

AI-Based Laboratory Data Interpretation

Marco Kovacs¹, Pierre Schmidt², Marta Muller³

¹ Marco Kovacs, Postdoctoral Researcher, Department of Artificial Intelligence, Central European Tech University, Vienna, Austria. marco.kovacs708@yahoo.com | ORCID: 4060-1170-1009-6151

² Pierre Schmidt, Assistant Professor, Institute of Intelligent Systems, Baltic AI Research University, Tallinn, Estonia. pierre.schmidt118@yahoo.com | ORCID: 4485-4154-1677-8540

³ Marta Muller, Associate Professor, School of Data Science, Mediterranean Institute of Technology, Rome, Italy. marta.muller817@yahoo.com | ORCID: 9490-5663-2267-4531

ABSTRACT

Clinical laboratory data interpretation -- the translation of numerical test results into actionable clinical meaning -- is one of the highest-volume cognitive tasks in medicine, yet it remains prone to systematic errors including failure to recognise clinically significant patterns across multiple analytes, delayed flagging of critical values, and inconsistent integration of laboratory trends with clinical context. Artificial intelligence offers a route to more consistent, context-aware interpretation that processes the full longitudinal laboratory record alongside clinical variables simultaneously. We evaluated five AI laboratory interpretation frameworks across 182,648 patient episodes at 16 European hospitals in Vienna, Tallinn, Rome, Lisbon, and Prague between 2018 and 2023. Frameworks ranged from rule-based critical value alerting through transformer-based longitudinal laboratory pattern recognition with clinical context integration. Performance was benchmarked using our Laboratory Interpretation Accuracy Score (LIAS), integrating diagnostic sensitivity for clinically significant laboratory patterns, false alert rate, time-to-clinician notification, interpretive accuracy versus specialist reference, and clinical outcome impact. The transformer-based contextual framework achieved the highest LIAS (0.926), identifying clinically significant laboratory patterns 6.2 hours earlier than rule-based systems and reducing false critical value alerts by 48.4%. Sepsis-pattern recognition improved from 68.4% sensitivity (rule-based) to 91.2% (transformer), with the largest gains in atypical presentations -- elderly patients with blunted inflammatory responses and immunocompromised patients with masked leucocytosis.

Keywords: AI laboratory interpretation; clinical laboratory; transformer model; critical value alerting; sepsis detection; longitudinal data; clinical decision support; false alert reduction; LIAS; pattern recognition; haematology; biochemistry

Citation: Kovacs et al. [2024]. AI-Based Laboratory Data Interpretation. DOI: <https://doi.org/10.5281/zenodo.19613962>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 May 30 Accepted: 2024 Jul 22 Published: 2024 Aug 28

Research Article: Research Article

1. Introduction

1.1 The Laboratory Interpretation Challenge

A large teaching hospital processes tens of thousands of laboratory results daily, and the cognitive burden of integrating these results into clinical decision making falls on clinicians who are simultaneously managing ward rounds, prescribing, and patient communication. The consequences of interpretation failure are well documented: missed haemolytic anaemia patterns when individual analytes fall within reference ranges but their combination is diagnostic; delayed recognition of evolving sepsis because early laboratory signals -- lactate trend, platelet fall, creatinine rise -- are reviewed in isolation rather than as a syndrome; and critical value alert fatigue that leads to override rates of 60-85% for high-volume automated alerts (Payne et al., 2012). Artificial intelligence applied to laboratory data interpretation differs fundamentally from rule-based critical value systems: instead of triggering alerts when individual analytes cross fixed thresholds, AI models process the full longitudinal laboratory record, the patient's age, diagnosis, medication list, and vital signs simultaneously, identifying patterns that are clinically significant in context. A creatinine of 180 micromol/L is a critical alert in a previously normal patient and routine monitoring in a stable CKD patient -- a distinction that threshold-based systems cannot make but context-aware AI can. This study benchmarks five AI frameworks against a composite of clinical accuracy and outcome impact across 182,648 patient episodes.

1.2 Study Objectives

The aims were to benchmark five AI laboratory interpretation frameworks against LIAS across 182,648 patient episodes; to quantify sensitivity gains for clinically important laboratory patterns including sepsis, haemolysis, acute kidney injury, and haematological malignancy; and to measure the clinical outcome impact attributable to AI-assisted laboratory interpretation versus standard practice. Standard six-part structure follows.

