# Exploiting EHRs using natural language processing to enable research in emergency medicine: a protocol for a study on hospitalization rates

Vicky Rubini[1], Franco Aprà[2], Giulia Irene Ghilardi[3], Jacek Górka[4], Katarina Hricova[5], Isaac John[6], Zora Lazúrová[7], Peter Mitro[7], Giovanni Nattino[3], George Notas[8], Chiara Pandolfini[3]*, Giovanni Porta[9], Gregor Prosen[10], Pankaj Sharma[6], Matej Strnad[10] and Guido Bertolini[3]

[1]Università Statale di Milano, Milan, Italy, [2]Ospedale San Giovanni Bosco, Turin, Italy, [3]Department of Medical Epidemiology, Istituto di Ricerche Farmacologiche Mario Negri IRCCS, Milan, Italy, [4]Center for Intensive Care and Perioperative Medicine, Uniwersytet Jagielloński, Kraków, Poland, [5]Hospital – Nemocnica AGEL Kosice-Saca, Košice, Slovakia, [6]Ashford and St Peter's Hospitals NHS Foundation Trust, Chertsey, United Kingdom, [7]1st Cardiology Clinic, Univerzita Pavla Jozefa Šafárika v Košiciach, Košice, Slovakia, [8]Department of Emergency Medicine, University of Crete School of Medicine and 7th Health Region of Crete, Heraklion, Greece, [9]Ospedale Santa Maria delle Grazie – Pozzuoli, Naples, Italy, [10]Center for Emergency Medicine, Univerzitetni Klinicni Center, Maribor, Slovenia

Increasing demands on emergency departments (EDs) call for optimized decision-making processes to improve patient outcomes and resource allocation. Overcrowding is a significant issue, and the propensity of EDs to hospitalize patients is a key contributing factor to limiting in-patient bed availability, with inappropriate decisions negatively impacting healthcare quality and costs. In this setting research in emergency medicine to improve these difficulties is challenging. The main obstacles are the large volume of cases handled, the paucity of staff availability, and the resulting lack of time to dedicate to data entry. Furthermore, the electronic health record (EHR) systems currently used in EDs are not optimized for collection of data for research. Even retrospective data analyses cannot be performed due to the lack of robust data. Moreover, the EHR contains not only structured data but also abundant information in a free-text format which is challenging to use for research purposes. This protocol describes a study, the Use Case 1 study, which is part of the more general Horizon Europe eCREAM (enabling Clinical Research in Emergency and Acute-care Medicine) project. The study will test the reliability of an advanced natural language processing model set up in eCREAM to exploit EHRs by extracting robust, structured data to enable research in EDs. Specifically, the study will test the validity of the data extracted from the EHRs by addressing the issue of hospitalization rate. We will develop a predictive model to assess emergency department hospitalization rates, thereby enabling standardized comparisons across centers, ultimately leading to improved decision-making and reduced unnecessary hospital admissions.

Retrospective patient data from 2021 to 2023 from 30 centers across Europe will be analyzed, and multivariable models will be employed to predict hospitalization and adjust comparisons between centers. The results are expected to improve decision-making in these departments. More generally,

should the data extraction system prove valid, our results would serve as a practical demonstration that, despite the abundance of free-text data, EHRs can be exploited to conduct research in the emergency medicine field.

**Clinical trial registration:** Clinicaltrials.gov, identifier: NCT06354764.

# Introduction

The increasing burden on emergency departments (EDs) requires efficient decision-making processes to ensure optimal patient outcomes and resource utilization. Many factors influence performance, such as the aging population and the growing complexity of medical care, and in this context, the ED's decision to hospitalize patients plays an important role. The propensity to hospitalize patients contributes to ED and hospital overcrowding, with negative implications for patient care quality and healthcare costs (1–3). Admitting a patient in the absence of any real need also leads to unnecessary tests and treatments, exposing patients to the risk of adverse events, avoidable clinical errors, unmotivated emotional and physical distress, and increased costs for the patient and society (4–6). Regrettably, inappropriate hospitalizations are very common. It is estimated that the extent of this phenomenon varies between 16% and 55% in Europe (7–9). For these reasons, the ED rate of hospitalization is considered a primary performance indicator of these departments with a continuous pressure for its reduction. Conversely, an overly optimistic assessment of a patient's condition in the ED and a subsequent improper discharge can also have serious health consequences with potential lengthy litigation. Hence, improving the appropriateness of this decision is of paramount public interest to protect individual patients and reduce the waste of resources that characterizes today's system.

