OncovigIA: Artificial Intelligence for Early Lung Cancer Detection and Referral in a Chilean Public Hospital

Jose Peña, MD[1,2,3,4]; Sebastián Santana[5]; Juan Cristobal Morales, MD[2,3]; Natalie Pinto[2,6]; Mariano Suárez[5]; Carola Sánchez[3]; Juan Carlos Opazo, MD[2,7]; Rodrigo Villarroel, MD[2]; Claudio Montenegro, MD[2]; Bruno Nervi, MD[1,4]; and Richard Weber, PhD[4,5,8]

DOI https://doi.org/10.1200/CCI-25-00035

ABSTRACT

PURPOSE Lung cancer is a leading cause of death in Chile, where late-stage diagnoses and high mortality rates prevail. Here, we describe the development of *OncovigIA*, a novel digital tool powered by natural language processing that enhances the identification of potential lung cancer cases by surveilling computed tomography (CT) reports in a large public Hospital in Santiago, Chile.

MATERIALS AND METHODS We combined natural language processing and large language models with state-of-the-art machine learning techniques and approaches to treat unbalanced data sets and determine the best solution to implement in *OncovigIA*. Focusing on key sections of the reports and using various machine learning models, including a balanced Random Forest, the tool achieved high performance with 0.90 accuracy and 0.84 F1-score on the test set.

RESULTS When applied to 13,326 CT chest reports from 2022, it successfully identified 377 CTs of patients with suspected lung cancer previously undetected and not managed by the multidisciplinary local lung cancer team.

CONCLUSION This study underscores the potential of artificial intelligence in early cancer detection and highlights the importance of its integration into local health care ecosystems. By promptly increasing the number of patients referred for specialized management, the tool *OncovigIA* offers a promising path toward improving lung cancer survival rates in Chile and beyond. Moreover, this article provides avenues for its broader implementation, extending it to other cancer types and/or health care–related texts for continuous surveillance, aiming at the early referral and treatment of cancer in low-resource settings.

ACCOMPANYING CONTENT

Data Supplement

Accepted August 7, 2025
Published October 2, 2025

JCO Clin Cancer Inform
9:e2500035

INTRODUCTION

Lung cancer remains a leading cause of death worldwide and in Chile, where survival rates lag behind regional averages.[1,2] A key factor is late-stage diagnosis, exacerbated by the absence of national screening programs and high smoking rates.[3,4] Consequently, many patients present advanced disease or cancer-related complications, underscoring the urgent need for early detection.[5,6] These factors emphasize the need for active management of pulmonary nodules. These lesions must be followed appropriately and eventually studied invasively to rule out lung cancer. Their management options include delayed follow-up imaging or positron emission tomography-computed tomography, obtaining a biopsy, and treating the lesion by surgical resection or radiation therapy; the decision is made on a case-by-case basis by an interdisciplinary team specializing in lung cancer management.[7,8]

Artificial intelligence (AI) has shown promise in lung cancer diagnosis through computed tomography (CT) image analysis.[9] Although its application in lung cancer detection and diagnosis has revealed promising results,[10,11] challenges remain, demanding further validation, collaborative approaches, and adaptability to diverse contexts.[12]

In 2020, one of the largest public hospitals in Chile (Hospital Dr Sótero del Río) implemented a multidisciplinary lung cancer team to manage patients throughout their diagnostic and treatment procedures. Supported by a computer system, this team tracks patients with suspected or confirmed lung cancer (including suspicious lung nodules), providing them with education, follow-up, and a communication channel between professionals. In 2022, this model had already been applied to 245 patients (80-100 per year, approximately), subjectively improving care coordination between the different agents involved in their management. However, one of the main limitations to its widespread use is that it still receives a low number of referrals by health care professionals, partly because of the lack of fluent communication between primary care centers and hospitals and also because

CONTEXT

Key Objective

To develop and implement a system capable of accurately identifying suspected lung cancer cases from computed tomography (CT) radiology reports and facilitating timely referral within a large hospital setting.

Knowledge Generated

This study introduces a novel integration of natural language processing and machine learning techniques for the early identification of lung cancer from CT radiology reports. Furthermore, it demonstrates the system's potential to support timely and effective referrals in a resource-constrained public hospital setting with high patient volumes and extended waiting times. Preliminary findings indicate improvements in detection accuracy and a reduction in the time from diagnosis to treatment initiation.

Relevance (*P.-M. Putora*)

The presented natural language processing–based tool significantly enhances early lung cancer detection by analyzing CT reports, identifying previously undetected cases with the potential to improve patient outcomes.*

*Relevance section written by *JCO CCI* Associate Editor Paul-Martin Putora, MD, PhD, MA.

of the heavy workload faced by specialized medical personnel in charge of referring suspected cases (eg, emergency physicians, radiologists, etc).

