WHITEPAPER

Nym’s Autonomous Medical Coding Engine

A technical overview of Nym’s autonomous medical coding engine and its benefits for healthcare revenue cycle management.

www.nym.health

CONTENTS

About Nym

Nym transforms revenue cycle management by automating the medical coding process. Our multi-specialty autonomous medical coding engine assigns medical charge codes in seconds and with zero human intervention, enabling hospitals, health systems, and provider groups to improve quality, accelerate payment cycles, and reduce coding-related costs. Today, our engine is powering automation in revenue cycle management and bringing accuracy and speed to medical billing, processing over 5.5 million charts annually in more than 250 healthcare facilities across the US.

See Nym in action

Executive Summary

While other industries have successfully begun to leverage the latest advancements in artificial intelligence, healthcare revenue cycle management continues to rely on inefficient manual processes that burden healthcare staff and cost providers billions annually. This is mainly due to the complexity of clinical language in patient medical records which, until recently, was one of the biggest barriers to the adoption of emerging technologies in healthcare revenue cycle management.

Nym has broken down this barrier. Our interdisciplinary team of physicians, computational linguists, engineers, and AI experts have developed innovative clinical language understanding (CLU) technology that understands the structured and unstructured clinical data in patient records and provides an entirely new way of analyzing and decoding medical records into clear, concise, and actionable information.

Nym currently leverages this CLU technology to fully automate medical coding, a process that has remained largely manual and error-prone for decades and which contributes significantly to the $256.6 billion the U.S. spends annually on "administrative complexity." Combining our innovative CLU technology with rules-based medical coding ontologies, Nym's industry-leading autonomous medical coding engine translates provider notes within patient charts into medical codes in seconds with over 95 percent accuracy and absolutely zero human intervention. This enables healthcare providers to improve coding quality, reduce coding-related costs by up to 35 percent, accelerate payment cycles, and drive countless other operational and financial benefits across medical coding and the broader revenue cycle management process.

This whitepaper provides a brief overview of the challenges facing the medical coding industry and the barriers to technology adoption in revenue cycle management before diving into Nym's differentiated approach to autonomous medical coding and the benefits that our engine drives for hospitals, health systems, and provider groups across the US.

Challenges to Technology Adoption in Medical Coding

Healthcare revenue cycle management (RCM) is the process of tracking and optimizing the financial lifecycle of a healthcare organization, from patient registration and appointment scheduling to insurance claims processing and payment collection. Despite its importance, many processes within RCM remain extremely inefficient and costly. One of the most inefficient processes is medical coding, a mid-revenue cycle process that involves the assignment of medical codes to patient records for reimbursement purposes.

Historically, medical coding has relied on certified medical coders to manually review patient records and select the most appropriate codes based on the information in the patient record. This approach is not only extremely time-consuming (especially for higher-complexity encounters such as surgeries or inpatient visits) but also costly, with medical coding being identified as a significant contributor to the $256.6 billion spent each year on administrative complexity (1, 2). Today, a shortage of medical coders, thin operating margins, and seemingly endless guideline updates have exacerbated the challenges associated with manual medical coding and have revenue cycle leaders turning to technology for a solution (3, 4).

However, developing technology that can drive meaningful improvements to medical coding efficiency, cost, and accuracy is much easier said than done due largely to the complexity of clinical language.

The language used by physicians to document a patient's medical history, procedure notes, and other written sections of the medical record is full of linguistic nuance (negation, subjectivity, etc.), medical terminology, abbreviations, context-dependent information, and duplicate content that present challenges for traditional natural language processing (NLP) technology.

This is illustrated by computer-assisted coding (CAC), a category of medical coding software that leverages traditional NLP to identify key terms in medical documentation and suggest medical codes based on those terms. The suggested codes must then be validated by certified medical coders before the encounter is sent to billing.

Today, a shortage of medical coders, thin operating margins, and seemingly endless guideline updates have exacerbated the challenges associated with manual medical coding and have revenue cycle leaders turning to technology for a solution.