One effective, validated approach to induce improvements in complex decision-making processes within organizations is peer-to-peer comparison (10, 11). In this context, a standardized and methodologically rigorous comparison of the hospitalization rate of EDs must adjust for possible between-center differences. The first step to reach this goal is to develop a predictive model that accurately estimates each patient's probability of being hospitalized, based on clinical conditions and contextual factors. Such a model would make it possible to calculate, for each ED, the expected hospitalization rate. The comparison of the observed hospitalization rate in an individual ED with the expected one derived from the model provides a robust method of comparing the department with the average performance, taking into account the characteristics of the patients treated and the conditions under which the ED operated. In other words, the predictive model represents the benchmark against which each ED is evaluated.

However, collecting data for research and improvement in the ED is challenging. Electronic health records (EHRs) in use in EDs represent an important source of information but are not optimized for research, hindering data-driven improvement research. Although these records contain structured data, they also contain abundant information in a free-text format hindering their use for research purposes.

The eCREAM (enabling Clinical Research in Emergency and Acute-care Medicine) project, a European program involving 11 partners in 8 countries coordinated by the Laboratory of Clinical Epidemiology of the Istituto di Ricerche Farmacologiche Mario Negri IRCCS in Italy, seeks to address this obstacle by enabling the use of ED EHRs for research. The main aims of eCREAM (see the article "The format of emergency department electronic health records in Europe. The European initiative and the eCREAM proposal" in this issue for details) are to develop new technical solutions to extract reliable clinical data from structured and unstructured information contained in different electronic patient files, to test the exploitation of the databases created in two relevant use cases, and to FAIRify (i.e., make data Findable, Accessible, Interoperable, and Re-usable) the established databases for clinicians, researchers, health policymakers and citizens. The two use cases that will be used to test the data represent real-life situations relevant to the ED and involve the concrete application of the data to address them. The first of the two use cases, Use Case 1 (UC1), is the focus of this study protocol and involves the assessment of ED propensity to hospitalize patients. The second use case involves the development of dashboards to be used by healthcare personnel, policymakers, and citizens to improve the quality of ED care (Use Case 2). Notably, being centered around the efficient reuse of health data for research and innovation, eCREAM is fully aligned with the European Health Data Space (EHDS) initiative (12).

The general aim of UC1 is to test the solution adopted to extract data from EHRs, i.e., the eCREAM language model (13) based on natural language processing (NLP), and to characterize the propensity of different EDs to hospitalize patients. Because of the wide heterogeneity of ED patients and because UC1 aims to test the methodological framework to extract valuable data rather than to provide a comprehensive assessment of the propensity to hospitalize, the analysis will focus on two subgroups of patients, i.e., those presenting to the ED either with dyspnea or following a transient loss of consciousness (TLoC). The study will also investigate the association between the adjusted hospitalization rate of different EDs and a short-term clinical outcome (i.e., 30-day survival) to test the possibility of also extracting data from administrative records.

In this framework, the primary objectives of the study are:

1) to create two separate databases (one for each of the two subgroups) containing the information considered necessary to study the propensity to hospitalize these patients, to assess their 30-day mortality, and to evaluate the accuracy of the extracted information;
2) to develop two multivariable models that predict the probability that patients presenting to the ED with dyspnea (first model) or after a TLoC (second model) will be admitted to the hospital;
3) to provide the participating EDs with an adjusted comparison of the hospitalization rates for the patients with selected symptoms, to enable an improvement in the quality of care.