Here, we present the development of a digital tool to enhance early detection and referral of suspected lung cancer by integrating AI into specialized cancer case management and monitoring chest CT reports. We aim to pioneer massive surveillance of lung nodules and suspected lung cancer linked to specialized care in Chile.

MATERIALS AND METHODS

Data Description

First, we present the data available for this study. Then, we address the particular problem of labeling that we encountered because of low data quality.

Chest CT reports are structured in three main sections: clinical history and diagnostic hypothesis, study findings, and conclusions.

Clinical History

This section summarizes the patient's relevant clinical history, possible symptoms, and any recent relevant procedures or surgeries. It also includes the reason for obtaining the chest CT scan. This information is crucial for radiologists to interpret the findings and make eventual patient care or follow-up recommendations.

Findings

This section describes specific findings on the chest CT scan images. Findings are typically divided into two categories: normal findings and abnormal findings. In general terms, normal findings are expected and do not indicate any medical problems, while abnormal findings are not expected and may indicate a problem. Radiologists thoroughly describe abnormal findings, including size, location, and specific features if applicable (Fig 1).

Conclusions

This section summarizes the most critical findings of the CT study and eventually provides recommendations for patient care. Radiologists may also discuss differential diagnosis (a list of possible causes for abnormal findings) and could recommend follow-up tests or procedures if necessary.

Data Labeling

Like most hospitals in Chile, Hospital Sótero del Río has no centralized system for tracking lung cancer cases. Thus, there is no unique and accurate way to know how many people are diagnosed with lung cancer annually. Additionally, independent systems that do exist (eg, pathology studies) are often incomplete or inaccurate. Furthermore, we aim not only to identify definitive or confirmed lung cancers but also to find people with lung nodules who need tomographic follow-up or other specific management to rule out the diagnosis.

One challenge of lung cancer is that many cases are not confirmed (eg, by biopsy), particularly in low- and middle-income countries, where access to cancer screening, diagnostic workups, and medical appointments is limited. In this scenario, manual labeling of chest CT reports can facilitate the identification of underrepresented lung cancer cases by reviewing scans from patients whose cancers were either not confirmed or inadequately investigated, including those who required follow-up studies.

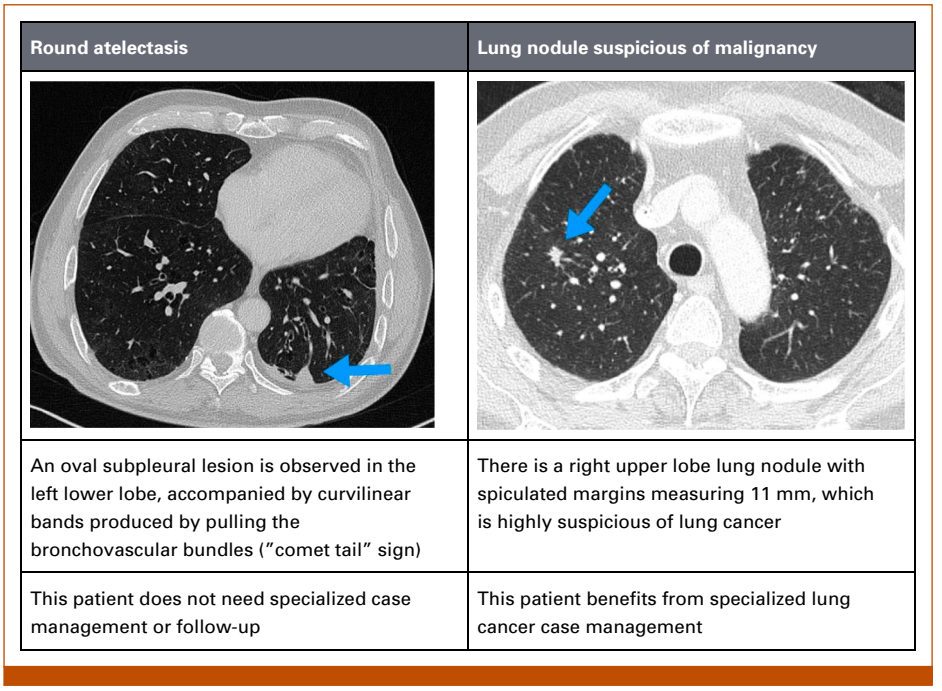

FIG 1. Examples of abnormal chest computed tomography findings and their descriptions by radiologists.