While CAC is an improvement upon manual coding, the NLP technology powering these solutions mostly uses less sophisticated approaches like synonym search and dictionaries. This approach does not enable CAC solutions to understand the context of the patient encounter or concepts such as negation and temporality. As a result, the suggested codes may not accurately reflect the patient encounter, thereby hindering the expected efficiency gains. As noted in a 2019 literature review, "productivity impacts [of CAC systems] vary widely, depending on the specific deployment.

Some studies reported a drop in productivity when medical coders were forced to validate, and frequently eliminate, a large number of suggested codes" (5).

Today, advancements in AI, specifically in large language models (LLMs), hold promise for improving the accuracy of medical coding automation solutions. However, a recent study evaluated the medical code–generating performance of available LLMs (GPT-3.5, GPT-4, Gemini Pro, and Llama2-70b Chat) and found that no model had an exact match rate for generated ICD-9-CM, ICD-10-CM, and CPT codes above 50 percent (the accuracy standard in medical coding is 95 percent or higher). The researchers concluded that LLMs alone do not have the level of sophistication required to accurately assign codes to patient encounters, and are therefore unsuitable for medical coding as they exist today (6).

That's where autonomous medical coding comes in. The latest advancement in medical coding technology, autonomous medical coding solutions leverage multiple subfields of artificial intelligence (AI) such as natural language processing, machine learning, and deep learning to instantly and accurately assign codes to patient encounters. These solutions are designed specifically for medical coding, with the goal of coding patient encounters and sending them directly to billing without the need for certified coders to validate the solution's output.

Autonomous medical coding has been very well-received by health systems, hospitals, and other healthcare organizations, and interest continues to increase year after year. A recent survey revealed that around 60 percent of healthcare organizations either use autonomous coding or plan to.

Additionally, the study found that half of the respondents who plan to incorporate the technology intend to adopt a solution within six to 12 months (7).

The following sections provide a detailed overview of Nym's differentiated approach to autonomous coding, including a breakdown of Nym's medical coding engine and the countless benefits that it drives for health systems, hospitals, and provider groups.

By truly automating medical coding with an autonomous coding solution, healthcare organizations can expect to see significant improvements in coding speed, quality, cost, and operational efficiency.

Nym's Autonomous Medical Coding Engine

Nym's autonomous medical coding engine can be broken down into four components: integration, parsing, clinical language understanding, and medical code assignment. Throughout these components, the engine leverages a combination of proprietary ML algorithms based on large language models (LLMs) and rules-based clinical and medical coding ontologies which enable accurate, configurable, and scalable autonomous medical coding (see page 10 for a high-level overview of Nym's engine architecture).

Integration

The first component of Nym's engine is integration, which refers to the process used by the engine to retrieve the data in patient charts from the customer's EHR.

Nym currently leverages the latest advancement in healthcare interoperability, Fast Healthcare Interoperability Resources (FHIR), to retrieve relevant patient data from EHRs quickly, securely, and in real-time, regardless of the format in which the data is stored.

With a FHIR integration, Nym can use the same integration process across any health system, removing the need for additional engine configurations and reducing the IT lift associated with traditional implementation processes.
For customers who integrate with Nym via FHIR, the patient chart is represented in a FHIR file. For customers who are not yet integrated with FHIR, the patient chart is represented in a PDF file.

Nym currently leverages the latest advancement in healthcare interoperability, Fast Healthcare Interoperability Resources (FHIR), to retrieve relevant patient data from EHRs quickly, securely, and in real-time, regardless of the format in which the data is stored.

Parsing

The parsing component of Nym's engine is responsible for taking the patient chart, whether represented in a FHIR file or PDF file, and developing an understanding of the overall chart structure.

This means identifying the different sections of a patient chart, including narrative sections written by providers (e.g. history of present illness, procedure notes, etc.) as well as sections containing structured clinical data (e.g. medications, lab results, etc.).

To accomplish this, Nym's engine leverages logic and visual-based engines enhanced with content-aware ML models designed specifically for each major EHR (Epic, Oracle Cerner, etc.). The content-aware ML models have linguistic understanding capabilities that enable the engine to quickly identify sections based on the content, just like a medical professional can quickly tell the difference between the HPI and discharge summary by quickly looking at the content of the text.

