specifications.
2. Develop clear, specific, and measurable objectives and system success criteria for the CPC-related and patent search AI tools.
3. Revise project management policies and practices to encourage increased engagement, as appropriate, with end users and others affected by the system, such as the public, earlier in the project lifecycle.
4. Create an AI-specific risk management plan (or revise existing policies and procedures to reflect the same) that requires a system-specific risk management approach, including identifying, documenting, analyzing, and managing risks.
5. Implement AI-specific risk management practices for the next-generation CPC AI tool and the patent search AI tool.

## Summary of Agency Response and OIG Comments

On March 24, 2025, we received USPTO's response to our draft report. In its response, USPTO concurred with the report's finding and recommendations and described actions it has taken, or will take, to address them. USPTO also provided one technical comment. We considered USPTO's comment and revised the final report as appropriate. USPTO's complete response is included in this report as appendix B.

Overall, we are pleased that USPTO concurred with our recommendations and look forward to reviewing its proposed corrective action plan.

## Appendix A: Objective, Scope, and Methodology

The objective of our audit was to determine whether USPTO had an effective governance structure and processes in place to manage its AI tools.<sup>32</sup>

Our audit scope focused on USPTO's organizational and system-level governance processes and practices for its AI tools; we considered information from 2019, when USPTO implemented its first AI tool, through October 2024.

To accomplish our objective, we did the following:

- We reviewed applicable laws, government frameworks, industry best practices, as well as USPTO policies and procedures relevant to AI oversight, including:
  - Pub. L. No. 116-283, Div. E, National Artificial Intelligence Initiative Act of 2020
  - Executive Order 13960, *Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government*, December 2020
  - GAO, *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities*, GAO-21-519SP, June 2021
  - National Institute of Standards and Technology, *AI Risk Management Framework*, January 2023
  - Software Engineering Institute, *CMMI for Acquisition*, November 2010
- We interviewed officials from USPTO's Offices of the Chief Information Officer, Information Technology for Patents, and Human Resources.
- We gathered data from AI.gov on USPTO AI use cases.
- From a universe of six AI tools, we judgmentally selected two AI systems, the CPC AI and the patent search AI tools, as case studies to evaluate USPTO system-level governance practices.
- We evaluated system-level governance practices for both AI tools related to system goals and objectives, roles and responsibilities, stakeholder involvement, system specifications, and transparency. To evaluate these practices, we reviewed and analyzed contract documents, system design documents, project charter, and other internal documents.
- We sent an online survey to 8,805 patent examiners, primary patent examiners, and supervisory patent examiners in August 2024 to assess whether their perspectives (as end users or system stakeholders) were collected and considered for both AI tools we reviewed, and to gather data on examiners' experiences using the tools. The survey included both close-ended questions and free-text fields. We received responses from

---

<sup>32</sup> During the entrance meeting with USPTO in March 2024, we discussed our plan to use GAO's *AI Accountability Framework* as criteria for this audit. USPTO did not express any concerns about the framework's applicability.

1,822 examiners.<sup>33</sup> The results cannot be generalized to all examiners due to the survey's low response rate.

- We evaluated organizational-level governance practices related to risk management and workforce needs. To do this, we reviewed risk management frameworks, system security plans, IT capacity and funding documents, recruitment incentive plans, and other internal documents.

We assessed the internal controls significant to the audit objective. This included an assessment of the five internal control components: control environment, risk assessment, control activities, information and communication, and monitoring.<sup>34</sup> We also assessed the underlying principles of internal controls. The team identified internal control weaknesses during this audit and proposed recommendations to address them.

In addition, we assessed the reliability of a computer-processed AI system list, created and maintained in Excel. To assess the list's reliability, we (1) obtained and compared the data to the public information on AI.gov, (2) worked with USPTO to address any discrepancies and ensure completeness, (3) reviewed the list for any obvious errors, omissions, or duplications, and (4) traced key fields in the data to source documents. We determined that the data was sufficiently reliable to support the findings and conclusions in this report.