The study's secondary objective is to assess whether the adjusted admission rate is associated with the patients' 30-day survival.

This study protocol describes the UC1 study, expanding on specific aspects that are particularly relevant to this special issue, namely the importance of EHRs, and their unexploited potential, for secondary research.

# Methods and analysis

## Study design

This is an observational, retrospective study that analyzes data from multiple EDs across Europe. The data will be automatically extracted by the eCREAM NLP system so that no additional burden will be placed on ED staff or patients for the data collection. An initial, minor, effort on the part of a few clinicians is necessary for developing the system and involves annotating a subset of clinical notes and filling in a virtual case report form from a subset of notes.

## Selection of subjects

The study will involve 30 EDs from six European countries (Italy, Poland, Slovakia, Slovenia, Switzerland, and the UK). It was ensured that centers of different sizes and heterogeneous contexts would participate. All patients who arrived at the participating EDs between January 2021 and December 2023 will be eligible for data collection.

A limited amount of data will be collected on all eligible patients. This restricted data collection will be aimed primarily at estimating the level of ED crowding, which is expected to significantly affect hospitalization decisions, but also at evaluating the average boarding time of patients to be hospitalized, the flow of incoming patients, and other conditions characterizing the context where the decision to hospitalize is taken. Conversely, an extended data collection will be carried out on two subgroups of interest for UC1: adult patients presenting with dyspnea and those presenting because they experienced TLoC. Patients referred to specialists outside the ED will be excluded from the study because these cases are not severe and the reasoning behind any possible hospital admission of these patients would be very different from that of patients who are visited in the ED (although we don't expect any

dyspnea or TLoC cases to be sent elsewhere). For dyspnea and TLoC cases, the necessary data will be identified from the triage section of the EHRs and from the final diagnoses, limiting the analysis to epilepsy or syncope-related TLoC cases.

The inclusion criteria for the patients eligible for the extended data collection (i.e., including all demographic and clinical factors) are:

- adult patients who presented to the ED with dyspnea. These patients will be selected by analyzing the triage section of the EHRs;
- adult patients who presented to the ED following TLoC. To reduce the complexity of such a vast group of patients, the analysis will be restricted to patients who experienced TLoC due to epilepsy or syncope. These patients will be selected by analyzing the EHRs' triage section and the final ED diagnosis.

## Study duration

The eCREAM UC1 study will last 24 months. The first 12 will be used to install the eCREAM platform, which will contain the different modules responsible for data extraction, in the various participating EDs and to retrieve the required clinical and administrative information, check its quality, and make any necessary adjustments. The following 12 months will be devoted to data analysis, research reports production, and dissemination of results.

## Data collection

Three sets of information influence the decision to admit or discharge a patient with dyspnea or TLoC and will therefore be collected during UC1 (please see Supplementary material for the full list of the variables):

- *Structural and organizational data* (e.g., the availability and size of the observation unit in the ED, or if point of care devices such as a blood gas analyzer or an ultrasound machine, are available in the ED) were collected at the ED level through a questionnaire administered to the ED staff or department head, and do not include patient data. These data were collected on a voluntary basis by a single clinician from each center at the beginning of the study and were not collected during routine ED service.
- *Transitory organizational data* (e.g., real-time crowding level in the ED) will be estimated based on the restricted data collected on all patients who arrived at the ED. In this framework, the only variables collected on all patients will be those necessary to calculate the number and type of patients in the ED at any given time, and will involve data such as date and time of arrival at the ED, and triage code.
- *Demographic and clinical data on patients with dyspnea or TLoC at ED arrival.* An international panel of experts identified the variables in the extended data collection needed to predict the hospitalization of these patients through a

Delphi-modified method and literature review. We therefore refer to this list of variables as the eCREAM virtual case report form (vCRF). As stated previously, the associated data will be obtained automatically directly from the EHRs in use in the ED. Examples of these data are presence of chronic pulmonary disease, active neoplasia, and various lab tests.

