

The Intersection of GenAI and PDF Editing

White Paper

TUNGSTEN
AUTOMATION

The Intersection of GenAI and PDF Editing

Exploring the Power and Potential of GenAI and PDF Editing

A technological revolution is underway, and the possibilities seem endless and exciting. With the rise of generative artificial intelligence, we're gaining innovative ways to interact with computers and data. This technology has many potential applications but could be valuable in document-based workflows. Some everyday tasks could become substantially simpler through GenAI and a PDF editor with integrated functionality.

Users now have new ways to accomplish tasks and apply creative thinking to problem-solving. Though it might not yet be the famous Star Trek computer, we're closer than ever, thanks to how GenAI naturally processes user queries to perform simple and complex tasks. **Tungsten Power PDF** now includes the option to bring your own GenAI token to access your preferred model's power.

The new functionality opens many doors and creates exciting opportunities. This white paper provides some background on this technology and its place in document editing. We'll highlight Power PDF's integration capabilities and explore how it can transform PDF editing workflows.

A Quick Overview of Generative AI Technology

Understanding the complex technologies behind today's generative AI models could fill volumes. Many have written lengthy books and are still analyzing the complex details. Understanding what makes this technology work highlights its potential and manages expectations.

Several elements contribute to the creation of a generative AI model:

- Neural networks replicate human modes of cognition to process and analyze data through **machine learning**.
- Deep learning uses extensive data sets to create the frameworks necessary for AI training and development.
- Natural language processing (NLP) enables models to interpret and respond in natural language as a human would answer questions.



Generative AI can take many forms, but in each case, it accepts an input and uses its training to generate an output. In the case of large language models (LLMs) such as ChatGPT, that usually means text input and output. We've also seen the rise of other forms of generative AI, which can interpret text or images to output pictures, audio, and video. It's proving to be **a helpful tool for many types of workflow automation** today.

Generative AI often creates new or novel responses, but it is not a form of *true* artificial intelligence. All GenAI relies on the data used to train its models, and GenAI models can still produce mistakes or falsehoods in some circumstances.

Even so, today's models can increasingly understand complex instructions to work alongside humans and accomplish many tasks better. The ability to phrase requests in natural language and to receive prompt results opens the door to straightforward, faster work.

Let's consider how such technology can integrate with **standard document workflows** by spotlighting our solution, Tungsten Power PDF.

Tungsten Power PDF: A Solution Built for Business

PDFs are among the most common file types used in business for documents. The reasoning is easy to understand. These files are often lightweight and highly portable, and their contents are displayed consistently regardless of the computer system. Even mobile devices are better equipped to handle PDFs today.

The right software to create and interact with these files is essential. Power PDF Advanced features all the abilities that business users need in a convenient and familiar ribbon-style interface. Licensing is flexible and includes a one-time payment option for a permanent license to the current version. Adoption is quick, and learning easy—your staff can use Power PDF to:

- Create full-text PDF documents using industry-leading optical character recognition that generates better outcomes and requires fewer corrections.
- Convert documents to and from the PDF format, such as Excel to PDF, PDF to Word and many more combinations.
- Enforce advanced security features, including password-based encryption and document access controls.
- Integrate eSignature functionality to create and sign or send documents for signature via email for rapid, secure authorizations.
- Collaborate in real-time with other business users in your organization to streamline document creation, editing and finalization.

As we continue to develop and improve Power PDF, we always look for new ways to open the door to better business functionality. With the rise of GenAI, the need to integrate an opportunity to leverage such models became clear. Power PDF Advanced now has innovative capabilities in the AI space.

Integrating Power PDF With GenAI Capabilities

The decision to add GenAI capabilities to Power PDF came only after carefully considering how to achieve a safe, reliable implementation. Many organizations rightfully have concerns about software with built-in AI capabilities and the potential risks to private, proprietary and sensitive data. Those risks include the potential for leaks and concerns about proprietary data making its way into LLM training data.