We also assessed the reliability of a computer-processed examiner list maintained in Excel. The data was retrieved from the Patent Application Locating and Monitoring system, effective August 7, 2024. The data's use was limited to identifying recipients for survey distribution. To assess the list's reliability, we (1) performed reasonableness tests of the data, (2) traced the list to National Finance Center system records,<sup>35</sup> (3) discussed the data with knowledgeable officials, and (4) obtained the system's ATO. We determined that the data was sufficiently reliable to support the findings and conclusions in this report.

We conducted our audit from March 2024 through February 2025 under the authority of the Inspector General Act of 1978, as amended (5 U.S.C. §§ 401–424), and Department Organization Order 10-13, as amended October 21, 2020.

We conducted this performance audit in accordance with generally accepted government auditing standards. Those standards require that we plan and perform the audit to obtain sufficient, appropriate evidence to provide a reasonable basis for our finding and conclusion based on our audit objective. We believe that the evidence obtained provides a reasonable basis for our finding and conclusion based on our audit objective.

---

<sup>33</sup> We received 1,829 completed surveys but removed the surveys of seven respondents who were not patent examiners, primary patent examiners, or supervisory patent examiners.

<sup>34</sup> GAO. September 2014. *Standards for Internal Control in the Federal Government*, [GAO-14-704G](#), 7–8. Accessed February 2024.

<sup>35</sup> The National Finance Center provides human resource services for the Department.



## Appendix B: Agency Response

USPTO's response to our draft report begins on the next page.




## UNITED STATES PATENT AND TRADEMARK OFFICE

UNDER SECRETARY OF COMMERCE FOR INTELLECTUAL PROPERTY AND  
DIRECTOR OF THE UNITED STATES PATENT AND TRADEMARK OFFICE

March 24, 2025

MEMORANDUM FOR: Kevin D. Ryan  
Acting Assistant Inspector General for Audit and Evaluation

FROM: Coke Morgan Stewart   
Acting Under Secretary of Commerce for Intellectual Property and  
Acting Director of the United States Patent and Trademark Office

SUBJECT: Response to Draft Report, "USPTO Should Improve Governance  
to Promote Effective Oversight of Its Artificial Intelligence Tools"

### Executive Summary

We appreciate the effort you and your staff made in reviewing the United States Patent and Trademark Office's (USPTO) oversight of its artificial intelligence (AI) tools. Your office found that the USPTO should strengthen key organizational and system-level governance practices. We believe that the USPTO's AI-related practices have been and continue to be among the most robust in the federal government. However, the USPTO recognizes that additional measures, including those detailed in your office's recommendations and in further guidance to be set forth by President Trump's Administration, can bolster the USPTO's readiness to adopt and govern this rapidly evolving technology. While the USPTO has taken many positive steps to promote safe and efficient use of AI as a critical technology to support the agency's mission, we acknowledge that more can be done to improve our AI governance and practices.

The USPTO views AI as a critical technology to further the agency's mission and enhance the efficiency of agency operations. The agency is proud to be at the forefront of federal AI innovation, having implemented AI capabilities for critical agency needs since 2019. The USPTO's portfolio of AI tools has expanded considerably since then, with eight use cases meeting the criteria for disclosure in the 2024 federal AI Inventory, and many more use cases in exploratory stages across the USPTO's business units. With thousands of active users, this portfolio of tools is delivering, and will continue to deliver, value to USPTO personnel and public stakeholders each and every day.

In addition to leaning forward into AI innovation across the enterprise, the USPTO has created an agency AI Innovation and Governance Council, is continuing to develop internal policies and procedures for the development and risk management of AI, has solicited employee feedback for AI-related information technology (IT) applications through dedicated channels, and has paved the way for agency use of frontier technologies from leading U.S. AI research labs. The agency has been nationally recognized for its leadership in the AI domain, with agency personnel having

received numerous awards, including the Department of Commerce Gold Medal, the Presidential Rank Award, the Samuel J. Heyman Service to America Medal, the WashingtonExec Pinnacle Award for Artificial Intelligence, the FutureEdge 50 Awards, and the CIO 100 Award.

