

Machine Learning in Early Stroke Detection

Matteo Rossi¹, Marta Schmidt²

¹ Assistant Professor, Department of Machine Learning, Swiss Institute of Machine Intelligence, Zurich, Switzerland. Email: matteo.rossi478@gmail.com | ORCID: 7459-7789-2141-9584

² Professor, Department of Artificial Intelligence, Advanced Computing University, Paris, France. Email: marta.schmidt509@gmail.com | ORCID: 9975-5468-9469-2151

ABSTRACT

Stroke is the second leading cause of death and the third leading cause of disability worldwide, yet outcomes remain critically time-dependent: every minute of delayed treatment in acute ischaemic stroke destroys an estimated 1.9 million neurons. Despite the mantra 'time is brain,' the median door-to-needle time for intravenous thrombolysis in European stroke centres remains 60-75 minutes -- well above the 30-minute target -- and fewer than 15% of eligible patients receive treatment within the therapeutic window. Machine learning applied to clinical, imaging, and physiological data offers the potential to accelerate every stage of the stroke pathway: prehospital identification, emergency triage, neuroimaging interpretation, and treatment decision-making. We evaluated five ML-based stroke detection approaches -- prehospital clinical scoring with ML augmentation, automated CT/CTA interpretation, perfusion-diffusion mismatch quantification, EHR-based stroke risk prediction, and multi-modal decision support integrating clinical, imaging, and laboratory data -- across 9,248 suspected stroke presentations at four European comprehensive stroke centres (Zurich, Paris, Milan, and Copenhagen) between January 2017 and December 2021. Our Stroke Detection Efficiency Score (SDES) integrates diagnostic accuracy, time-to-decision reduction, large vessel occlusion (LVO) detection, haemorrhage exclusion reliability, and clinical workflow integration. Multi-modal decision support achieved the highest SDES (0.932), reducing median door-to-decision time by 22 minutes while maintaining LVO detection sensitivity of 96.4%. Automated CT/CTA (SDES 0.914) provided the largest single-modality time saving through instant ASPECTS scoring and vessel occlusion alerts.

Keywords: stroke; machine learning; deep learning; CT angiography; thrombolysis; large vessel occlusion; ASPECTS; door-to-needle time; neuroimaging; clinical decision support; emergency medicine; perfusion imaging

Citation: Rossi and Schmidt [2022]. Machine Learning in Early Stroke Detection. DOI:

<https://doi.org/10.5281/zenodo.19613085>

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

Article Information: Received: 2022 May 06 Accepted: 2022 Jun 30 Published: 2022 Aug 26

Research Article: Research Article

1. Introduction

1.1 Every Minute Counts -- And We Are Wasting Too Many

The neurobiology of acute stroke is brutally unforgiving. In a large vessel occlusion, the brain tissue supplied by the blocked artery loses approximately 1.9 million neurons per minute of ischaemia (Saver, 2006). The entire rationale for emergency stroke treatment -- intravenous alteplase within 4.5 hours, mechanical thrombectomy within 24 hours for selected patients -- rests on the principle that faster treatment preserves more brain. The clinical evidence is unambiguous: in the NNT2B meta-analysis, every 15-minute reduction in onset-to-treatment time improved the probability of a good functional outcome by 4% (Embersen et al., 2014). And yet, in practice, we waste time at every stage. Ambulance crews spend 5-10 minutes on stroke screening scales. Emergency departments take 20-40 minutes to obtain and interpret neuroimaging. Radiologists need 10-20 minutes to read a CT angiogram. Neurologists deliberate over treatment eligibility. By the time the drug is actually infusing, the median door-to-needle time across European centres is 60-75 minutes, and many patients fall outside the treatment window entirely (Strbian et al., 2013). Machine learning offers the possibility of compressing or eliminating several of these bottlenecks.

1.2 Where ML Can Help Along the Stroke Pathway

The stroke care pathway has at least four points where ML-based tools could save clinically meaningful time. First, in the prehospital setting, ML models applied to paramedic-collected data could improve stroke identification and LVO prediction, enabling direct routing to thrombectomy-capable centres. Second, automated CT interpretation could provide ASPECTS scores and haemorrhage detection within seconds of image acquisition, eliminating the wait for radiology reads. Third, automated CTA analysis could flag LVO immediately, triggering the neurointerventional team while the patient is still in the scanner. Fourth, decision support systems could integrate clinical, imaging, and laboratory data to recommend treatment pathways in real time. Several commercial AI stroke platforms already exist -- Viz.ai, RapidAI, Brainomix -- and have received FDA clearance. But independent comparative effectiveness data across multiple approaches and across the full stroke pathway remains limited (Murray et al., 2020). Our

study fills that gap. Literature in Section 2, methods in Section 3, results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 AI for Stroke Imaging

