

Hospital Readmission Predictive Modeling

Sofia Schmidt¹, Nina Petrov²

¹ Assistant Professor, Institute of Intelligent Systems, Mediterranean Institute of Technology, Rome, Italy. Email: sofia.schmidt974@yahoo.com | ORCID: 6187-2252-1887-3944

² Professor, Institute of Intelligent Systems, Mediterranean Institute of Technology, Rome, Italy. Email: nina.petrov893@yahoo.com | ORCID: 5734-1261-4808-2257

ABSTRACT

Unplanned hospital readmissions within 30 days of discharge affect 12-20% of hospitalised patients across European health systems, costing an estimated EUR 15-25 billion annually and serving as a widely used marker of healthcare quality. Despite two decades of modelling efforts, the discriminative performance of readmission prediction has plateaued at C-statistics of 0.60-0.70 -- barely better than chance -- raising questions about whether readmissions are fundamentally unpredictable or whether existing models simply use the wrong data. We analysed 124,682 hospitalisations across three Italian academic hospitals (Rome, Milan, and Naples) from January 2016 to December 2021, comparing five predictive modelling approaches: administrative claims models, clinical regression models using structured EHR data, NLP-enhanced models incorporating discharge summary text, temporal deep learning models using time-series EHR data, and multi-modal ensemble models integrating all data sources. We evaluated prediction for all-cause 30-day readmission across four high-volume conditions: heart failure (HF), pneumonia, COPD exacerbation, and hip/knee arthroplasty. Our Readmission Prediction Utility Score (RPUS) integrates discrimination, calibration, net reclassification improvement, actionability of identified risk factors, and implementation feasibility. The multi-modal ensemble achieved the highest RPUS (0.916), with C-statistic 0.78 for all-cause readmission -- a meaningful advance over administrative models (C-stat 0.62). NLP-enhanced models (RPUS 0.892) contributed the largest single-source improvement by capturing social determinants, functional status, and care plan adherence barriers from unstructured clinical notes that coded data systematically misses.

Keywords: hospital readmission; predictive modelling; machine learning; natural language processing; electronic health records; heart failure; COPD; discharge planning; healthcare quality; deep learning; clinical decision support; risk stratification

Citation: Schmidt and Petrov [2022]. Hospital Readmission Predictive Modeling. DOI: <https://doi.org/10.5281/zenodo.19612993>

Copyright: © 2022 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2021 Dec 26 Accepted: 2022 Feb 24 Published: 2022 May 03

Research Article: Research Article

1. Introduction

1.1 The Readmission Problem Nobody Has Solved

Hospital readmissions are one of those problems that everyone agrees is important but nobody seems able to fix. The numbers are remarkably stable across countries and decades: roughly 12-20% of hospitalised patients bounce back within 30 days, and about half of those readmissions are considered potentially preventable (van Walraven et al., 2011). In the US, the Hospital Readmissions Reduction Program (HRRP) began financially penalising hospitals for excess readmissions in 2012, and while this generated enormous attention and investment, the actual impact on readmission rates has been modest -- about a 1-2 percentage point reduction nationally, with ongoing debate about whether even that reflects genuine quality improvement or coding changes (Wasfy et al., 2017). In Europe, readmission rates remain stubbornly high despite various national quality improvement programmes. The fundamental challenge is that readmissions are driven by a complex tangle of clinical factors (disease severity, comorbidity burden, medication management), system factors (discharge planning quality, post-discharge follow-up availability), and social factors (social isolation, health literacy, housing stability, medication affordability) that no single intervention addresses comprehensively (Donze et al., 2016).

1.2 Why Prediction Has Been So Difficult

The modelling challenge mirrors the clinical complexity. The HOSPITAL score (Donze et al., 2013) -- one of the most widely validated readmission predictors -- achieves a C-statistic of only 0.65-0.72, which is genuinely poor for a binary classification task. The CMS readmission models used for HRRP penalties perform even worse (C-stat 0.60-0.63), partly because they deliberately use only administrative claims data to ensure fairness across hospitals with different EHR capabilities (Kansagara et al., 2011). Several meta-analyses have examined why models underperform. Kansagara et al. (2011) reviewed 26 readmission prediction models and found none with C-statistic above 0.72. Artetxe et al. (2018) reviewed ML approaches and found marginal improvements (C-stat 0.68-0.75) over traditional logistic regression. The emerging consensus is that the barrier is not algorithmic sophistication but data richness: the social and behavioural determinants that drive many readmissions are

simply not captured in structured EHR fields. Discharge summaries, nursing notes, and social work assessments contain critical information about a patient's ability to manage their care at home, but this information exists as free text, invisible to models trained on coded data. Our study tests whether NLP and multi-modal integration can break through this ceiling. Section 2 reviews the literature, Section 3 describes methods, results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Administrative and Clinical Models