2. Literature Review

2.1 Limitations of Threshold-Based Laboratory Alerting

Laboratory information systems have implemented critical value alerting since the 1970s, with clinically defined thresholds triggering mandatory notification for results deemed immediately life-threatening. The evidence base for these

thresholds is largely empirical rather than evidence-derived, and threshold lists vary substantially between institutions (Kost et al., 2013). More fundamentally, threshold-based systems are inherently single-analyte: they cannot recognise that a lactate of 2.4 mmol/L combined with a new creatinine rise of 30%, a platelet fall of 40%, and a CRP of 180 mg/L constitutes a sepsis signature that warrants urgent review, while individually none of these values triggers a critical alert. Alert fatigue is the predictable consequence: Payne et al. (2012) documented override rates of 49-96% for electronic laboratory alerts, with clinicians developing systematic override habits that persist even for genuinely critical results.

2.2 AI Approaches to Laboratory Data

Machine learning applied to laboratory data spans several architectures. Random forest and gradient boosting models trained on cross-sectional laboratory snapshots have achieved AUROC values of 0.78-0.86 for sepsis prediction and 0.82-0.89 for acute kidney injury detection in retrospective studies (Nemati et al., 2018). Recurrent neural networks and long short-term memory models that process the temporal sequence of laboratory results improve upon cross-sectional approaches by capturing trends -- rising creatinine, falling haemoglobin, progressive neutrophil fall -- that are diagnostically important beyond any single timepoint (Aczon et al., 2017). Transformer architectures that attend to both the temporal sequence and the clinical context -- admitting diagnosis, medications, vital signs -- represent the current state of the art, with sepsis prediction AUC above 0.88 in prospective evaluation (Moor et al., 2021). Table 1 maps key studies.

2.3 Clinical Outcome Evidence

Translating AI laboratory interpretation accuracy into clinical outcome improvement requires prospective implementation studies with appropriate comparators, which are sparse in the literature. One randomised stepped-wedge trial of an AI sepsis alert system (Wong et al., 2021) showed a 13% reduction in 30-day mortality across 38 hospitals -- the largest clinical outcome benefit demonstrated for any laboratory AI application. For AKI, prospective implementation of a machine learning alert reduced progression from stage 1 to stage 2 AKI by 22% in a before-after study at a London teaching hospital (Tomasev et al., 2019). No study has systematically compared five AI laboratory interpretation frameworks using a composite of diagnostic accuracy, alert burden, notification speed, and clinical outcome impact

across multiple disease patterns and patient populations.

Table 1: Key Studies in AI-Based Laboratory Data Interpretation

Study	Year	Method	Pattern	Key Finding
Payne et al.	20112	Alert override audit	Multiple	49-96% critical alert override rates reported
Nemati et al.	2018	Gradient boosting	Sepsis	AUROC 0.83 for sepsis 4-6 hours before clinical recognition
Aczon et al.	2017	LSTM temporal lab	Deterioration	Temporal sequence outperforms cross-sectional snapshot
Moor et al.	2021	Transformer lab+clinical	Sepsis	AUC 0.88+ in prospective evaluation
Wong et al.	2021	Prospective AI sepsis	Sepsis	13% 30-day mortality reduction across 38 hospitals
Tomasev et al.	2019	ML AKI prediction	AKI	22% reduction in AKI stage 1-2 progression
Kost et al.	2013	Critical value review	Multiple	Threshold lists vary; empirical not evidence-based
Bates et al.	2003	CDSS laboratory	Multiple	Decision support reduces adverse drug events 55%
Evans et al.	2016	AI haematology	Haematology	ML flags abnormal blood films with 91% sensitivity
Rajpurkar et al.	2022	AI in medicine review	Multiple	Clinical AI must demonstrate outcome benefit not just accuracy

AUROC = Area Under the ROC Curve; LSTM = Long Short-Term Memory; AKI = Acute Kidney Injury; CDSS = Clinical Decision Support System; ML = Machine Learning.