After Nym's engine has identified the different sections of the chart, it consolidates the narrative sections and structured sections and organizes the information in a "Standard Nym Chart" which can then be processed by the engine's CLU component.

Clinical Language Understanding

Nym has developed innovative CLU technology that leverages proprietary ML models and rules-based clinical ontologies to accurately translate unstructured clinical data in patient records into a comprehensive narrative of the patient encounter.

The following sections provide a step-by-step breakdown of Nym's CLU technology, from Semantic Representation to Narrative Reconstruction.

Step 1. Semantic Representation

During this phase, Nym's engine breaks down every sentence in the chart documentation into syntactic elements (objects, nouns, verbs, etc.). This information is used to create a semantic representation that shows the relationships and logical structure of each sentence. Nym's engine then identifies the medical concepts (procedures, symptoms, etc.) within each sentence.

The semantic representation phase is powered by LLMs trained to understand natural language and trained on large amounts of medical data to extract medical concepts.

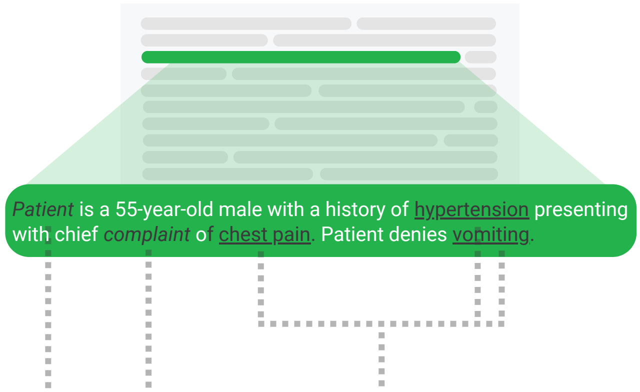

SYNTACTIC ELEMENTS
Nym's engine identifies that *"patient"* is the agent of the action *"complains."*

MEDICAL CONCEPTS
Nym's engine identifies "hypertension" "chest pain" and "vomiting" as the medical concepts in the sentence.

Figure 1. Semantic Representation

Step 2. Context Extrapolation

During this phase, Nym's engine assesses the contextual elements associated with each medical concept to understand whether concepts are negated, whether they are subjective or objective, whether a concept is a current, ongoing, or past concern (temporality), and other contextual information. This enables Nym's engine to develop an accurate understanding of each medical concept in the documentation as well as which medical concepts in the documentation should be considered for medical coding.

The context extrapolation phase is powered by multiple LLMs, each of which Nym has specifically trained to understand complex linguistic concepts like subjectivity, negation, and temporality. Each LLM uses linguistic rules to validate quality and accuracy, which is critical since coding can sometimes be associated with things the provider considered that were not necessarily true.

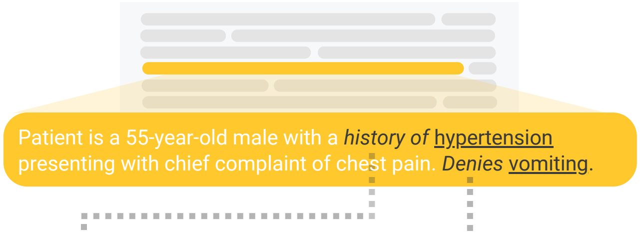

TEMPORALITY
The phrase "*history of*" indicates that the concept "hypertension" is a chronic condition (therefore it may not pertain to the current visit). The engine leverages ontologies described in the following sections to differentiate between chronic and acute conditions based on related symptoms, exams, and other factors from the documentation.

NEGATION
Because the word *"Denies"* precedes the concept "vomiting," the engine knows that "vomiting" is a negated concept (the patient did not vomit despite the entity appearing in the documentation).

Figure 2. Context Extrapolation - Negation and Temporality

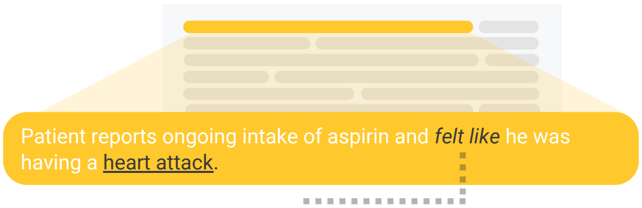

SUBJECTIVITY VS OBJECTIVITY
Because the phrase "felt like" precedes the concept "heart attack," the engine would consider "heart attack" to be a subjective concept and therefore, it would not be taken into consideration for medical code assignment.
On the other hand, an objective concept (e.g. a final diagnosis reported by the provider) would considered for medical coding.