The CMS readmission models (Krumholz et al., 2011) for HF, pneumonia, and AMI use administrative claims variables -- age, sex, and a curated set of comorbidity codes from the index admission and 12-month lookback -- achieving C-statistics of 0.60-0.63. These models were intentionally simple to avoid penalising hospitals that serve sicker patients, but their poor discrimination means they cannot meaningfully identify high-risk individuals for targeted intervention. The HOSPITAL score (Donze et al., 2013) added clinical variables -- haemoglobin at discharge, sodium, procedure during stay, number of admissions in the prior year, length of stay, and discharge from an oncology service -- and achieved C-stat 0.72 in the derivation cohort, though external validations consistently report 0.65-0.68. The LACE index (van Walraven et al., 2010) uses length of stay, acuity, comorbidity (Charlson), and ED visits, achieving similar performance. The recurring finding across all these models is that coded clinical data captures disease burden but misses the psychosocial context that determines whether a patient can actually manage their condition at home.

2.2 Machine Learning and NLP Approaches

ML models have been explored extensively but with mixed results. Rajkomar et al. (2018) used deep learning on the full longitudinal EHR (diagnoses, medications, lab values, vital signs, nursing flowsheets) and achieved C-stat 0.75-0.77 for readmission -- a genuine improvement over traditional models but still far from the 0.80+ threshold that clinicians typically consider useful. Huang et al. (2019) applied NLP to discharge summaries using ClinicalBERT and found that adding text-derived features improved C-statistic from 0.68 to 0.74 for HF readmission -- the largest single-source improvement in the literature. The extracted features that mattered most were

mentions of social isolation, difficulty with medications, unstable housing, and caregiver absence -- precisely the factors that structured EHR data does not capture. Shameer et al. (2017) reviewed 64 ML readmission models and found that ensemble methods (random forests, gradient boosting) consistently outperformed single algorithms. Table 1 summarises the key evidence.

2.3 The Gap

No published study has compared administrative, clinical, NLP, temporal deep learning, and multi-modal approaches head-to-head across multiple conditions in a single cohort using a comprehensive metric that includes not just discrimination but also calibration, reclassification, actionability, and implementation feasibility. That is the contribution of this work.

Table 1: Key Studies in Hospital Readmission Prediction

Study	Year	Approach	Conditions	Key Finding
Krumholz et al.	2011	CMS claims model	HF, PNA, AMI	C-stat 0.60-0.63
Donze et al.	2013	HOSPITAL score	All-cause	C-stat 0.72 derivation, 0.65-0.68 external
van Walraven et al.	2010	LACE index	All-cause	C-stat 0.68, widely validated
Kansagara et al.	2011	Systematic review	All-cause	26 models, none > 0.72
Rajkumar et al.	2018	Deep learning EHR	All-cause	C-stat 0.75-0.77
Huang et al.	2019	Clinical BERT NLP	Heart failure	C-stat 0.68 to 0.74 with text
Shameer et al.	2017	ML review	All-cause	Ensembles best (C-stat 0.68-0.75)
Artetxe et al.	2018	ML review	All-cause	Marginal improvement over logistic reg.

Study	Year	Approach	Conditions	Key Finding
Wasfy et al.	2017	HRRP evaluation	US national	1-2% readmission reduction only
Donze et al.	2016	Risk factors review	All-cause	Social factors largely unmeasured

HF = Heart Failure; PNA = Pneumonia; AMI = Acute Myocardial Infarction; CMS = Centers for Medicare and Medicaid Services; HRRP = Hospital Readmissions Reduction Program.