Hence, bearing in mind our target population, manual labeling of chest CT reports was necessary to create an accurate data set for training our models. To be labeled as positive, reports had to describe either a suspicious image (such as a lung mass), a single (or multiple) lung nodule(s), or other images in need of follow-up studies according to the radiology report (such as a subsequent chest CT in 3-6 months). Of note, cases were considered positive for suspicion irrespective of whether there was another primary tumor described in the same study (eg, a person being treated for colorectal cancer with a sigmoid tumor develops a suspicious lung nodule). Cases were considered negative in case of definitive metastatic disease to the lungs originating elsewhere. The concept behind this strategy was to exclusively detect lung cancer or its precursor lesions regardless of the simultaneous presence of other cancers, but to avoid the detection of overt metastatic lung disease from other primary locations. Initially, we labeled 1,190 cases, of which 180 were considered suspicious, and 1,010 were not, leading to an imbalanced data set. These CT reports corresponded to 2021, so we could evaluate our model later in cases from 2022. Table 1 shows the resulting data set.

AI Approaches

Models

We evaluated several machine learning models for identifying potential lung cancer cases from CT reports. These included Decision Trees (DT),[13] Random Forest (RF),[14] XGBoost,[15] Bidirectional Encoder Representations from Transformers (BERT),[16] and Long Short-Term Memory (LSTM).[17] DT is a supervised machine learning algorithm that recursively segments data into homogeneous subsets by applying sequential binary splits on the basis of feature-based decision criteria, generating a hierarchical tree–like predictive model for classification or regression tasks. RF is an ensemble learning method that constructs multiple DT and aggregates their predictions through majority voting or averaging to enhance predictive performance and reduce overfitting. XGBoost is a gradient-boosting decision tree algorithm optimized for computational efficiency and predictive accuracy. BERT and LSTM represent deep neural network architectures designed to capture complex contextual and sequential relationships within textual data.[18] We performed feature extraction on the basis of expert knowledge. During an extensive knowledge acquisition phase, where we observed an expert analyzing several chest CTs, we noticed that they use mainly the first and third of the three sections mentioned in the Section "Data description," that is, "Clinical history" and "Conclusions." The extracted features were used as input variables for the respective models. Model

TABLE 1. Available Data (CT year 2021)

Type of Data	Number of Cases
Labeled as "suspected lung cancer"	180
Labeled as "not suspected lung cancer"	1,010
Not labeled	11,071
Total	12,261

Abbreviation: CT, computed tomography.

architectures and implementation details are elaborated in the Data Supplement.

For hyperparameter optimization, we implemented different approaches. We used a standard grid search methodology for XGBoost, RF, and DT models to identify optimal parameter configurations systematically. Because of the significant computational resources required for training BERT, we constrained our hyperparameter exploration to three carefully selected configurations and reported results from the best-performing variant. All traditional machine learning models (XGBoost, RF, and DT) were implemented using the scikit-learn library. By contrast, we leveraged a pretrained Spanish language version available in the Flair library for the BERT model. This pretrained model was fine-tuned on our data set of CT reports to adapt it to the specific medical terminology and contextual nuances before being deployed as a classifier.

Text Representation

Machine learning models require effective text representation to bridge unstructured data with learning algorithms. We used two primary approaches:

1. Term Frequency-Inverse Document Frequency (TF-IDF): For traditional models (RF, DT), we used TF-IDF to create structured matrices representing word importance.[19] This technique quantifies the significance of each term within documents relative to the entire corpus.
2. Word Embeddings: neural networks (BERT, LSTM) use word embeddings, which capture semantic relationships between words in a vector space.[18] These dense vector representations encode contextual meaning, allowing the models to process linguistic nuances more effectively.

For the XGBoost, RF, and DT models, we generated a TF-IDF matrix using the scikit-learn library. We selected 12,000 features for the representation to balance computational efficiency with information retention. Subsequently, we applied standard grid search for hyperparameter tuning to optimize model performance.

Class Imbalance

The data set exhibited a class imbalance, with 84.9% of reports classified as "Not suspected lung cancer." We used several strategies to address this: (1) training RF, DT, and XGBoost models on the original, unmodified data set; (2) balancing the data via undersampling the majority class and oversampling the minority class to achieve a 1:1 ratio; and (3) fine-tuning a pretrained BERT model (although the imbalance limited its performance).