The eCREAM platform equipped with the eCREAM NLP system to automatically extract the values of the variables listed in the vCRF from the clinical notes and to estimate the transitory organizational data will be installed in each participating hospital.

## Data processing

The processing of data collected in the context of UC1 has been planned with the specific aim of protecting the patients' privacy to the highest possible level. Figure 1 provides a simplified representation of the data processing steps. Two separate data flows will be generated for the restricted and extended data collections, with the latter being performed on patients arriving at the ED for dyspnea or TLoC. Data will be sent to a database on the Central eCREAM platform which will be used for all the analyses in the project. After appropriate data curation the data will also be transferred to the Medical Informatics Platform (MIP) database where a complex system developed within the EBRAINS infrastructure (www.ebrains.eu) will share the data with the broad research community for analyses while meeting GDPR and national regulations for privacy protection. Further details on data processing and on the MIP can be found in the UC1 study protocol (14).

By adopting this data processing plan only data necessary to pursue the study objectives will be transmitted from the participating centers to the study coordination center. The study will, therefore, comply with the principle of data minimization as set out in the GDPR. Prior to starting data analyses the coordinating center will assess the internal consistency of the information collected.

## Data analysis

Different methods will be used to achieve the objectives of the study.

To assess the reliability of the extracted dataset (primary objective 1) we will randomly sample a subset of the ED visits from the database of the Central eCREAM platform and will verify its accuracy. The whole vCRF will be evaluated, including both variables sourced by structured datasets and variables generated by the NLP algorithms. Experienced ED physicians will be asked to complete the vCRF for those same sampled records and the resulting data will be compared to those extracted automatically through the eCREAM system. The human-generated data will be considered as the gold standard. The performance of the eCREAM extraction will be primarily evaluated in terms of overall accuracy. For all binary variables of the vCRF, we will also evaluate the results in terms of sensitivity (i.e., among the patients with the presence

of a factor, the proportion of patients for which the factor has been flagged as present) and specificity (i.e., among the patients without the factor, the proportion of patients for which the factor has been flagged as not present). This will be done in order to separately assess the reliability of the eCREAM system to retrieve present clinical conditions and to report their absence.
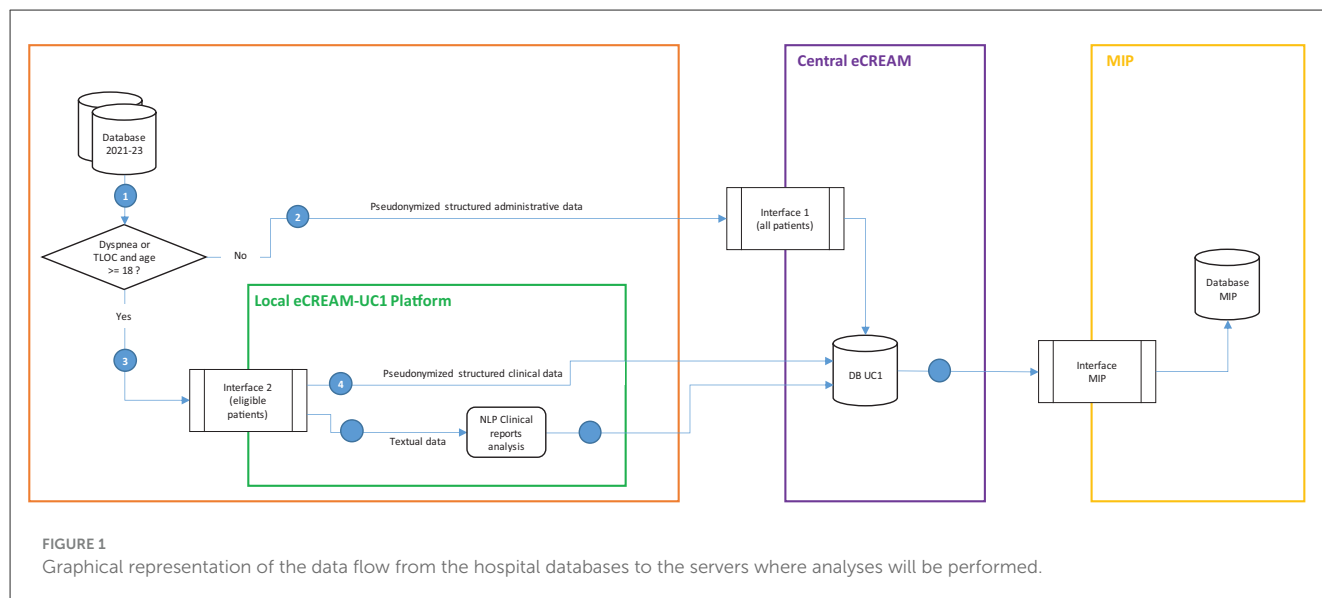
To achieve primary objective 2, we will develop a multivariable model estimating the probability of hospitalization based on patient-specific factors (e.g., medical conditions), transitory-specific factors (e.g., ED crowding), hospital-specific factors (e.g., availability of medical resources), and societal-specific factors (e.g., presence of healthcare services). All factors collected will be considered as candidate values for the multivariable model. The type of model to be used will need to address different peculiarities of the data source, including the presence of missing values and the heterogeneity of ED presentations that results in heterogeneous associations between factors and probability of hospitalization. In this context, decision trees represent a broad family of statistical models that appropriately address these two data characteristics. We will explore the use of different models from this family, including Classification and Regression Trees (CARTs), Boosted Trees, Random Forests, and Bayesian Additive Regression Trees (BARTs). We will select the specific algorithm with the best performance which will be evaluated in terms of discrimination (with the area under the ROC curve) and calibration (using the Calibration belt). In particular, the development of a model that achieves acceptable calibration (nonsignificant test associated with the Calibration belt) and adequate discrimination (area under the ROC curve > 0.8), i.e., good performance estimating the probability of patients being hospitalized, will support the reliability of the information extracted from the EHR of the participating centers for the purposes of this research project.

