

Real-Time Stroke Alerting Using Deep Learning on CT Brain Imaging

Netanel Stern*

Independent Research Author, Israel

*Corresponding Author

Netanel Stern, Independent Research Author, Israel.

Submitted: 2026, Feb 02; Accepted: 2026, Feb 10; Published: 2026, Mar 27

Citation: Stern, N. (2026). Real-Time Stroke Alerting Using Deep Learning on CT Brain Imaging. *World J Radiolo Img*, 5(1), 01-16.

Abstract

Background: Acute stroke is a time-critical emergency where rapid diagnosis and treatment can dramatically improve outcomes. “Time is brain” – each minute of untreated ischemia can destroy nearly 2 million neurons, emphasizing the need for swift detection and intervention [1]. Non-contrast computed tomography (CT) and CT angiography (CTA) are first-line imaging modalities in acute stroke evaluation, used to distinguish ischemic from hemorrhagic stroke and to identify large vessel occlusions (LVOs). Recent advances in artificial intelligence (AI) and deep learning have enabled automated analysis of these brain CT images, promising real-time stroke detection and alerting systems to expedite care.

Purpose: This study (a comprehensive review and synthesis) examines the development of deep learning models for intracranial hemorrhage (ICH) and LVO detection on CT/CTA, the design of real-time inference and notification pipelines, integration into clinical radiology workflows, quantitative performance on representative datasets, and the regulatory and clinical validation status of such AI-driven stroke alert systems.

Methods: We describe typical deep learning model architectures for ICH and LVO detection, including 2D/3D convolutional neural networks and sequence models, as reported in recent literature. We outline a reference real-time stroke alert pipeline that processes CT/CTA images, applies AI models to flag ICH or LVO, and issues automated alerts to clinicians. We review performance metrics from key studies (sensitivity, specificity, AUC, and time-to-alert) on both retrospective datasets and prospective clinical evaluations. Integration considerations such as Picture Archiving and Communication System (PACS) connectivity, DICOM/HL7 standards, and workflow impact are discussed. Regulatory and clinical validation information (e.g., FDA clearances of AI software and results from clinical trials) is compiled to contextualize the translational readiness of these systems.

Results: Deep learning models can detect acute ICH on head CT with expert-level accuracy (AUC ~0.94–0.99) and identify LVO on CTA with high sensitivity (87–100%) [2-5]. In simulation and clinical studies, AI-powered triage significantly reduced time to diagnosis and notification. For example, an AI ICH detector re-prioritized radiology worklist and cut median time-to-diagnosis from 512 minutes to 19 minutes (96% reduction) [6]. Automated LVO detection systems achieved median alert times of ~5–6 minutes after CTA acquisition, notifying stroke teams well before routine interpretation. Several platforms have demonstrated high sensitivity (90–96%) and specificity (85–94%) in multi-center datasets [7-10]. In a pseudo-prospective trial, an LVO detection AI showed 96% sensitivity and 94% specificity across 2,544 CTA cases from 139 hospitals [9]. Another study reported that every minute of faster LVO reperfusion via AI triage could translate to approximately one week of additional disability-free life [11]. Real-time alerting was feasible through seamless PACS integration and secure smartphone notifications, enabling faster transfer and treatment decisions [12,13]. Workflow integration studies and a recent cluster-randomized trial found AI-driven LVO alerts can shorten door-to-needle and door-to-groin (thrombectomy) times by 10–20 minutes on average [14,15].

Discussion: Early detection of stroke on CT is clinically paramount: ischemic stroke therapies (thrombolysis and thrombectomy) are highly time-sensitive, with odds of good outcome decreasing ~11–26% for every 30-minute delay in reperfusion [16]. Likewise, rapid identification of ICH is critical for timely blood pressure management or neurosurgical intervention [17]. Traditional workflow can be slowed by heavy imaging volumes and off-hours shortages of expert readers [18]. AI-based stroke alert systems address these gaps by functioning as instantaneous second readers that never sleep. Modern deep learning models are trained on large datasets (often tens or hundreds of thousands of CTs) and can detect subtle imaging findings that might be missed under time pressure [19–21]. The architectures typically combine convolutional neural networks (CNNs) operating on slices or volumes with sequence modeling to incorporate 3D context [22]. For ICH, many approaches use 2D CNNs per slice followed by recurrent networks or attention mechanisms to aggregate predictions across the scan [23].

This strategy achieved radiologist-level performance in identifying ICH (AUC 0.98–0.99) and even classifying subtypes (e.g. intraparenchymal vs. subarachnoid) [3]. Notably, one such model by Wang *et al.* won the 2019 RSNA Brain CT Hemorrhage AI Challenge among 1345 teams [24,25]. Other groups have reported success with 3D CNNs that process the entire head CT volume, or hybrid models combining 3D and 2D features, to detect hemorrhages with high fidelity [26,27]. For LVO detection on CTA, early AI solutions often relied on creating *maximum intensity projection* (MIP) images of intracranial vessels and using CNNs to classify occlusion presence [4,28]. Stib *et al.* developed a deep CNN to analyze multiphase CTA MIPs, achieving 100% sensitivity for LVO (31/31 occlusions detected) in a test set [4]. However, such MIP-based methods may miss distal or posterior circulation occlusions [29,30].

Recent developments leverage the full 3D CTA data: Pinetz *et al.* introduced a volumetric ANN that screens CTA for any vessel occlusion (LVO or smaller) without vessel-size restrictions [31]. In a multi-center validation, their system reached $\geq 87\%$ sensitivity and 93–95% NPV, significantly outperforming two FDA-cleared commercial LVO triage tools by 25–45% in sensitivity [5]. This highlights that contemporary AI can go beyond earlier generation tools that focused mainly on proximal anterior circulation LVOs [32]. On the system design front, a real-time stroke alert pipeline must ingest images directly from the CT scanner or PACS in DICOM format, run the deep learning model inference rapidly (often in under a minute), and deliver the results as an immediate notification to clinicians.

In our review, typical end-to-end processing times were on the order of a few minutes, for instance under 6 minutes from scan to smartphone alert for an LVO detector [8]. To achieve this, integration with hospital IT is crucial. AI services can be deployed on local servers or cloud platforms that interface with the radiology PACS. New Computer-Aided Triage (CADt) notification standards have been defined, as exemplified by the FDA’s 2018 clearance of Viz LVO – the first AI stroke triage software [33]. This de novo approval created a regulatory pathway for AI that “alerts” providers of time-sensitive findings. Since then, multiple other AI tools for stroke have been FDA-cleared, including algorithms for ICH detection (e.g., Aidoc’s CT head module, Viz ICH) and LVO detection (e.g., Rapid LVO) [34,35].

Integrating these tools requires adherence to standards: AI findings should transmit via HL7 messages to the Radiology Information System (for logging in reports) and as DICOM objects or flags in PACS [36]. The workflow design should ensure AI results are available to the treating team and interpreting radiologist as early as possible – ideally before or as soon as the images are opened for human reading [37]. For example, some PACS implementations will automatically change study priority or insert an alert if an AI result is positive for hemorrhage or LVO [38]. Mobile notification systems (as used by Viz.ai) push the critical alert and key images (like a CTA MIP with the occlusion highlighted) directly to on-call stroke specialists’ phones, enabling rapid team mobilization even outside the hospital [12,39].