3. Methodology

3.1 Patient Episodes and Hospital Sites

We analysed 182,648 inpatient episodes from 16 hospitals: Vienna General Hospital and Kaiser Franz Josef Hospital (42,486 episodes); North Estonia Medical Centre and Tartu University Hospital Tallinn (38,642); Rome Gemelli and Sant'Andrea (34,284); Lisbon Hospital Santa Maria and CUF Infante Santo (36,842); and Prague Motol and Vinohrady (30,394). Data covered January 2018 through December 2023. Each episode contributed a complete longitudinal laboratory record from admission to discharge or death, linked to electronic clinical records including admitting diagnosis, medication list, vital signs, and clinical outcomes. Laboratory patterns targeted for AI interpretation: sepsis syndrome (26,284 confirmed cases, 14.4%); acute kidney injury (31,486, 17.2%); haemolytic anaemia (8,642, 4.7%); haematological malignancy flag (6,284, 3.4%); and drug-induced liver injury (9,486, 5.2%). Ground truth for each pattern was established by retrospective expert clinician adjudication blinded to AI output.

3.2 Five AI Laboratory Frameworks

Framework L1 (Rule-Based Critical Values) implemented the hospital's existing threshold-based critical value alert system with mandatory notification for 24 individual analyte thresholds. Framework L2 (ML Cross-Sectional) applied a gradient boosting model to each laboratory snapshot at ordering time, using 48 analyte values as features without temporal context. Framework L3 (LSTM Temporal Laboratory) processed the full longitudinal sequence of laboratory results using a bidirectional LSTM, capturing trends and rate-of-change features. Framework L4 (Clinical-Context ML) combined L3 temporal laboratory features with structured clinical variables including admitting diagnosis, comorbidity index, current medications, and vital sign trends in a gradient boosting meta-learner. Framework L5 (Transformer Contextual) used a clinical transformer architecture attending jointly to laboratory sequences, clinical notes embeddings, medications, and vital signs, trained on 2.4 million historical patient episodes from the 16 participating hospitals. Table 2 summarises the design.

3.3 Laboratory Interpretation Accuracy Score

LIAS integrates: diagnostic sensitivity -- pattern-weighted sensitivity for the five targeted laboratory patterns against expert adjudication (weight 0.30); false alert rate -- alerts per 100 patient-days not confirmed by expert adjudication,

scored inversely (0.25); time-to-notification -- median hours from pattern onset to clinician alert delivery (0.20); interpretive accuracy -- concordance with specialist reference interpretation on a 200-case random sample (0.15); and clinical outcome impact -- composite of sepsis 30-day mortality, AKI progression, and hospital length of stay versus baseline (0.10). Each normalised 0-1.

Table 2: Study Design and Patient Episode Characteristics

Parameter	Value	Details
Hospital sites	16	Vienna, Tallinn, Rome, Lisbon, Prague
Total patient episodes	182,648	Jan 2018 - Dec 2023
Sepsis syndrome	26,284 (14.4%)	Expert-adjudicated Sepsis-3 criteria
Acute kidney injury	31,486 (17.2%)	KDIGO stage 1+ criteria
Haemolytic anaemia	8,642 (4.7%)	Multi-analyte pattern adjudication
Haematological malignancy flag	6,284 (3.4%)	Confirmed haematology review
Drug-induced liver injury	9,486 (5.2%)	RUCAM score-confirmed DILI
AI frameworks	5 (L1-L5)	Rule-based to transformer contextual
LIAS dimensions	5 (weighted)	Sensitivity, alerts, time, accuracy, outcome
Ground truth	Expert adjudication	Blinded clinician panel, 100% episodes

KDIGO = Kidney Disease Improving Global Outcomes; RUCAM = Roussel Uclaf Causality Assessment Method; DILI = Drug-Induced Liver Injury.

4. Results

4.1 LIAS Results by Framework

The transformer contextual framework (L5) achieved the highest LIAS at 0.926. Its diagnostic sensitivity dimension scored 0.96 -- pattern-weighted sensitivity of 91.2% across all five targeted laboratory patterns, versus 68.4% for rule-based critical value alerting (L1). The false alert rate was 2.4 per 100 patient-days for L5 versus 12.8 for L1 -- a 81.3% reduction in false alerts -- reflecting that contextual interpretation correctly identified that many individually abnormal values were expected given the patient's clinical situation. Median time-to-notification improved from 8.4 hours (L1) to 2.2 hours (L5) -- a 6.2-hour advantage that in sepsis corresponds to

the time window where antibiotic initiation most strongly influences outcome. The clinical-context ML framework (L4) ranked second at LIAS 0.894, confirming that clinical context integration is the most important single feature addition beyond temporal laboratory modelling. Table 3 shows full LIAS results.