Figure 3. Context Extrapolation - Subjectivity vs Objectivity

Step 3. Ontological Linking

By this phase, Nym's engine has a list of all the medical concepts in the clinical documentation and a contextual understanding of each concept. The next step is to link these concepts to entities in Nym's clinical ontologies.

Nym's clinical ontologies consist of a collection of entities (diagnosis, symptoms, medications, etc.) and describe the relationships between those entities.

By linking concepts to ontological entities, Nym's engine is able to develop an understanding of how the numerous medical concepts in the clinical documentation relate to one another

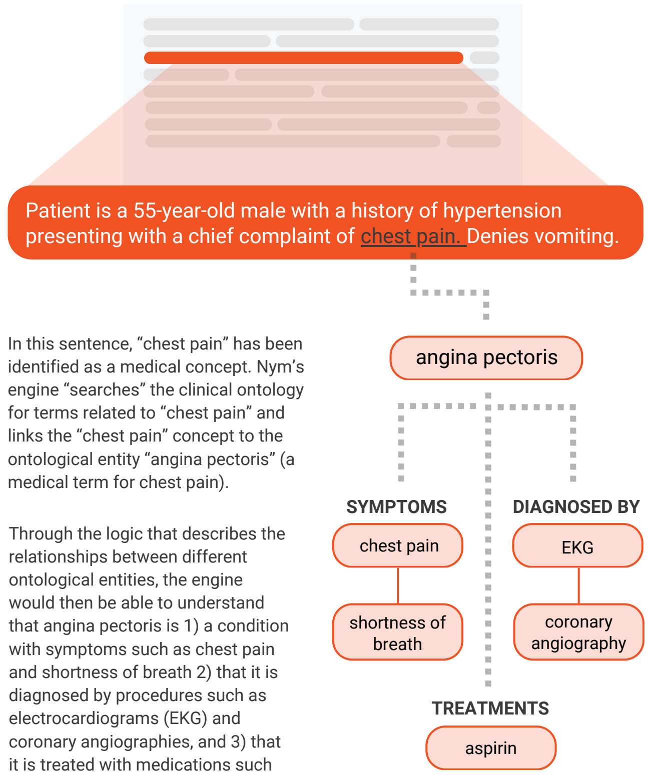

In this sentence, "chest pain" has been identified as a medical concept. Nym's engine "searches" the clinical ontology for terms related to "chest pain" and links the "chest pain" concept to the ontological entity "angina pectoris" (a medical term for chest pain).

Through the logic that describes the relationships between different ontological entities, the engine would then be able to understand that angina pectoris is 1) a condition with symptoms such as chest pain and shortness of breath 2) that it is diagnosed by procedures such as electrocardiograms (EKG) and coronary angiographies, and 3) that it is treated with medications such as aspirin.

Figure 4. Ontological Linking

Step 4. Narrative Construction

In the last phase of the CLU component, Nym's engine links together the medical concepts using the syntactic, contextual, and ontological information gained during the previous phases. The output is a comprehensive clinical picture of the patient encounter to which medical codes can be applied.

The Narrative Construction phase is powered by a combination of LLMs and clinical ontologies.

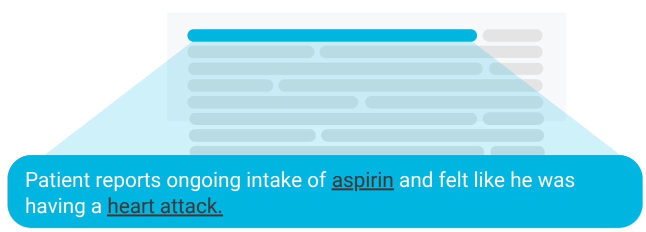

The medical concepts in this sentence are "aspirin" and "heart attack." Taking into account the syntactic, contextual, and ontological information it gained during the previous three phases, Nym's engine understands that the patient is taking aspirin as a preventative measure to reduce the risk of heart attack, which is also the reason for the patient's visit. This same process is replicated for the other medical concepts throughout the clinical documentation and pieced together to clearly and accurately describe the patient encounter.