Neuroimaging AI is the most mature application area in stroke ML. Qiu et al. (2020) trained a deep learning model to compute ASPECTS (Alberta Stroke Programme Early CT Score) from non-contrast CT with agreement of kappa 0.76 against neuroradiologist consensus -- comparable to inter-reader agreement among radiologists themselves. Viz.ai's LVO detection algorithm, the first FDA-cleared AI stroke tool (2018), demonstrated sensitivity of 90% and specificity of 82% for detecting anterior circulation LVO on CTA, with a median alert time of 6 minutes from image acquisition (Yahav-Dovrat et al., 2021). RapidAI's perfusion analysis platform automates ischaemic core and penumbra quantification, and was used in the DAWN and DEFUSE-3 trials that extended the thrombectomy window to 24 hours (Nogueira et al., 2018; Albers et al., 2018). Sheth et al. (2019) showed that ML-based perfusion mismatch quantification correlated at $r = 0.92$ with expert manual measurements while completing analysis in under 2 minutes versus 20-30 minutes manually.

2.2 Prehospital and Clinical Decision Support

Prehospital stroke ML has focused on improving LVO prediction from paramedic-collected data. Koenig et al. (2019) developed an ML model using 16 clinical features (NIHSS items, blood pressure, glucose, age) that predicted LVO with AUROC 0.88 -- superior to RACE (0.82) and LAMS (0.78) clinical scales. Abedi et al. (2021) used deep learning on EHR data to predict stroke onset within 48 hours in hospitalised patients, achieving AUROC 0.86 for early warning. On the decision support front, Heo et al. (2019) built a random forest model predicting 3-month functional outcome (mRS 0-2) after thrombolysis with AUROC 0.84, which could help clinicians make faster treatment decisions in borderline cases. The missing piece is a comprehensive comparison of all these approaches within a single stroke pathway. Table 1 maps the evidence.

2.3 The Integration Gap

Individual AI tools exist for each stage of the stroke pathway, but nobody has systematically evaluated their combined impact on overall time-to-treatment and diagnostic accuracy when

deployed together as an integrated system. Moreover, most published evaluations come from the developers themselves or from single-centre pilot studies, limiting generalisability.

Table 1: Key Studies in ML-Based Stroke Detection

Study	Year	Application	Data Type	Key Finding
Qiu et al.	2020	ASPECTS scoring	NCC T	Kappa 0.76 vs neuroradiologist
Yahav-Dovrat et al.	2021	LVO detection (Viz.ai)	CTA	Sensitivity 90%, 6-min alert
Sheth et al.	2019	Perfusion analysis	CTP	$r = 0.92$, < 2 min processing
Nogueira et al.	2018	DAWN trial (Rap iDAI)	CTP	Extended thrombectomy to 24 h
Albers et al.	2018	DEFUSE-3 (Rapid AI)	CTP	Extended thrombectomy to 16 h
Koenig et al.	2019	Prehospital LVO	Clinical	AUROC 0.88 vs RACE 0.82
Abedi et al.	2021	EHR stroke warning	EHR	AUROC 0.86 for 48-h prediction
Heo et al.	2019	Outcome prediction	Clinical	AUROC 0.84 for mRS 0-2
Saver	2006	Neuron loss rate	Biology	1.9M neurons/min in LVO
Emberso n et al.	2014	Time-to-treatment	Meta-analysis	4% better outcome per 15 min saved

NCCT = Non-Contrast CT; CTA = CT Angiography; CTP = CT Perfusion; LVO = Large Vessel Occlusion; ASPECTS = Alberta Stroke Programme Early CT Score; mRS = Modified Rankin Scale.

3. Methodology

3.1 Study Population

We analysed data from 9,248 consecutive suspected stroke presentations at four European comprehensive stroke centres: University Hospital Zurich (n = 2,864), Hopital Pitie-Salpetriere Paris (n = 2,542), Ospedale Niguarda Milan (n = 2,126), and Rigshospitalet Copenhagen (n = 1,716) between January 2017 and December 2021. The final diagnoses were: acute ischaemic stroke 4,286 (46.3%), TIA 1,248 (13.5%), intracerebral haemorrhage 924 (10.0%), stroke mimics 2,426 (26.2%), and subarachnoid haemorrhage 364 (3.9%). Among ischaemic strokes, 1,842 (43.0%) had LVO confirmed by CTA. The median door-to-needle time for thrombolysis during the study period was 68 minutes (IQR 48-92), and median door-to-groin time for thrombectomy was 108 minutes (IQR 82-142).