3. Methodology

3.1 Cohort and Data Sources

We extracted data from the EHR systems of three Italian academic hospitals: Policlinico Gemelli (Rome, 52,486 hospitalisations), Ospedale Niguarda (Milan, 42,128), and Azienda Ospedaliera dei Colli (Naples, 30,068) from January 2016 to December 2021. We focused on four high-volume, high-readmission conditions: heart failure (n = 28,462, readmission rate 22.4%), pneumonia (n = 32,186, rate 14.8%), COPD exacerbation (n = 26,842, rate 19.6%), and elective hip/knee arthroplasty (n = 37,192, rate 5.8%). Readmission was defined as any unplanned acute inpatient admission within 30 days of discharge from the index hospitalisation, to any hospital (not just the discharging one), captured through regional health information exchanges. Structured EHR data included demographics, diagnoses (ICD-10), procedures, medications, laboratory values, vital signs, and prior utilisation. Unstructured data included discharge summaries, nursing notes, and social work assessments, totalling 4.2 million clinical text documents.

3.2 Five Modelling Approaches

Model R1 (Administrative Claims) used logistic regression on age, sex, DRG, Elixhauser comorbidity flags, and 12-month prior utilisation -- mirroring the CMS methodology. Model R2 (Clinical EHR) used XGBoost on 86 structured variables including labs at discharge, medication count, LOS, vital sign trajectories, and functional status codes. Model R3 (NLP-Enhanced) added features extracted from discharge summaries using a ClinicalBERT model fine-tuned on Italian clinical text: 24 NLP-derived variables capturing social determinants (lives alone, caregiver availability), functional barriers (mobility

limitations, cognitive impairment), medication concerns (polypharmacy burden, affordability), and discharge plan complexity. Model R4 (Temporal Deep Learning) used a transformer encoder on the entire temporal sequence of structured EHR events during hospitalisation (labs, vitals, medications, orders), capturing clinical trajectory patterns. Model R5 (Multi-Modal Ensemble) fused R2, R3, and R4 through a stacking meta-learner. Table 2 presents the study setup.

3.3 The RPUS Score

The Readmission Prediction Utility Score integrates: discrimination via C-statistic (weight 0.25), calibration via Brier score (0.15), net reclassification improvement over the administrative model (0.20), actionability -- whether identified risk factors suggest specific interventions (0.25), and implementation feasibility rated by a hospital informatics panel (0.15). Temporal splitting used 2016-2019 for training, 2020 for validation, and 2021 for testing. The 2020-2021 split also tests robustness to the COVID-19 pandemic, which disrupted hospitalisation patterns substantially.

Table 2: Study Design and Cohort Summary

Parameter	Value	Details
Hospitals	3 Italian academic centres	Rome, Milan, Naples
Hospitalisations	124,682	Jan 2016 - Dec 2021
Heart failure	28,462 (readmit 22.4%)	Highest readmission rate
Pneumonia	32,186 (readmit 14.8%)	Including COVID pneumonia 2020-21
COPD	26,842 (readmit 19.6%)	Acute exacerbation
Arthroplasty	37,192 (readmit 5.8%)	Elective hip and knee
Clinical text documents	4.2 million	Discharge summaries + nursing notes
Models compared	5 (R1-R5)	Admin through multi-modal ensemble
RPUS dimensions	5 (weighted)	C-stat, calibration, NRI, actionability, feasibility
Temporal split	2016-19 / 2020 / 2021	Train / validation / test

COPD = Chronic Obstructive Pulmonary Disease; NRI = Net Reclassification Improvement; COVID = SARS-CoV-2 pneumonia included in pneumonia cohort from 2020.

4. Results

4.1 Multi-Modal Ensemble Breaks the 0.75 Barrier

The multi-modal ensemble (R5) achieved the highest RPUS at 0.916, with C-statistic 0.78 for all-cause 30-day readmission on the 2021 test set. This represents a substantial and clinically meaningful advance over the administrative model (R1, C-stat 0.62) and over the structured clinical model (R2, C-stat 0.72). The NLP-enhanced model (R3) contributed the largest single improvement, pushing C-statistic from 0.72 (R2 alone) to 0.76 -- a 0.04 gain entirely attributable to information extracted from unstructured text. The most predictive NLP features were: patient lives alone (OR 1.68 for readmission), documented medication affordability concern (OR 1.82), no outpatient follow-up arranged at discharge (OR 1.94), and cognitive impairment noted in nursing assessment (OR 2.12). None of these are captured in any structured EHR field at our institutions. The temporal deep learning model (R4) achieved C-stat 0.74, showing that the trajectory of clinical variables during hospitalisation (improving vs declining lab trends, vital sign stability) adds predictive value beyond point-in-time discharge snapshots. Table 3 has the full RPUS breakdown.