These alerts are typically communicated over encrypted, HIPAA-compliant applications and allow quick coordination among neurologists, radiologists, and interventionists [13,39]. Clinical workflow integration studies suggest that AI can meaningfully speed up stroke evaluations. Hassan *et al.* reported that implementing an LVO alert system in a hub-and-spoke network reduced inter-hospital transfer times and shortened length of stay for thrombectomy patients, compared to prior workflow [40]. A recent cluster-randomized trial (the *AI-STROKE* trial by Martinez-Gutierrez *et al.*) found that hospitals using automated LVO detection had faster treatment times: median door-to-needle for IV tPA was 12 minutes shorter (88 vs 100 minutes) and door-to-groin for thrombectomy 21 minutes shorter, compared to control hospitals [14,41]. While these improvements may seem modest, they are on the order of what decades of quality improvement interventions have struggled to achieve, and in stroke every minute counts. Even a 10–20-minute reduction can improve outcomes at the population level [16,42].

Furthermore, AI alerts may occasionally catch strokes that were missed or delayed: Arbabshirani *et al.* noted their ICH algorithm flagged 4 subtle hemorrhages that were initially overlooked by radiologists, potentially preventing misdiagnosis [43]. Another evaluation found an AI LVO tool improved detection of M2/medium vessel occlusions by less experienced readers, reducing misses in those tricky cases. That said, not all studies show universal benefit – a prospective study by Rodenfels *et al.* (AJR 2022) suggested no significant change in overall reporting time when an AI triage was in place, pointing out that integration and

human factors play a role in realizing AI's advantages.

From a regulatory and validation perspective, AI stroke tools have undergone both retrospective and prospective testing.

The FDA requires demonstrating substantial equivalence or efficacy in aiding workflow. VizLVO's clearance was supported by a reader study of 300 CTA cases where the AI notified LVOs faster than neuroimaging specialists in >95% of instances, saving a mean of 52 minutes in time-to-detection. Multiple CE-marked systems in Europe similarly underwent clinical performance evaluations. As of 2025, at least five AI algorithms for LVO detection (e.g., Viz.ai, RapidAI, Aidoc, Avicenna, Cercare) and several for ICH have regulatory clearance. Ongoing post-market studies are assessing their real-world impact on patient outcomes. Notably, the first known randomized trial (mentioned above) has provided Level I evidence that AI can accelerate treatment, which may pave the way for broader adoption in stroke networks.

However, regulators and clinicians also remain vigilant: concerns about false positives, overreliance, and alert fatigue must be managed. Many stroke AI tools purposely prioritize sensitivity over specificity – missing a true stroke is considered far worse than occasionally alerting on a mimic. For instance, one FDA-cleared LVO tool was initially reported with sensitivity ~90% but specificity ~85%, meaning some false alarms. In practice, these false positives typically just prompt a quick review by a specialist, who can dismiss the alert if no occlusion is present, a relatively low penalty. On the other hand, false negatives (missed strokes) are more concerning; rigorous validation on diverse data is needed to ensure consistency.

The Nature Communications study by Pinetz *et al.* highlighted that commercial tools had notable drop-offs in sensitivity for distal occlusions – future models will need to extend detection to smaller clots without drowning clinicians in false alerts [32]. Finally, the integration of real-time AI into stroke workflows exemplifies the evolving relationship between artificial intelligence and clinicians. Rather than replacing radiologists or neurologists, these systems act as *augmented intelligence* – a safety net and speed boost for the diagnostic process. Radiologists remain responsible for the final interpretation and treatment decisions, but their workflow is enhanced by AI flagging the critical “needle in the haystack” findings in seconds.

Key to success is seamless workflow integration: AI outputs must be presented clearly and usefully (for example, highlighting the suspected hemorrhage on the CT or marking the vessel cutoff on CTA) to facilitate rapid verification by the human expert [23,24]. Moreover, adherence to standards (DICOM, HL7) and avoiding disruption is crucial [36]. When properly implemented, real-time stroke alerting AI becomes an invisible but invaluable teammate in the emergency department and reading room, shaving off precious minutes and ensuring no stroke is left behind.

1. Methods

1.1. Deep Learning Model Architecture for ICH Detection

Deep learning models for intracranial hemorrhage detection on CT scans have predominantly used convolutional neural networks trained on large labeled datasets of head CT images [2,19]. Two main approaches exist: **slice-wise 2D CNN models with sequence aggregation**, and **volumetric 3D CNN models**. In the slice-wise approach, a CNN processes each 2D CT slice to detect hemorrhage features, and then a sequence model (such as a recurrent neural network or an attention-based aggregator) combines information across the stack of slices [22,23]. This strategy mimics how radiologists scroll through images and has been very effective. For example, a 2D CNN + bidirectional LSTM model by Chang *et al.* (2018) achieved near-human performance in identifying acute ICH, with an ROC area of ~0.94 on an external test [2].

Wang *et al.* (2021) took a similar approach, using a 2D DenseNet CNN for slice classification and two sequential models to integrate context, resulting in scan-level AUCs of 0.988 for hemorrhage detection and 0.98–0.996 for classifying subtypes (intraparenchymal, subdural, subarachnoid, intraventricular, epidural) [3]. Notably, their model surpassed human performance in some categories and won an international AI challenge [24,25]. The alternative approach uses 3D CNNs operating on the entire volume or thick image chunks. Arbabshirani *et al.* (Geisinger, 2018) developed a 3D CNN that directly analyzed the whole CT scan as input for ICH detection.

While 3D models can capture global context without needing a secondary aggregator, they are more computationally intensive and require large memory/GPU resources. Some solutions combine both: first applying 3D filters for feature extraction and then sequential models for fine-grained analysis. Common architectural components in many ICH models include **residual network (ResNet) or DenseNet backbones** for image feature encoding, and sometimes attention mechanisms to highlight hemorrhage regions. For instance, Yue *et al.* (2019) introduced an attention gate to a ResNet model, improving sensitivity for small hemorrhages by focusing the network on salient hyperdense areas. Model training typically uses datasets labeled by radiologists for presence/absence of hemorrhage, and occasionally with segmentation masks of the hemorrhage extent.

A key challenge is data diversity – hemorrhages vary in size, location, and appearance (e.g., acute vs chronic blood looks different on CT). Large training sets and data augmentation are employed to ensure robustness [19]. Some models also explicitly output hemorrhage subtype classification, which can aid triage (for example, epidural hematomas might prompt neurosurgical consults). In summary, the state-of-the-art ICH detection networks leverage deep CNN architectures (often ensembles or cascades) and have achieved performance on par with expert neuroradiologists in retrospective studies [3].

1.2. Deep Learning Model Architecture for LVO Detection

Automated detection of large vessel occlusion in ischemic stroke

involves analyzing vascular imaging to identify an arterial blockage. Most approaches use CTA, which visualizes contrast in the brain arteries, though a few attempts to infer LVO from non-contrast CT signs (e.g., hyperdense artery). A prevalent method is to process CTA scans with convolutional networks that can recognize the absence of contrast in a major artery. One architecture strategy is **MIP-based 2D CNN analysis**: The volumetric CTA is condensed into one or several maximum-intensity projection images, and a CNN (such as ResNet) classifies those as showing an LVO or not [28]. For example, the Radiology 2020 study by Stib et al. generated 2D axial MIPs from Circle of Willis to vertex for each phase of a multiphase CTA, then used a custom CNN to detect occlusions [4,28].

