GUIDE

The Developer's Guide to Speech to Text for Medical Transcription

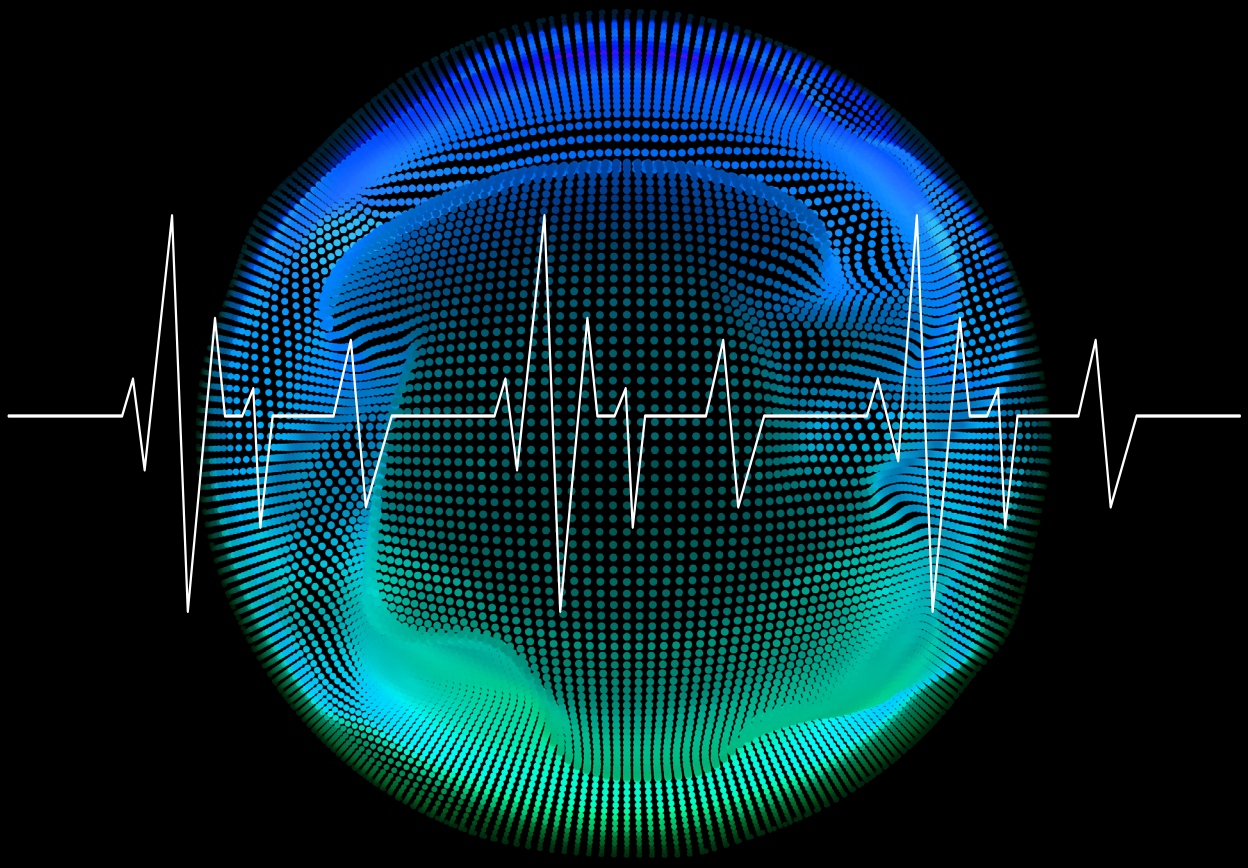

Deepgram

Table of contents

Executive Summary

Medical transcription is essential in healthcare, yet traditional methods are often slow, error-prone, and increase provider fatigue, creating a need for more advanced solutions. Deepgram's Nova-2 Medical speech-to-text (STT) API was designed to meet these challenges, leveraging powerful AI models to provide fast, accurate, and cost-effective automatic speech recognition (ASR) capabilities that developers can easily integrate to build transcription solutions tailored for healthcare.

This whitepaper examines how Nova-2 Medical can be used to enhance healthcare documentation. With **16% higher word recall rates (WRR)** for medical terms and **11% improved word error rate (WER)** over previous versions, Nova-2 Medical outperforms leading competitors by 42.8% on average. Key features include real-time transcription, rapid processing speeds (up to 40x faster than other models), and affordable pricing at only $0.0043 per minute. Built with strict adherence to HIPAA regulations, the model integrates seamlessly with healthcare systems and is customizable for specialized terminology.

Nova-2 Medical transforms workflows in clinical documentation, telemedicine, patient sessions, and medical research, helping healthcare organizations improve accuracy, efficiency, and patient care while reducing costs. Its scalable, secure architecture is a reliable tool for healthcare providers, streamlining operations and supporting various applications, from compliance checks to virtual agents.

Nova-2 Medical's performance leads to better patient care and more efficient operations. Healthcare applications that incorporate Deepgram's Nova-2 Medical model can see a 30% drop in the number of transcription errors, speeds that are up to 40 times faster than other options, and big cost savings up to 5 times more affordable compared to alternative solutions.

Introduction

The advent of automatic speech recognition (ASR), also known as speech-to-text (STT) or transcription, has ushered in a new era of human-computer interaction. Speech recognition and understanding are critical components shaping the future of not only the digital interfaces we use daily but also fundamentally redefining the way businesses, customers, and individuals interact.

This transformation transcends professional environments, influencing our personal spaces, and reshaping relationships with family, friends, and trusted advisors, including healthcare providers and clinicians.

ASR technology, already a familiar tool in our everyday lives, is continuously evolving. Powered by deep neural networks, the next generation of ASR tech offers unrivaled processing speed and transcription accuracy.

As an independent healthtech developer or a corporate team, you may wonder: how can I incorporate advanced medical transcription APIs to enhance my product and refine my business strategy?

Traditional medical transcription is unsuitable for most uses

Medical transcription that is both accurate and completed promptly is essential to providing quality healthcare. It improves patient outcomes by facilitating better communication between healthcare providers, reducing the likelihood of medical errors, and accurately documenting vital patient information.

However, traditional transcription methods, which frequently depend on manual processes, face several challenges:

- **High Costs:** Outsourcing transcription services or maintaining in-house teams can be expensive.
- **Turnaround Time Delays:** Manual transcription is time-consuming, leading to delays in accessing critical patient information.
- **Accuracy Concerns:** Human errors are inevitable, potentially leading to misdiagnosis, inappropriate treatments, and medico-legal complications that could impact patient care.

- **Compromised Patient Interaction:** The administrative burden of documentation can detract from valuable patient-provider face time, impacting the quality of care.

Fix them with Deepgram's AI-powered Nova-2 Medical

Since its inception in 2015, Deepgram has consistently defined the state of the art in automatic speech recognition (ASR) technology, setting new standards for speed and accuracy in the industry year after year after year. Our customers rely on our API-first platform, integrating our ASR and other voice AI capabilities into their products and workflows to deliver solutions across a broad spectrum of diverse use cases.

Deepgram's flagship ASR model, Nova-2, is widely recognized for its versatility across various domains where text conversion from human speech is required. Organizations like **NASA** use Nova-2 use Nova-2 to transcribe communications between the International Space Station (ISS) and mission control. **Five9** also uses Nova-2 to improve customer service through real-time transcription and audio analysis.

Recognizing that general-purpose models may not fully meet the need of specialized fields like healthcare, we developed Deepgram Nova-2 Medical. Nova-2 Medical is a finely-tuned model that is highly proficient at transcribing healthcare-specific language such as:

Medical Terminology: Symptoms, diagnoses, treatments, medications, and clinical jargon.