4.2 Condition-Specific Performance

Performance varied substantially by condition. Heart failure showed the highest C-statistics across all models (R5 C-stat 0.82), reflecting the well-characterised clinical predictors of HF decompensation -- BNP trajectory, sodium at discharge, diuretic response, and functional class. COPD was second (R5 C-stat 0.78), with exacerbation history and inhaler adherence barriers (captured by NLP) as the strongest predictors. Pneumonia was harder (R5 C-stat 0.74), partly because the 2020-2021 test period included COVID pneumonia patients whose readmission patterns differed from pre-pandemic pneumonia. Arthroplasty was hardest to predict (R5 C-stat 0.70) because the low baseline readmission rate (5.8%) means most readmissions are driven by rare, unpredictable surgical complications rather than systematic risk factors. Table 4 presents condition-specific results.

4.3 Reclassification and Clinical Impact

The NRI analysis showed that R5 correctly reclassified 18.4% of patients compared to R1 -- meaning that nearly one in five patients would be managed differently (more intensive discharge planning for upward-reclassified patients, fewer unnecessary resources for downward-reclassified ones). Among patients reclassified from intermediate to high risk by R5, the observed readmission rate was 28.6%, compared to 14.2% for those remaining intermediate -- confirming that the reclassification identified genuinely higher-risk individuals. In practical terms, if these high-risk patients received intensive transitional care (home visits, telephone follow-up, medication reconciliation), and if such programmes reduce readmissions by 20-30% as the literature suggests, the model could prevent an estimated 640-960 readmissions per year across the three hospitals.

Table 3: Readmission Prediction Utility Score (RPUS) by Model

Model	C-Statistic	Calibration	NRI	Actionability	Feasibility	RPUS
R5: Multi-Modal	0.94	0.90	0.94	0.94	0.76	0.916
R3: NLP-Enhanced	0.90	0.88	0.90	0.92	0.82	0.892
R4: Temporal DL	0.88	0.86	0.86	0.82	0.78	0.858
R2: Clinical EHR	0.84	0.86	0.82	0.84	0.92	0.852
R1: Admin Claims	0.72	0.78	Ref.	0.68	0.96	0.768

$RPUS = 0.25 \times C\text{-stat} + 0.15 \times \text{Calibration} + 0.20 \times NRI + 0.25 \times \text{Actionability} + 0.15 \times \text{Feasibility}$. NRI = Net Reclassification Improvement vs R1. Ref. = reference.

Table 4: C-Statistic for 30-Day Readmission by Condition and Model

Model	Heart Failure	Pneumonia	COPD	Arthroplasty	All-Cause
R1: Admin Claims	0.66	0.60	0.64	0.58	0.62
R2: Clinical EHR	0.76	0.70	0.74	0.66	0.72

Model	Heart Failure	Pneumonia	COPD	Arthroplasty	All-Cause
R3: NLP-Enhanced	0.80	0.74	0.76	0.68	0.76
R4: Temporal DL	0.78	0.72	0.74	0.68	0.74
R5: Multi-Modal	0.82	0.74	0.78	0.70	0.78

All values on 2021 temporal test set. Pneumonia includes COVID-19 pneumonia cases admitted in 2021. COPD = Chronic Obstructive Pulmonary Disease.




Figure 1: RPUS by Predictive Model




Figure 2: C-Statistic by Condition -- Admin vs Multi-Modal




Figure 3: Top NLP-Derived Risk Factors (Odds Ratio for Readmission)




Figure 4: Multi-Dimensional Comparison -- Top 3 Models

5. Discussion

5.1 The Text Is Where the Signal Lives

The single most important finding of this study is that NLP extraction from discharge summaries produced a larger improvement in discrimination (delta C-stat +0.04) than any other individual data source. This is not because clinical text contains exotic biomarkers or obscure laboratory values -- it is because text captures the human context of illness that structured data systematically ignores. A discharge summary that says 'patient lives alone, limited mobility, daughter visits weekly but works full-time, patient expressed concern about affording medications' encodes critical readmission risk information that exists nowhere in the structured EHR. The OR of 2.12 for 'cognitive impairment noted in nursing assessment' is particularly telling: this is arguably the strongest single readmission risk factor we identified, yet cognitive impairment is rarely coded as a diagnosis unless it is the primary reason for admission. The implication for the field is clear: readmission prediction will not substantially improve until models routinely incorporate clinical text.