Figure 5. Narrative Construction

If the engine does not have the highest degree of confidence with a medical concept extraction, contextual element, or other insight derived during the parsing or CLU phase, the chart is dropped by Nym's engine and returned for manual coding. The lack of the engine's confidence can be due to documentation errors, ambiguity in the text, missing information, or can be due to one of the algorithms within Nym's engine not having high confidence in its output.

Medical Code Assignment

Nym's innovative CLU technology enables Nym's engine to accurately translate the unstructured data in the patient record into a comprehensive clinical narrative. The medical coding component of the engine is up next, taking into account standard coding guidelines and customer-specific configurations to apply the most accurate, compliant medical codes to the patient encounter.

Standard Coding Guidelines

Similar to how Nym's engine leverages clinical ontologies to understand the links between different medical concepts, it also uses medical coding ontologies to link medical concepts to actual medical codes. The medical coding ontologies take into account the latest standard coding guidelines put in place by organizations like the Centers for Medicare and Medicaid (CMS), the American Medical Association (AMA), and Managed Care Plans.

When coding guidelines are updated, which can happen on an annual basis (e.g. ICD, CPT) or quarterly basis (e.g. NCCI), Nym's team makes the necessary updates to the medical coding ontologies, thereby ensuring continuous compliance. Additionally, Nym uses the latest data from the Food and Drug Administration (FDA) to ensure that all medication-related information is always up-to-date.

Custom Configuration

At Nym, we realize that health systems, hospitals and provider groups can have different interpretations of certain standard coding guidelines and therefore have different coding "philosophies."

For these reasons, Nym builds a layer of customer-specific, rules-based logic that the engine references in addition to the standard medical coding ontologies.

Multi-Specialty Capabilities

Nym's autonomous medical coding engine currently supports six different specialties and service lines across both inpatient and outpatient care, including emergency medicine, radiology, outpatient surgery, outpatient visits, inpatient professional services, and urgent care.

To support these specialties, Nym's team of medical doctors and medical coding and compliance auditors have built a clinical ontology that encompasses entities (diagnosis, symptoms, etc.) from all specialties and describes the relationships between them. The ability to leverage a single comprehensive Ontology for all specialties allows Nym to scale its CLU across specialties in an efficient and comprehensive manner.

Nym currently supports six different specialties and service lines across both inpatient and outpatient care, including emergency medicine, radiology, outpatient surgery, outpatient visits, inpatient professional services, and urgent care.

This differentiated approach to autonomous medical coding enables Nym to build capabilities for new specialty areas without needing to develop an entirely new engine. As a result of this approach, Nym is able to provide a truly multi-specialty, scalable, autonomous medical coding solution to our customers.