Physician Dictation: Accurately capturing physician notes, discharge summaries, and other medical reports.

Patient-Physician Interactions: Transcribing consultations, ensuring comprehensive documentation of patient histories and treatment plans.

Nova-2 Medical enables healthcare application providers to accurately capture and document vital patient information. This has improved patient care and communication among healthcare professionals, reduced errors, and increased operational efficiency.

A preview of this guide

This guide has the following structure:

An overview of the challenges associated with traditional medical transcription.

An introduction to Deepgram Nova-2 Medical and its underlying technology.

A detailed exploration of Nova-2 Medical's key features and how it addresses these issues through its advanced capabilities and domain-specific training.

Real-world use cases that demonstrate the impact of Nova-2 Medical in healthcare settings.

A comprehensive comparison of the different, leading AI-powered medical transcription tools on the market, including Deepgram's Nova-2 Medical.

A conclusion summarizing the value proposition and encouraging further action.

An appendix addressing frequently asked questions.

After reading this guide, you will have a comprehensive understanding of why Deepgram Nova-2 Medical is the best option available to fulfill your medical transcription needs.

The High Cost of Inefficient Medical Transcription

Healthcare providers face numerous obstacles with traditional medical transcription, affecting patient care and operational efficiency. These challenges arise in various settings, whether during in-person consultations, telemedicine appointments, or over-the-phone conversations.

The following are some of the most significant obstacles:

- **Decreased patient engagement due to time spent on documentation.**
- **Time-consuming manual documentation.**
- **Risk of errors and inaccuracies.**
- **Inefficiencies and delays.**

Let's explore each challenge in detail.

Decreased patient engagement due to time spent on documentation

Documentation demands often force healthcare providers to split their attention between patient interactions and record-keeping.

According to a 2024 study, 80% of physician respondents agreed that "the effort or time required for me to complete documentation tasks impedes patient care."[1] This staggering figure contributes to inefficiencies and impedes patient engagement.

Dividing a physician's focus between documentation and patient care can compromise the quality of care, as clinicians may feel pressured to rush consultations, which, of course, could result in patients feeling less engaged and valued.

Recent research has shown a significant increase in EHR workload for primary care physicians over the past few years (Arndt et al., 2024)[2].

The study found that:

Primary care physicians spent a median of 36.2 minutes on the EHR per patient visit. This is significantly more than typical face-to-face time with patients, often around 15-20 minutes per visit.

Physician burnout and reduced productivity

Healthcare professionals, particularly physicians, can spend nearly half their workday on transcription tasks, which pulls them away from patient care. This administrative burden not only reduces the time available for patient care but also contributes to physician burnout, a growing concern in the healthcare industry.

According to the Medscape Physician Compensation Report for 2023, physicians spend an average of 15.5 hours per week on paperwork and administrative tasks (Kane, 2023)[3]. This report, which surveyed physicians across 23 specialties, provides a more recent and comprehensive look at the administrative burden doctors face.

The following are some key findings from this report:

- ***Physicians spend the most time on paperwork and administrative tasks, averaging 19 hours per week.***
- ***Several specialties, including critical care, internal medicine, nephrology, neurology, and oncology, spend an average of 18 hours per week on these tasks.***
- ***Family medicine physicians devote approximately 17 hours weekly to paperwork and administrative tasks[3].***

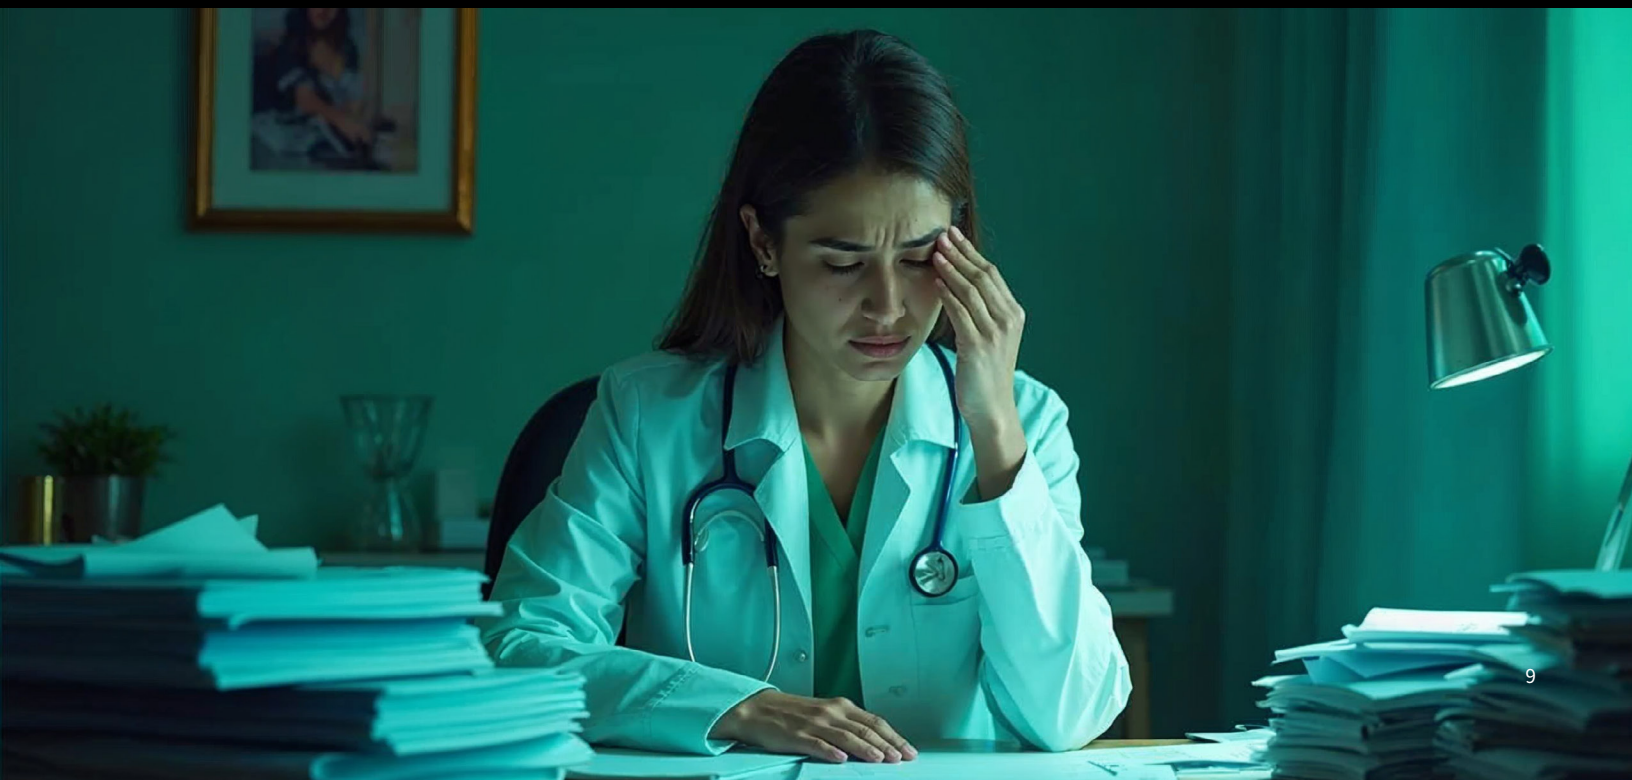

Compromised patient safety due to a high risk of errors and inaccuracies

Manual transcription poses a high risk of error. Fatigue, time constraints, and the complexity of medical terminology can result in errors that may lead to misdiagnosis, inappropriate treatments, or even medico-legal complications.