Deployment

We deployed the models as a web application using Streamlit in Python, accepting CT reports (as text or spreadsheets) as input. The platform displays cancer risk probabilities and provides recommendations, highlighting cases where risk is detected. Clinician workflow analysis revealed a focus on "Clinical History" and "Conclusions," leading us to develop a refined model using only these sections, which improved performance. The most accurate models were trained on balanced data, excluding the "Findings" section. All the results reported are for the test set; the entire data set is summarized in Table 2.

As shown in Table 2, the RF model provides the highest F1-score and accuracy. It was trained using balanced data and exclusively relied upon the Clinical History and Conclusions extracted from the chest CT reports. Notably, because of the inherent data imbalance within our problem, models trained on nonbalanced data encompassing the complete report yielded an average F1-score of 0.7. Upon rebalancing the data set and incorporating the entirety of the report, these models

TABLE 2. Performance of the Best Model Under Different Data Treatments

Data	Model	Recall	Precision	F1-Score	Accuracy
Not balanced/full report	RF	0.48	0.84	0.77	0.89
Not balanced/full report	DT	0.48	**0.91**	0.78	**0.90**
Not balanced/full report	XGBoost	0.57	0.83	0.81	**0.90**
Not balanced/full report	Balanced RF	**0.93**	0.34	0.62	0.66
Not balanced/full report	Balanced Bagging Classifier	**0.93**	0.43	0.71	0.76
Not balanced/full report	BERT	0.10	0.50	0.55	0.87
Balanced/full report	RF	0.68	0.62	0.78	0.87
Balanced/full report	DT	0.59	0.51	0.72	0.82
Balanced/full report	XGBoost	0.73	0.70	0.82	0.89
Balanced/clinical history + conclusions	RF	0.75	0.73	**0.84**	**0.90**
Balanced/clinical history + conclusions	DT	0.59	0.58	0.74	0.85
Balanced/clinical history + conclusions	XGBoost	0.70	0.72	0.82	**0.90**

NOTE. Bold numbers indicate the best value per column.
Abbreviations: BERT, bidirectional encoder representations from transformers; DT, decision tree; RF, random forest.

achieved an average F1-score of 0.77. Subsequently, after the rebalancing process and restricting the model inputs to solely the Clinical History and Impressions, the models achieved an average F1-score of 0.8. This outcome signifies a noteworthy improvement of 0.1 in the F1-score by applying our modeling approach.

RESULTS

For external validation, we evaluated our best-performing model on an independent data set of chest CT reports from 2022, completely separate from the training data. This approach allowed us to assess the model's generalizability to new, unseen data collected in a different period. Our specialized clinical team reviewed the validation results to ensure clinical relevance and accuracy. This temporal validation strategy helps confirm that our model can maintain performance when applied to newly generated reports and is not merely overfitting to patterns specific to the 2021 data set.

Our model detected 1,208 suspicious reports (9.06% out of 13,326 reports from 2022) that were evaluated by the clinical team, confirming 621 (51.40%) as correctly labeled by the algorithm for suspected lung cancer (or lung nodules worth a particular study or follow-up); of those, 151 (12.5%) cases were already confirmed lung cancer cases and were being treated accordingly. Interestingly, 461 (38.16%) of the correctly labeled instances were not being treated for lung cancer, nor had they been referred to the lung cancer team. Therefore, those cases are new findings from the lung cancer team's perspective and could benefit from their specialized management. Figure 2 shows the results after the evaluation by the clinical team of 1,208 cases reported as suspicious by the algorithm.

We suggest the following validation plan to prove the impact of our system, OncovigIA, over time. During its first 6 months of operation in Hospital Dr Sótero del Río, we will record all chest CTs and their reports weekly. We will apply OncovigIA and submit cases of suspected of lung cancer to the hospital's case management team for evaluation and follow-up, measuring the following indicators: true-positive rate, false-positive rate, number of patients under follow-up to rule out lung cancer, time from detection of suspicious cases to confirmation of disease, and time from detection of suspicious cases to treatment initiation. We will compare these indicators with those recorded for a similar period in the year before OncovigIA entered into operation.

DISCUSSION

Our work is preceded by the growing relevance of data analysis technologies for the early detection of cancer.[5] Regarding case management, multidisciplinary teams are considered the gold standard of cancer care. They result in better clinical processes and outcomes for patients with cancer, with evidence of improved survival and positive consequences for patients in multiple dimensions.[20] The model proposed in this report considers an AI tool to foster referral to such multidisciplinary teams. We agree with a recent review underscoring the need for interdisciplinary engagement to ensure choosing an appropriate case use, a