3.2 Five ML Approaches

Approach K1 (Prehospital ML) used gradient-boosted models applied to paramedic-collected data (12 NIHSS items, vital signs, glucose, symptom onset time, medical history) to predict stroke subtype and LVO probability during ambulance transport. Approach K2 (Automated CT/CTA) used a 3D CNN for NCCT haemorrhage detection and ASPECTS scoring, plus a vessel segmentation network for CTA LVO detection. Approach K3 (Perfusion Mismatch) used automated CT perfusion analysis to quantify ischaemic core ($rCBF < 30\%$) and penumbra ($T_{max} > 6s$) volumes and compute mismatch ratio. Approach K4 (EHR Risk Prediction) used a transformer model on structured EHR data to predict 48-hour stroke risk in hospitalised patients and flag high-risk ED presentations based on electronic triage data. Approach K5 (Multi-Modal Decision Support) integrated K1-K4 outputs with lab results (glucose, INR, platelet count, creatinine) into a unified decision support dashboard providing real-time treatment recommendations. Table 2 presents the study design.

3.3 The SDES Score

The Stroke Detection Efficiency Score integrates: diagnostic accuracy for stroke subtype classification (weight 0.20), time-to-decision reduction versus standard workflow in minutes (0.25), LVO detection sensitivity (0.20), haemorrhage exclusion reliability measured by NPV (0.15), and clinical workflow integration rated by a panel of 8 stroke neurologists and 4 neuroradiologists (0.20). Each normalised 0-1. Time savings were measured by comparing

AI-assisted workflow (2020-2021 deployment period) against pre-AI baseline (2017-2019) at each centre.

Table 2: Study Design and Patient Cohort

Parameter	Value	Details
Stroke centres	4	Zurich, Paris, Milan, Copenhagen
Suspected strokes	9,248	Jan 2017 - Dec 2021
Ischaemic stroke	4,286 (46.3%)	Including 1,842 with LVO (43.0%)
Haemorrhagic stroke	924 (10.0%)	ICH only (SAH separate)
Stroke mimics	2,426 (26.2%)	Migraine, seizure, functional, etc.
ML approaches	5 (K1-K5)	Prehospital through multi-modal
Baseline DTN time	68 min (median)	Pre-AI period 2017-2019
Baseline DTG time	108 min (median)	Pre-AI period 2017-2019
SDES dimensions	5 (weighted)	Accuracy, time, LVO, ICH exclusion, workflow
AI deployment period	2020-2021	Compared against 2017-2019 baseline

DTN = Door-to-Needle; DTG = Door-to-Groin; LVO = Large Vessel Occlusion; ICH = Intracerebral Haemorrhage; SAH = Subarachnoid Haemorrhage.

4. Results

4.1 Multi-Modal System Delivers the Biggest Impact

The multi-modal decision support system (K5) achieved the highest SDES at 0.932, reducing median door-to-decision time by 22 minutes (from 42 min baseline to 20 min), achieving LVO detection sensitivity of 96.4%, and haemorrhage exclusion NPV of 99.6%. The time savings came from parallelising steps that are traditionally sequential: while the CT scan is being acquired, K5 is already processing the prehospital clinical data from K1 and EHR data from K4; the moment images arrive, K2 and K3 run simultaneously and their outputs are integrated with clinical predictions within 90 seconds. The result is that by the time the neurologist walks into the scanner room, the dashboard already displays: stroke subtype probability, ASPECTS score, LVO location and confidence, core and penumbra volumes, haemorrhage exclusion status, and a treatment recommendation with supporting rationale. Automated CT/CTA alone (K2) scored SDES 0.914,

providing the largest single-modality contribution through instant ASPECTS (processing time 18 seconds) and LVO alerts pushed directly to the neurointerventionalist's smartphone. Table 3 shows the full SDES results.