Additionally, issues like illegible handwriting and incomplete records further disrupt the continuity of care. According to a study in the Journal of Patient Safety, hospitals cause between 210,000 and 400,000 deaths annually due to preventable harm (James, 2013)[4].

The Journal of the American Medical Informatics Association published a 2019 study examining the rate of manual transcription errors in outpatient point-of-care testing[5].

This study found that:

- ***260 out of 6,930 (3.7%) manual entries differed from their interfaced result.***
- ***37 of the 260 (14.2%) errors were discrepant by more than 20% and included potentially dangerous mistranscriptions.***
- ***An additional 37 (14.2%) errors were due to the inclusion of non-numeric characters.***
- ***The study concluded that clinically significant discrepancies for clinic-entered point-of-care results occurred at a rate of approximately 5 per 1,000 results (0.5%).***

Operational bottlenecks that cause inefficiencies and delays

Traditional transcription methods are often inefficient and slow. The multi-step process, including dictation, transcription, and review, can delay document turnaround, impacting critical decision-making and slowing the overall care process.

Delayed access to medical records can hinder timely diagnosis and treatment plans, potentially compromising patient outcomes. These challenges highlight the need for a more efficient, accurate, and streamlined approach to medical transcription.

We developed the Nova-2 Medical model to alleviate these issues and transform how healthcare practitioners approach medical documentation.

Addressing Medical Transcription Challenges with Artificial Intelligence (AI)

At Deepgram, we've collaborated with leading healthtech providers such as Lyrebird, Five9, Tortus, NOCD, Clinical Notes AI, and Phonely to tackle critical challenges in medical transcription, including automating EHR and SOAP—or subjective, objective, assessment, and plan—notes, transcribing patient sessions, and managing patient intake processes.

These partnerships helped us identify the unique demands of the healthcare sector and inspired us to develop a solution that addresses these challenges with accuracy and efficiency.

While our general-purpose Nova-2 model[6] has delivered exceptional results in various industries, we recognized the need for a specialized automatic speech recognition (ASR) model that excels in the medical field. This led us to create the Deepgram Nova-2 Medical model, specifically developed for the healthcare industry.

Introducing Deepgram's Nova-2 Medical model

Deepgram's Nova-2 Medical[7] is a speech-to-text (STT) model specifically engineered for the healthcare sector. It delivers high accuracy in medical transcription while maintaining the speed and cost-effectiveness of the original Deepgram Nova-2.

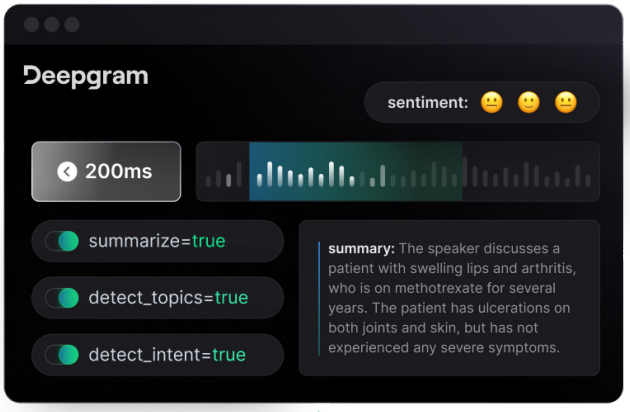

By combining Deepgram Nova-2 Medical with Deepgram's Audio Intelligence, tasks such as generating summaries of clinical encounters can be efficiently performed.

This model accurately captures medical terminology, symptoms, diagnoses, treatments, and medications so that healthcare professionals receive clear, actionable transcripts.

The Nova-2 Medical model not only maintains the original Nova-2's remarkable speed and cost-effectiveness, but it also adapts to handle the complexities of medical language. Like all of our other voice AI models, Nova-2 Medical is accessible through a single, powerful API, enabling developers to quickly integrate its capabilities into their existing applications.

How Deepgram transforms medical transcription with Nova-2 Medical

The Deepgram Nova-2 Medical model uses the same core speech-to-text (STT) process as the original Nova-2 model but with enhancements tailored to the healthcare domain.

The audio data, whether capturing pre-recorded conversations or real-time interactions between healthcare professionals and patients, goes through the following stages before transforming into text:

1. **Audio Preprocessing:** The system enhances sound quality and filters out background noise to capture even complex medical conversations clearly.
2. **Feature Extraction:** Next, the system extracts the important features like pitch, frequency, intensity, duration, and formants as phonemes and creates a spectrogram that the model can process.
3. **Acoustic Model:** The acoustic model analyzes the spectrogram to identify phonemes—the fundamental units of sound in spoken language.
4. **Language Model:** The language model uses its understanding of medical terminology and context to convert the identified phonemes into coherent words.
5. **Decoding and Post-Processing:** The system combines the individual words to generate a complete and accurate transcription, ready for integration into healthcare workflows.

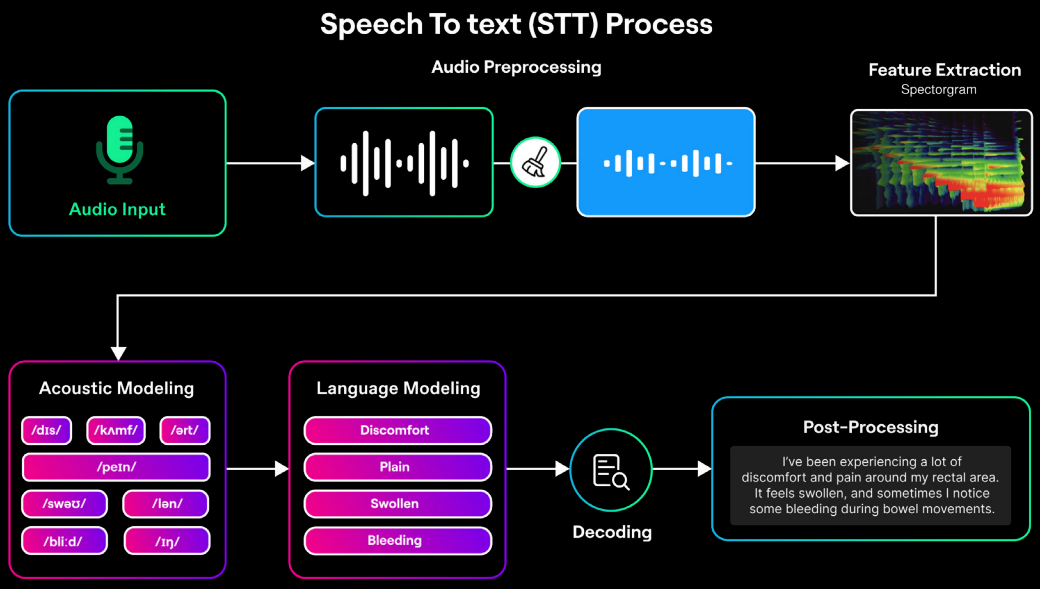

An illustration of the medical transcription process of Deepgram Nova-2 Medical (audio input → feature extraction → acoustic and language modeling → transcript).

At its core, we built Nova-2 Medical on an encoder-decoder Transformer architecture designed specifically for high-accuracy STT—or speech-to-text—tasks. The two key components are:

- **Acoustic Model:** The encoder processes and encodes the audio into acoustic embeddings to capture essential sound features such as phonetic and tonal elements.
- **Language Model:** The decoder converts these sound embeddings into correct medical text with lots of context for high-quality transcription.