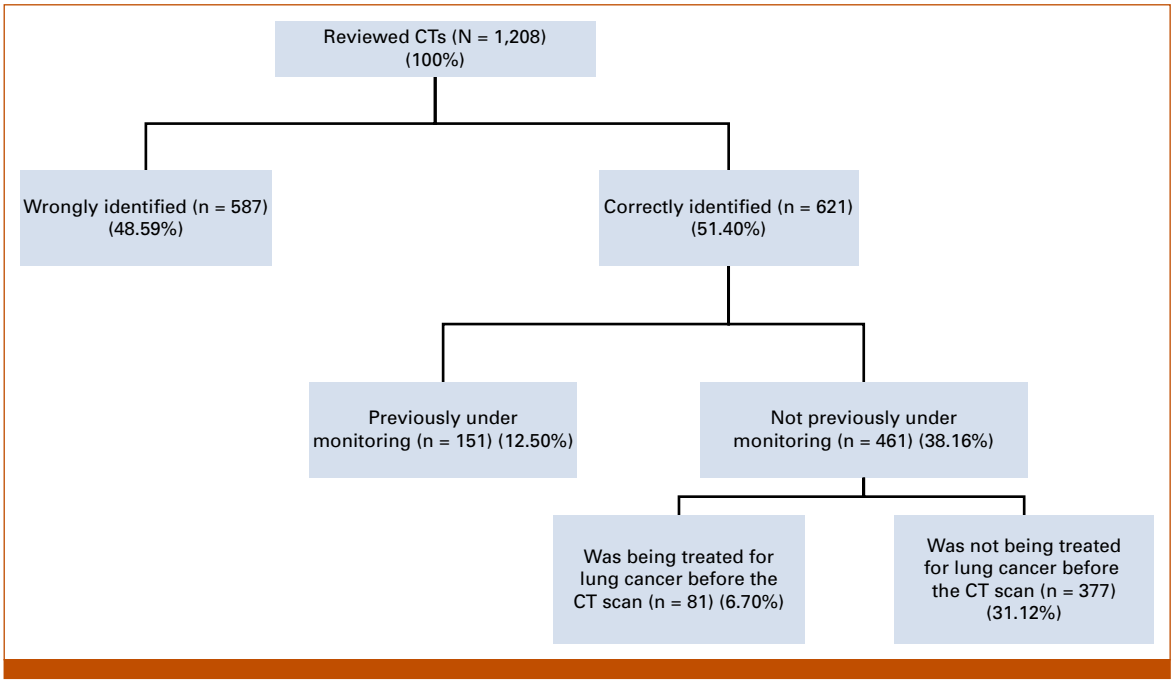

FIG 2. Outcome of the reviewed cases. CT, computed tomography.

robust development and testing phase before its adoption into health care systems, and the requirement of a local ecosystem to promote the growth of these initiatives.[21] Bearing all those factors in mind, despite the highly reported diagnostic accuracy of AI in the detection of lung cancer on imaging studies,[12] we chose a text-based approach at this first stage for considering it more attractive, given its potential applicability to the large number of documents that are produced in a health care system daily that could be actively surveilled (eg, biopsy reports, outpatient visits records, etc), but also given the current state of the local health care system and its already ongoing models of care.

Beyond classification, OncovigIA is designed to actively enhance early referral and case management within a multidisciplinary workflow, addressing practical challenges in a low-resource health care system. It identifies confirmed lung cancer cases and flags high-risk nodules requiring further follow-up, something many previous NLP-based approaches have not explicitly targeted; see, for example, Mithun et al.[22] Additionally, our system is tailored to the Chilean health care context, where access to structured cancer screening programs is limited. It moves beyond retrospective classification by ensuring real-time case tracking and referral facilitation. A key distinction is that OncovigIA was externally validated on real-world clinical data after deployment. By contrast, Mithun et al[22] relied on retrospective validation using CT reports from 2014 to 2016, without testing the model in a live clinical setting.

From the clinical standpoint, it is important to highlight several areas. First, this strategy's preventive approach provides surveillance for suspected lung cancer and lung nodules that require subsequent studies to rule it out. This is particularly relevant in our context—very common in countries with similar conditions—where there is no national Lung Cancer Screening program, patients struggle to find medical appointments, and the overwhelming majority of patients with lung cancer are diagnosed in advanced stages of their illness. Second, although it roughly failed half of the time, the model detected an elevated number of CTs to be studied or followed up. We can safely assume there will be several hundred cases to manage each year (500-600 CTs approximately), which demands a more systemic approach to develop a strategy that effectively enables their timely management. This should include designing a currently nonexistent information system or service that could serve as an interface and promote patient education and empowerment, incorporating patients more actively in pursuing appropriate suspected cancer management.

We plan to scale this intervention in a comprehensive process as shown in Figure 3.