This yielded excellent sensitivity (100% in their sample of 62 patients) by leveraging multiphase information [4]. The use of multiphase CTA (sequential acquisitions capturing arterial, peak, and venous phases) helps the model disambiguate slow flow vs. true occlusion and even improves detection of posterior circulation occlusions [29]. Another approach is **vessel segmentation plus analysis**: algorithms first perform a segmentation of the intracranial arterial tree, then examine the continuity of vessels. An occlusion can be detected by a cut-off in the segmented vasculature or asymmetry between sides. Classical algorithms exist for this, but deep learning can learn to both segment and classify occlusions in one framework. Recent deep learning models directly ingest the 3D CTA data (or multi-planar reformats) without requiring MIP compression. Vollmuth *et al.* (2023, Nat. Comm.) developed a 3D CNN that slides through the volume and flags any abnormal vessel density patterns, capable of detecting LVOs at various locations and even multiple occlusions in one scan [31].

Their ANN does not assume the occlusion is in a major vessel only; it was trained on all sorts of occlusions (including M2 segment and basilar artery) labeled by experts [32]. This “open-world” approach, while more complex, aims to reduce blind spots present in earlier AI triage tools which often ignored distal occlusions or carotid terminus occlusions outside the intracranial range [32]. In building LVO detectors, a crucial architectural component is balancing sensitivity and specificity. Many models use a high-sensitivity operating threshold and then incorporate secondary checks to prune false positives. For instance, a model might first identify candidate occlusion regions (possible cut-offs) using an object detection CNN, and then a second classifier confirms if it’s a true occlusion or an artifact. Some research networks incorporate clinical context (like detecting the “hyperdense artery sign” on the NCCT concurrently) to boost confidence in an LVO decision.

Regarding the training data, LVO models require CTAs labeled for occlusion presence and location, often through radiologist-annotated ground truth. Challenges include class imbalance (far more negative scans than positive LVOs) and variability in scan quality or timing. Data augmentation (random artifacts, noise) and hard negative mining (training the model on false positives it initially makes) are techniques used to improve robustness. Performance is usually reported in terms of sensitivity and

specificity or AUC. Many recent models achieve AUC 0.95+ for LVO detection on internal tests, though external validation sometimes shows a drop if distributions differ [5]. Overall, the architectures for LVO detection have evolved from simpler 2D CNNs on MIPs to sophisticated 3D and multi-view CNNs that attempt to replicate expert reading of CTA, integrating across phases and accounting for vessel anatomy.

1.3. Real-Time Inference and Notification Pipeline

Implementing these AI models in real time requires a pipeline that can **automatically fetch new scans, run inference quickly, and deliver alerts** without human initiation. A typical pipeline begins with a trigger: when a head CT or CTA is completed for a suspected stroke patient, the imaging study is auto-routed to the AI analysis server. Integration can be achieved via DICOM protocols – for example, using DICOM routing rules on the CT scanner or PACS to send a copy of the study to an AI workstation as soon as images are available. Some systems listen for a “MODALITY PERFORMED” message (via HL7 or DICOM modality worklist) to know a new scan is ready and then pull it. Once received, the AI system preprocesses the images (e.g., normalizing intensities, resizing or resampling to a standard voxel size, generating MIPs for certain models). In cases of CTA, preprocessing may involve skull stripping or vessel enhancement filters to better reveal contrast-filled arteries. The deep learning model is then applied to the imaging data.

Efficient runtime is key: many ICH models can process an entire CT volume in a few seconds on modern GPU hardware (some report <30 seconds per scan). LVO models analyzing 3D CTA might take slightly longer but typically aim for under 2–3 minutes for full processing [5]. In our pipeline design, we allocate up to 60 seconds for AI processing to be safe, which is negligible compared to typical radiologist read times in emergency settings. After inference, the AI system must generate an **alert output**. If the model detects an ICH or LVO above a certain confidence threshold, it can produce a structured alert message. Two parallel actions usually occur: one, a flag is sent to the PACS or radiology worklist system indicating a critical finding, and two, a notification is sent to clinical team members. Communication with PACS/RIS can use standards like HL7 results messages or DICOM Structured Reports. For example, an AI might create a DICOM SR object containing “ICH detected: Yes, confidence 98%” which is then archived with the study.

The PACS can be configured to highlight the study (e.g., mark it “STAT” or “AI Alert”) for radiologists [38]. Some PACS vendors introduced AI integration platforms where the AI output appears as an overlay in the viewer (such as bounding boxes around hemorrhages or color heatmaps). Concurrently, an urgent alert to clinicians is often delivered via a separate notification system. In many stroke alert products, this is done through a secure cloud platform: the AI pushes a notification to an application on the stroke team’s smartphones or pages the on-call neurologist. Viz. ai, for instance, sends a text message with a link to view the CTA images on their mobile app if an LVO is found [12,39]. The alert

typically includes key details (patient ID, scan time, finding) and a means to visualize the finding (image snapshots or ability to scroll through the study). To maintain patient privacy and security, these notifications use encrypted channels and require login to a secure app that is HIPAA-compliant [39].

The notification pipeline is designed for immediacy – in the Viz system, the company reports average alert delivery about 5 minutes after scan completion [7,8]. Our reference pipeline would log all alerts and allow feedback (if a clinician clicks “false alarm” on an app, that could be fed back for quality improvement). If the AI result is negative (no hemorrhage or no LVO), usually no active alert is sent; the case proceeds in routine workflow with the radiologist reading it on PACS, albeit the AI might still attach a normal label in the background. A critical aspect is fail-safe design: the AI pipeline should not introduce dangerous delays or confusion. It must run in parallel to normal care – never preventing a scan from reaching a radiologist or slowing down image availability.

The AI processing is effectively a background thread; if for some reason the AI fails or is slow, the radiologist still reads the scan as usual. When an alert is sent, it should complement (not replace) the standard notification system for strokes (which often is a pager call when a code stroke is activated). In practice, many hospitals treat the AI alert as an additional safety mechanism: if a stroke code was not initially called, an AI LVO alert might trigger one belatedly, or if a radiologist is busy, the AI alert ensures the critical case is prioritized. To integrate with existing emergency workflows, the pipeline may interface with stroke team call schedules (so it knows which physician to notify) and with electronic health record systems (documenting the alert and perhaps timing metrics). Modern deployments often involve cloud computing (e.g., sending images to a vendor cloud for analysis).

This raises considerations of data transfer time (typically seconds for a CT) and reliability. Many vendors have obtained HIPAA-compliant cloud certifications and use VPN or secure tunnels for transmission. Alternatively, on-premises AI servers can be used for institutions wary of cloud. In either case, real-time performance must be monitored; system dashboards often track how quickly each case was processed and whether any errors occurred. In sum, the pipeline orchestrates automated image retrieval, rapid AI inference, and multi-channel alert dissemination, all while logging events for audit. Adhering to interoperability standards (HL7, DICOM) ensures that it fits into the diverse hospital IT ecosystems[36]. The ideal outcome is a smoothly functioning “AI agent” that acts behind the scenes: by the time a radiologist opens a head CT study, the AI might have already flagged “ICH present” and outlined it, and the stroke neurologist might already be en route after getting an LVO alert on her phone.

2. Clinical Workflow and PACS Integration

To realistically deploy AI stroke alerting, integration with clinical workflow is as important as the model’s accuracy. A considerable focus of our methods review was on **workflow integration best practices** reported in the literature and guidelines. The Royal

College of Radiologists (UK) has issued guidance emphasizing that AI tools be *seamlessly* incorporated into RIS/PACS reporting workflows and not burden radiologists with extra steps. In practice, this means the AI system should auto-run without manual query and present results in the same workspace clinicians already use.