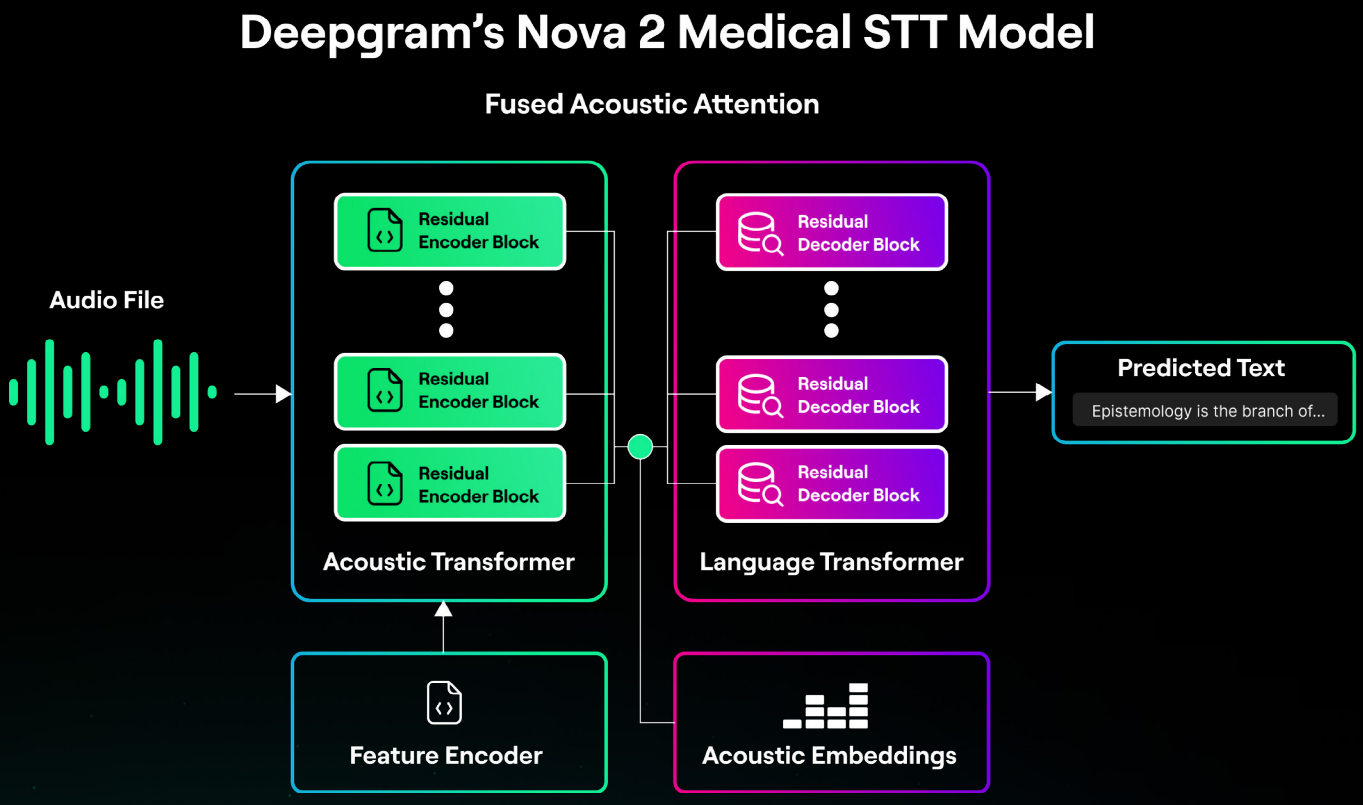

The Deepgram Nova-2 Medical architecture is based on the Transformer model. It includes an encoder (acoustic model) that converts audio into embeddings and a decoder (language model) that transforms those embeddings into text.

What Makes Deepgram's Nova-2 Medical Model Unique?

The Nova-2 Medical model from Deepgram has features made to meet the needs of healthcare providers. Built for accuracy, efficiency, and security, Nova-2 Medical transforms how medical professionals manage documentation.

Here's how Nova-2 Medical stands out:

1. We have developed a highly accurate speech-to-text (STT) model.
2. There is compliance and security.
3. Real-time transcription with exceptional speed.
4. A low-cost API with huge savings.
5. We prioritize integration and scalability.
6. The process involves customization and deployment.

Highly accurate speech-to-text (STT) model (42.8% WER ⬇ 20.5% WRR ⬆)

Nova-2 Medical also achieves an 11% improvement in overall word error rate (WER) for pre-recorded (batch) transcription compared to the previous version. It outperforms leading medical alternatives like Rev AI, Speechmatics, or Assembly AI by 42.8% on average.

With Nova-2 Medical, you will achieve outstanding accuracy for medical terminology without sacrificing general speech recognition performance. This creates a balanced and reliable solution for transcribing clinical interactions while enabling use cases such as customer support bots, QA/compliance analysis, and virtual clinical agents.

