

WHITEPAPER

FilevineAI

Introduction

Why Invest in FilevineAI?

There's a new AI tool advertised every day. So why does FilevineAI deserve your attention?

Because we are uniquely great at legal AI. In fact, we are better at this than any other company.

Why? **Because we have the data.**

From the first days of Filevine, we have been obsessed with gathering everything from each legal matter into one location. We have poured all our energy into building a system good enough to house this much data.

This goes beyond documents. Every single deadline, text, note, task, fax, and email associated with a legal matter sits in Filevine. There is not a company in the world that has the kind of unstructured legal data Filevine has.

And something extremely precious rests within that data: reasoning. Each line shows legal professionals thinking through a situation and figuring out the best next steps to take. That data exists as natural language — and with the latest AI technology, we can finally surface that reasoning to help legal professionals take action.

This paper will give you all the technical details of how we do this. But the foundation rests here: Filevine is the best company with the best opportunity to make the biggest difference in the history of legal tech — ever.

We hope you make history with us.

Best,

Ryan Anderson
CEO and Co-founder of Filevine



The Filevine Advantage

Filevine's AI-powered platform is transforming the practice of law — not by replacing human expertise, but by amplifying it, creating a future where legal professionals can focus on what matters most: delivering justice and serving their clients.

Filevine's deployment of AI is unique in a number of ways:

- **It operates directly within your system of record.** While our competitors charge you per file and require data export, Filevine works with your data where it lives, providing real-time insights across all your documents, notes, and project data.
- **It works across your entire dataset.** Unlike point solutions that handle single tasks, Filevine's AI can engage with the complex relationships between documents, notes, calendar entries, and project data to provide more intelligent, context-aware results.
- **It is trained specifically on legal and medical expertise.** Generic AI tools are trained on vast amounts of data across all subject matters — some of it inaccurate or unprofessional. FilevineAI's proprietary embeddings model is trained specifically on trusted legal and medical sources, greatly reducing the risks of errors and hallucinations.
- **It protects your data.** Because you're operating within one system, you can benefit from Filevine's industry-leading security measures and agreements, without the risk associated with less vetted software.
- **It gives you answers in real time.** While other tools require you to wait days for processing, Filevine can give you answers immediately.

These benefits combined have made Filevine a leader in the field of legal AI.



Understanding Filevine AI Technology

Filevine uses a proprietary embeddings model trained specifically for legal and medical contexts.

This custom embedding model enables you to use sophisticated searching beyond the standard simple keyword matching. For instance: when you search for “cancer” it doesn’t just find that word. It also incorporates related terms like “leukemia,” “melanoma,” and “glioblastoma.” By using this model we also allow your firm to handle massive documents efficiently in cases, without the typical fatigue that hits a legal worker after a long day of reading.

Search Term: “Cancer”

Traditional Search	Vector Based Search
Pre- Cancer	Cancer
Pre- Cancerous	Cancers
Cancer	Cancerous
Cancers	Leukemia
Cancerous	Melanoma
	Glioblastoma

Our Data Science, Product, Security, and Engineering teams focus on daily updates to improve accuracy and expand functionality, all while prioritizing integrating these solutions directly into your daily workflows within Filevine.



How Does Filevine Minimize Bias in Data?


Filevine follows established Data Science standards and practices to determine the sample sizes required for each training or validation set created. It leverages statistical methods to ensure that our data set is not biased or incomplete in representation. This governance policy is managed and audited by the VP of Data Science for compliance with information Security and Internal Audit oversight.

Where Does Filevine Get Its Training Data?

Filevine currently has two data sources in which we source data for model training.

The first is **publicly available data**, which Filevine crawls and retrieves from public domains (such as open-source legal sites). These domains are vetted for biases, accuracy, and purpose. Filevine uses statistics, sampling, and feature creation to ensure that data sources are representative and complete. Filevine also ethically sources data from public arenas — we do not use any third party data source to train (such as The Wall Street Journal). Filevine uses this data to find relationships and patterns and provide a legal and medical knowledge set for our models.

The Filevine Difference: Data and Search in one machine language



The second source is **anonymized Filevine user data**, where Filevine removes all personally identifiable information and metrics. These data sets are curated for specific models like our LeadAI model, in which Filevine recommends the likelihood of a lead to convert to client status through the LeadDocket product. Users can opt out of this, in which case Filevine will not include their data in these anonymized datasets. This is further ensured by filter measures built into the code that ensure no information from opt-out users is included in our training data.

Once the models are trained with anonymized data, Filevine purges our production systems of the training data so that only the trained model runs in Production/Security and DevOps-provisioned environments that handle user data. Filevine also trains the model in a production environment so that limited contributors can access it and to ensure non-user data does not mix with user data. All processes and steps of the Data Science team fulfill the requirements of our SOC 2 Type II and HIPAA compliance.

Which AI Models Are Used?

Filevine selects the model that will perform best for the task at hand. We use a mixture of traditional, natural language processing, and generative models, depending on the task that Filevine is solving. Filevine tests which methodology will be the most appropriate given several factors, including, but not limited to: data requirements due to geographic or high-security requirements, speed, cost, and, most importantly, accuracy. Filevine uses standard precision and recall tests and guardrail evaluators to determine the accuracy of our models' decisions.

Filevine monitors the performance of our models by continuing to monitor and provide new validation set metrics (precision/recall).