In the future, the developed tool (*OncovigIA*) will periodically monitor several massive sets of health care–related documents (eg, CT reports, pathology reports, etc) and deliver suspicious cases to the clinical team to initiate their management, with a specialized nurse in charge of the process of expert medical evaluations and diagnostic and treatment procedures. A traceability tool will aid the navigating nurse record and update the patient's current cancer stage while monitoring the necessary steps for each patient; these could include confirmation studies in cases of a very high probability of cancer (eg, biopsy, surgery, etc). However, most lung nodules where lung cancer is not probable enough at the time of the index CT will lead to later follow-up studies (eg,

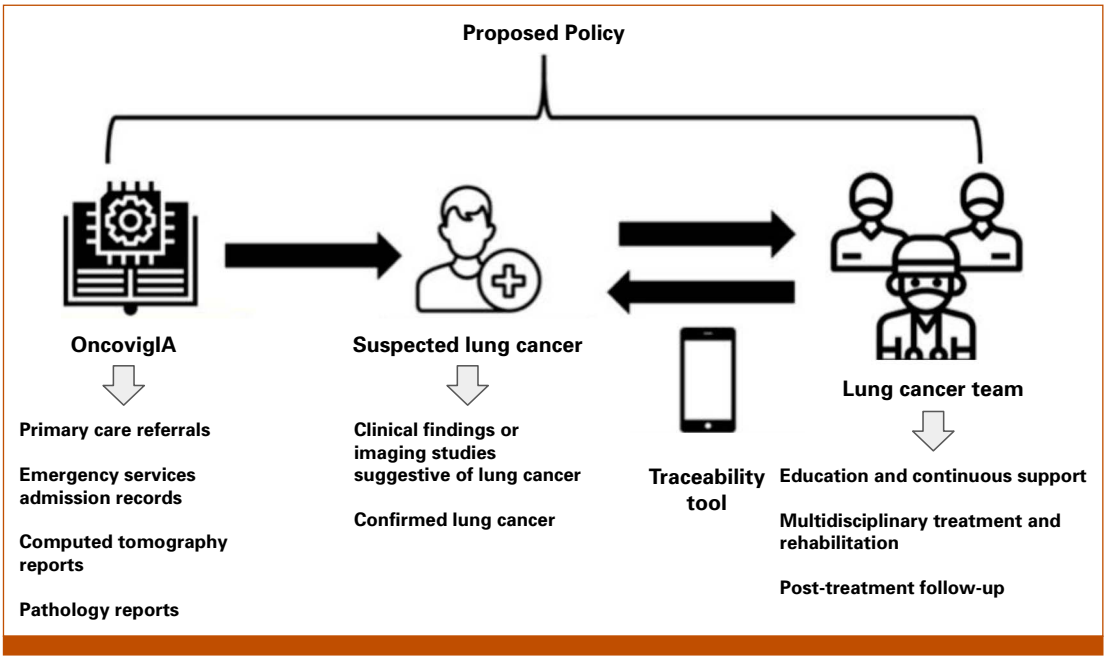

FIG 3. Proposed policy for early lung cancer detection, diagnosis, and treatment.

3-6 months after). We expect this strategy to help allocate scarce resources, shorten the diagnosis and treatment times for patients with cancer, and hopefully improve their survival rate.

At the same time, implementing the previously described process will provide additional information to monitor the developed models effectively: We will check the system's classification and react accordingly in the case of errors of both types, retraining the respective models, particularly during the first months of their application. Although the current models are classifying satisfactorily, we anticipate several changes that could harm their accuracy over time. Modifying the current protocol to describe CT scans would lead to different input vectors for the predictive models. In that case, standardization of the respective reports could ensure systematically adequate input vectors. For permanent changes (new technology revealing novel findings in CT scans, improved protocols, etc), the respective models could be retrained, or alternative models could provide better results; see Table 2, where different models are compared. Monitoring the model's performance over time could also establish insights regarding the feature's importance and provide feedback for the protocols or specific terminology radiologists should follow when describing CT scans.

In the literature, machine learning is noted for its high effectiveness in predicting various types of cancer, including breast, brain, lung, liver, and prostate cancers.[23] Extending this lung cancer detection tool to other types of cancer entails several challenges, but it is feasible with specific adjustments and considerations using our approach. First, clinical, radiologic, and other relevant data must be collected to detect the particular cancer. These data may include biopsy reports, imaging studies reports, electronic health records, and other medical documents. Each type of cancer may display unique characteristics in the data that are crucial for its detection, necessitating the development of new features. Furthermore, machine learning models and algorithms designed to detect novel cancer types must be tailored accordingly, which involves tuning hyperparameters, selecting relevant features, and training specific models. Collaboration with clinical oncology experts and other medical professionals is vital for the success of these initiatives.