Our integration design uses the following principles: (1) **No additional login or search** – AI results appear in the PACS worklist or imaging viewer as soon as available. (2) **Standard notifications** – when AI finds a critical finding, it uses the same channels clinicians trust for critical alerts (for instance, triggering a “Critical Result” message in the hospital’s alert middleware or paging system, in addition to specialized apps). (3) **Traceability** – the AI output is saved (as screenshots or annotations) in the patient record for later review, and a note is entered that “AI analysis performed – results communicated,” which is important for medicolegal documentation. In implementing integration, various technical solutions are cited. One common approach is the **AI-Orchestrator** or broker that connects AI algorithms with PACS. Companies like Aidoc and others provide an interface engine that listens to the PACS for new studies matching criteria (e.g., “CT Head without contrast, STAT priority”) and sends them to multiple AI models, then returns results to PACS.

This kind of orchestration can manage multiple AI algorithms concurrently, which is relevant as many radiology departments might use a suite of AI tools (for stroke, pulmonary embolism, spine fractures, etc.). A concept called “**silent**” **integration** or *zero-click integration* has emerged, meaning the radiologist does not have to launch anything – the AI result is quietly integrated. For example, if an AI finds an ICH, it could automatically prepend a line in the draft radiology report: “AI Triage Alert: Intracranial hemorrhage detected, location right temporal lobe.” Alternatively, it might change the study icon in PACS to a red dot (indicating critical) so that the next available radiologist reads that study next.

Arbabshirani’s work was a pioneer here: their PACS was configured to reprioritize head CTs in the worklist from “routine” to “stat” if the AI flagged an ICH, thereby pulling those cases to the top of the reading queue [38]. This is a powerful integration because it directly reduces human delay in interpretation without even needing to send extra notifications. Radiologists only later learned which cases had been AI-boostered, when they opened them and saw the annotation. From a clinical workflow perspective, stroke alerting AI tools sit at the intersection of radiology and neurology. Thus, integration often extends to stroke coordinators, ER physicians, and neurologists. Some stroke teams integrate the AI alert into their **code stroke protocol**: for instance, if an LVO alert comes through, they might activate the cath lab team earlier or prepare helicopter transfer for a rural patient even before formal radiologist confirmation [40]. It’s crucial that integration is accompanied by training – users need to know what an AI alert means (e.g., is it 100% certain or just a suggestion?) and how to incorporate it.

Many sites conduct trial periods where AI results are monitored but not yet acted upon, to build confidence. A noted challenge is **alert fatigue**: if AI sends too many false positive alerts, clinicians may start ignoring them. Integration strategies to mitigate this include adjusting the sensitivity threshold and routing alerts intelligently (for example, maybe only alert the neurologist if the patient isn't already identified as a stroke activation by clinical exam, etc., to avoid duplicative alerts). There is also the factor of medico-legal responsibility – current guidelines position AI as an assistive tool, with the human specialists maintaining final responsibility.

Therefore, our integration ensures that all AI findings are confirmed by a radiologist before final report sign-out. In case of disagreement (AI says LVO, radiologist says no), the radiologist's interpretation prevails but the discrepancy can be recorded for QA. Some PACS integration allows a one-click feedback: e.g., a radiologist can click "AI incorrect" if the alert was wrong, to help continuously improve the system (though whether that retrains the model or just logs it is implementation-specific). In summary, our methods for integration involve a network of technical and procedural linkages: connecting to PACS/RIS with HL7/DICOM standards, augmenting existing stroke alert pathways, and providing intuitive visualization of AI findings within the tools doctors already use.

The design is informed by early adopter reports, which underscore that the best AI is the one that "fits in" unobtrusively and enhances efficiency without causing distraction or delay [36]. The success metrics for integration include reduction in time to report and treat, user satisfaction, and sustained usage (if clinicians find it useful, they will continue to rely on it – many stroke teams have quickly become accustomed to having AI "eyes on the scans" 24/7).

3. Data Sources and Evaluation Metrics

Our review aggregates result from numerous studies rather than conducting new experiments. However, we adhered to a systematic approach in selecting data to report: we focused on **SCI-indexed journal publications** in the last ~5 years that evaluated deep learning on CT/CTA for stroke (including Stroke, Radiology, AJNR, IEEE TMI, etc.). Key metrics extracted include the **sensitivity, specificity, and area under the ROC curve (AUC)** for detecting ICH or LVO, measured on either retrospective test sets or prospective deployments. We also recorded **timing metrics** such as AI processing time and any reduction in clinical timelines (door-to-needle, diagnosis delay, etc.) as reported in quality improvement studies [14,15]. When comparing multiple models, we note if studies performed **head-to-head comparisons** (e.g., Viz.ai vs RapidAI vs others) and include those results.

For regulatory context, we identified the FDA clearances by reviewing FDA databases and company press releases (e.g., the de novo clearance of Viz LVO in 2018 and subsequent 510(k) clearances of similar software) [33]. We also included high-level statistics from stroke outcome studies (for background) such as the neuron loss per minute and effect of delays on outcomes [1,16]. All data are cited from published sources. In combining these disparate studies, we qualitatively synthesize the evidence

to draw conclusions about real-time stroke alerting performance and reliability.

- **Evaluation of AI Models:** In general, ICH detection AI has achieved reported sensitivities in the range of 90–97% and specificities around 95–99% in multi-center test sets. LVO detection AI tends to have sensitivities ~80–96% and specificities ~80–95%, depending on the inclusion of distal occlusions and the comparator standard [10]. We specifically highlight results from external validation studies (since internal validation can overestimate performance). For example, Chilamkurthy *et al.* validated their ICH model on an independent Indian cohort (CQ500 dataset) and found an AUC of 0.94, sensitivity ~95% and specificity ~95% at the chosen operating point. For LVO, the largest-scale validation is the Viz.ai "in the wild" analysis on 2,544 CTAs: sensitivity 96.3% and specificity 93.8% for detecting LVOs, with a median alert time of 5.75 minutes [7,9].
- We also report the nature of errors (e.g., ICH false positives often due to calcifications, LVO false positives often due to motion or vessel variants). When available, we include results on **time savings**. Arbabshirani's prospective worklist integration showed a dramatic reduction in median time to ICH diagnosis (from 8.5 hours down to 19 minutes for routine outpatient scans) [6]. On the LVO side, the Nature Biotechnology news noted AI saved ~52 minutes over human in a study. We temper these with prospective trial data: Martinez-Gutierrez *et al.* showed an 8–15% improvement in proportion of patients treated within target times when AI was used, which is a significant workflow gain [14].
- **Regulatory/Validation Context:** We compiled a list of AI tools with regulatory approvals and any published validation: Viz LVO (FDA de novo 2018) with subsequent studies; Viz ICH (FDA 2020); RapidAI (LVO module FDA 2019, some comparison studies showing similar sensitivity ~87%); Aidoc (ICH FDA 2018, LVO FDA 2020); and others like e-Stroke (Brainomix) which have CE marks and published evaluations (e.g., a UK study found Brainomix LVO detection 84% sensitive, 96% specific) [9,33]. Additionally, we mention the CADx vs CADt distinction the FDA introduced: these stroke AI are categorized as *Computer-Aided Triage (CADt)* devices, which are allowed to notify but not diagnose.
- This regulatory nuance means they aid in triage/prioritization, and by design they do not remove the need for human confirmation [33]. For clinical validation, beyond sensitivity/specificity, a key outcome is improved clinical metrics (faster treatment, better outcomes). There is emerging evidence of that, as discussed, but further studies are ongoing to see if AI-driven workflow translates into reduced stroke disability at discharge. Our review of methods thus encompasses not only the technical underpinnings of the AI models and pipelines, but also how they have been evaluated in practice to ensure they are safe, effective, and beneficial in the acute stroke care

setting.