Architecture of Nym's Autonomous Medical Coding Engine

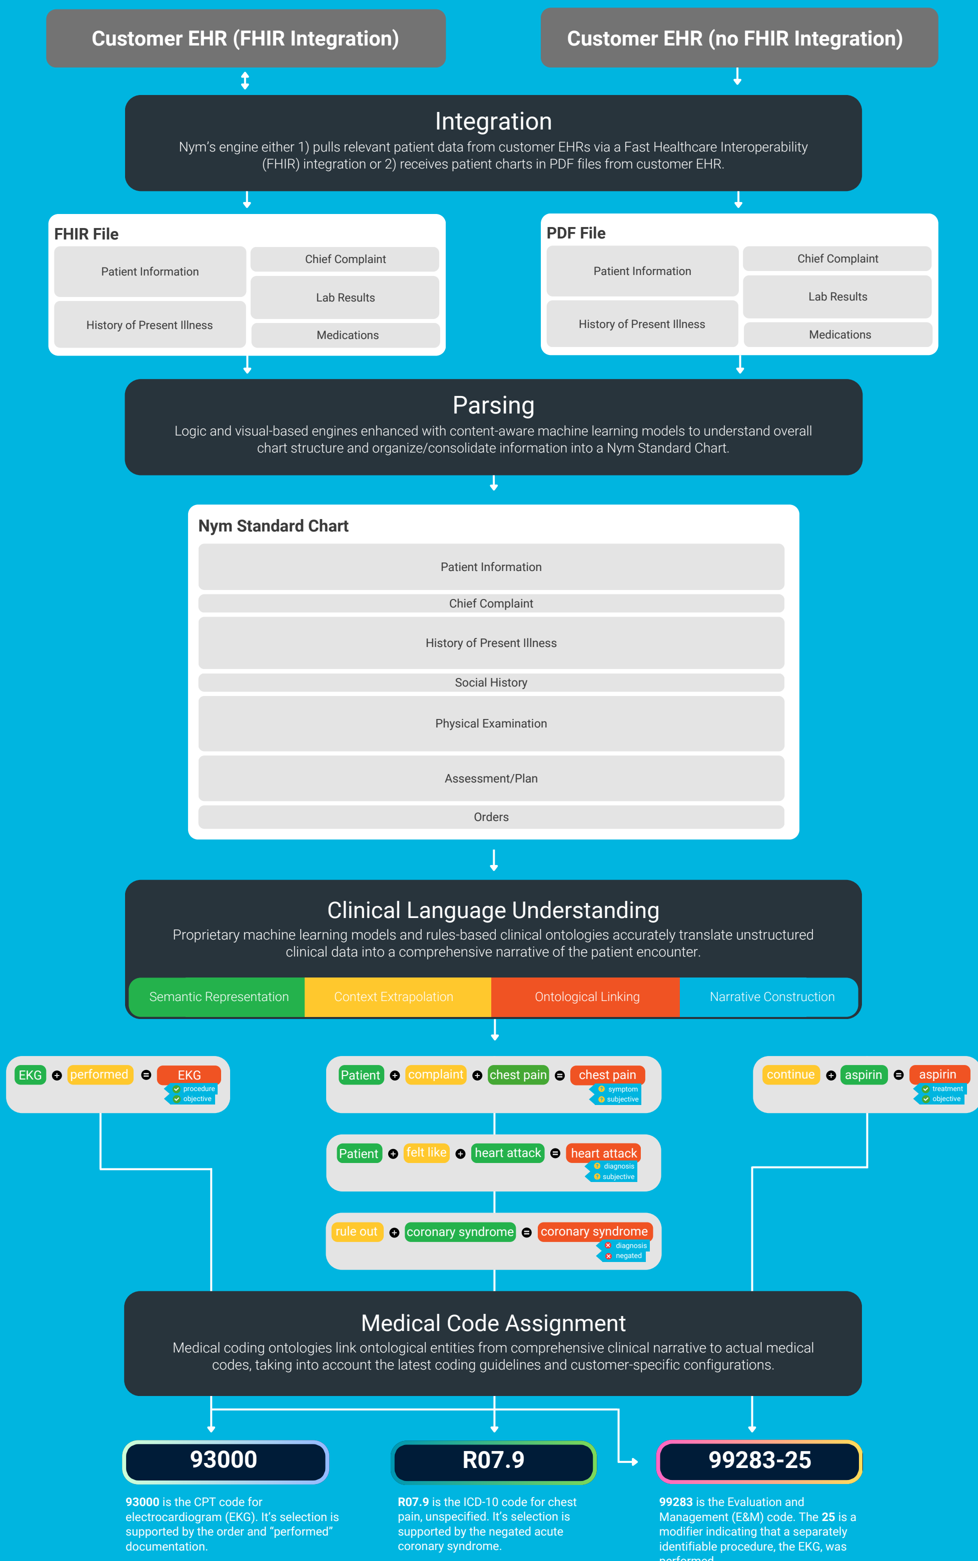

Benefits of Autonomous Coding with Nym

The underlying technology described in the previous sections enables Nym's autonomous medical coding engine to deliver exceptional coding quality and compliance, true configurability to internal coding guidelines, transparency into every coding decision made by the engine, and other operational and financial benefits.