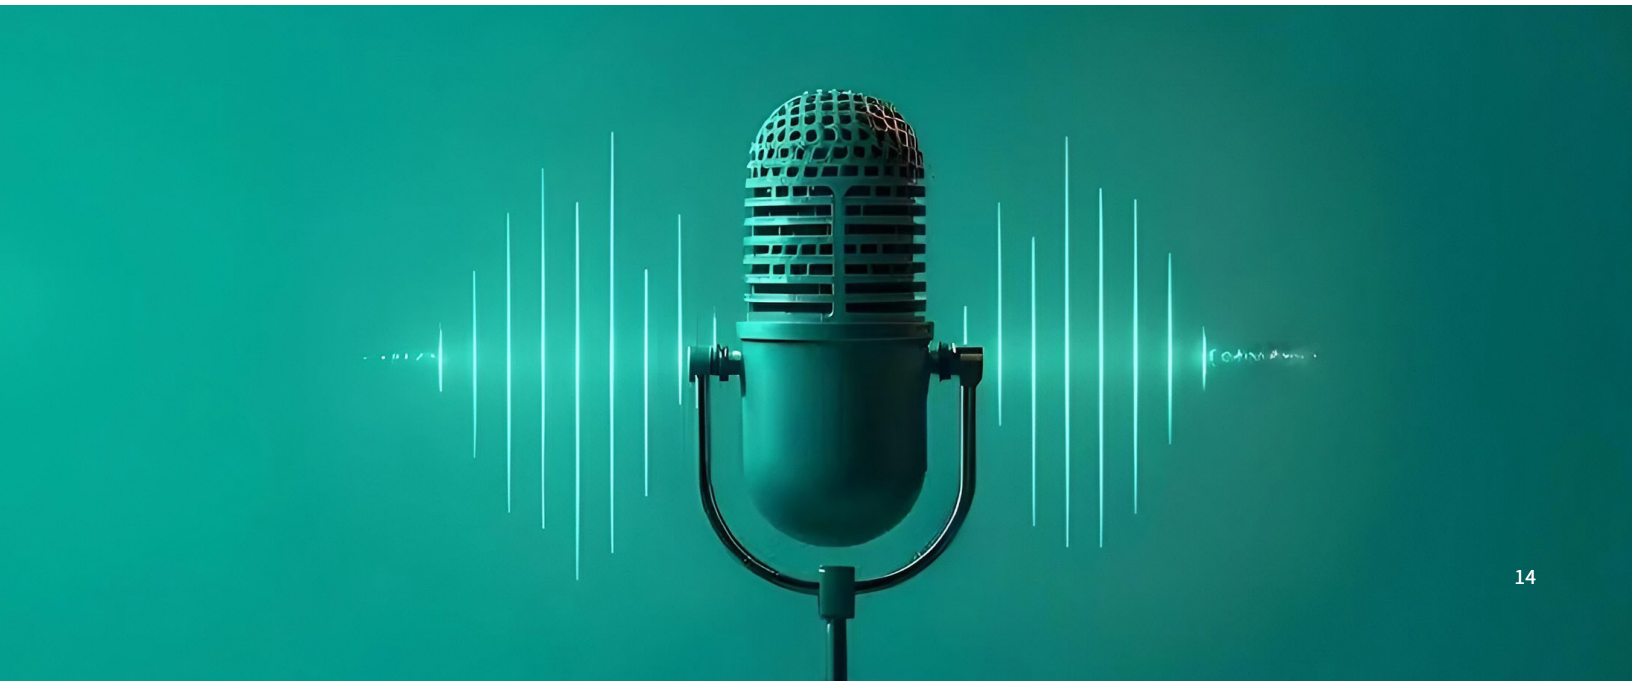

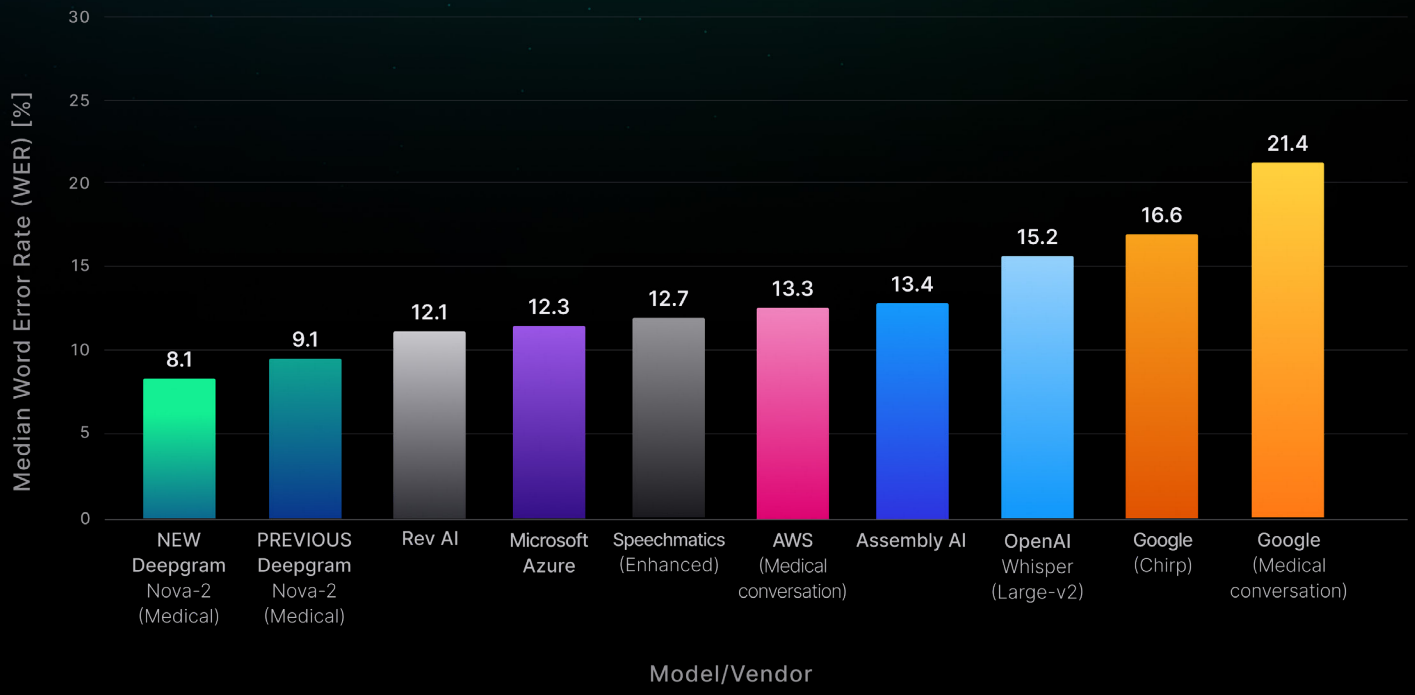

We present the median file word error rate (WER) for pre-recorded English transcription across all benchmarked medical domain test sets.

Nova-2 Medical ensures that healthcare professionals can record and capture correct and complete medical information. It has a 16% higher word recall rate (WRR) for medical terms than the previous Nova-2 medical model and a 20.5% average improvement over its leading competitors.

WRR measures the percentage of words in the ground truth text correctly predicted or matched (i.e., true positives), with higher scores indicating better accuracy.

> *Deepgram's speech-to-text has been far and away the best we've seen for medical transcription accuracy. Plus, their approach makes it simple to add in new, uncommon words outside of their training domain." **-Will Bodewes, CEO, Phonely AI***

HIPPA compliance and patient data security

The Nova-2 Medical model ensures patient safety and privacy during the transcription process by adhering to the highest medical documentation standards, including HIPAA compliance.

Independent audits and certifications of Deepgram's secure infrastructure and data handling procedures ensure that healthcare organizations can adhere to legal and regulatory requirements while protecting sensitive data.

Real-time transcription and instant access to patient data

The Nova-2 Medical model can transcribe in real-time, which lets doctors and nurses get patient information right away during appointments or procedures. Because they don't have to do it by hand, doctors can spend more time talking to patients and less time writing things down.

Real-time transcription also helps people make decisions more quickly and correctly by giving them instant access to relevant patient data.

Another advantage is that for clinics where patients are skeptical about recording their data, real-time audio intelligence with Nova-2 Medical can help healthcare IT practitioners build systems that process sessions on the fly and discard the recordings right after.

> *If you want to build a real-time voice application these days, I don't see anything else being used other than Deepgram."* **– Lukas Wolf, Co-Founder, Sonia**

Improved efficiency through faster transcription and documentation speeds

Nova-2 Medical's transcription rate is approximately 5 to 40 times faster than competing models, with a median inference time of just 29.8 seconds per hour of audio. It is also one of the only speech recognition models on the market capable of meeting real-time application demands.

Our quick processing feature reduces the administrative burden on healthcare professionals. It allows them to document patient interactions more efficiently and focus on delivering care.