The USPTO is committed to continued innovation and sound governance in the agency's use of AI. In particular, we will align our practices with direction and guidance on AI set forth by the Trump Administration. The Administration has already acted swiftly to remove barriers to U.S. leadership in AI, and has indicated an intent to issue new and revised guidance on various AI-related matters. As these documents are issued, the USPTO will take appropriate measures to ensure that agency AI practices are consistent with the directives and recommendations therein.

The USPTO concurs with OIG's recommendations as explained in our detailed responses below. Our technical comments are also attached.

### **OIG Recommendations**

OIG recommendation that the Under Secretary of Commerce for Intellectual Property and Director of the United States Patent and Trademark Office direct the Commissioner for Patents and the Chief Information Officer to take the following actions:

- 1. Revise policies and practices to ensure that future AI systems develop measurable objectives and system success criteria during project planning and ensure that such objectives trace to system specifications.*

#### **USPTO Response:**

The USPTO concurs with this recommendation. Development on the USPTO's early AI tools commenced prior to the completion of the USPTO's transition to a product-based organization using the Agile framework for software development. Since then, the USPTO has continued to refine and mature its Agile software development processes. In addition, the USPTO updated its Capital Planning and Investment Control (CPIC) Guide to support Agile Development, Security, and Operations (DevSecOps) processes. The latest update to the CPIC Guide was made on November 6, 2024, to reflect further refinements with USPTO IT Planning and Quarterly Review processes and more clearly aligning IT work with agency strategic objectives and business priorities, among other enhancements.

The USPTO acknowledges that we can continue to improve how we establish, monitor, and refine measurable objectives and success criteria, including with respect to AI systems. The USPTO is committed to following its DevSecOps process to ensure that objectives and success criteria are clearly captured in Epics, which define the IT work and are reviewed and verified as "done" by the Product Owner. This Agile process takes the place of the Requirements Traceability Matrix used in traditional waterfall style projects.

2. *Develop clear, specific, and measurable objectives and system success criteria for the CPC-related and patent search AI tools.*

USPTO Response:

The USPTO concurs that the agency can improve our existing processes for development of objectives and system success criteria.

For CPC-related AI tools, the USPTO has been focused on refining metrics for success in alignment with particular use cases, with a primary use case now focused on ensuring high-quality CPC classification that supports accurate document retrieval during prior art search, which benefits both examiners and the public. As part of this use case, a team of subject matter experts at the USPTO has planned to incorporate user feedback on classification search strategies that can ultimately help develop appropriate measurable objectives that relate to the intended use. This includes internal feedback through focus sessions and surveys as well as public feedback through appropriate forums. The USPTO has also explored potential A/B preference testing to understand auto class performance compared to contractor performance with respect to assessing the suitability of the auto classification tool for prior art search against appropriate measurable objectives. This would allow the USPTO to more precisely determine the impact of utilizing AI-based classification to replace the current non-AI (contractor)-based classification solution. The teams focused on this effort will continue efforts to solicit feedback and set appropriate objectives that are clear, specific, and measurable.

For Patent Search AI Tools, the system was developed with metrics included at the time of development that focused on how examiners leveraged the tools in order to make patentability decisions. The USPTO is actively working to ensure current and future systems that are developed include these metrics, as well as ever-evolving best practices, with a focus on system success criteria. This will continue to ensure quality tools from early in the development process. The development teams, for public end users as well as examiners, will continue to work with data experts to ensure that metrics derived from the system are consistent with current best data practices and methods. In this way, teams will maintain the capability to make the best development decisions throughout the development process.

3. *Revise project management policies and practices to encourage increased engagement, as appropriate, with end users and others affected by the system, such as the public, earlier in the project lifecycle.*

USPTO Response:

The USPTO concurs that the agency can improve on current policies and practices in order to increase existing engagement with system end users. The agency has already taken active steps toward engaging users earlier in development in other projects that can be used as a blueprint to continue to expand user engagement in future early-stage AI efforts.

