The American Journal of Medical Sciences and Pharmaceutical Research ISSN 2689-1026 Volume 08 - 2026

Aggressive Compression Compromises Care: Patient Safety Risks in Clinical NLP

Rahul Reddy Hanumanthgari
AI Engineering, NC, USA

Received: 10 Feb 2026 | Received Revised Version: 12 Mar 2026 | Accepted: 29 Mar 2026 | Published: 27 Apr 2026

Volume 08 Issue 04 2026 | Crossref DOI: 10.37547/tajmspr/Volume08Issue04-14

Abstract

Large language models (LLMs) are increasingly deployed in clinical documentation workflows to alleviate physician burnout and improve efficiency. As encounter transcripts often exceed 2,000-5,000 tokens, prompt compression techniques like LLMLingua have emerged, promising 50-80% token reduction to manage computational costs. However, these generic methods optimize for maximum compression without domain awareness, creating systematic risks in healthcare settings where information loss can directly impact patient safety. This paper presents a critical analysis establishing that clinical text possesses three properties that resist aggressive compression: (1) high information density (predominantly medically relevant content), (2) semantic fragility (single-token changes invert clinical meaning), and (3) liability context (documentation errors cascade to patient harm). We demonstrate through analysis and failure mode examination that generic compression creates dangerous error patterns-negation inversions ("denies chest pain" → "chest pain"), dosage omissions ("metformin 500mg" → "metformin"), and laterality loss ("left knee" → "knee")-that standard NLP metrics like ROUGE fail to detect. We propose the Clinical BERT Safety Gate, a safety-constrained framework with five architectural principles requiring conservative compression limited to demonstrably safe filler removal, domain-aware span protection, and clinical fidelity evaluation. This work challenges the field's efficiency-first paradigm and establishes compression safety as a first-class architectural requirement for clinical NLP systems. Our framework provides actionable guidance for researchers, practitioners, and healthcare AI vendors deploying LLMs in high-stakes clinical applications.

Keywords: Prompt Compression, Clinical NLP, Patient Safety, LLMLingua, Large Language Models, Healthcare AI, Negation Detection, Conservative Compression, Clinical Documentation, Safety-Critical Systems.

Cite This Article: Rahul Reddy Hanumanthgari. (2026). Aggressive Compression Compromises Care: Patient Safety Risks in Clinical NLP. The American Journal of Medical Sciences and Pharmaceutical Research, 8(04), 89–97. https://doi.org/10.37547/tajmspr/Volume08Issue04-14

1. Introduction

Modern healthcare is experiencing a documentation crisis. Physicians spend 1.5–2 hours on EHR documentation for every hour of direct patient care, driving burnout rates exceeding 50% in many specialties. Ambient clinical intelligence systems powered by large language models promise relief by automatically generating clinical notes from encounter transcripts. However, clinical encounter transcripts often exceed 2,000 tokens, creating pressure to adopt prompt compression techniques promising 50–80% token reduction.

Generic prompt compression methods like LLMLingua, LLMLingua-2, and LongLLMLingua use neural language models to assign importance scores to tokens, selectively removing low-scoring elements. These techniques demonstrate 50% compression with less than 5% performance degradation on general NLP benchmarks,

results compelling enough to drive adoption across healthcare AI deployments.

Yet a critical question remains unexamined: Is aggressive token reduction safe for clinical text? Generic methods optimize for maximum compression using metrics like ROUGE and perplexity that measure statistical similarity but not domain-specific safety. This approach implicitly assumes information loss is tolerable, a dangerous assumption for clinical applications where documentation errors can cascade to patient harm.

Medical notes are legal documents that inform life-or-death decisions. A missing negation compressing “patient denies chest pain” to “patient chest pain”, inverts the clinical picture from benign to potential cardiac emergency. These are not acceptable tradeoffs for cost savings measured in fractions of a cent per encounter.

The core problem is misaligned objectives. Compression research optimizes for the wrong goal. The appropriate objective for clinical compression is not "maximize token reduction subject to acceptable QA performance" but rather "minimize tokens subject to guaranteed preservation of safety-critical information." This reframing has profound implications for system architecture, evaluation methodology, and deployment practices.