In conclusion, in this report, we present a novel system called *OncovigIA*, an AI-driven tool for the early detection of suspected lung cancer from chest CT reports in a hospital in Santiago de Chile, designed to improve referrals to a specialized multidisciplinary team. Applied to 13,326 cases, *OncovigIA* achieved promising results, with F1-score of 0.8, accuracy of 0.9, and precision of 0.73 in detecting relevant cases. Of note, it found many patients needing specialized lung cancer case management not previously referred by their treating physicians. Considering the local clinical context, implementing *OncovigIA* has the potential to increase timely referrals, expedite treatment, and ultimately improve patient outcomes in this low-resource setting.

Future efforts will focus on adapting this system to other hospitals and extending it to other cancer types and medical documents. The first of these two challenges is based on the already developed feature extraction process (see the Models section in this article), since CT reports are standardized across Chile. The hospital where we developed our tool is one of the largest in the country, serving a representative segment of the country's population. It might be necessary, however, to recalibrate *OncovigIA* in some specific private hospitals since patients' characteristics and access to health care could differ from those we had available, resulting, for example, in a higher proportion of early cancer diagnosis compared with the Dr Sótero del Río population. This fine-tuning could be performed easily by recalibrating the model's parameters using the respective cases.

A new feature extraction process is necessary to extend *OncovigIA* to other cancer types since the triggers in the respective texts are different. Our experience from the development of *OncovigIA* assures the efficient construction of such new systems for cancer detection.

Although the system OncovigIA has demonstrated strong performance, its limitations must be acknowledged. The model was trained on a relatively small data set of 1,206 labeled cases, which affects its ability to generalize and contributes to misclassifications, particularly false positives. To enhance its performance, future work will focus on expanding the data set, incorporating more advanced data augmentation techniques, refining classification thresholds, and exploring alternative feature representations beyond TF-IDF. Additionally, although OncovigIA serves as a valuable screening tool, it is not intended as a standalone diagnostic system but as an assistive technology to identify potential cases for further expert review. Addressing these challenges will be key to improving the system's robustness and ensuring its successful integration into clinical workflows.

AFFILIATIONS

[1]Departamento de Hematología-Oncología, Escuela de Medicina, Pontificia Universidad Católica de Chile, Santiago, Chile
[2]Hospital Dr Sótero del Río, Santiago, Chile
[3]Servicio de Salud Metropolitano Sur Oriente, Santiago, Chile
[4]Centro para la Prevención y el Control del Cáncer (CECAN), Santiago, Chile
[5]Departamento de Ingeniería Industrial, Universidad de Chile, Santiago, Chile
[6]Escuela de Enfermería, Pontificia Universidad Católica de Chile, Santiago, Chile
[7]Departamento de Medicina Interna, Escuela de Medicina, Pontificia Universidad Católica de Chile, Santiago, Chile
[8]Instituto Sistemas Complejos de Ingeniería (ISCI), Santiago, Chile

CORRESPONDING AUTHOR

Richard Weber, PhD; Social Media Handle: @Richard_Weber; e-mail: richard.weber@uchile.cl.

SUPPORT

Supported by a grant from Merck Sharp & Dohme (I.A.) LLC (Chile), a subsidiary of Merck & Co, Inc, Rahway, NJ; by the Agencia Nacional de Investigación y Desarrollo (ANID) FONDAP 152220002 (CECAN), ANID PIA/PUENTE AFB230002 (ISCI), ANID BECAS/DOCTORADO NACIONAL 21211812, and by Chilesincancer Foundation.

AUTHOR CONTRIBUTIONS

Conception and design: Jose Peña, Sebastián Santana, Juan Cristobal Morales, Natalie Pinto, Mariano Suárez, Carola Sánchez, Rodrigo Villarroel, Richard Weber
Financial support: Bruno Nervi
Administrative support: Juan Cristobal Morales, Claudio Montenegro, Bruno Nervi
Provision of study materials or patients: Juan Cristobal Morales, Carola Sánchez
Collection and assembly of data: Jose Peña, Sebastián Santana, Juan Cristobal Morales, Natalie Pinto, Carola Sánchez, Juan Carlos Opazo
Data analysis and interpretation: Jose Peña, Sebastián Santana, Juan Cristobal Morales, Natalie Pinto, Mariano Suárez, Carola Sánchez, Juan Carlos Opazo, Claudio Montenegro, Bruno Nervi, Richard Weber
Manuscript writing: All authors
Final approval of manuscript: All authors
Accountable for all aspects of the work: All authors