For CPC-related AI tools, the USPTO has been focused on how to better incorporate feedback from internal users as well as external users throughout the lifecycle of development in an effort



to identify appropriate measurable objectives. For internal users, this includes gathering information from users early in the process through focus sessions and/or surveys or other appropriate avenues. Similarly, the USPTO intends to continue engaging the public on CPC auto-classification through venues such as the Patent Public Advisory Committees (PPAC) as well as other avenues. It is anticipated that this engagement with users will provide a framework for future discussions and data results that can ultimately demonstrate the effectiveness of the CPC-related AI tools and ultimately boost confidence in the system.

For Patent Search AI Tools, the USPTO will continue to engage internal end users at multiple levels from the beginning of development. This includes having product teams with embedded examiners as team members, including quarterly and biweekly collaborative planning sessions. To further improve engagement with internal end users, the USPTO will continue our robust practice of surveying end users, ensuring feedback is received prior to development beginning. To enhance our engagement with public end users, surveys have been added to our externally facing search website so that feedback can be gathered prior to development starting. In addition, the USPTO has developed a process to gather feedback from the public by using the Patent and Trademark Resource Centers and the Public Search Facility to test newly developed features prior to deployment. This practice will also be used for future Patent Search AI Tool features in order to further engage end users. The USPTO also intends to continue engaging the public on Patent Search AI Tools through the PPAC and through updates on the USPTO's Artificial Intelligence webpage.

Finally, the USPTO has created dedicated channels for early-stage stakeholder engagement on new AI initiatives. In late 2024, the USPTO launched an agency-wide AI Ideation Community through which employees can submit AI-related ideas and suggestions, as well as engage in discussion on the broader portfolio of such ideas. During a month-long pilot period, the USPTO received over 100 submissions pertaining both to specific USPTO business units and to the enterprise as a whole. These employee submissions are being methodically reviewed by technical and business leadership at the agency, and they serve as yet another valuable signal of USPTO end user needs.

4. *Create an AI-specific risk management plan (or revise existing policies and procedures to reflect the same) that requires a system-specific risk management approach, including identifying, documenting, analyzing, and managing risks.*

#### USPTO Response:

The USPTO concurs with the recommendation that it is advisable to adopt AI-specific risk management practices. As part of the agency's ongoing efforts to promote AI accountability and the responsible use of AI systems, the agency has established governance structures and processes to manage, operate, and oversee the implementation of AI systems.<sup>1</sup> In January 2024,

---

<sup>1</sup> The GAO report, *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities*, highlights the importance of setting up an AI governance structure at both the organizational and system levels, with establishing an AI-specific risk management plan as one of the many subcomponent practices recommended for agencies to take when creating such a governance structure. See GAO, *Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities*, GAO-21-519SP, June 2021.

the agency stood up the USPTO AI Governance Council (AIGC), supported by working group efforts such as the AI and Emerging Technologies (AI/ET) Working Group. The AIGC brings together internal stakeholders across the agency from all business units to discuss current agency AI efforts, to include external stakeholder impact and overall agency risk in using AI systems. Moreover, the AIGC is working toward developing an AI user policy to assist in managing risk across the agency for the internal use of AI systems. These policy efforts began in October 2024 in consideration of the publication of the NIST AI Risk Management Framework and Office of Management and Budget (OMB) guidance on the governance and risk management of AI. While the agency is still in the early stages of building its AI governance structure and processes, the agency has nonetheless started work on creating an AI-specific risk management framework. In May 2024, the AIGC and AI/ET Working Group began work to develop and scope an agency-specific AI and intellectual property risk management framework that is still in progress.

In addition to the USPTO's AI-specific accountability and risk management efforts, the agency also engages in significant IT and enterprise-wide risk management practices that account for many AI system risks, even in the absence of an adopted AI-specific risk management framework. Many of those practices and procedures are detailed below.