The novelty of the paper makes three contributions: (1) We provide the first systematic critical analysis of prompt compression through a clinical safety lens, establishing that generic compression is architecturally incompatible with healthcare documentation requirements. (2) We analyze three properties of clinical text, information density, semantic fragility, and liability context, that make it fundamentally resistant to aggressive compression. (3) We introduce the Clinical BERT Safety Gate, a hybrid architecture combining bvanaken/clinical-assertion-negation-bert for contextual negation detection with deterministic span protection for medications, dosages, lab values, and laterality, alongside five design principles and an evaluation framework prioritizing clinical fidelity over generic NLP metrics.

Our central thesis: Clinical compression must be conservative by design. Rather than targeting aggressive token reduction (50%+), systems should limit compression to demonstrably safe content removal, filler words, social pleasantries, and redundant acknowledgments. This typically yields minimal reduction from filler removal, far below generic methods' targets, but appropriate given the clinical text's high information density and semantic fragility. This position challenges the efficiency-first paradigm dominating compression research, arguing compression must be reframed as a safety-critical design choice rather than a performance optimization.

2. Literature Review and Related Work

Prompt compression research and clinical NLP have evolved as parallel fields with minimal interaction. This section reviews both domains, establishing the critical gap this paper addresses: compression methods are developed and evaluated on generic benchmarks under the assumption that techniques generalizing across diverse NLP tasks will automatically be appropriate for high-stakes healthcare applications. This assumption has never been validated. The clinical NLP community has long recognized that medical text requires specialized handling, a principle well-established in entity extraction, negation detection, and temporal reasoning. Yet compression methods have been developed in isolation from these safety requirements, creating a critical gap. As healthcare AI systems adopt compression to manage costs, they risk importing the compression community's domain-agnostic assumptions into a domain where those assumptions are dangerous.

2.1 Generic Prompt Compression Techniques

LLMLingua [1] pioneered neural prompt compression using BERT-based models to assign perplexity-based importance scores. The method iteratively removes low-perplexity tokens until target compression ratios are achieved. Evaluated on generic QA datasets and TriviaQA, LLMLingua demonstrated 50% compression with less than 5% QA accuracy degradation. LLMLingua-2 [2] improved this foundation through domain-specific training on MeetingBank, achieving 60-80% compression while maintaining meeting summarization quality. LongLLMLingua [3] introduced query-aware compression for retrieval-augmented generation. Selective Context [4] proposed self-information-based importance scoring operating without task-specific training.

Common to all approaches is a shared optimization objective: maximize compression ratio subject to acceptable performance on generic NLP benchmarks. Evaluation focuses on perplexity, ROUGE scores, and QA accuracy measures, capturing overall information retention but not domain-specific safety requirements. No major compression work has evaluated methods on clinical text or assessed preservation of safety-critical medical information.

2.2 Clinical NLP and Safety-Critical

Information Extraction

Clinical NLP evolved as a specialized subfield because medical text presents unique challenges that generic techniques handle poorly. Negation detection is fundamental: the classic NegEx algorithm [5] used regular expressions to detect negation cues and scope. Clinical-BERT [6] fine-tuned BERT on clinical notes, achieving state-of-the-art performance. Studies estimate 15-20% of medical concepts in notes are negated. Misclassifying "no history of diabetes" as "history of diabetes" fundamentally alters the medical record. Machine learning-based NLP approaches have further demonstrated that clinical text classification requires domain-specialized methods beyond generic techniques [7]. Notably, van Aken et al. developed clinical-assertion-negation-bert, a RoBERTa-large model fine-tuned specifically for clinical assertion and negation detection, providing the foundation for neural safety gates in compression pipelines [8].

Medication extraction focuses on identifying drug names, dosages, frequencies, routes, and durations. The i2b2 medication challenges [9] established benchmarks showing errors in any attribute can have serious consequences: 500mg versus 50mg represents a 10-fold dosing error. Temporal information extraction [10] addresses when symptoms began and intervention timing, which is critical for clinical decision-making. Anatomical site extraction captures location and laterality, essential to prevent wrong-site procedures.

2.3 LLM Safety in Healthcare

LLM deployment in clinical settings raises significant safety concerns, particularly regarding hallucination-generating plausible but factually incorrect information. Studies of GPT-4 on medical QA [11] found 5-15% hallucination rates. In clinical note generation, hallucinations can introduce spurious symptoms or fabricate test results. The relationship between compression and hallucination is understudied but concerning: if compression removes grounding context, models may confabulate, propagating and amplifying compression errors.