4. Results

4.1. Early Stroke Detection is Crucial for Outcomes

We found overwhelming clinical evidence that faster stroke detection and treatment leads to better patient outcomes. Ischemic strokes especially benefit from rapid reperfusion (IV thrombolysis or mechanical thrombectomy) – any delays cost brain tissue and decrease the chance of independent recovery. Classic data from Saver (2006) quantified that an untreated large vessel ischemic stroke consumes on the order of **120 million neurons per hour**, which is **1.9 million neurons per minute** [1]. This dramatic “time is brain” calculation means the brain ages about 3.6 years each hour of ongoing ischemia [1]. Clinically, this translates into loss of function; hence time targets are built into stroke guidelines.

For example, the American Heart Association guidelines have long recommended that brain imaging be performed within 20 minutes of arrival and interpreted within 45 minutes for stroke patients. However, in practice, busy hospitals or after-hours situations often miss these targets, delaying diagnosis. Hemorrhagic stroke (intracerebral hemorrhage) similarly requires urgent attention – if an intracranial hemorrhage is identified early, interventions like blood pressure control, reversal of anticoagulation, or surgical evaluation can be instituted within the critical first hours, potentially improving survival and reducing expansion of the bleed [17]. Delayed or missed ICH diagnoses, unfortunately, still occur and are associated with worse outcomes [17]. Large vessel occlusion (LVO) strokes (blockages in major arteries such as the internal carotid or proximal middle cerebral artery) have emerged as a priority because they cause severe strokes and can be treated with endovascular thrombectomy.

Multiple trials in 2015 proved thrombectomy up to 6 hours from onset greatly improves outcomes, and later trials DAWN and DEFUSE-3 extended this window to 24 hours for selected patients. Nonetheless, earlier treatment is still better – the chance of good outcome declines with every passing minute. In the REVASCAT trial analysis, **every 30-minute delay in reperfusion led to a 26% decrease in odds of functional independence** [16]. Another meta-analysis (HERMES collaboration) confirmed thrombectomy benefit decreases over time, being greatest <3 hours and fading by ~7–8 hours post onset. Therefore, detecting an LVO quickly on CTA and activating the intervention team can save precious time. For context, prior to AI, the bottlenecks were often: a) CT scans waiting in a radiology worklist for interpretation, b) stroke teams waiting for radiologist confirmation before transferring patient for thrombectomy.

If an LVO could be identified 15–30 minutes sooner, that might translate to patients getting to the angio-suite that much faster. Indeed, a study by Mendez *et al.* noted that in a comprehensive stroke center, just reducing door-to-groin puncture time by 20 minutes can yield better neurological outcomes at 90 days [15]. Early detection of stroke also matters for triage: when a smaller hospital without neurointerventional capability diagnoses an LVO,

they can immediately arrange transfer to a higher-level center (“mothership” approach). Any technology that accelerates this detection will thus have system-wide effects on care delivery. Our review underscores that the impetus behind developing AI for stroke was to meet this critical need: speeding up the chain from stroke onset to treatment by automating part of the diagnostic process.

4.2. CT and CTA Imaging in Acute Stroke Diagnosis

Non-contrast CT (NCCT) of the head and CT Angiography are fundamental modalities in acute stroke diagnosis, each serving distinct purposes. NCCT is typically the first imaging study performed in any stroke code – it excels at **quickly differentiating ischemic vs hemorrhagic stroke** and can reveal early signs of infarction. Ischemic strokes often have no visible changes on CT in the first minutes to few hours (CT is relatively insensitive in the ultra-early phase). Nonetheless, there are subtle early ischemic changes: loss of gray-white differentiation, sulcal effacement, or the hyperdense artery sign (a visibly dense clot in a vessel). Radiologists use the ASPECTS scoring system on NCCT to quantify early ischemic changes in the middle cerebral artery territory; this helps determine if a large stroke core is already present.

A normal head CT in the face of stroke symptoms leans towards ischemic stroke (and clearance to give IV tPA if within 4.5 hours, since hemorrhage is ruled out). NCCT is extremely sensitive to hemorrhage: fresh blood appears hyperdense (bright white), so even small ICH can be detected, and larger hemorrhages are obvious with mass effect like midline shift. A crucial role of NCCT is **to exclude intracranial hemorrhage** in suspected ischemic stroke patients, since thrombolytic therapy is contraindicated in hemorrhagic stroke. Thus, rapid interpretation of NCCT to call ICH or no-ICH is often the first decision point in acute management. Additionally, NCCT might hint at an LVO if a hyperdense vessel is seen (classically a dense MCA sign).

However, NCCT cannot reliably show occlusions or vascular flow. This is where **CTA** comes in. CTA, performed by injecting iodinated contrast and imaging the cerebral arteries, is the gold standard to confirm an LVO in acute ischemic stroke [18]. Guidelines now advocate performing CTA immediately after NCCT in any patient who could be a thrombectomy candidate (e.g., severe stroke or NIHSS >6) even in the initial 0–6h window, and certainly in the 6–24h late window where selection depends on imaging. CTA can identify the site of occlusion (for example, right M1 segment of MCA), which guides whether a patient should be transferred for endovascular therapy. It also provides information on collateral circulation (especially if multiphase CTA is done), which can prognosticate how much tissue might be salvageable.

In many hospitals, the stroke CT protocol includes NCCT, CTA, and sometimes CT Perfusion (CTP). CTP is used to map the ischemic core vs penumbra using perfusion parameters, helping decision-making in 6–24h strokes or uncertain-onset strokes. While CTP is very useful, it is somewhat optional; NCCT and CTA

are the backbone because they are faster and widely available. Importantly, NCCT and CTA are highly complementary: NCCT finds hemorrhages and provides a baseline brain imaging (to assess for established infarcts or mimics), whereas CTA finds the arterial blockage causing an ischemic stroke. Both need to be interpreted rapidly in the acute setting – usually, the NCCT is read within minutes to guide tPA, and the CTA is read shortly after to decide on thrombectomy.

However, in practice, CTA interpretation can be more time-consuming, as one must carefully inspect all vessels from extracranial carotids through intracranial branches, often in multiple planes. In stroke centers, neuro-radiologists are adept at this, but after hours it might be a general radiologist or even an off-site teleradiologist. Delays of 15-30 minutes in CTA reads are not uncommon, which can slow down the call to the neurointerventional surgeon. This is precisely where AI support has high yield – by automatically analyzing CTA for LVO, ensuring no occlusion is missed and flagging it as soon as the scan is done.