For example, the agency follows the guidance outlined in the NIST Special Publication (SP) series as one of many methods of overseeing its IT risk management, as mandated by OMB Circular A-130, to include complying with and integrating many aspects of its voluntary cybersecurity framework, privacy framework, and risk management framework. The methods and concepts in the NIST SPs are universal and can be applied to any IT application, to include AI systems. For example, NIST publishes two methods that work in concert and that the agency follows for managing risk of its IT applications: the NIST risk management framework (SP 800-37) and the appropriate list of security and privacy controls that go along with it (SP 800-53 series).

In July 2024, NIST also published its AI risk management framework, which the agency has evaluated and started to incorporate aspects of into its IT procurement risk management. For example, the agency's IT Acquisition Checklist, provided by the Department of Commerce, which is used to check for multiple areas of IT risk, was updated in August 2024 to include checking for AI systems or applications. This step created an additional risk management practice for the agency to consider an AI system's security, privacy, safety, and rights-impacting potential prior to solicitation or acquiring the AI application. Through the AIGC, starting in October 2024, the agency began developing a more robust process for ensuring proper stakeholders were involved in this IT acquisition risk management process for AI systems to better consider the potential safety and rights-impacting implications, along with the security and privacy implications that are already baked into the process.

Many AI-specific risks intersect with areas such as privacy and cybersecurity, and the USPTO has robust practices in place to ensure that AI systems meet applicable standards and requirements with respect to those areas. For instance, through the use of the IT Acquisition Checklist process, as outlined above, the privacy team is notified about acquisitions (including those of AI systems or capabilities) to conduct the necessary privacy checks prior to the AI collecting, using, storing, or disseminating of Personally Identifiable Information. Additionally,



the IT Checklist will ensure the appropriate privacy clauses are included within the contract language to hold the contractor accountable and ensure the privacy rights of the public and USPTO employees. To ensure the integrity of all AI systems acquired, the USPTO, through the IT Checklist, does a review to determine if a supply chain risk assessment needs to be done. A supply chain risk assessment, if required, would be a collaboration with the Department of Commerce Supply Chain Risk Management Program to evaluate the risk of the purchase and enable leadership to make an informed risk-based decision about the acquisition.

Moreover, the agency's implementation of the NIST Risk Management Framework system security and privacy assessment and authorization activities provides a structured process for managing security and privacy risk. The outcome of this process provides USPTO senior leaders and executives with the necessary information to make informed risk management and system authorization decisions about the systems supporting their missions and business functions throughout the lifecycle of those systems. Additionally, continuous monitoring activities, to include regular vulnerability and compliance scanning and annual assessments, provide situational awareness of the security and privacy posture of USPTO systems on an ongoing basis to assist with risk management efforts for all IT. Overall, system security and privacy assessment and authorization activities facilitate the connection of risk management processes at the USPTO (system, mission/business unit, and organization).

As part of the USPTO's implementation of the NIST Risk Management Framework, and in compliance with requirements established by the Federal Information Security Modernization Act of 2014 (FISMA 2014), the USPTO determines the levels of information security and privacy appropriate to protect USPTO information and information systems, to include those systems with AI components and capabilities, and performs initial and periodic testing and evaluation of information security and privacy controls and techniques to ensure they are effectively implemented. The USPTO also adheres to the Federal Risk and Authorization Management Program (FedRAMP) when authorizing the use of commercial cloud service offerings by the USPTO. FedRAMP is conducting an analysis to determine the impact to security and privacy controls with the introduction of AI systems and components within FedRAMP authorization boundaries, and upon publication of related standards and guidance from FedRAMP, the USPTO will incorporate the necessary changes into applicable USPTO policies and procedures.

The USPTO also implements requirements established by OMB Memorandum 22-18, Enhancing the Security of the Software Supply Chain through Secure Software Development Practices, dated September 14, 2022; OMB Memorandum M-23-16, Update to Memorandum M-22-18, Enhancing the Security of the Software Supply Chain through Secure Software Development Practices, dated June 9, 2023; and Department of Commerce Procurement Memorandum 2024-08 (Revised), dated September 9, 2024, to help obtain the necessary information from software producers in order to make risk-based decisions about procuring software.