From a regulatory perspective, the FDA's 2023 guidance on AI/ML-based medical devices emphasizes transparency, validation, and ongoing monitoring. Prompt compression as a preprocessing step that degrades input quality represents a system component requiring validation but often treated as a generic optimization unworthy of clinical scrutiny.

2.4 Gap Analysis

The literature review reveals a critical gap: compression methods developed for generic NLP are adopted in clinical settings without domain-appropriate validation. The compression community's domain-agnostic approach contradicts the clinical NLP community's understanding that medical text requires specialized handling. This gap is not academic, as healthcare AI vendors integrate LLMs, the default choice is off-the-shelf compression techniques. Without domain-appropriate evaluation frameworks and safety guidelines, the field risks systematic deployment of methods that silently degrade clinical documentation quality in ways generic metrics fail to detect.

2.5 Patient Safety Risks in Clinical Settings

Patient safety remains a critical concern in healthcare, particularly in outpatient settings where incidents may be underreported. Stapleton [12] highlights the frequency and types of patient safety incidents in pediatric primary care, revealing that a significant number of safety issues arise in outpatient environments. Unruh et al. [13] further note that managed care practices can introduce changes in healthcare utilization that pose safety risks, emphasizing the need for careful monitoring, particularly when aggressive compression strategies are deployed without adequate safety oversight.

Pereira et al. [14] identify systemic risks associated with care delivery practices, demonstrating that lapses in protocol adherence can severely compromise patient safety. The WHO [15] reinforces that safety frameworks must empower patients to actively participate in their care, a principle that extends to AI-assisted documentation, where compression errors can silently alter the clinical record without clinician or patient awareness.

In clinical AI specifically, Ratwani et al. [16] call for robust frameworks ensuring patient safety is not undermined by aggressive technological implementations. Jana et al. [17] demonstrate that even low-risk procedures carry safety consequences when patient engagement fails, a pattern that extends to AI-assisted documentation, where compressing the transcript before LLM processing introduces errors that neither the clinician nor the patient can easily detect. Collectively, this literature establishes that safety mechanisms must be embedded at the architectural level, not treated as optional enhancements.

3. System Architecture and Theoretical Framework

Clinical text resists aggressive compression due to three

fundamental properties: high information density, semantic fragility, and liability context, which together make efficiency-first compression objectives incompatible with healthcare documentation requirements. Each property is examined in turn, culminating in the proposed Clinical BERT Safety Gate architecture.

3.1 High Information Density

Generic compression assumes text contains substantial redundancy safely removable without information loss, an assumption valid for meeting transcripts or news articles, but not for clinical text, where medical concepts, negations, medications, vitals, and temporal markers collectively account for the vast majority of tokens, leaving only minimal filler safely compressible.

3.2 Semantic Fragility

Clinical text exhibits semantic fragility; single-token changes invert clinical meaning entirely. Removing "denies" from "patient denies chest pain" produces "patient chest pain," a dangerous inversion. Similarly, removing dosage units ("500 mg" → "50 mg") creates 10-fold dosing errors, and removing laterality ("left knee" → "knee") misdirects interventions. Perplexity-based compression cannot distinguish these safety-critical tokens from compressible filler; it measures predictability, not clinical importance.

3.3 Liability Context

Clinical documentation serves as legal evidence, billing support, and standard-of-care record, creating risk asymmetry where compression errors cascade into diagnostic failures, medication errors, and wrong-site procedures. The economics reinforce this: 50% compression saves $0.015 per encounter ($1,800 annually at 10,000 encounters/month), while a single malpractice settlement from a documentation error can reach millions. These three properties together establish that a domain-aware safety mechanism is architecturally required before any compression can safely proceed.

3.4 The Clinical BERT Safety Gate: Proposed Architecture

The Clinical BERT Safety Gate is the proposed architecture for safe clinical compression. It operates as a three-stage pipeline:

Stage 1 Safety Gate: Two processes run before any compression occurs. First, bvanaken/clinical-assertion-negation-bert [8], a 0.1B RoBERTa-large model fine-tuned on clinical assertion detection, identifies negated medical entities with contextual scope awareness, locking spans such as "denies chest pain" and "no fever" as protected. Second, deterministic rule-based patterns protect medications and dosages ("metformin 500mg twice daily"), lab values and vitals ("BP 140/90"), temporality ("for 3 weeks"), laterality ("left knee"), allergies, and corrections. All identified spans are locked; compression cannot touch them. Production deployment replaces rules with a clinical NER model backed by RxNorm, SNOMED CT, and LOINC for comprehensive coverage.