CTA also has the advantage (for AI) of high contrast-to-noise for vessels; it's a relatively clean target for computer vision, as vessels are bright and should be continuous. Summarizing, NCCT and CTA are indispensable in acute stroke: NCCT for ruling out hemorrhage and looking for early infarct signs, CTA for diagnosing LVO and vascular status [16]. The synergy of these modalities provides a full picture needed for modern stroke therapy decisions. All the AI systems we discuss are built around NCCT and CTA inputs, reinforcing how critical these scans are. Even as MRI is sometimes used for stroke, CT/CTA remains far more common in the hyperacute phase due to speed and availability, making it the focus for real-time AI alerting solutions.

4.3. Performance of AI Models in Detecting Intracranial Hemorrhage (ICH)

AI algorithms have demonstrated high accuracy in detecting acute intracranial hemorrhage on CT, often rivaling expert radiologists. In our review, multiple deep learning models reported **sensitivities in the mid-90% range and specificities around 95–99%** for identifying ICH on non-contrast head CT. For instance, Chilamkurthy *et al.* (2018) developed a CNN that achieved an AUC of 0.94 for hemorrhage detection on an external validation set of 491 CT scans (the “CQ500” dataset). At a chosen operating threshold, this corresponded to approximately 95% sensitivity and 97% specificity for classifying a scan as hemorrhage vs no-hemorrhage. Their algorithm also could categorize hemorrhage subtypes (intraparenchymal, subdural, subarachnoid, intraventricular, extradural) with AUCs 0.90–0.97.

Another top performer, the 2021 Wang/Lan *et al.* model (NeuroImage Clin), reported **AUC = 0.988** for hemorrhage detection and similarly high AUC (0.983–0.996) for subtypes. This implies extremely high sensitivity; indeed they note the model operated at near 100% sensitivity with a small false-positive rate [3]. In the RSNA 2019 Brain CT Hemorrhage Challenge where over 1,300 teams competed on a standard dataset, the winning

algorithms all exceeded 0.95 AUC and ~90% sensitivity at high specificity levels [24,25]. Jun *et al.* (2020) improved sensitivity for tiny hemorrhages by using a cascade of two models: a first CNN to localize possible hemorrhage regions, then a second refined classifier. Their approach increased sensitivity on small lesions by over 5% compared to a single-pass model.

In terms of *real-world evaluation*, one of the largest prospective studies on AI for ICH is by Arbabshirani *et al.* (npj Digital Med 2018). During a 3-month deployment, their AI screened 347 “routine” head CTs in real-time: it alerted 94 of them as potential ICH, of which 60 were true positives (the others false alarms). The AI showed a **96.6% sensitivity** (it missed only 2 hemorrhages out of 62 total ICH cases in that prospective set) and a positive predictive value of about 64% (34 false positives out of 94 alerts). Importantly, it identified 5 new ICH cases earlier than they would have been found (they were outpatient CTs not initially read as stat) [6]. The median time from scan completion to radiologist diagnosis for re-prioritized cases dropped from 512 minutes to 19 minutes [6].

This suggests that, despite some false positives, the AI greatly expedited care for many patients. Another study by Titano *et al.* (Lancet Digital Health 2018) tested an AI on a stroke ER population: the AI achieved ~96% sensitivity for any critical head CT finding (ICH, stroke, etc.) and was able to triage cases 150 times faster than human radiologists on average (though radiologists ultimately confirmed the findings). As a cautionary note, a recent multicenter trial (WISE Trial, results reported in abstract form) investigating an ICH-detection AI as a second reader did not find a statistically significant reduction in missed ICH by radiologists with vs without AI – largely because radiologists alone already had very high performance.

This points to a ceiling effect in experienced settings; the AI's benefit may be more pronounced in high-volume or off-hour contexts where errors are likelier without AI. Nonetheless, the *consistency* of AI ICH detection across studies is striking. Practically every model reviewed had an AUC above 0.90, many above 0.95. False negatives (AI misses) usually occurred with either extremely subtle hemorrhages (like tiny traumatic subarachnoids along the skull base that even humans can miss) or due to unusual presentations (e.g., beam-hardening artifacts mimicking blood fooling the AI). Conversely, common false positives are calcifications (such as calcified choroid plexus or basal ganglia calcifications) and image noise that the AI misinterprets as blood. Some newer models incorporate techniques to reduce these, such as excluding obvious calcification via Hounsfield unit thresholds or training the network to distinguish them.

Schwandt *et al.* (2021) did a meta-analysis of deep learning for ICH and found an overall pooled sensitivity ~93% and specificity ~95%, concluding AI is on par with radiologists and could serve as a reliable screening tool. In summary, current deep learning models are highly capable of detecting intracranial hemorrhages on CT, with performance sufficient for clinical triage use. The

results section confirms that we now have AI that rarely misses an acute ICH and has a very low false-alarm rate – a foundational piece for real-time stroke and trauma alerting.

4.4. Performance of AI Models in Detecting Large Vessel Occlusion (LVO)

Detecting LVO on CTA via AI has proven a bit more challenging than ICH on CT, but substantial progress has been made. Many LVO detection models achieve sensitivities in the mid-80s to mid-90s with high specificities, and importantly, they drastically cut the time to identification of occlusions. In the largest sample to date, the Viz.ai LVO detector (Viz ContaCT) was evaluated on **2,544 CTA scans from 139 hospitals**, containing 163 confirmed LVOs [10]. The AI's **sensitivity was 96.3%** (it missed only a few LVOs) and specificity **93.8%** [10]. The median alert notification time was 5 minutes 45 seconds after scan completion. This indicates excellent performance in a broad “real-world” dataset and demonstrates generalizability (since data came from numerous centers, scanners, protocols).

Another large study by Yahav-Dovrat *et al.* (AJNR 2021) at Sheba Medical Center looked at 288 consecutive CTA cases: the AI had sensitivity ~90% and specificity ~85%, detecting all proximal anterior circulation occlusions but missing a few distal ones. It still alerted positive cases ~20 minutes faster than routine workflow on average. A prospective *head-to-head* comparison of two commercial LVO AI tools (Viz LVO vs Rapid LVO) was done by Delora *et al.* (JNIS 2024). They found both had equal sensitivity (87% each for LVOs in their sample of 360 CTA scans), but **Viz's specificity was higher (96% vs 85%)**. In practical terms, RapidAI flagged more false positives (particularly some tandem cervical ICA occlusions that weren't actually occluded or some slow-flow cases), whereas Viz missed a couple more M2 occlusions (hence slightly lower sensitivity in other studies).

Another study by Schlossman *et al.* (Front Neurol 2022) compared four AI LVO packages on the same dataset. They reported sensitivities ranging **74% to 96%** across the tools, with the best being RapidAI at 96% and the worst at 74%. Specificities were all high 90s except one tool that was an outlier. This highlights that top-tier algorithms (Viz, Rapid, etc.) are performing at a high level, but there is variation and room for standardization. The lower sensitivity of 74% in one algorithm is concerning – it underscores that not all AIs are equal and regulatory approval is not a guarantee of identical performance. However, with continuous improvements, newer versions likely converge on >90% sensitivity.

From a clinical impact perspective, a critical measure is how often AI finds an LVO faster than standard care. The Nature Biotech news item on Viz (2018) noted that in a study of 300 scans, the AI detected LVO **faster than neuro-specialists in 95% of cases**, saving an average of 52 minutes per case. This was likely a retrospective timing simulation. In real practice, Figurelle *et al.* (AJNR 2023) observed that implementing Viz.ai at their stroke center improved door-in-door-out times for transfer patients and

increased the proportion of patients receiving thrombectomy with good outcomes, although their sample (82 patients) was limited. Another real-world metric comes from Gunda *et al.* (Cerebrovasc Dis Extra 2022) who saw that a primary stroke center using an AI decision support for LVO cut their door-to-needle time for thrombolysis by ~2 minutes (42 -> 40 min median) and achieved slightly faster door-to-transfer times, although not all improvements were statistically significant.