The adjusted comparison of the hospitalization rate across EDs (primary objective 3) will be based on the probabilities of hospitalization estimated by the developed multivariable model. Specifically, the expected center-specific hospitalization rate will be calculated by summing these probabilities over all patients visiting a single ED. The ratio between observed and expected rates, which will be referred to as standardized hospitalization rate (SHR), represents an adjusted indicator to measure the propensity of the ED to hospitalize patients. Based on the binomial distribution, the SHR will be provided with the corresponding 95% confidence interval.

Finally, we will use another multivariable logistic regression model to assess the association of the adjusted admission rate with 30-day survival (secondary objective). We will consider the 30-day survival as the model's response and include, as predictors, the SHR estimated for primary objective 3 and other relevant prognostic factors. The coefficient of the SHR in the model will provide information on the association between the ED-specific propensity to hospitalize patients and the patient outcomes, after adjusting for patient-specific prognostic factors.

## Sample size

To achieve primary objective 1 and verify the accuracy of the database extracted from the EHR, ED physicians will need

FIGURE 1
Graphical representation of the data flow from the hospital databases to the servers where analyses will be performed.

to complete the vCRF on a random sample of 384 records. This is the necessary sample size to obtain, for each extracted variable, a 95% confidence interval of the accuracy with semiwidth narrower than 5% in the worst-case scenario, i.e., when the accuracy is 50%.

For primary objective 2, following the recommendation of 10 events and non-events per variable for binary outcome models, at least 200 hospitalized and 200 discharged patients from the ED are needed to develop a model with 20 predictors, which is around three times higher than that of the models recently published in the literature. Assuming hospitalization rates of 20% and 50% for dyspnea and TLoC, at least 400 and 1,000 patients are needed, respectively.

To address primary objective 3, setting the type-I error to 5% and power to 90%, to be able to detect as significant an SHR of 1.2 (i.e., an observed hospitalization rate 20% higher than expected), the study will need at least 577 and 1,442 patients with dyspnea and TLoC, respectively, for each ED.