5.2 From Prediction to Prevention

Prediction without action is a waste of everyone's time. The advantage of NLP-derived risk factors over traditional clinical variables is that they point directly toward specific interventions. If the model identifies 'lives alone' and 'no follow-up arranged' as key risk drivers for a particular patient, the obvious intervention is to arrange a home visit within 48 hours of discharge and schedule an outpatient appointment before the patient leaves the hospital. If 'medication affordability' is flagged, a social work consult to explore assistance programmes is indicated. This is why actionability is weighted so heavily (0.25) in our RPUS -- a model that achieves high C-statistic but identifies

only unmodifiable risk factors (age, comorbidity count) is less useful than a slightly less discriminating model that flags actionable social determinants.

5.3 Limitations

Several limitations apply. First, the NLP model was developed on Italian clinical text; adaptation to other languages requires language-specific fine-tuning of the BERT model. Second, the 2020-2021 test period includes substantial COVID-19 disruption, which may limit generalisability to post-pandemic practice. Third, we measured 30-day readmission to any hospital, but inter-hospital data sharing is imperfect in Italy and some readmissions to non-network hospitals may have been missed. Fourth, we did not implement the model in clinical practice, so the estimated impact on readmission prevention is modelled rather than observed.

6. Conclusion

6.1 What This Study Shows

Multi-modal readmission prediction integrating structured EHR data, NLP-derived social and functional features, and temporal clinical trajectories achieves C-statistic 0.78 -- a clinically meaningful advance over administrative models (0.62) and structured clinical models (0.72). NLP extraction from discharge summaries provides the largest single improvement (+0.04 C-stat) by capturing social isolation, functional barriers, medication concerns, and care plan gaps that coded data misses. The model correctly reclassifies 18.4% of patients versus administrative baselines, with potential to prevent 640-960 readmissions annually across three hospitals through targeted transitional care. Heart failure shows the best prediction (C-stat 0.82) while arthroplasty remains hardest (C-stat 0.70).

6.2 What Comes Next

Three priorities follow. First, a stepped-wedge cluster-randomised trial deploying the R5 model as a real-time discharge planning tool, measuring impact on 30-day readmission rates, patient satisfaction, and cost savings. Second, adaptation and validation of the NLP pipeline in other European languages and health systems to test generalisability. Third, incorporation of post-discharge data streams -- patient-reported outcomes via smartphone apps, home vital sign monitoring, and pharmacy refill data -- to enable dynamic risk updating after discharge rather than a single point-of-discharge prediction. All code,

NLP model weights, and the RPUS calculator are available from the corresponding author on request.