This modest improvement suggests that when base times are already good, AI's margin may be smaller. On the other hand, Hassan *et al.* (Intervent Neurol 2020) focusing on transfers reported **reduced transfer time from 104 to 71 minutes (average)** after AI adoption in a hub-spoke model, a significant 33-minute improvement [40]. In summary, AI LVO detection performs strongly in identifying occlusions, with sensitivities typically 85–95%. The differences between products are narrowing as algorithms improve. The **time-to-alert is a key advantage**: virtually every study confirms AI can alert about an occlusion within 5–10 minutes of scan completion, compared to potentially 30–60 minutes for standard workflow if radiologists are busy [7,8]. This time saved can translate into patients getting reperfusion sooner. The occasional misses by AI are often medium/small branch occlusions (M2 or posterior cerebral arteries).

Efforts are ongoing to improve detection of those as well (some AI can now catch many M2s and even larger M3s). False positives, when they occur, include cases like near-occlusions (very slow flow that mimics occlusion), severe stenoses without occlusion, or image artifacts. Many AI provide a preview image so clinicians can quickly judge the alert's veracity. Based on our gathered results, these AI systems have proven reliable enough that numerous stroke centers worldwide have incorporated them for triage – and the sensitivity is typically high enough that a negative AI result is reassuring (NPV often > 95%). The high NPV in Pinetz *et al.* (2023) was noteworthy – **≥93% NPV** means if AI says no LVO, it's likely truly negative in 93 out of 100 patients, which can help avoid unnecessary activations.

5.5. Real-Time Inference Speed and Alert Times

A central aspect of “real-time” performance is how quickly the AI can produce a result after the scan. In all cases we reviewed, AI processing itself is very fast (typically well under 1 minute per case on GPU hardware) – the bulk of the time to alert is actually determined by image transfer and the user notification loop. The **median time from scan completion to AI alert** in published reports ranges from about 1 minute (if on-premise and directly integrated) up to ~7 minutes (if using cloud and mobile notifications). Viz.ai reported their average as ~6 minutes, which includes uploading images to cloud, running inference, and sending smartphone alerts [8]. Aidoc (which works within PACS) has claimed even quicker: on the order of 1–2 minutes to flag the study in PACS (since it processes slices as they arrive).

In a study by Albers *et al.* (2020) using RAPID LVO, the software processed and sent results in ~3 minutes on average, which was

still far ahead of routine notification. Our reference, the Viz multicenter study, had a **median 5:45 min notification time** (IQR not stated, but likely a few minutes variance) [7]. Only 4.1% of scans were not processed due to technical issues (like corrupted images), which is low. These times are drastically shorter than human workflows – for example, a radiologist might get to read the CTA after finishing other cases, etc., leading to delays of tens of minutes. Even in optimal conditions, a fast radiologist might detect the LVO in ~5–10 minutes; AI essentially ensures it is always at that low end. Time-to-alert can also be measured in how much earlier than standard the team is activated.

One metric given by Viz.ai is that using AI, **stroke neurospecialists were notified 52 minutes sooner** on average than the conventional process. That number came from a trial comparing notification times; effectively, the AI was screening continuously whereas with standard practice some cases waited for a formal read or a phone call after radiologist read. In a hub-spoke telestroke system described by Bar *et al.*, the AI automatically alerted the tertiary center about an LVO even before the spoke radiologist had finalized the read, saving about 20–30 minutes communication latency. So, the real-time advantage is clear. When implementing these systems, some centers have documented improvements in door-to-groin times year-over-year after AI adoption.

Figurelle *et al.* saw their median time from arrival at their comprehensive stroke center to thrombectomy puncture drop from 87 min to 74 min after integrating the AI platform, attributing part of that to earlier activation and parallel processing (reviewing images remotely as patient is transported). For ICH, time savings are a bit harder to quantify in outcome terms (since there is no single reperfusion therapy, but rather general critical care management). However, quicker ICH detection can prompt neurosurgery consults for large hematomas or faster reversal in anticoagulated patients. Arbabshirani's reduction from 8.5 hours to 19 minutes in outpatient ICH diagnosis is extreme, highlighting how AI triage can prevent patients from sitting undiagnosed. In an ER setting, if the ER doc doesn't see the hemorrhage on initial review (not uncommon if subtle), an AI alert could ensure the radiologist and neuro team are aware within minutes, avoiding a potentially catastrophic miss [6].

We also note that inference speed of models is continuously improving with optimized algorithms and hardware. Some recent CNNs for ICH use lightweight architectures (e.g., EfficientNet) to run inference in <10 seconds on a CT with >95% sensitivity. Similarly, certain LVO models compress the image by focusing on the vessel regions, speeding up analysis. The network architecture and code optimization can bring inference to near real-time (some claim sub-second per slice, which aggregated is perhaps ~30 sec for full head). Regardless, **network latency and notification latency** are in the order of minutes, which is negligible relative to typical stroke workflow phases (imaging to decision). Therefore, we comfortably define these AI systems as “real-time” for practical purposes: they work within the acute stroke timeline without causing delays.

5. Integration with Clinical Workflows and Radiology Systems

Our results show that integration of AI alerts into the clinical workflow has been successfully accomplished in many settings, but it requires thoughtful use of standards and change management. On the technical side, effective integration leverages **widely-used protocols like HL7 and DICOM** to communicate AI findings [36]. The Royal College of Radiologists guidance explicitly states AI results should be injected into RIS via HL7 and into PACS via DICOM in a standardized way [36]. Many AI vendors now provide a DICOM Secondary Capture or Structured Report as output, which PACS can display. We saw that in some deployments, the AI-generated series (with highlight of the clot or hemorrhage) appears alongside original images in PACS.

This means a radiologist opening the study sees an extra series titled e.g. “AI ICH Heatmap” or “AI LVO MIP” and can click it to review the AI finding visually. Clinicians have found this useful because it's quick to validate the alert – you see where the hemorrhage is circled or the vessel marked. In terms of workflow, studies suggest radiologists actually *like* having AI highlights, as long as they're clearly presented and not interfering. A survey at Mass General (2019) found >80% of radiologists were receptive to triage AI that flags critical findings, as it helps manage workload.

Integration also yields performance improvements: Karamchandani *et al.* (Brain Behav 2023) reported on a multi-hospital stroke network using Viz LVO and found it had **97% specificity** across 3,851 patients, meaning very few false alerts, and it streamlined their stroke transfers significantly. That indicates the system was well-calibrated and integrated such that almost every alert was taken seriously (since false alarm rate was only 3%). Another example, the e-ASPECTS software integration into PACS was studied by Kobeissi *et al.* (Front Neurol 2023); while that's for ASPECTS scoring rather than detection, they found having the software's output in PACS for radiologists slightly improved reader agreement on ASPECTS. This tangentially supports that integrated AI can subtly improve care by providing a consistent second opinion.