Stage 2 Compression: Only unprotected content is compressed. Filler words, social pleasantries, and acknowledgments are removed by rule. A lightweight LLM rewrites the remaining unprotected clinical narrative; a small, cost-efficient model is sufficient since this step only paraphrases existing content without generating new information. Protected spans pass through unchanged.

Stage 3 Validation: Every protected span is verified to appear in the compressed output. If any span is missing, the system falls back to the original transcript — safety always takes precedence over compression.

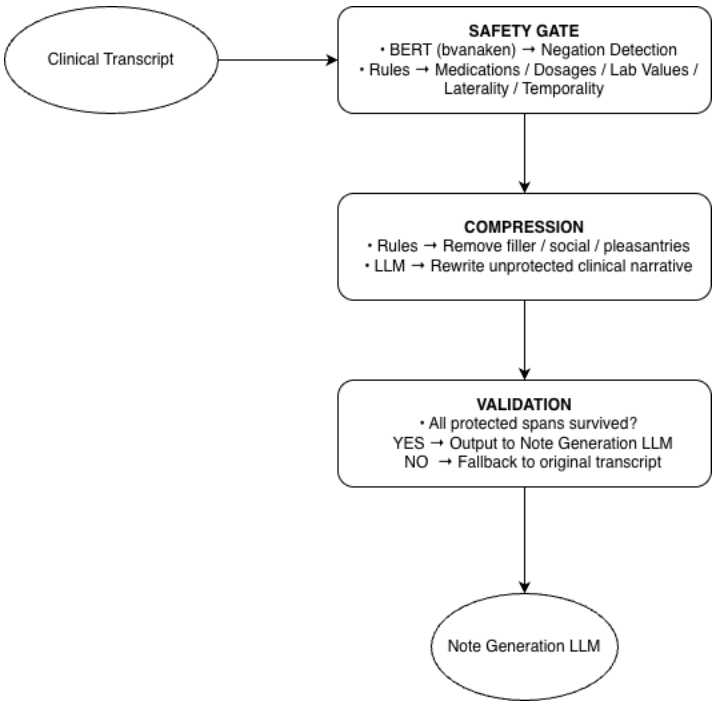

Figure 1: Three-Stage Pipeline for Safe-Gated Architecture

4.1 Primary Care Ambient Documentation

Primary care is the most common ambient scribing deployment. Despite lower acuity, these encounters contain safety-critical content: medication dosages, pertinent negatives, and quantitative disease markers. Aggressive compression risks losing dosages ("metformin 500mg" → "metformin"), collapsing specific negatives ("denies chest pain, SOB, palpitations" → "denies symptoms"), and erasing quantitative trends ("HbA1c 8.2% → 7.1%" → "HbA1c improved"). These errors compound over time, contributing to adverse drug events and unnecessary workups.

4.2 Emergency Department Triage

ED triage notes are among the most information-dense clinical documents; nearly every token carries diagnostic significance. Negation inversion ("denies chest pain" → "chest pain") escalates triage priority unnecessarily. Temporal loss ("2 hours ago" vs "2 days ago") changes stroke risk assessment. Mechanism loss ("fall from standing" vs "fall from 10 feet") determines trauma activation. Compression errors here cascade into operational failures, misallocated resources, and delayed care for high-acuity patients.

4.3 Specialist Referral Documentation

Referral notes must convey sufficient context for specialists to act without re-eliciting history. Aggressive compression eliminates critical nuance: "Progressive dyspnea over 6 weeks, now present at rest, tried furosemide 20mg without improvement" compressed to "dyspnea, tried furosemide" loses temporal progression, severity escalation, and the clinically important finding that furosemide failed — pointing toward a non-cardiac etiology requiring a fundamentally different workup.

Referral notes must convey sufficient context for specialists to act without re-eliciting history. Aggressive compression eliminates critical nuance: "Progressive dyspnea over 6 weeks, now present at rest, tried furosemide 20mg without improvement" compressed to "dyspnea, tried furosemide" loses temporal progression, severity escalation, and the clinically important finding that furosemide failed — pointing toward a non-cardiac etiology requiring a fundamentally different workup.