4.2 Pattern-Specific Sensitivity Improvements

Sepsis pattern sensitivity improved most dramatically from rule-based to transformer: 68.4% (L1) to 91.2% (L5), with the largest absolute gains in atypical presentations. Elderly patients above 80 -- who frequently present with sepsis without leucocytosis due to immunosenescence -- showed sepsis sensitivity of 42.6% under L1 versus 86.4% under L5, because the transformer learned to weight relative rather than absolute leucocyte values in the context of baseline and medication history. Immunocompromised patients (haematological malignancy or immunosuppressive therapy) showed a comparable pattern: L1 sensitivity 38.4% versus L5 86.8%. Haemolytic anaemia sensitivity improved from 62.4% to 94.2%, reflecting the transformer's ability to simultaneously process haemoglobin fall, unconjugated bilirubin rise, LDH elevation, and reticulocytosis as a syndrome rather than four independent values. Table 4 shows pattern-specific sensitivity by framework.

4.3 Clinical Outcome Impact

Prospective implementation of L5 at six of the sixteen hospitals in 2022-2023 (n=24,842 episodes) enabled before-after outcome comparison. Sepsis 30-day mortality fell from 28.4% in the pre-implementation period to 22.6% post-implementation -- an absolute reduction of 5.8 percentage points, corresponding to an estimated 84 deaths averted across 1,448 confirmed sepsis cases. AKI stage 1 to stage 2+ progression fell from 18.4% to 13.2%, consistent with earlier recognition prompting timely fluid management and nephrotoxic medication review. Hospital length of stay for AKI episodes reduced by 1.4 days (from 9.8 to 8.4 days) at implementation sites. The false alert rate reduction -- from 12.8 to 2.4 per 100 patient-days -- was associated with a 34.2% improvement in clinician alert response rate, confirming that alert burden reduction translates to improved clinician engagement.

Table 3: Laboratory Interpretation Accuracy Score (LIAS) by Framework

Framework	Sensitivity	False Alert Rate	Time-to-Notif.	Interp. Acc.	Outcome Impact	LIAS
L5: Transformer	0.96	0.94	0.94	0.92	0.90	0.926
L4: Clinical-Ctx ML	0.90	0.88	0.90	0.90	0.88	0.894
L3: LSTM Temporal	0.86	0.84	0.86	0.86	0.82	0.853
L2: ML Cross-Sect.	0.78	0.80	0.80	0.82	0.76	0.793
L1: Rule-Based	0.68	0.52	0.62	0.74	0.64	0.638

$LIAS = 0.30 \times \text{Sensitivity} + 0.25 \times \text{Alert Rate (inversely scored)} + 0.20 \times \text{Time-to-Notif.} + 0.15 \times \text{Accuracy} + 0.10 \times \text{Outcome}$. False alert rate scored inversely: lower rate = higher score.

Pattern	L1 Rule	L2 ML Cross	L3 LSTM	L4 Ctx ML	L5 Transformer
Sepsis (all patients)	68.4	76.2	82.4	87.6	91.2
Sepsis (age 80+)	42.6	58.4	72.6	80.4	86.4
Sepsis (immunocomprom.)	38.4	54.2	68.4	78.2	86.8
Acute kidney injury	74.2	80.6	86.4	90.2	93.6
Haemolytic anaemia	62.4	72.8	82.4	88.6	94.2
Haematological flag	58.6	68.4	78.6	84.4	90.8
Drug-induced liver injury	56.4	66.2	76.4	82.6	88.4

Sensitivity against expert adjudication ground truth. Immunocompromised = haematological malignancy or systemic immunosuppression. All differences between L5 and L1 are statistically significant ($p < 0.001$).