Power PDF will not have its own generative AI model within the software. Instead, it can integrate with a wide range of third-party LLM models. Power PDF allows you to bring your own token and integrate the ability to send queries securely to your preferred model. You remain in control of your data in line with the Terms of Service governing the model you choose.

Some key features of this integration include:

- Capability for defining URL endpoints for interacting with LLMs and receiving output responses.
- Options for defining specific authentication methods for controlling access and only submitting valid queries.
- Mass configuration capabilities via group policy options and a specific customization kit. Effortlessly set up multiple users across your organization and maintain adequate controls.

The only limitations to what you can achieve with your queries are those imposed by the model you choose. The easy-to-configure integration in Power PDF ensures you can swap to another model when required.

How Can GenAI Reshape Document Editing Workflows?

What can you achieve with that functionality once you connect generative AI and a PDF editor? This integration opens the door to making complex tasks more accessible and to building human-in-the-loop systems that keep automated tasks accountable for reliability. However, we see several integral use cases at the intersection of PDF editing and GenAI.

Ask questions and get answers with AI document querying

The first considerable usage we anticipate as a source of value is the capability to conduct a live Q&A with an LLM about the contents of a document. Depending on its contents, a PDF may be dozens or even hundreds of pages long. Those contents could also be complex; besides text, they may include images, charts, graphs and other representations of information. It's not always easy to find what you want, even if you generally know what you want to locate.

With an integrated LLM, you can ask the model questions and let it use the document to generate its answers. This feature can save considerable time, especially when working on expansive documents.

Some examples of how you might query an LLM about a document:

- Identifying the existence of sensitive or personally identifiable information in the document.
- Isolating specific subjects within a large document, such as, “Does this document contain any information related to future pricing trends?”
- Explaining concepts or ideas laid out in a document.
- Defining unfamiliar terms in a document.
- Calculating answers based on data in tables.
- Check whether specific formatting has been adequately set up or verify that the document contains particular information.

Ultimately, the only limit on the questions you could ask about a document are those imposed by your model or your imagination. You could issue a query about an entire document or simply a selection of text. **Performing a wide range of tasks with a PDF becomes much simpler when you can instruct an LLM** in natural terms about what you want to accomplish.

This Q&A capability also ties in with another primary use case.

Summarize and distill complex documents

The ability to **summarize a PDF with AI** has an immense potential value. Knowledge workers often have many tasks on their plate most days. Finance, law, and insurance employees will likely encounter information-dense documents during the day. Consider a customer support agent for an insurance company handling queries about an individual’s insurance policy.

Working independently, the agent might take several minutes or even require a callback to provide the answer. With GenAI, you could ask a model to summarize the policy as it relates to the elements of the customer’s query. The same tactic could work for a sales agent. An LLM can quickly provide all the critical facts about a policy and talking points to support closing a sale.

This example is brief. AI summarization of documents has many potential applications. Consider a lawyer trying to understand whether an item of case law applies to a situation they’re handling for a client. An AI could quickly summarize the key findings in a court opinion or the other cases cited in the decision. That shortcut can slash the time it takes for the lawyer to develop a case or reach conclusions about questions central to a suit.

With a PDF summarizer, extracting the most critical information from a document takes seconds, not minutes or hours. Streamline research, save time and money, and equip teams with better access to the knowledge stored across all your documents.

Accurately translate text with speed

AI document translation is perhaps the most powerful and robust use case for GenAI within a PDF environment.

It can take, on average, at least one to two days to translate a document of 1,000 words. If you add a review by a second translator, the duration can extend to three days. The longer the document, the more time it will take to produce a translation. Turnaround times for long or complex documents can sometimes stretch to a week or more. These speeds can seem unacceptably sluggish in today’s fast-paced and globally interconnected economy.

GenAI has already triggered a rapid transformation across the translation industry. With **integration directly into Power PDF**, you can provide a token for an LLM purpose-trained to deliver more accurate document translations. The results of such translation tools speak for themselves.