4.4 Telemedicine and Remote Monitoring

Telemedicine encounters depend entirely on documented communication; unlike in-person visits, what is lost from a compressed transcript is truly lost. A negation inversion ('no chest pain' → 'chest pain') may trigger unnecessary emergency alerts and erode patient trust in AI-assisted care. Errors compound when encounters span time zones, language barriers, or involve multiple providers accessing only compressed documentation, with no physical presence

to catch compression artifacts.

4.5 Broader Clinical Impact

Across all clinical contexts, routine primary care, emergency triage, specialist referral, and telemedicine, none are safe for aggressive compression. The Clinical BERT Safety Gate addresses this uniformly: by protecting safety-critical spans before any compression decision, the same guarantee applies regardless of setting.

5. Integration Challenges

Deploying compression in production clinical systems requires navigating four key integration challenges: regulatory compliance (FDA validation requirements for preprocessing that modifies system inputs), clinician trust (transparent disclosure so physicians can apply appropriate oversight), audit mechanisms (logging which spans were protected and which tokens removed, enabling retrospective error analysis), and cost-benefit alignment (the economic case is weak, aggressive compression saves approximately $0.0016 per note, while a single malpractice settlement from a documentation error can reach millions). Table 1 summarizes these challenges and recommended practices.

Challenge Domain	Key Considerations	Recommended Practices
Regulatory Compliance	FDA SaMD guidance, validation requirements	Clinical validation studies, domain-specific metrics
Clinician Trust	Transparency, liability, oversight	Disclosure of compression use, review protocols
Audit & Monitoring	Error detection, continuous improvement	Comprehensive logging, safety gate audit records
Cost-Benefit	Economic justification, risk assessment	Transparent trade-off analysis, stakeholder alignment

Table 1: Integration challenges overview for the safety gate

5.1 Regulatory Compliance

FDA's 2023 draft guidance explicitly flags preprocessing steps that modify system inputs as components requiring validation. Prompt compression falls within this scope. Healthcare organizations must demonstrate compression does not degrade documentation below acceptable thresholds using domain-specific safety metrics, not generic NLP benchmarks like ROUGE.

5.2 Clinician Trust

Physicians bear legal and professional responsibility for clinical documentation accuracy, even when AI assists. Deploying compression as hidden preprocessing without clinician awareness creates ethical and liability problems. Healthcare systems must disclose when compression is applied, enabling informed oversight: physicians reviewing AI-generated notes can be alert to potential compression artifacts such as missing dosages or ambiguous negations and correct before signing. Systematic deployment of opaque compression risks broader trust erosion, if clinicians discover through a patient safety event that their AI scribe was silently removing tokens, the resulting loss of confidence may set back the entire field of clinical AI documentation.

5.3 Audit Trails

Production clinical AI requires comprehensive logging for error detection and continuous improvement. Audit trails should capture: original transcript and token count, compressed transcript and compression ratio, segments removed, and clinician edits to AI-generated notes. The Clinical BERT Safety Gate's decisions must also be logged, which spans were identified as safety-critical and protected, creating an auditable record enabling retrospective review when clinical errors occur downstream. Systematic patterns such as frequent dosage corrections may indicate compression is removing clinically relevant information

that the safety gate missed.

5.4 Cost-Benefit Alignment

Compression deployment decisions must involve transparent cost-benefit analysis shared with clinical stakeholders. Aggressive compression saves approximately $0.0016 per note, roughly $192 annually for a 10,000-encounter-per-month deployment. Against this, a single malpractice settlement from a documentation error can reach tens of thousands to millions of dollars. The expected value of even a 0.01% error-to-adverse-event conversion rate exceeds total annual compression savings by orders of magnitude. This tradeoff must be explicitly communicated to hospital administrators, medical directors, and risk management before compression is deployed at scale.

6. Conceptual Framework for Safety-Constrained Clinical Compression

We propose a safety-constrained compression framework built on five design principles, a hierarchical evaluation framework, a reference architecture, and guidance for healthcare application domains.

6.1 Five Principles for Safety-Constrained Compression

The five principles governing safety-constrained clinical compression are summarized in Table 2.