AUTHORS' DISCLOSURES OF POTENTIAL CONFLICTS OF INTEREST

The following represents disclosure information provided by authors of this manuscript. All relationships are considered compensated unless otherwise noted. Relationships are self-held unless noted. I = Immediate Family Member, Inst = My Institution. Relationships may not relate to the subject matter of this manuscript. For more information about ASCO's conflict of interest policy, please refer to www.asco.org/rwc or ascopubs.org/cci/author-center.
Open Payments is a public database containing information reported by companies about payments made to US-licensed physicians (Open Payments).

No potential conflicts of interest were reported.

ACKNOWLEDGMENT

All authors' disclosure reports can be found at the ASCO Disclosure Management System.

REFERENCES

1. Siegel RL, Miller KD, Wagle NS, et al: Cancer statistics, 2023. CA Cancer J Clin 73:17-48, 2023
2. World Health Organization, International Agency for Research on Cancer: Rates: Cancers by Population. Chile, Globocan, 2018. http://gco.iarc.fr/today/data/factsheets/populations/152-chile-fact-sheets.pdf
3. Allemani C, Matsuda T, Di Carlo V, et al: Global surveillance of trends in cancer survival 2000–14 (CONCORD-3): Analysis of individual records for 37 513 025 patients diagnosed with one of 18 cancers from 322 population-based registries in 71 countries. Lancet 391:1023-1075, 2018
4. OECD: Health at a Glance: OECD Indicators. Paris, France, OECD Publishing, 2019
5. Okma KG, Tenbensel T, Crivelli L (eds): Health Reforms Across the World: The Experience of Twelve Small and Medium-sized Nations with Changing Their Healthcare Systems. Singapore, World Scientific, 2020
6. Valvani A, Martin A, Devarajan A, et al: Postobstructive pneumonia in lung cancer. Ann Transl Med 7:357, 2019
7. Vitale C, D'Amato M, Calabrò P, et al: Venous thromboembolism and lung cancer: A review. Multidiscip Respir Med 10:28-29, 2015
8. Mazzone PJ, Lam L: Evaluating the patient with a pulmonary nodule: A review. JAMA 327:264-273, 2022
9. Villena F, Pérez J, Lagos R, et al: Supporting the classification of patients in public hospitals in Chile by designing, deploying and validating a system based on natural language processing. BMC Med Inform Decis Mak 21:208-211, 2021
10. Goetz LH, Schork NJ: Personalized medicine: Motivation, challenges, and progress. Fertil Steril 109:952-963, 2018
11. Chiu HY, Chao HS, Chen YM: Application of artificial intelligence in lung cancer. Cancers 14:1370, 2022
12. Liu M, Wu J, Wang N, et al: The value of artificial intelligence in the diagnosis of lung cancer: A systematic review and meta-analysis. PLoS One 18:e0273445, 2023
13. Breiman L, Friedman J, Olshen RA, et al: Classification and Regression Trees (ed 1). New York, Chapman and Hall/CRC, 1984
14. Breiman L: Random forests. Mach Learn 45:5-32, 2001
15. Chen T, He T, Benesty M, et al: Xgboost: Extreme gradient boosting. R package version 0.4-2. 1:1-4, 2015
16. Devlin J, Chang M-W, Lee K, et al: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies. Minneapolis, Minnesota, Volume 1, pp, 4171-4186
17. Hochreiter S, Schmidhuber J: Long short-term memory. Neural Comput 9:1735-1780, 1997
18. Wang B, Wang A, Chen F, et al: Evaluating word embedding models: Methods and experimental results. APSIPA Trans Signal Inf Process 8:e19, 2019
19. Salton G, Buckley C: Term-weighting approaches in automatic text retrieval. Inf Process Manag 24:513-523, 1988
20. Taberna M, Gil Moncayo F, Jané-Salas E, et al: The multidisciplinary team (MDT) approach and quality of care. Front Oncol 10:85, 2020
21. Koh DM, Papanikolaou N, Bick U, et al: Artificial intelligence and machine learning in cancer imaging. Commun Med 2:133, 2022
22. Mithun S, Jha AK, Sherkhane UB, et al: Development and validation of deep learning and BERT models for classification of lung cancer radiology reports. Inform Med Unlocked 40:101294, 2023
23. Zhang B, Shi H, Wang H: Machine learning and AI in cancer prognosis, prediction, and treatment selection: A critical approach. J Multidiscip Healthc 16:1779-1791, 2023