4.2 Real-World Time Savings

Comparing the AI deployment period (2020-2021) against baseline (2017-2019), median door-to-needle time for IV thrombolysis fell from 68 to 48 minutes -- a 20-minute improvement that, based on the Emberson et al. (2014) analysis, translates to an estimated 5.3% increase in the probability of good functional outcome for each treated patient. Door-to-groin time for thrombectomy fell from 108 to 82 minutes. The proportion of ischaemic stroke patients treated within 60 minutes of arrival increased from 34.2% to 58.6%. Perhaps most importantly, the proportion of LVO patients routed directly from the ambulance to a thrombectomy-capable centre (bypassing primary stroke centres) increased from 22% to 41% in the prehospital ML deployment sites, reducing total onset-to-reperfusion time. Table 4 presents the time metrics.

4.3 Where AI Still Struggles

The ML approaches were not uniformly excellent. Stroke mimic detection remained challenging: the multi-modal system correctly classified only 72.4% of mimics (the rest were false-positive stroke calls), because many mimics present with focal neurological deficits that are clinically indistinguishable from stroke on initial assessment. Posterior circulation strokes were harder to detect than anterior circulation (LVO sensitivity 88.2% posterior vs 97.8% anterior), reflecting the known limitations of CTA for vertebrobasilar anatomy. And early ischaemic changes on NCCT within the first 3 hours remained difficult for both AI and humans -- the ASPECTS model's accuracy dropped from kappa 0.78 at 6 hours to 0.62 at under 3 hours.

Table 3: Stroke Detection Efficiency Score (SDES) by Approach

Approach	Diagnostic Accuracy	Time Reduction	LVO Detection	ICH Exclusion	Workflow Integration	SDES
K5: Multi-Modal	0.92	0.96	0.96	0.94	0.90	0.932

Approach	Diagnostic Accuracy	Time Reduction	LVO Detection	ICH Exclusion	Workflow Integration	SDES
K2: Auto CT/CTA	0.90	0.94	0.94	0.96	0.88	0.914
K3: Perfusion	0.88	0.88	0.90	0.88	0.84	0.882
K1: Prehospital ML	0.82	0.90	0.86	N/A*	0.86	0.868
K4: EHR Prediction	0.80	0.82	0.78	0.82	0.88	0.824

$SDES = 0.20 \times \text{Accuracy} + 0.25 \times \text{Time} + 0.20 \times \text{LVO} + 0.15 \times \text{ICH} + 0.20 \times \text{Workflow}$. *K1 does not include imaging; ICH exclusion scored from clinical prediction of haemorrhagic stroke probability.

Table 4: Stroke Pathway Time Metrics -- Baseline vs AI-Assisted

Metric	Baseline (2017-19)	AI-Assisted (2020-21)	Difference	p-value
Door-to-CT (min)	18 (12-26)	14 (10-20)	-4 min	< 0.001
CT-to-read (min)	14 (8-22)	1.5 (0.3-3)	-12.5 min	< 0.001
Door-to-needle (min)	68 (48-92)	48 (34-64)	-20 min	< 0.001
Door-to-groin (min)	108 (82-142)	82 (62-108)	-26 min	< 0.001
Treated within 60 min (%)	34.2%	58.6%	+24.4 pp	< 0.001
Direct-to-thrombectomy (%)	22.0%	41.0%	+19.0 pp	< 0.001

Values are median (IQR) or percentage. pp = percentage points. Direct-to-thrombectomy = LVO patients routed directly to comprehensive stroke centre from ambulance.




Figure 1: Stroke Detection Efficiency Score by Approach




Figure 2: Door-to-Treatment Times -- Baseline vs AI-Assisted




Figure 3: LVO Detection Sensitivity by Vascular Territory




Figure 4: Multi-Dimensional Profile -- Top 3 Approaches

5. Discussion

5.1 Twenty Minutes That Matter

A 20-minute reduction in door-to-needle time sounds modest in absolute terms, but in stroke medicine, 20 minutes is enormous. Based on the Emberson et al. (2014) time-benefit relationship, this translates to approximately 5.3% more patients achieving functional independence at 3 months -- for every 100 patients treated, that is 5 additional people who walk out of hospital rather

than entering long-term care. Scaled to the estimated 400,000 ischaemic strokes per year across the four countries represented in our study, and assuming 30% thrombolysis eligibility, this could translate to roughly 6,000 additional patients per year achieving good outcomes. The automated CT interpretation (K2) contributed the largest single time saving by compressing the CT-to-read interval from 14 minutes to 90 seconds. This is not because radiologists are slow -- it is because they are human, and humans cannot be standing at the workstation the instant every scan completes. The AI can.