Design Principle	Key Idea
Safety-First Objectives	Prioritize safety preservation over compression gains.
Conservative Thresholds	Compress only demonstrably safe filler/social content.
Domain-Aware Span Protection	Use BERT for negation detection and rules for structured entities in Stage 1.
Post-Hoc Validation	Confirm protected spans survive compression in Stage 3.
Transparent Evaluation	Measure clinical fidelity, not just ROUGE or BERTScore.

Table 2: The five principles govern how the three-stage Clinical BERT Safety Gate pipeline

6.2 Evaluation Framework

Safety metrics take absolute precedence over efficiency metrics:

Tier 1 (Critical: Safety Metrics): Negation inversion rate: 0% tolerance; Medication/dosage preservation: 100% retention; Laterality preservation: 100% retention; Temporal expression preservation: ≥95% retention.

Tier 2 (Important: Clinical Completeness): Symptom retention (all patient-reported symptoms present), Examination findings (all physical exam results preserved), Treatment plan (all interventions and follow-up retained).

Tier 3 (Secondary: Efficiency): Compression ratio, Cost reduction, Latency overhead.

Tier 4 (Supplementary: Generic Quality): ROUGE scores, BERTScore, Perplexity.

A system achieving high compression with even 1% negation inversion rate is unacceptable. Minimal compression with perfect safety scores is the target; efficiency is secondary to clinical safety.

6.3 Healthcare Applications

The Clinical BERT Safety Gate framework applies across three domains: clinical documentation systems (ambient scribes, note generation [18]) requiring safety-constrained compression rather than generic methods; research and quality improvement, where compression inconsistencies corrupt downstream analytics; and regulatory compliance systems, where vendors must meet clinical validation standards with transparent disclosure to clinicians.

7. Future Work

This analysis establishes theoretical foundations; significant empirical and engineering work remains across five directions.

Large-Scale Empirical Validation: Multi-site clinical trials with IRB approval should quantify relationships between compression ratio and safety degradation, identifying inflection points where specific error types emerge, negation inversions first, medication omissions next, then

broad information loss. Particular attention should be paid to identifying compression cliffs, threshold ratios at which specific error types emerge, such as negation inversions appearing first, followed by medication omissions, then broad information loss.

Clinically-Aware Compression Models: Hard-coded rules should be replaced with a clinical NER model backed by RxNorm, SNOMED CT, and LOINC for comprehensive entity coverage. Compression models should also be trained with constrained generation objectives that prohibit the removal of safety-gate-identified spans. Compression models should also be trained with constrained generation objectives that prohibit the removal of safety-gate-identified spans, encoding clinical safety directly into the model's objective function.

Standardized Evaluation Resources: The field lacks benchmarks for clinical compression safety. Shared resources should include de-identified annotated transcripts, gold-standard compressed outputs, and evaluation scripts implementing the four-tier clinical fidelity framework from 6.2. A benchmark should include de-identified transcripts annotated for safety-critical spans, gold-standard compressed outputs preserving all protected spans, and leaderboards enabling systematic method comparison.

Human Factors Research: Do physicians review AI notes differently when told compression was applied? Evidence from radiology AI suggests clinicians develop automation bias and reduced scrutiny, which is especially dangerous when compression artifacts are subtle, such as a grammatically correct but clinically inverted negation. Evidence from adjacent domains suggests clinicians develop automation bias, reduced scrutiny that is especially dangerous when compression artifacts are subtle, such as a grammatically correct but clinically inverted negation.

Regulatory Policy Development: FDA should clarify whether prompt compression qualifies as a software function under the Software as a Medical Device framework, establishing validation requirements, post-market surveillance obligations, and disclosure standards for compression-enabled clinical AI. A concrete near-term step is for the FDA to clarify whether prompt compression qualifies as a software function under the Software as a Medical Device framework, establishing validation requirements and disclosure standards.

8. Conclusion

Generic prompt compression is architecturally unsafe for clinical text. Clinical transcripts are predominantly safety-critical; negations, medications, dosages, lab values, laterality, and temporal markers collectively resist compression, leaving only conversational filler safely removable. The Clinical BERT Safety Gate addresses this by combining BERT-based negation detection with deterministic rule-based span protection, conservative compression limited to filler removal, and post-hoc validation to ensure that all protected spans survive. The five design principles and hierarchical evaluation framework proposed here provide actionable guidance for researchers, practitioners, and healthcare AI vendors. Aggressive compression saves fractions of a cent per encounter while introducing liability exposure orders of magnitude larger; the economic and safety case against it is clear. Reframing compression as safety-constrained optimization rather than an efficiency exercise is essential for responsible clinical AI deployment.