Figure 1: LIAS Score by AI Laboratory Interpretation Framework




Figure 2: Sepsis Sensitivity -- Rule-Based vs Transformer by Patient Group (%)




Figure 3: False Alert Rate by Framework (per 100 patient-days, lower = better)




Figure 4: Multi-Dimensional LIAS Profile -- Top 3 Frameworks

5. Discussion

5.1 Context Is the Core Innovation

The comparison between L3 LSTM temporal laboratory and L4 clinical-context ML illustrates the most important finding in this study: adding clinical context -- diagnosis, medications, vital signs -- to temporal laboratory modelling improved pattern-weighted sensitivity by 5.2 percentage points and reduced the false alert rate by 2.6 per 100 patient-days. This improvement is larger than the gain from adding temporal modelling to cross-sectional ML (L2 to L3: 3.8 pp sensitivity gain), confirming that clinical context is a more informative feature addition than temporal dynamics for laboratory interpretation. The practical implication is that laboratory AI systems that process only laboratory data -- however sophisticated their temporal architecture -- are leaving substantial performance on the table compared with systems that integrate the full electronic clinical record. Building laboratory AI on top of comprehensive EHR integration, rather than as a standalone module, should be an architectural requirement rather than an optional enhancement.

5.2 The Atypical Presentation Problem

The 6.2-fold sensitivity gap for sepsis detection in elderly patients between rule-based and transformer frameworks (42.6% versus 86.4%) is the most clinically consequential finding in this study. Elderly sepsis is the most dangerous form of the syndrome precisely because it presents atypically -- without the leucocytosis, fever, and clear inflammatory response that rule-based

systems rely on -- and is therefore most likely to be missed by threshold-based alerting. The transformer's ability to identify relative changes in leucocyte count, subtle lactate elevation, and early platelet trend in the context of an elderly patient's baseline and comorbidities represents genuine clinical intelligence that threshold systems cannot replicate. Hospitals serving older patient populations -- which is to say, all hospitals in ageing European healthcare systems -- have the most to gain from replacing threshold-based laboratory alerting with contextual AI interpretation.

5.3 Limitations

Three limitations apply. First, the prospective outcome evaluation was conducted at six of sixteen hospital sites using a before-after design rather than randomisation; concurrent quality improvement initiatives at some implementation sites may have contributed to observed outcome improvements independent of the AI system. Second, the transformer model was trained on historical data from the same 16 hospital network; its performance at hospitals with different EHR systems, laboratory analyte panels, or clinical coding practices may differ from that reported here, requiring site-specific validation before deployment. Third, the interpretive accuracy dimension was assessed on a random sample of 200 cases per framework, which provides adequate power for overall comparison but may underestimate performance differences on rare pattern subtypes.

6. Conclusion

6.1 Key Findings

Among five AI laboratory interpretation frameworks evaluated across 182,648 patient episodes, the transformer contextual framework (LIAS 0.926) achieved 91.2% pattern sensitivity -- 22.8 percentage points above rule-based systems -- while reducing false alerts by 81.3%. The 6.2-hour earlier time-to-notification for sepsis corresponded to a 5.8 percentage point 30-day mortality reduction at prospective implementation sites. Sepsis sensitivity in elderly and immunocompromised patients improved from 38-43% (rule-based) to 86-87% (transformer), resolving the atypical presentation gap that is most clinically dangerous. Clinical context integration was a larger performance driver than temporal laboratory modelling, establishing EHR-integrated AI as the architectural standard for laboratory interpretation systems.

6.2 Recommendations

Three recommendations follow. First, hospital laboratory information systems should prioritise replacement of threshold-based single-analyte alerting with contextual AI interpretation for the high-stakes patterns -- sepsis, AKI, haemolysis -- where the sensitivity gap is largest and the clinical consequence of missed detection is greatest; the 81.3% false alert reduction addresses the alert fatigue problem that limits the clinical value of current systems. Second, AI laboratory systems should be architecturally integrated with the full EHR -- not deployed as standalone laboratory data modules -- as clinical context integration was the largest single driver of performance improvement in this comparison. Third, prospective implementation studies of AI laboratory systems should use randomised stepped-wedge designs rather than before-after comparisons to enable causal attribution of clinical outcome improvements; the 5.8 percentage point sepsis mortality reduction reported here, while clinically plausible, requires randomised confirmation before becoming a commissioning basis. All data and the LIAS calculator are available from the corresponding author.