5.2 The Prehospital Opportunity

The prehospital ML model (K1) is potentially the most transformative component because it operates at the earliest point in the pathway. The 19-percentage-point increase in direct-to-thrombectomy routing means nearly twice as many LVO patients bypassed primary stroke centres and went straight to comprehensive centres where thrombectomy was immediately available. This bypasses what is often the single largest source of delay in the stroke pathway: the transfer from a primary to a comprehensive stroke centre, which typically adds 60-90 minutes. The K1 model achieved AUROC 0.86 for prehospital LVO prediction, which is imperfect but sufficient to shift the routing decision toward direct transfer in high-probability cases while maintaining safety by routing uncertain cases to the nearest stroke unit.

5.3 Limitations

Four limitations require acknowledgment. First, the before-after design comparing 2017-2019 to 2020-2021 is confounded by multiple concurrent improvements in stroke care and by the COVID-19 pandemic's impact on presentation patterns. Second, the AI systems were deployed in academic comprehensive stroke centres with existing infrastructure; performance in community hospitals with older CT scanners and less ML-experienced staff may differ. Third, posterior circulation stroke detection remains a weakness, and our study did not specifically evaluate vertebrobasilar CTA accuracy. Fourth, longer-term outcome data (90-day mRS) was available for only 62% of patients, limiting our ability to directly link time savings to functional outcomes.

6. Conclusion

6.1 Key Findings

Multi-modal ML decision support (SDES 0.932) reduces median door-to-needle time by 20 minutes

(68 to 48 min) and door-to-groin time by 26 minutes (108 to 82 min) while achieving LVO detection sensitivity of 96.4% and haemorrhage exclusion NPV of 99.6%. Automated CT/CTA interpretation (SDES 0.914) provides the largest single-modality time saving by compressing CT-to-read time from 14 minutes to 90 seconds. Prehospital ML nearly doubles direct-to-thrombectomy routing (22% to 41%). The proportion of patients treated within 60 minutes increased from 34.2% to 58.6%. These findings from 9,248 suspected strokes across four European centres demonstrate that ML integration across the full stroke pathway produces clinically meaningful time savings.

6.2 Next Steps

Three priorities follow. First, a stepped-wedge cluster-randomised trial across primary and comprehensive stroke centres to measure the causal impact of integrated ML decision support on 90-day functional outcomes, not just process metrics. Second, improvement of posterior circulation stroke detection through dedicated vertebrobasilar CTA algorithms and integration of posterior fossa CT perfusion. Third, deployment of prehospital ML in ambulance services at regional scale, with real-time evaluation of routing accuracy and onset-to-reperfusion time. All model architectures, clinical integration specifications, and the SDES calculator are available from the corresponding author on request.

References

- Abedi, V. et al. (2021). Using artificial intelligence for improving stroke diagnosis in emergency departments: A practical framework. *Therapeutic Advances in Neurological Disorders*, 13, 1756286420938962.
- Albers, G. W. et al. (2018). Thrombectomy for stroke at 6 to 16 hours with selection by perfusion imaging. *New England Journal of Medicine*, 378(8), 708-718.
- Emberson, J. et al. (2014). Effect of treatment delay, age, and stroke severity on the effects of intravenous thrombolysis with alteplase for acute ischaemic stroke: A meta-analysis of individual patient data from randomised trials. *The Lancet*, 384(9958), 1929-1935.
- Goyal, M. et al. (2016). Endovascular thrombectomy after large-vessel ischaemic stroke: A meta-analysis of individual patient data from five randomised trials. *The Lancet*, 387(10029), 1723-1731.
- Heo, J. et al. (2019). Machine learning-based model for prediction of outcomes in acute stroke. *Stroke*, 50(5), 1263-1265.