In summary, to address the study's primary objectives, the most stringent requirement regards primary objective 3, which requires data collection on about 600 and 1,500 patients presenting to the ED for dyspnea and loss of consciousness, respectively, from each of the 30 participating centers. In a previous analysis that involved a mid-sized (about 60,000 patients/year) Italian ED (Maggiore Policlinico Hospital, Milan) over the years 2017–2019, the proportion of patients with dyspnea and TLoC was 3.5% and 2.5%, respectively, which correspond to about 2,100 and 1,500 patients/year (data not published). Since we will also include smaller EDs (even those with about 30,000 patients/year), and in order to be conservative, we plan to retrospectively analyze eligible patients visiting the participating EDs over 3 years. This means that we will recruit at least 3,150 patients with dyspnea and 2,250 patients with TLoC in each participating center. This sample size will be ample to assess the accuracy of the data collection, develop the predictive models, and compute the ED-specific indicators on a meaningful group of

patients. Should participating centers not enroll enough patients in either of the two groups we will extend the retrospective data collection for an appropriate amount of time to achieve the desired target.

# Discussion

The principal aim of eCREAM's UC1 is to improve ED decision-making processes by using EHR data through advanced IT systems and data analytics. Its impact will be significant and multifaceted.

In the short term, the study will provide healthcare professionals with rigorous data on the variability of the hospitalization rate between EDs and the staff of each participating ED with an adjusted evaluation of the center-specific propensity to hospitalize patients. Such analyses will be the first implementation of a feedback system that enables healthcare professionals to make more informed decisions, reduce inappropriate hospital admissions and discharges, and, ultimately, improve patient outcomes. Additional decision-making domains and processes, such as those related to hospital management decisions regarding e.g., high boarding rates in hospitals, may be impacted by this feedback as well.

In the medium-to-long term, the impact of the study will be 2-fold. First, the use case will provide the ED research community with one of the largest and richest databases of patients presenting to the ED with dyspnea or TLoC. If the results of this study are positive, the data could provide answers to many other research questions as well. The database, which will be shared through the MIP infrastructure to maximize its secondary use, will enable a wide range of analyses including comparative effectiveness research studies to assess the effect of treatment, interventions, and policies. Second, UC1 will define the methodological basis for the automatic extraction of accurate data from existing EHRs

and, thus, create a sustainable framework for clinical research in emergency medicine. In the future, such a framework will be exploitable for a wide range of purposes, including the generation of a data source for new research projects, quality assessment programs, and interactive management dashboards. The overall impact of these efforts is expected to enhance the quality of emergency care, reduce healthcare costs, and improve patient safety and satisfaction.

## Ethics statement

The studies involving humans were approved by the Ethics Committee Comitato Etico Territoriale (CET) Interaziendale, AOU Città della Salute e della Scienza di Torino, Italy. Written informed consent to participate in this study was not required from the participants or the participants' legal guardians/next of kin in accordance with the national legislation and the institutional requirements.

## Author contributions

VR: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Writing – original draft, Writing – review & editing. FA: Conceptualization, Investigation, Methodology, Writing – review & editing. GG: Conceptualization, Methodology, Project administration, Writing – review & editing. JG: Conceptualization, Investigation, Methodology, Writing – review & editing. KH: Conceptualization, Investigation, Methodology, Writing – review & editing. IJ: Conceptualization, Funding acquisition, Investigation, Methodology, Writing – review & editing. ZL: Conceptualization, Investigation, Methodology, Writing – review & editing. PM: Conceptualization, Investigation, Methodology, Writing – review & editing. GNa: Conceptualization, Formal analysis, Investigation, Methodology, Supervision, Writing – review & editing. GNo: Conceptualization, Investigation, Methodology, Writing – review & editing. CP: Conceptualization, Methodology, Project administration, Writing – review & editing. GPo: Conceptualization, Investigation, Methodology, Writing – review & editing. GPr: Conceptualization, Investigation, Methodology, Writing – review & editing. PS: Conceptualization, Funding acquisition, Investigation, Methodology, Writing – review & editing. MS: Conceptualization, Investigation, Methodology, Writing – review & editing. GB: Conceptualization, Funding acquisition, Methodology, Project administration, Supervision, Writing – review & editing.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declare that no Gen AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/femer.2025.1558444/full#supplementary-material