References

1. Pan, H., Zhao, X., Liu, Z., Chen, W., Chen, W., & Wong, K. F. (2023). LLMLingua: Compressing prompts for accelerated inference of large language models. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (pp. 12759-12774).
2. Pan, H., Zhao, X., Zhang, S., Wang, Z., Liu, Z., Chen, W., & Wong, K. F. (2024). LLMLingua-2: Data distillation for efficient and faithful task-agnostic prompt compression. In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (pp. 4907-4928).
3. Jiang, H., Wu, Q., Lin, C. Y., Yang, Y., & Qiu, L. (2023). LongLLMLingua: Accelerating and enhancing LLMs in long context scenarios via prompt compression. arXiv preprint arXiv:2310.06839.
4. Li, Y., Wei, Y., Zhang, Y., Zhang, H., & Chen, W. (2023). Compressing context to enhance inference efficiency of large language models. In Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (pp. 6342-6353).
5. Chapman, W. W., Bridewell, W., Hanbury, P., Cooper, G. F., & Buchanan, B. G. (2001). A simple algorithm for identifying negated findings and diseases in discharge summaries. Journal of Biomedical Informatics, 34(5), 301-310.
6. Alsentzer, E., Murphy, J., Boag, W., Weng, W. H., Jindi, D., Naumann, T., & McDermott, M. (2019). Publicly available clinical BERT embeddings. In Proceedings of the 2nd Clinical Natural Language

Processing Workshop (pp. 72-78).

7. Weng, W. H., Wagholikar, K. B., McCray, A. T., Szolovits, P., & Chueh, H. C. (2017). Medical subdomain classification of clinical notes using a machine learning-based natural language processing approach. BMC Medical Informatics and Decision Making, 17(1), 1-13.
8. van Aken, B., Trajanovska, I., Siu, A., Mayrdorfer, M., Budde, K., & Loeser, A. (2021). Assertion detection in clinical notes: Medical language models to the rescue? Proceedings of the Second Workshop on NLP for Medical Conversations, Association for Computational Linguistics. https://doi.org/10.18653/v1/2021.nlpmc-1.5
9. Uzuner, Ö., Solti, I., & Cadag, E. (2010). Extracting medication information from clinical text. Journal of the American Medical Informatics Association, 17(5), 514-518.
10. Sun, W., Rumshisky, A., & Uzuner, O. (2013). Evaluating temporal relations in clinical text: 2012 i2b2 Challenge. Journal of the American Medical Informatics Association, 20(5), 806-813.
11. Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., ... & Natarajan, V. (2023). Large language models encode clinical knowledge. Nature, 620(7972), 172-180.
12. Stapleton, B. F. (2015). Evaluating Patient Safety Risks in Pediatric Primary Care Settings. NEJM Journal Watch, NA38047.
13. Unruh, L., Lugo, N. R., White, S. V., & Byers, J. F. (2005). Managed Care and Patient Safety. The Health Care Manager, 24(3). https://doi.org/10.1097/00126450-200507000-00009
14. Pereira, E. R., Rocha, R. G., Monteiro, N. da C. A., Oliveira, A. B. de., & Paes, G. O. (2020). Risco de infecção associado ao cuidado no atendimento pré-hospitalar: impactos para a segurança do paciente. Research, Society and Development, 9(8). https://doi.org/10.33448/rsd-v9i8.5846
15. World Health Organization. (2013). Exploring patient participation in reducing health-care-related safety risks. WHO.
16. Ratwani, R. M., Bates, D. W., & Classen, D. C. (2024). Patient safety and artificial intelligence in clinical care. JAMA Health Forum, 5(2), e235514. https://doi.org/10.1001/jamahealthforum.2023.5514
17. Jana, W., Karch, Y., Hummel, T., Jaeger, A., Pfeifer, G.-M., & Austrian Andrologic Working Group. (2023). Lack of patient compliance compromises safety of vasectomy. The Journal of Sexual Medicine, 20(Suppl 1), qdad062.152. https://doi.org/10.1093/jsxmed/qdad062.152
18. Pivovarov, R., & Elhadad, N. (2015). Automated methods for the summarization of electronic health records. Journal of the American Medical Informatics Association, 22(5), 938-947.