- Koenig, M. A. et al. (2019). Machine learning prehospital large vessel occlusion prediction from clinical assessment. *Stroke*, 50(Suppl_1), A134.
- Murray, N. M. et al. (2020). Artificial intelligence to diagnose ischemic stroke and identify large vessel occlusions: A systematic review. *Journal of NeuroInterventional Surgery*, 12(2), 156-164.
- Nogueira, R. G. et al. (2018). Thrombectomy 6 to 24 hours after stroke with a mismatch between deficit and infarct. *New England Journal of Medicine*, 378(1), 11-21.
- Qiu, W. et al. (2020). Machine learning for detecting early infarction in acute stroke with non-contrast-enhanced CT. *Radiology*, 294(3), 638-644.
- Saver, J. L. (2006). Time is brain -- quantified. *Stroke*, 37(1), 263-266.
- Sheth, S. A. et al. (2019). Machine learning-enabled automated determination of acute ischemic core from computed tomography angiography. *Stroke*, 50(11), 3093-3100.
- Strbian, D. et al. (2013). Thrombolytic therapy in acute ischemic stroke: Outcome for every delay of treatment by 10 minutes. *Stroke*, 44(3), 838-840.
- Turc, G. et al. (2019). European Stroke Organisation guidelines for intravenous thrombolysis in acute ischaemic stroke. *European Stroke Journal*, 4(3), 192-201.
- Yahav-Dovrat, A. et al. (2021). Evaluation of artificial intelligence-powered identification of large-vessel occlusions in a comprehensive stroke center. *American Journal of Neuroradiology*, 42(2), 247-254.
- Powers, W. J. et al. (2019). Guidelines for the early management of patients with acute ischemic stroke. *Stroke*, 50(12), e344-e418.
- Campbell, B. C. V. et al. (2019). Imaging selection in ischemic stroke: Feasibility of automated CT-perfusion analysis. *International Journal of Stroke*, 10(1), 51-54.
- Mouridsen, K. et al. (2020). Artificial intelligence applications in stroke. *Stroke*, 51(8), 2573-2579.
- Bivard, A. et al. (2017). Perfusion CT in acute stroke: A comprehensive analysis of infarct and penumbra. *Radiology*, 284(1), 120-128.
- Hassan, A. E. et al. (2020). Early experience utilizing artificial intelligence shows significant reduction in transfer times and length of stay in a hub and spoke model. *Interventional Neuroradiology*, 26(5), 615-622.
- Berkhemer, O. A. et al. (2015). A randomized trial of intraarterial treatment for acute ischemic stroke. *New England Journal of Medicine*, 372(1), 11-20.

Declarations

Funding

Supported by the Swiss National Science Foundation (grant 320030_204418) and the French National Research Agency (ANR-19-CE45-0020). No funder influenced study design, data analysis, or publication decisions.

Conflict of Interest

Both authors declare no competing interests. Neither has received consulting fees, research funding, or equity from Viz.ai, RapidAI, Brainomix, or any other AI stroke technology company.

Data Availability Statement

De-identified stroke pathway time metrics and ML model performance data are available from the corresponding author on reasonable request. Neuroimaging data cannot be shared due to patient privacy regulations. All model architectures and the SDES calculator are available on request.

Ethical Approval

Approved by the Cantonal Ethics Committee of Zurich (BASEC 2019-01648), CPP Ile-de-France VI (2019-A02186-51), Comitato Etico Milano Area 3 (N. 312-2019), and De Videnskabsetiske Komiteer Region Hovedstaden (H-19048762). The retrospective component received consent waiver; the prospective AI deployment received institutional quality improvement approval.

Appendix A

SDES Calculation and Model Architecture Specifications

This appendix details the SDES scoring methodology, the CNN architecture for CT/CTA analysis, the prehospital gradient-boosted model features, and the multi-modal fusion network design.

Part I -- SDES Dimensions

- 1a. Diagnostic Accuracy: Macro-F1 for stroke subtype (ischaemic, haemorrhagic, mimic, TIA)
Against neuroradiologist + neurologist consensus
- 1b. Time Reduction: (baseline time - AI time) /
baseline time Higher reduction = higher
score
- 2a. LVO Detection: Sensitivity for anterior + posterior
LVO on CTA Against interventional
neuroradiologist ground truth
- 2b. ICH Exclusion: Negative predictive value for
haemorrhage on NCCT Must exceed 99% for
clinical acceptability
- 3a. Workflow Integration: Expert panel rating (8
neurologists, 4 neuroradiologists) Ease of
integration into existing stroke protocols
- 3b. $SDES = 0.20 * Accuracy + 0.25 * Time + 0.20 * LVO$
 $+ 0.15 * ICH + 0.20 * Workflow$

Part II -- Model Specifications

- 4a. K2 CT/CTA: 3D ResNet-50, input 512x512x32
slices, ASPECTS regions segmented 18-sec
processing, 0.3-sec LVO alert
- 4b. K1 Prehospital: XGBoost, 28 features (12 NIHSS
+ vitals + history) Runs on tablet, < 5 sec
inference
- 5a. K3 Perfusion: U-Net segmentation + rCBF/Tmax
thresholding Core and penumbra volumes in
45 sec
- 5b. K5 Fusion: Late fusion attention network, clinical
(64d) + imaging (256d) + lab (32d) embeddings
..... Dashboard update in 90 sec from scan
completion