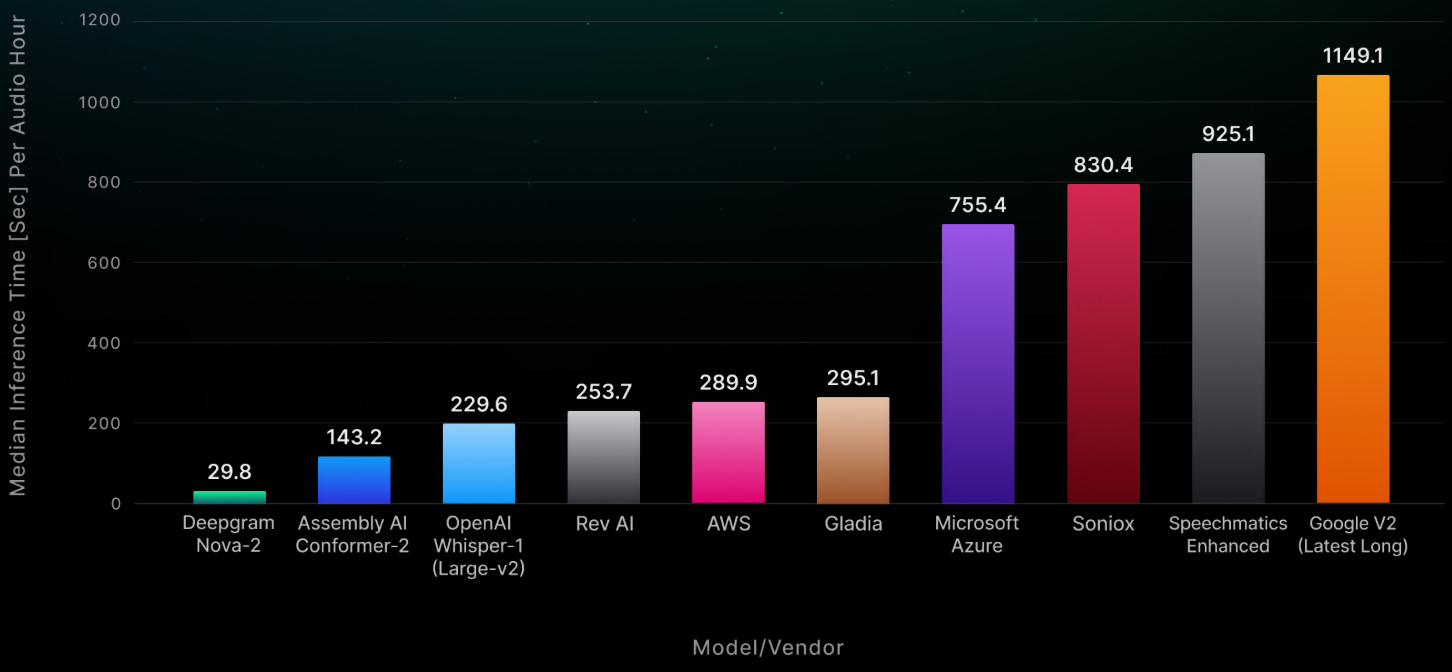

Deepgram Nova-2 Medical transcribes one hour of audio in 29.8 seconds, while its closest competitor is four times slower.

Nova-2 Medical has no trade-off between accuracy, speed, and cost, which makes it the top choice for streaming voice applications.

> *Deepgram's Nova-2 medical model can transcribe entire phone conversations with precision and speed that rivals human transcribers. Thanks to Nova-2, we've delivered unparalleled ROI to our customers in the healthcare sector."*
> **– Antoni Rosinol, CEO, StackAI**

Low-cost API with huge savings compared to other solutions

Nova-2 is three to five times more affordable than other full-service transcription providers, with prices starting at just $0.0043 per minute. For a healthcare organization transcribing 1,000 hours of audio per month, this translates to a potential cost savings of $12,000 to $20,000 annually.

Our cost-effective API enables healthcare organizations to cut transcription costs while maintaining high-quality outputs. This makes it an ideal choice for institutions seeking both affordability and performance.

Scalable API that integrates with existing systems (EHR, automation workflows, etc.)

The Nova-2 Medical API integrates with existing healthcare workflows with minimal disruptions and configurations required. The API can handle rising or falling transcription volumes as the organization's needs change.

Whether a small clinic processes 100 hours of audio per month or a large hospital network deals with 10,000 hours, Nova-2 easily adapts to meet evolving demands.

Customize Nova-2 Medical for your use case

Deepgram provides custom model training for specific medical use cases, enhancing accuracy for rare or domain-specific terminology.

Healthcare professionals can improve the accuracy of uncommon keywords (e.g., new drug names) with rapid custom model training services that boost the Nova-2 medical model's already impressive, out-of-the-box performance.

In a pilot study, a Nova-2 model custom-trained for oncology practice had a 98.5% success rate, while the generic model only had a 92% success rate.

> We have been using the Nova-2-medical model for the last few months and have had a very positive experience. Having accurate medical speech recognition makes a significant difference in the notes that doctors create using our AI Scribe. We are definitely impressed with its ability to recognize medical terms and the blazing processing speed of our audio files.
>
> **– Gerardo Guerra Bonilla, CEO, Chartnote.**

Deploy Nova-2 Medical as a managed service or within your self-hosted environment (on-prem or VPC)

Nova-2 offers flexible deployment options, allowing organizations to choose between a managed service or secure self-hosting on-premises or in a virtual private cloud (VPC) to meet operational and security requirements.

Next, we will examine the various healthcare applications of Nova-2, providing examples from various companies.

Nova-2 Medical Transforms Healthcare Workflows: Use Cases

The Nova-2 Medical model from Deepgram has Accurate and efficient transcription is crucial in the healthcare sector for improving patient care, streamlining clinical workflows, and advancing medical research. We specifically tailored the Nova-2 Medical model to meet these diverse needs, providing a range of applications to improve healthcare outcomes.

Here's how to use Nova-2 Medical effectively in a variety of scenarios:

Efficient and Accurate Clinical Documentation

Deepgram's Nova-2 Medical model streamlines the creation of essential clinical documents such as patient notes, discharge summaries, and other medical reports.

Healthcare providers spend significantly less time on paperwork due to the model's accurate transcription of complex medical terms such as symptoms, diagnoses, treatments, and medications.

Implementing Nova-2 Medical for clinical documentation in a large hospital network led to a 40% reduction in paperwork time, enabling healthcare providers to concentrate more on patient care. Its high accuracy captures critical details accurately, which enhances the reliability of medical records.

> *We refer to ourselves as a 'therapeutic scribe.' We are using Deepgram to speak to the actual clinician so that they can have a conversation. It says, 'describe your session for me'. They'll describe what happened in their session, as well as how their client presented any assessments they did. The AI will then pose questions, and once the clinician responds, it will use the transcripted information to create a note or document.*
>
> **– Ross Young, CEO of Clinical Notes AI**

StackAI cut clinical application errors by 24% and saved 67% for a client

Conducting call center quality assurance is traditionally a tedious, manually intensive process that involves call monitoring, call scoring, and using analytics software to track qualitative metrics related to agent sentiment and adherence to call scripts and SOPs (standard operating procedures).

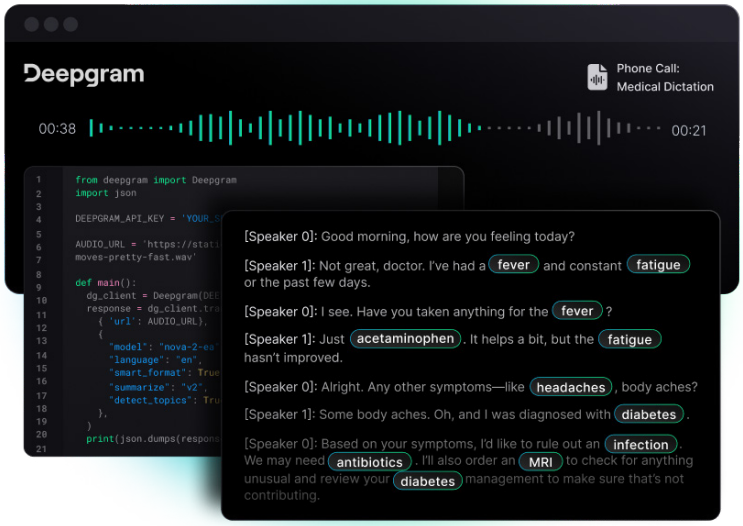

By combining Deepgram Nova-2 Medical with Deepgram's Audio Intelligence, tasks such as generating summaries of clinical encounters can be efficiently performed.

In healthcare domains, these tasks can be especially problematic as conversations often involve protected health information (PHI) that requires strict handling to ensure compliance with the Health Insurance Portability and Accountability Act (HIPAA).

StackAI recently helped their client deploy a HIPAA-compliant AI solution using the Nova-2 Medical model to automate their QA process. With Nova-2 Medical, they **reduced their error rate by 24% and increased call processing by nearly 7 times, resulting in an overall cost savings of 67%** and providing exactly the type of transformative outcome McKinsey[8] and others have predicted AI can deliver.

Nova-2 Medical improves virtual consultations in telemedicine

As virtual consultations become more common, accurate and real-time transcription has become essential. The Nova-2 Medical model enables healthcare professionals to generate transcripts during telemedicine sessions, offering immediate access to patient information.