References

- Artetxe, A. et al. (2018). Predictive models for hospital readmission risk: A systematic review of methods. *Computer Methods and Programs in Biomedicine*, 164, 49-64.
- Donze, J. D. et al. (2013). Potentially avoidable 30-day hospital readmissions in medical patients: Derivation and validation of a prediction model. *JAMA Internal Medicine*, 173(8), 632-638.
- Donze, J. D. et al. (2016). Causes and patterns of readmissions in patients with common comorbidities: Retrospective cohort study. *BMJ*, 354, i3428.
- Goldstein, B. A. et al. (2017). Opportunities and challenges in developing risk prediction models with electronic health records data. *Journal of the American Medical Informatics Association*, 24(1), 198-208.
- Guo, A. et al. (2020). Machine learning models for hospital readmission prediction: A systematic review. *BMC Medical Informatics and Decision Making*, 20(1), 56.
- Huang, Y. et al. (2019). ClinicalBERT: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*.
- Kansagara, D. et al. (2011). Risk prediction models for hospital readmission: A systematic review. *JAMA*, 306(15), 1688-1698.
- Krumholz, H. M. et al. (2011). An administrative claims measure suitable for profiling hospital performance based on 30-day all-cause readmission rates among patients with heart failure. *Circulation: Cardiovascular Quality and Outcomes*, 4(2), 243-252.
- Leppin, A. L. et al. (2014). Preventing 30-day hospital readmissions: A systematic review and meta-analysis of randomized trials. *JAMA Internal Medicine*, 174(7), 1095-1107.
- Rajkomar, A. et al. (2018). Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1(1), 18.
- Shameer, K. et al. (2017). Predictive modeling of hospital readmission rates using electronic medical record-wide machine learning. *JMIR Medical Informatics*, 5(3), e28.
- van Walraven, C. et al. (2010). Derivation and validation of an index to predict early death or unplanned readmission after discharge from hospital to the community. *CMAJ*, 182(6), 551-557.
- van Walraven, C. et al. (2011). Proportion of hospital readmissions deemed avoidable: A systematic review. *CMAJ*, 183(7), E391-E402.
- Wasfy, J. H. et al. (2017). Readmission rates after passage of the Hospital Readmissions Reduction Program. *Annals of Internal Medicine*, 166(5), 324-331.
- Zuckerman, R. B. et al. (2016). Readmissions, observation, and the Hospital Readmissions Reduction Program. *New England Journal of Medicine*, 374(16), 1543-1551.
- Dharmarajan, K. et al. (2013). Diagnoses and timing of 30-day readmissions after hospitalization for heart failure, acute myocardial infarction, or pneumonia. *JAMA*, 309(4), 355-363.
- Kripalani, S. et al. (2014). Reducing hospital readmission rates: Current strategies and future directions. *Annual Review of Medicine*, 65, 471-485.
- Futoma, J. et al. (2015). A comparison of models for predicting early hospital readmissions. *Journal of Biomedical Informatics*, 56, 229-238.
- Frizzell, J. D. et al. (2017). Prediction of 30-day all-cause readmissions in patients hospitalized for heart failure. *JAMA Cardiology*, 2(2), 204-209.
- Coleman, E. A. et al. (2006). The care transitions intervention: Results of a randomized controlled trial. *Archives of Internal Medicine*, 166(17), 1822-1828.

Declarations

Funding

Supported by the Italian Ministry of Health (Ricerca Finalizzata RF-2019-12372146) and institutional research funds from the Mediterranean Institute of Technology. No funder influenced study design, data analysis, or publication decisions.

Conflict of Interest

Both authors declare no competing interests. Neither has received funding from any EHR vendor, health IT company, or hospital system with a financial stake in readmission reduction programmes.

Data Availability Statement

De-identified structured EHR extracts and NLP model outputs are available from the corresponding author on reasonable request, subject to ethics approval. Raw clinical text cannot be shared due to re-identification risk. All Python and R analysis code, the ClinicalBERT fine-tuning pipeline, and the RPUS calculator are available on request.

Ethical Approval

Approved by the Comitato Etico Policlinico Gemelli (Prot. 0036482/21), Comitato Etico Milano Area 3

(N. 246-2021), and Comitato Etico Università Federico II Napoli (N. 184/21). Informed consent was waived for retrospective analysis of de-identified EHR data.

Appendix A

RPUS Calculation and NLP Feature Definitions

This appendix details the RPUS scoring methodology, the 24 NLP-derived features with definitions and extraction methods, and sensitivity analyses.

Part I -- RPUS Dimensions

1a. Discrimination: C-statistic on 2021 temporal test set Higher = better separation of readmitted vs not

1b. Calibration: 1 - Brier score Hosmer-Lemeshow $p > 0.05$ required

2a. NRI: Categorical NRI vs R1 at 10% and 20% thresholds R1 = reference

2b. Actionability: Expert panel rating (5 hospitalists, 5 discharge planners) Do identified risk factors suggest specific interventions?

3a. Feasibility: Informatics panel rating on data availability, compute, and integration 0-1 scale

3b. $RPUS = 0.25 \times C\text{-stat} + 0.15 \times \text{Calibration} + 0.20 \times NRI + 0.25 \times \text{Actionability} + 0.15 \times \text{Feasibility}$

Part II -- Key NLP Features

4a. Social isolation: 'lives alone', 'no family support', 'socially isolated' OR 1.68 for readmission

4b. Medication affordability: 'cannot afford', 'medication cost', 'financial concern' OR 1.82

5a. Follow-up gap: 'no outpatient appointment', 'follow-up not arranged' OR 1.94

5b. Cognitive impairment: 'confused', 'disoriented', 'memory difficulty', nursing MMSE note OR 2.12