## References

1. Stang AS, Crotts J, Johnson DW, Hartling L, Guttmann A. Crowding measures associated with the quality of emergency department care: a systematic review. *Acad Emerg Med*. (2015) 22:643–56. doi: 10.1111/acem.12682

2. Ouyang H, Wang J, Sun Z, Lang E. The impact of emergency department crowding on admission decisions and patient outcomes. *Am J Emerg Med*. (2022) 51:163–8. doi: 10.1016/j.ajem.2021.10.049

3. McKenna P, Heslin SM, Viccellio P, Mallon WK, Hernandez C, Morley EJ. Emergency department and hospital crowding: causes, consequences, and cures. *Clin Exp Emerg Med*. (2019) 6:189–95. doi: 10.15441/ceem.18.022

4. Cliff BQ, Avanceña ALV, Hirth RA, Lee SYD. The impact of choosing wisely interventions on low-value medical services: a systematic review. *Milbank Q*. (2021) 99:1024–58. doi: 10.1111/1468-0009.12531

5. Ciapponi A, Fernandez Nievas SE, Seijo M, Rodríguez MB, Vietto V, García-Perdomo HA, et al. Reducing medication errors for adults in hospital settings. *Cochrane Database Syst Rev*. (2021) 11:CD009985. doi: 10.1002/14651858.CD009985.pub2

6. Shrank WH, Rogstad TL, Parekh N. Waste in the US health care system: estimated costs and potential for savings. *JAMA*. (2019) 322:1501–9. doi: 10.1001/jama.2019.13978

7. Tavakoli N, Hosseini Kasnavieh SM, Yasinzadeh M, Amini M, Mahmoudi Nejad M. Evaluation of appropriate and inappropriate admission and hospitalization days according to Appropriateness Evaluation Protocol (AEP). *Arch Iran Med*. (2015) 18:430–4.

8. Pileggi C, Bianco A, Di Stasio SM, Angelillo IF. Inappropriate hospital use by patients needing urgent medical attention in Italy. *Public Health*. (2004) 118:284–91. doi: 10.1016/j.puhe.2003.06.002

9. Lombardi A, Tesoriere A, D'Amici P, Salvi PF, Puzzovio A, Di Paola M. [Appropriate hospital utilization in emergency surgery: application of the appropiatness evaluation protocol]. *G Chir*. (2000) 21:369–72.

10. Navathe AS, Volpp KG, Bond AM, Linn KA, Caldarella KL, Troxel AB, et al. Assessing the effectiveness of peer comparisons as a way to improve health care quality. *Health Aff*. (2020) 39:852–61. doi: 10.1377/hlthaff.2019.01061

11. Winickoff RN, Coltin KL, Morgan MM, Buxbaum RC, Barnett GO. Improving physician performance through peer comparison feedback. *Med Care*. (1984) 22:527–34. doi: 10.1097/00005650-198406000-00003

12. *Proposal for a Regulation of the European Parliament and of the Council on the European Health Data Space* (2022). Available online at: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52022PC0197 (Accessed August 10, 2025).

13. Bertolini G, Ghilardi GI, Pandolfini C, Nattino G, Lavelli A, Moretti F. *Study Protocol - NLP-DeVal: Development and Validation of a Natural Language Processing Tool to Enable Clinical Research in Emergency and Acute Care Medicine: Retrospective Cohort Study*. Zenodo (2024).

14. Bertolini G, Banzi R, Catania F, Lavelli A, Ghilardi GI, Pandolfini C, et al. *Study Protocol - Propensity to Hospitalize Patients From the ED in European Centers an Observational Retrospective Quality-of-Care Study*. Zenodo (2024).