This capability not only improves communication between healthcare providers and patients but also guarantees the real-time capture of all relevant details, thereby enhancing the overall quality of care.

Sonia Health Offers Fast and High-Precision Virtual Interactions with Patients.

Sonia is a mobile application that delivers AI-driven voice therapy sessions on demand and at a fraction of the cost of today's traditional services. The application uses Nova-2 Medical to provide the precision, speed, and cost efficiency required to deliver affordable, high-quality mental health support to anyone in need.

It puts a virtual therapist you can talk to at a moment's notice right in your back pocket.

Sonia is a fully AI-powered cognitive behavioral therapist (CBT).
Image Source: Soniahealth.com

Patient session transcription for comprehensive patient records

As noted earlier, the Nova-2 Medical model accurately transcribes patient sessions and captures conversations between healthcare providers and patients. This feature improves follow-up care by recording patient histories and detailed recommendations, providing valuable reference points for future treatment decisions.

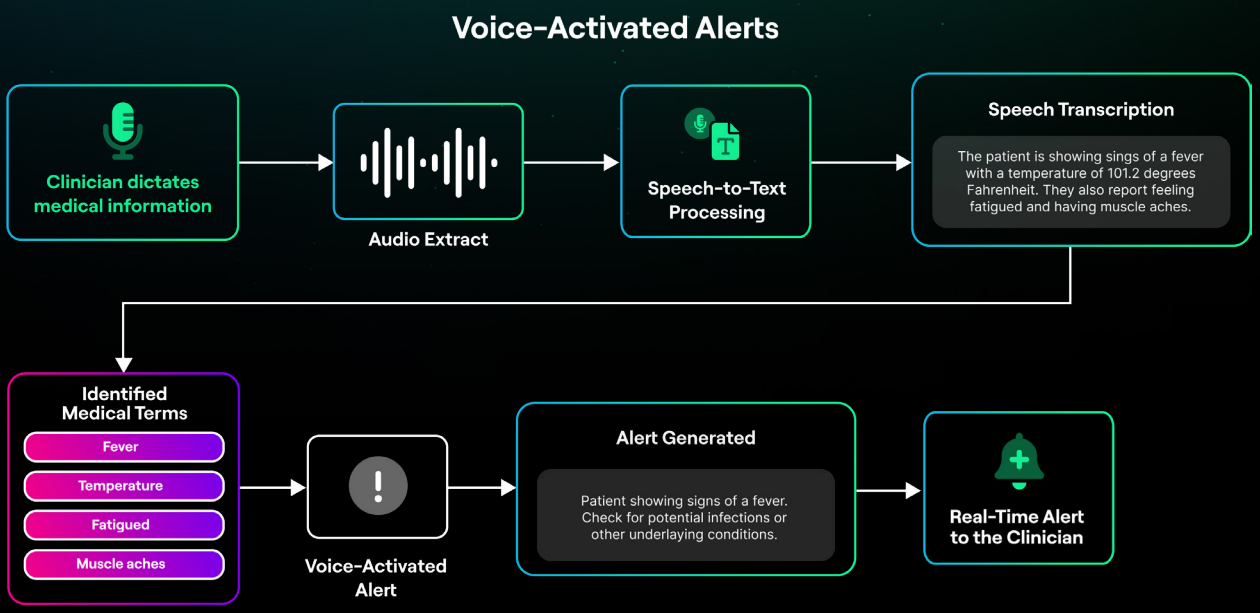

Deepgram Nova-2 Medical helps clinicians improve sessions by monitoring keywords like allergies, medication changes, and drug interactions. Clinicians receive instant alerts when these keywords occur, improving patient care and decision-making.

Lyrebird Health reduces burnout with patient session documentation.

Deepgram's Nova-2 Medical Model powers Lyrebird Health's medical documentation platform, helping practitioners like Dr. Carr interact better with patients. Their application manages every aspect of medical documentation.

Nova-2 Medical listens, learns, and generates notes instantly, reducing physician stress and minimizing errors due to the model's high accuracy.

Accelerating Discoveries in Medical Research

Nova-2 Medical plays a pivotal role in medical research by transcribing and analyzing large volumes of audio or video data from research studies or clinical trials. By turning spoken content into text quickly and easily, researchers can analyze their findings more quickly and accurately, speeding up progress in medicine overall.

For example, a research team using Nova-2 Medical could process and transcribe over 500 hours of clinical trial data in half the time it would have taken with traditional transcription methods.

Now that you know what makes Nova-2 awesome and why healthcare organizations use it, let's see how it compares to other leading solutions on the market.

Comparison to Other Medical Transcription Solutions

Deepgram Nova-2 Medical is the leading medical transcription service, but how does it compare to other options in the industry? This comparison examines its performance alongside Google Cloud Speech-to-Text, Amazon Transcribe Medical, Microsoft Azure Speech Service, Speechmatics, Rev AI, and more, focusing on key areas such as accuracy, speed, and cost.

The table below provides a detailed look at how each service fares in these critical areas.

Comparison to Other Medical Transcription Solutions

Feature/ Model	Deepgram Nova-2 Medical	Google Cloud	AWS	Microsoft Azure	Whisper Large v2	Speechmatics	AssemblyAI	Rev AI	Google Chirp
Accuracy	16% improvement in medical terminology WRR; 42.8% better WER than competitors	WER of 21.4%.	13.3% WER lower than Deepgram Nova-2 Medical by 5.2%	12.3% WER lower than Deepgram Nova-2 Medical by 4.2%	WER of 15.2%.	12.7% WER lower than Deepgram Nova-2 Medical by 4.6%	13.4% WER lower than Deepgram Nova-2 Medical by 5.3%	12.1% WER lower than Deepgram Nova-2 Medical by 4%	WER of 16.6%.
Speed	5 to 40 times faster than competitors; median inference time of 29.8 seconds per hour of audio	40x slower than Deepgram Nova-2 Medical.	10x slower than Deepgram Nova-2 Medical.	25x slower than Deepgram Nova-2 Medical.	Slower compared to specialized models; real-time performance can vary	30x slower than Deepgram Nova-2 Medical	5x slower than Deepgram Nova-2 Medical.	Slower than Deepgram Nova-2 Medical.	Slower than Deepgram Nova-2 Medical.
Cost	Starting at $0.0043 per minute; highly cost-effective compared to others	Pay-as-you-go pricing; costs can accumulate based on usage	Pay-as-you-go pricing; free tier available for limited use	Pricing varies by usage; generally higher than Deepgram	Open-source, free to use but requires infrastructure setup	3x more expensive than Deepgram Nova-2 Medical	3x more expensive than Deepgram Nova-2 Medical	5x more expensive than Deepgram Nova-2 Medical	4x more expensive than Deepgram Nova-2 Medical
Integration	Seamless integration with EHR systems; flexible deployment options (cloud/ on-premises)	Integrates well with Google Cloud services; requires S3 for audio storage	Integrates with AWS services and EHR systems easily	Integrates with Azure services but may require more setup for healthcare applications	Requires custom integration depending on deployment scenario; may need additional tools for EHR integration	Unified API simplifies integration across platforms	API designed for easy integration into existing workflows	API available for integration into various applications	Integrates well within Google's ecosystem
Customization	Custom model training available for specific medical terms and jargon	Limited customization options for specific medical specialties	Customizable for specific medical fields like cardiology or neurology	Some customization options are available but less flexible than Deepgram	Open-source allows for extensive customization but requires technical expertise to implement effectively	Offers custom dictionaries	Customizable models based on user needs	Customization options available	Limited customization options

The table shows that Deepgram Nova-2 Medical outperforms the competition across all listed criteria. In terms of accuracy, it outperforms the most accurate models in WER by over 4%. It is also more than 5% faster than the quickest competing models.