For third-party LLMs, Filevine has contracts in place for 0-day retention and 0 storage of any data being sent to the third-party LLM.

At this time, Filevine uses validation sets to ensure we have completeness and accuracy in AI decisions. Filevine has listed Third Party Data Subprocessors here in the [DPA](#).



How Private Is My Data?

For any model trained on anonymized data, FilevineAI interacts with specific features from user data that are anonymized by the user and not shared with other organizations. While certain data points like geographic location or severity of injury may be presented, they do not contain client names or other sensitive data.

For Filevine's generative models, the models do interact with the full document context which contains personally identifiable information (PII). Filevine generative models have a 0-day retention period, meaning after the data is processed for that specific limited session ID, it is purged from the system. This helps ensure there are no memory or usage leaks across users due to code and session limitations. The model does not manage any authentication, authorization, or session completion tasks — Filevine's code ensures there can be no loophole in the system to bypass these requirements.

Is Filevine's AI Reliable?

Filevine leverages generative AI models that use Filevine user's data in context to provide insights, extract data, and generate narrative. It protects the reliability of its responses through many interlocking processes, including:

- **Protecting the models from unvetted data.**

Filevine does not search the web or use third party opinions to inform the prompts, the only context informing the prompt is your secure Filevine data. Its models are never injected with public or unknown data, and never run automatically without human checkpoints in the loop. When the user inputs information and documents manually through the collection and custom field sections, that data is considered legitimate and safe. Filevine also uses the output from our AIFields extractions (like Medical Chronologies) to be used as input context for our DemandsAI or DraftingAI products. In every case, the data comes from a trusted source.



- **Limiting sessions to each user or organization**

We limit sessions to each user/organization so there is no memory or data leakage across users or organizations. This is all defined in our code base; no unauthorized user will ever see another user's generative output.

- **Leveraging Retrieval-Augmented Generation (RAG) systems.**

Filevine also leverages enhanced retrieval systems that inject specific and secure context into our generative prompts. RAG is managed by unique organization, user, and project level authorization requirements to ensure there is no bleed-over of data between users.

- **Utilizing secure prompt management systems.**

Filevine uses a prompt management system (Humanloop) to monitor the input and output of our generative systems to ensure there is no hallucination being returned.

- **Continual monitoring.**

Filevine guards against prompt injection or prompt hacking. Our Data Science team monitors the accuracy and completeness of our prompting libraries to ensure that there are no issues with the generative instructions and data source types required. Filevine also has evaluators on prompts to monitor and deny toxicity or ill-will prompts attempting to elicit information outside of the context intended.

- **Separating legal reasoning from general knowledge.**

For legal reasoning or decisions, Filevine does not consult the model's general knowledge. Instead, the response is informed by our robust training of legal context. Filevine provides proprietary instructions and system guidance to ensure that accuracy and completeness are met. Filevine's medical prompts return completeness and accuracy in the 90+% range depending on the data source, which far exceed the metrics of the 'do it yourself' GPT prompt metrics by over 50% on completeness. We put in the expertise and training on our end, so your teams can simply take advantage of the most innovative technology in a user-friendly way.



What Guardrails Are in Place?

Filevine uses our prompt management system guardrail evaluators to identify and remove any discriminatory output. We offer pre-determined prompts with strict guardrails, and all user-defined prompt functionality and engagement in Filevine are also subject to guardrails and coding/auth practices. These parameters limit the number of prompts and output that are rejected due to ill-will behavior or profanity.

Filevine has robust evaluator guardrails in the prompting system to ensure the output contains no hallucinations, ill-will, or hacking instructions. The level of these guardrails depends on the task. For instance, Filevine provides users with a robust pre-set prompt library in AIFields, while users can generate their own prompts in our DepoCopilot product.

How Does Filevine Support the Ethical Use of AI?

Filevine ensures statements of ethical values are communicated at the board level and down through the employee ranks. This applies to technology as well. These principles are audited annually as part of our SOC 2 Type II audit. Filevine also leverages [ABA opinions](#) and [model rules](#) to define and maintain ethical use of AI in our platform. Relevant ABA standards are [Opinion 512](#), Model Rule 1.1, Model Rule 1.5, Model Rule 1.6, Model Rule 7.1, Model Rule 5.1, and Model Rule 5.3.



Humans Always Make the Final Decisions

At the core of every AI tool Filevine creates is this principle: the final decision will always be made by a human.

Our tools make it easy for users to review the AI's decisions. Filevine provides the user several places to validate generative recommendations and output. Filevine encourages users to validate and use the output as a draft or starting place to ensure they can maximize their time and energy to quickly perform their jobs.

For example, in Filevine AIFields output, the user is given URL hyperlinks to the specific page to validate any data that has been cited. In Doc Review, users not only receive hyperlinks, but can also see a side-by-side review of the AI results with the source data. This ensures that users can quickly validate results and edit any inaccuracies if necessary. Filevine offers review and edit access — there is no 'auto-ship' process where Filevine automates human review away or automatically shares/publishes those AI generated outputs.



Building the Legal Technology of the Future

Filevine has already shown itself to be a leader in the field of legal technology — but we're only getting started.

In the years ahead, we will continue creating cutting-edge, AI-enhanced tools that make the practice of law better for legal professionals and their clients.

With every new development, we will continue to scrupulously protect your data, provide reliable results, and do everything possible to earn your trust.