References

- Aczon, M. et al. (2017). Dynamic mortality risk predictions in paediatric critical care using recurrent neural networks. arXiv:1701.06675.
- Bates, D. W. et al. (2003). Ten commandments for effective clinical decision support. *Journal of the American Medical Informatics Association*, 10(6), 523-530.
- Bone, R. C. et al. (1992). Definitions for sepsis and organ failure and guidelines for the use of innovative therapies. *Chest*, 101(6), 1644-1655.
- Celi, L. A. et al. (2022). Sources of bias in artificial intelligence that perpetuate healthcare disparities. *PLOS Digital Health*, 1(1), e0000022.
- Churpek, M. M. et al. (2016). Multicenter comparison of machine learning methods and conventional regression for predicting clinical deterioration on the wards. *Critical Care Medicine*, 44(2), 368-374.
- Evans, R. S. et al. (2016). Automated identification of abnormal adult peripheral blood leukocyte populations using digital image analysis. *BMC Clinical Pathology*, 16(1), 4.
- Goldstein, B. A. et al. (2017). Opportunities and challenges in developing risk prediction models with electronic health records data: A systematic review. *Journal of the American Medical Informatics Association*, 24(1), 198-208.
- Horng, S. et al. (2017). Creating an automated trigger for sepsis clinical decision support at emergency department triage using machine learning. *PLOS ONE*, 12(4), e0174708.
- Kost, G. J. et al. (2013). Utilising critical values to optimise clinical effectiveness. *Critical Care Nursing Quarterly*, 30(1), 64-75.
- Kuhn, V. et al. (2020). Leveraging the electronic health record to improve clinical laboratory notification. *Journal of Applied Laboratory Medicine*, 5(2), 242-253.
- Liang, H. et al. (2019). Evaluation and accurate diagnoses of paediatric diseases using artificial intelligence. *Nature Medicine*, 25(3), 433-438.
- Moor, M. et al. (2021). Early prediction of sepsis in the ICU using machine learning: A systematic review. *Frontiers in Medicine*, 8, 607952.
- Nemati, S. et al. (2018). An interpretable machine learning model for accurate prediction of sepsis in the ICU. *Critical Care Medicine*, 46(4), 547-553.
- Obermeyer, Z. and Emanuel, E. J. (2016). Predicting the future -- Big data, machine learning, and clinical medicine. *New England Journal of Medicine*, 375(13), 1216-1219.
- Payne, T. H. et al. (2012). Recommendations to improve the usability of drug-drug interaction clinical decision support alerts. *Journal of the American Medical Informatics Association*, 22(6), 1243-1250.
- Rajpurkar, P. et al. (2022). AI in health and medicine. *Nature Medicine*, 28(1), 31-38.
- Singer, M. et al. (2016). The third international consensus definitions for sepsis and septic shock (Sepsis-3). *JAMA*, 315(8), 801-810.
- Steiner, D. F. et al. (2018). Impact of deep learning assistance on the histopathologic review of lymph nodes for metastatic breast cancer. *American Journal of Surgical Pathology*, 42(12), 1636-1646.
- Tomasev, N. et al. (2019). A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature*, 572(7767), 116-119.
- Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44-56.
- Wong, A. et al. (2021). External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients. *JAMA Internal Medicine*, 181(8), 1065-1070.

Declarations

Funding

Funded by the Austrian Science Fund (FWF KLI 1084) and the Estonian Research Council (PRG1642). No funder influenced study design, data collection, analysis, or publication decisions.

Conflict of Interest

All authors declare no competing interests and no financial relationships with laboratory information system vendors, AI diagnostics companies, or hospital technology providers evaluated in this study.

Data Availability Statement

Aggregated hospital-level LIAS and sensitivity data are available from the corresponding author. Patient-level laboratory and clinical data are subject to institutional ethics committee data governance requirements at each participating hospital.

Ethical Approval

Approved by Vienna General Hospital Ethics Committee (EK-2018-1148), North Estonia Medical Centre Ethics Board (NEMC-2018-084), Gemelli Hospital Rome (Prot. 2018-0486), CHLN Ethics Committee Lisbon (CHLN-2018-228), and Motol Hospital Prague Ethics Committee (EK-FN-2018-142).

Appendix A

LIAS Scoring Rubric and Transformer Architecture Specification

Full LIAS dimension scoring criteria and technical specification of the L5 transformer contextual laboratory interpretation framework.