Specify your output language, select a portion of the PDF for translation, or submit the entire document. Where human translators may need days of intensive work, an LLM can produce a translated document in a fraction of that time. With Power PDF, you can modify or create a new document based on this translation.

It's important to note that this application perfectly demonstrates the need for human-in-the-loop systems. LLMs are highly advanced today and continuing to improve, but it is essential to remember that they are not infallible. Mistakes can occur—and language is considerably nuanced. A capable LLM might not detect every idiom or appropriately translate every metaphor.

GenAI doesn't mean replacing your translation team. Instead, you can provide AI-translated documents to translators for review, editing, and finalization. By shaving time off the front end of the process, you can instead focus a team's efforts on making the best final product possible.

The Benefits and Advantages Afforded by GenAI and PDF Integration

Each use case described above has its apparent advantages. A Q&A feature makes finding relevant data faster. Summarization makes knowledge more accessible and enables preventing bottlenecks in everyday workflows. Translation accelerates turnaround times and may improve final document quality with human assistance. However, there are broader benefits organizations may see as a result of leveraging these use cases.

Other advantages include:

- Cost savings. From less complex document workflows and support processes to a reduced reliance on expensive translators, GenAI can contribute to paying for itself. It could even eventually translate into more visible bottom-line impacts.
- Fewer errors in standard processes. With appropriate oversight of GenAI usage, teams can output higher-quality work.
- Faster and more responsive customer service. Both Q&A and summarization capabilities make core information more accessible to support providers. Better, faster translation can make your help materials available to more regions.
- Improved employee morale. With less time spent obtaining actionable information from documents, teams can achieve more and at higher quality while exercising their primary skills and creativity.

As you seek to unlock these benefits, it's important to remember to follow some accepted best practices during implementation.



Some Best Practices for Safe, Secure Work

There's no doubt that GenAI is a powerful and even transformative new tool. Using it safely and effectively in your organization is the key to unlocking its value. That's true for all use cases—not only when integrated into Power PDF. So, what are some of the practices you should keep in mind?

- Establish clear rules and guardrails regarding data privacy. Understand the Terms of Service provided by your LLM. Ensure that users only apply GenAI capabilities to the appropriate document types.
- Thoroughly train staff. Educate users on the responsible use of LLMs and what they can achieve when working on PDFs using this capability. Reinforce rules about security and privacy during training.
- Keep humans in the loop. Practice *trust but verify* with LLM outputs in high-value or high-risk situations. Reviewing translations, especially of longer documents, is essential for maintaining quality control.
- Remain vigilant about other technological best practices by maintaining a secure network and encrypting sensitive data where necessary.

By taking precautions and exercising care, you can tap into the incredible potential of current and future LLMs and remake your interaction with PDF files.

Explore What Your Teams Could Achieve Today

Ask for the location of specific types of data. Separate the main points from a lengthy document within moments. Translate crucial work documents to make them more accessible to audiences or colleagues worldwide. When you can integrate your preferred generative AI model into a proven solution for PDF editing, you begin to unlock the full potential of what's possible with documents today.

We're excited to add bring-your-own-token support to Power PDF, which will open the door to new and creative ways to manage your document processes. Combined with its advanced security features, collaborative capabilities and other valuable functionality, this capacity makes it the best choice for business users today.

GenAI and PDF editing go hand in hand. Consider how your teams could accomplish more daily with a PDF editing solution they can query naturally. [Learn more about adding Power PDF and these capabilities](#) to your organization today.



About Tungsten Automation

Tungsten Automation, formerly Kofax, is the global leader in intelligent automation solutions with a trusted legacy of nearly 40 years, with a team of 2,000+ employees in 40 countries, serving 25,000+ global customers. Our dedication to innovation and customer success has earned us industry recognition, including being named a leader in Intelligent Automation (IA), Intelligent Document Processing (IDP) and Process Orchestration by top analysts. We are trusted to help businesses gain unprecedented efficiencies and reduce costs through AI-powered workflow automations that propel their businesses into the future.

To learn more, visit www.TungstenAutomation.com

www.TungstenAutomation.com