Ensures Coding Quality

Coding Accuracy: Nym guarantees an accuracy rate of 95 percent or above for all charts coded by our engine. This ensures that healthcare providers are appropriately reimbursed for the services provided during the patient encounter and reduces the risk of coding-related denials and the associated revenue loss down the line. On the clinical side, accurate medical coding ensures that metrics like severity of illness and risk of mortality, risk adjustment factor scores, and quality measures such as Merit-Based Incentive Payment System (MIPS) are correctly reported.

Continuous Compliance: When new coding guidelines are released, Nym's team quickly assesses the updates and makes technical adjustments to the engine's CLU and medical coding components (where applicable) to reflect those updates. This ensures that Nym's customers are always compliant with the most recent guidelines, thereby reducing the risk of denials, payer takebacks post-audit, and concerns related to fraud, waste, or abuse of coding practices.

Configurability: Nym configures a unique version of our engine for every customer to align with internal guidelines and SOPs. This ensures that the coding output remains consistent during the transition from manual coding to autonomous coding.

Maximizes Audit Readiness

While other autonomous coding solutions rely heavily on machine learning for code assignment and may struggle to explain why certain medical codes were selected, Nym's engine leverages rules-based ontologies that provide complete transparency into coding rationale.

This is reflected in Nym's audit trails, which are produced for every patient encounter coded by the engine and include information such as the supporting documentation associated with each code assigned, links to the specific guidelines referenced by the engine during code selection, and much more. Nym's audit trails provide our customers with a comprehensive, actionable resource that they can use in the event of an audit, denial, or other matter related to compliance.

Accelerates Payment Cycles

Nym's engine codes patient encounters in a matter of seconds as opposed to the minutes it takes a medical coder to manually assign code (for example, discharged emergency department visits take an average of 32 minutes to manually code) (2). By leveraging Nym's engine, our customers significantly reduce the time in accounts receivable (A/R) by an average of 5 days and can resolve coding backlogs overnight.

Reduces Coding-Related Costs

Nym's engine reduces coding-related costs for our customers by up to 35 percent. These cost-savings are driven by many mechanisms, namely the reduction in the number of full-time medical coder hours required (FTE savings), the lowered risk of denials and associated revenue loss, and the decreased days in A/R that could lead to accumulated interest.

Enables Scalability

Nym's engine can easily support increases in patient volume, whether a result of seasonal increases, new facility expansions, etc., by simply increasing compute power. This is an improvement upon past solutions such as computer-assisted coding (CAC), which provide limited scalability since they still require coders to validate every code before an encounter is sent to billing.

References

1. Shrank, WH., Rogstad, TL., Parekh N. (2019). Waste in the US Health Care System: Estimated Costs and Potential for Savings. JAMA. 322(15):1501–1509. https://doi:10.1001/jama.2019.13978
2. Tseng P., Kaplan, R. S., Richman B. D., Shah M. A., Schulman K. A. (2018). Administrative Costs Associated With Physician Billing and Insurance-Related Activities at an Academic Health Care System. JAMA. 319(7):691–697. https://doi:10.1001/jama.2017.19148
3. Lubell, J. (2023, April 19). Addressing another health care shortage: medical coders. The American Medical Association. https://www.ama-assn.org/about/leadership/addressing-another-health-care-shortage-medical-coders
4. PwC Health Research Institute. (2023). Medical cost trend: Behind the numbers 2024. https://www.pwc.com/us/en/industries/health-industries/library/assets/pwc-behind-the-numbers-2024.pdf
5. Stanfill, M. H., & Marc, D. T. (2019). Health Information Management: Implications of Artificial Intelligence on Healthcare Data and Information Management. Yearbook of medical informatics, 28(1), 56–64. https://doi.org/10.1055/s-0039-1677913
6. Soroush, A. et. al., (2024). Large Language Models Are Poor Medical Coders — Benchmarking of Medical Code Querying. NEJM AI, 1(5). https://ai.nejm.org/doi/abs/10.1056/AIdbp2300040
7. https://www.hfma.org/technology/why-autonomous-coding-is-having-a-moment-in-healthcare/