In addition, the USPTO carefully evaluates its AI use cases with respect to AI-specific considerations. This applies to both AI capabilities built internally and those acquired from outside the agency (the IT Checklist includes a required review by AI subject matter experts for any acquisition involving AI/machine learning capabilities). As of 2024, all usage of new

frontier AI capabilities, including in USPTO commercial and government cloud environments, is off-by-default. Proposed AI capabilities are only enabled on a case-by-case basis after consultation between the requesting team and personnel within the USPTO's Office of the Chief Information Officer. Such consultation involves careful consideration of the specific details of the AI use case being contemplated and often results in the formulation of use case-specific requirements as a condition of development and/or deployment. Particularly with frontier technologies such as Generative AI, the USPTO has taken a steady, measured approach to studying such technologies in controlled lab/sandbox cloud environments in order to determine effective and appropriate uses of these capabilities.

Furthermore, several efforts are currently underway to mature the USPTO Enterprise Risk Management (ERM) program with an FY 2025 focus on marrying organizational performance management and ERM. USPTO performance measures are an important tool for monitoring risk. The USPTO has AI-related risk statements in its risk register. These risk statements were all scored as low severity (likelihood and impact) by the USPTO Risk Management Council in December 2024. The USPTO is in the progress of identifying where the USPTO can use AI to mitigate high-severity programmatic risks such as patent pendency. The risk profile created for each high-severity risk will detail the risk response initiatives (such as AI development, testing, and adoption in various areas), owners for each response, and measures used to monitor the success of the response.

The USPTO is now awaiting further Administration guidance before determining what additional steps are to be taken with respect to AI risk management. Prior to January 2025, the federal government's approach to AI risk management had been set forth in Executive Order 14110 and succeeding documents such as OMB M-24-10 and M-24-18, as well as the voluntary NIST AI RMF 1.0. The USPTO had taken steps to begin aligning its practices with these documents, as noted above. On January 20, 2025, President Trump identified Executive Order 14110 as a harmful Executive Order to be immediately revoked.<sup>2</sup> Subsequently, on January 23, 2025, President Trump directed the revocation of "certain existing AI policies and directives that act as barriers to American AI innovation," and as part of this directive called for the revision of OMB M-24-10 and M-24-18 to ensure consistency with Administration policy.<sup>3</sup> Accordingly, the USPTO will align our AI practices with these forthcoming documents and other Administration guidance such as the Action Plan being developed pursuant to Executive Order 14179.

*5. Implement AI-specific risk management practices for the next-generation CPC AI tool and the patent search AI tool.*

The USPTO concurs with this recommendation and notes several AI-specific technical practices, such as the profiling of individual operations in Patent Search AI model binaries, that have already been undertaken. As part of the USPTO's implementation of the NIST Risk Management Framework, and in compliance with requirements established by the FISMA 2014, in addition to AI-specific risk management practices, the USPTO will perform initial and periodic testing and

---

<sup>2</sup> Exec. Order No. 14148, Initial Rescissions of Harmful Executive Orders and Actions, 90 FR 8237 (Jan. 20, 2025).

<sup>3</sup> Exec. Order No. 14179, Removing Barriers to American Leadership in Artificial Intelligence, 90 FR 8741 (Jan. 23, 2025).

evaluation of information security and privacy controls and techniques to ensure they are effectively implemented for the CPC AI tool and Patent Search AI tool. The USPTO will align its AI practices, including with respect to the next-generation CPC AI tool and the Patent Search AI tool, with the forthcoming federal AI documents (described in response to Recommendation 4) as they are issued.

# REPORT

# FRAUD & WASTE ABUSE



## HOTLINE



Department of Commerce

**Office of Inspector General Hotline**

[www.oig.doc.gov](http://www.oig.doc.gov) | 800-424-5197