Additionally, Deepgram Nova-2 Medical offers seamless integration and full customization, advantages that most competitors lack. Despite its superior performance, it remains the most cost-effective medical transcription service, offering no trade-off between performance and price.

Empowering Healthcare Providers with AI-Powered Transcription

Deepgram Nova-2 Medical is at the forefront of transforming medical transcription, addressing the long-standing challenges that have burdened healthcare providers for years. Nova-2 Medical helps clinicians regain valuable time, improve patient care, and make operations more efficient by automating and streamlining documentation workflows.

More Time with Patients:

With real-time transcription and exceptional speed, Nova-2 Medical minimizes the administrative burden on healthcare professionals, allowing them to focus more on their patients and less on paperwork.

Improved Accuracy:

Nova-2 Medical's industry-leading accuracy guarantees the precise capture of critical patient information, thereby reducing errors and enhancing the reliability of medical records.

Reduced Physician Burnout:

Nova-2 Medical helps alleviate physician burnout by automating time-consuming transcription tasks, promoting better work-life balance and overall well-being.

Improved Operational Efficiency:

Deepgram's solution's seamless integration and scalability enable healthcare organizations to optimize their workflows and adapt to evolving demands.

What's Next? Embrace the Future of Medical Transcription

Deepgram Nova-2 Medical is more than a transcription tool; it's a catalyst for transforming healthcare documentation. Experience the power of AI-powered transcription and unlock new levels of efficiency, accuracy, and patient care.

We invite you to explore our **API Playground** or **sign up** to try the Nova-2 medical model firsthand. To get started, simply include **`model=nova-2-medical`** in your API calls. For more information, visit our **API Documentation.**

Join us in transforming healthcare with cutting-edge speech recognition technology tailored to your needs. To see how the Nova-2 medical model can integrate seamlessly into your system, contact us today!

About Us at Deepgram

Deepgram is a foundational AI company on a mission to understand human language. We've processed nearly 10,000 years of audio for trailblazing customers like Citi, TIAA, TalkDesk, Observe AI, Twilio, and Spotify, even transcribing NASA's radio communications between the ISS and Mission Control. We've served nearly 2 trillion enriched, computable words to our customers.

With just an API call, we provide any developer access to the most advanced speech AI transcription and understanding. Our models deliver the fastest, most accurate transcription alongside contextual features like summarization, sentiment analysis, and topic detection.

To learn more, visit **deepgram.com** or **contact us** to speak with one of our product experts.

Deepgram

Essential Building Blocks for **Voice AI**

Appendix A
Answers to Frequently Asked Questions (FAQs)

What's the difference between Deepgram Nova-2 and Deepgram Nova-2 Medical?
Deepgram Nova-2 is a general-purpose transcription model designed for a wide range of industries, while Deepgram Nova-2 Medical is specifically trained to recognize complex medical terminology and jargon.

Nova-2 Medical combines Nova-2's core capabilities with extensive training in medical terminology, clinical jargon, and physician dictation styles. This results in significantly higher accuracy for medical transcriptions, making it an ideal choice for clinical documentation, telemedicine, and research settings.

What does it mean for a transcription service to be HIPAA-compliant?
HIPAA compliance signifies that the transcription service adheres to the stringent guidelines set forth by the Health Insurance Portability and Accountability Act (HIPAA) in the United States. This ensures that all patient data, including Protected Health Information (PHI), is handled securely and confidentially throughout the transcription process, safeguarding patient privacy and maintaining regulatory compliance.

How does Deepgram Nova-2 Medical handle real-time transcription?
Deepgram Nova-2 Medical offers real-time transcription, converting spoken medical conversations into text instantaneously. This capability empowers healthcare providers to access patient information and generate documentation in real-time during consultations, procedures, or telemedicine sessions, improving clinical workflows and patient care.

What are the common applications of medical transcription?
- Medical transcription serves a variety of critical functions in the healthcare sector, including:
- Automating the creation of Electronic Health Records (EHRs) and SOAP notes.
- Streamlining patient intake processes by transcribing patient-provided information.
- Transcribing patient sessions and consultations for accurate and comprehensive documentation.
- Facilitating medical research by transcribing and analyzing audio/video data from studies and trials.

Can Deepgram Nova-2 Medical be customized for specific use cases?
Yes, Deepgram offers the flexibility to customize and fine-tune the Nova-2 Medical model to recognize specialized medical terminology or jargon unique to specific medical specialties, departments, or organizations. This ensures optimal accuracy and relevance for your specific transcription needs.

How fast is Deepgram Nova-2 Medical compared to other transcription services?
Deepgram Nova-2 Medical boasts exceptional speed, capable of processing transcriptions 5 to 40 times faster than many leading alternatives. This rapid turnaround time saves healthcare professionals valuable time, allowing them to focus more on patient care and less on administrative tasks.

In terms of accuracy, what sets Deepgram Nova-2 Medical apart?
Deepgram Nova-2 Medical delivers industry-leading accuracy, demonstrating a 16% improvement in Word Recall Rates (WRR) for medical terminology and a 42.8% better overall Word Error Rate (WER) than many competing transcription solutions. This ensures that your medical documentation is precise, reliable, and actionable, reducing the risk of errors in patient records.

What are the costs associated with using Deepgram Nova-2 Medical?
Deepgram Nova-2 Medical is a budget-friendly transcription solution. Pre-recorded audio costs $0.0043 per minute, making it affordable for healthcare organizations looking to improve transcription accuracy without incurring high costs.

How does Deepgram Nova-2 Medical integrate with existing healthcare systems?
Deepgram provides seamless integration through its user-friendly API, allowing the Nova-2 Medical model to effortlessly connect with existing Electronic Health Record (EHR) systems, telemedicine platforms, and other healthcare software. This flexibility ensures that Nova-2 Medical fits naturally into your organization's workflows, minimizing disruption to existing workflows and maximizing efficiency.

Appendix B

1 **"Payerchin, R. (June 10, 2024). Physicians say documentation burdens are impeding patient care." Medical Economics®.** https://www.medicaleconomics.com/view/physicians-say-documentation-burdens-are-impeding-patient-care

2 **Arndt, B. G., Micek, M. A., Rule, A., Shafer, C. M., Baltus, J. J., & Sinsky, C. A. (2024). More tethered to the EHR: EHR workload trends among academic primary care physicians, 2019-2023. The Annals of Family Medicine, 22(1), 12-18.** https://doi.org/10.1370/afm.3047

3 **Kane, L. (2023). Medscape physician compensation report 2023: The recovery. Medscape.** https://www.medscape.com/slideshow/2023-compensation-overview-6016341

4 **James, J. T. (2013). A new, evidence-based estimate of patient harms associated with hospital care. Journal of Patient Safety, 9(3), 122-128.** https://pubmed.ncbi.nlm.nih.gov/23860193/

5 **Mays, J. A., & Mathias, P. C. (2019). Measuring the rate of manual transcription error in outpatient point-of-care testing. Journal of the American Medical Informatics Association, 26(5), 433-437.** https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6351970/

6 **Josh Fox. (2024, June 27). Introducing Nova-2: The next generation of Deepgram's speech-to-text API.** https://deepgram.com/learn/nova-2-speech-to-text-api

7 **Josh Fox. (2024, June 12). Nova-2 Medical Model: The #1 Choice for Medical Transcription.** https://deepgram.com/learn/best-medical-transcription-with-nova-2-medical-speech-to-text-api

8 **Chui, M., Hall, B., Singla, A., Sukharevsky, A. (2021, December 8). The State of AI in 2021.** https://www.mckinsey.com/capabilities/quantumblack/our-insights/global-survey-the-state-of-ai-in-2021