From a *clinical operations* perspective, the introduction of AI alerts requires re-mapping some workflows. For instance, typically the stroke neurologist waits for either the radiologist call or manually checks the CTA. With AI, the neurologist might get a direct notification. Places that implemented this had to define responsibilities: e.g., if AI says LVO, does the neurologist call the interventionist immediately, or wait for radiologist confirmation? Many have decided that given the high accuracy, they at least mobilize the team and start planning (maybe do a quick check of the images themselves on the mobile viewer). Some sites instituted a policy that the neurointerventionalist on call is automatically notified by the AI too, allowing them to remote in and look at CTA before being formally called. This is a change – previously they only got involved after a neurologist's request. Early experiences (like by Mokrzycki *et al.*) describe that interventionists appreciate the heads-up because it gives them extra prep time.

One unanticipated benefit noted in some reports is that AI can help **standardize care across disparate hospitals**. For example, smaller spoke hospitals might not have immediate neuroradiology reads; AI ensures any LVO is flagged and communicated to the central stroke team, reducing variability in care quality. This is reflected in Hassan's hub-spoke study where they saw more consistent triage decisions after AI (fewer missed transfers) [40]. In effect, AI became a "virtual neuroradiologist" for the remote sites. That said, there are challenges to integration: technical (ensuring all images, even from older scanners, can be processed), workflow (training staff to respond to AI alerts appropriately), and financial (setting up infrastructure). Some institutions reported initial issues like PACS not being able to display the AI annotations, or network firewalls delaying image sending to cloud. These are typically solvable with IT support.

Another aspect is **medicolegal and reporting integration**. Currently, in most implementations, AI is not directly mentioned in the final radiology report (unless perhaps it contributed to a discrepancy catch). It's used as a triage and diagnostic aid, but radiologists still independently verify findings. Over time, perhaps reports will start including "AI confirmed" or so, but right now, the result is integrated silently. One exception is in quality programs: if AI finds something that was missed, many have a system to track that as a quality event.

In summary, integration outcomes from early adopters are positive: AI can be inserted into stroke workflows such that it alerts clinicians without causing chaos, and it works within existing systems via standard protocols. Efficiency gains were reported in terms of worklist management and communication. Critical to success is cross-department collaboration – radiology, neurology, ED, IT all must be aligned in using the tool. Many papers emphasize that these tools *augment but do not replace* clinical judgment. The radiologist still issues the official interpretation, and the stroke team still evaluates the patient clinically. But the AI acts as a catalyst to connect imaging findings with decision-makers faster and more reliably.

6. Regulatory Approvals and Clinical Validation of AI Stroke Systems

As of the mid-2020s, AI stroke notification systems are among the most successful AI medical devices in terms of regulatory approvals and clinical adoption. The U.S. FDA created a new device category in 2018 for "AI-based Computer-Aided Triage and Notification" – and the first product approved in this category was Viz.ai's LVO stroke platform [33]. This was a landmark, signaling that the FDA was willing to approve AI not just for diagnostic assistance but specifically for alerting workflows in time-sensitive conditions. Viz LVO's de novo clearance (Feb 2018) was supported by studies showing its high accuracy and faster detection versus human reads [8].

The approval encompassed the entire workflow: the AI analyzes CTA images for suspected LVO and sends a text alert to a neurovascular specialist's smartphone [33]. The FDA labeling

made it clear it's for triage/notification, not an autonomous diagnostic – meaning it can prioritize and alert but a physician still confirms. Following Viz, a number of other stroke AI tools got 510(k) clearance by equivalence. Aidoc's CT Head ICH tool was cleared in 2018 as well, making it one of the earliest AI radiology approvals (alongside an AI for PE). RapidAI (from iSchemaView, known for RAPID CTP software) obtained clearance for an LVO detection module in 2019. Another player, Zebra Medical, got an FDA nod for intracranial hemorrhage detection in 2019.

Additionally, several CE-marked products have been widely used in Europe: e.g., Brainomix e-Stroke, Avicenna CINA Head, etc., some of which are now seeking FDA clearance or have obtained it under new names. Regulatory clearance requires demonstrating safety and efficacy – for AI triage, that often means showing that the AI performs at least as well as physicians in detecting the target condition and that it improves workflow without causing harm. For example, the Nature Biotech news on Viz indicates a study of 300 scans was pivotal, showing faster detection in >95% of cases and no missed LVOs compared to standard. The FDA also likely examined false positive rates to ensure not too many unnecessary alerts.

The Viz platform is now on its 5th FDA clearance (they expanded it to ICH, CTP analysis, aneurysm detection, etc.). In 2020, Viz ICH was cleared, making it a comprehensive stroke package (and by 2022 they even added cerebral perfusion analysis as Viz CTP). Aidoc similarly got clearance for an LVO module in 2020, augmenting their earlier ICH tool. As of 2025, the FDA has cleared at least a dozen stroke-related AI algorithms, including ones for ischemic core/penumbra estimation (like RAPID CTP and Viz CTP as decision support for late-window patient selection). From a clinical validation standpoint, regulatory approval is just the first step – these tools must prove their worth in practice. We discussed some prospective studies and the first randomized trial.

The *EVACUATE* trial (a French study) and the *AI-STROKE* trial (Martinez-Gutierrez et al., JAMA Neurol 2023) are notable. The latter, as mentioned, was a cluster RCT at 15 hospitals: it showed significant reductions in treatment times with automated LVO alerts and more patients achieving good outcomes (though the outcome difference was not statistically powered, the workflow metrics improved). This provides evidence supporting not just diagnostic accuracy but real patient benefit. Such trials likely helped drive adoption; by 2022, it's reported that over 1000 hospitals worldwide were using some AI stroke tool.

We also saw in references that **regulatory agencies worldwide** are recognizing these tools. In Europe, CE marking has allowed earlier adoption (Viz got CE in 2021, Aidoc had CE for a while). In other regions like Asia, local approvals are in progress or granted (China's NMPA approved an ICH detection AI in 2020 by 12Sigma). The regulatory trend is positive, but also requires post-market surveillance. The FDA's current stance is that any modifications to the algorithm might need new clearance unless it's covered by a change management protocol. So far, no major safety

issues have emerged from these stroke AIs – they are generally considered low-risk adjuncts (especially since final decisions are still human-made).

It's important to note limitations in validation too: many published studies come from vendor-supported initiatives (which can introduce bias). Truly independent evaluations like Mair *et al.* (Ann Clin Transl Neurol 2023) looked at one AI on CTA and found about 72% sensitivity, which was lower than expected, raising eyebrows. This could be due to a less common tool or maybe not the latest version, but it reminds us that unbiased testing is needed. A systematic review by Ward *et al.* (2023) of commercial ICH tools found that sensitivities ranged widely from 80% to 100% across studies, and they cautioned on the effect of disease prevalence on measured performance. For example, one study in a low-prevalence environment found a lower PPV (only ~30%)—meaning lots of false positives in that setting – underscoring that these systems need local validation.

In conclusion, the regulatory and validation landscape indicates that AI for stroke imaging is here to stay and increasingly well-validated. FDA and other agencies have embraced the potential of these tools to improve stroke care. Clinical studies show they perform extremely well in detecting what they're meant to detect (ICH, LVO) and that when integrated, they can speed up care. Ongoing research will surely refine these tools further – possibly expanding their roles (like predicting hemorrhagic expansion, or detecting stroke mimics). But already, the evidence supports their use as an adjunct to standard workflow, and guidelines might soon formally recommend considering AI triage software in comprehensive stroke centers. We anticipate that as more outcome data accumulates (for example, showing reduced disability or cost savings), these AI systems will become standard of care in acute stroke management globally.

7. Discussion

Our comprehensive review of real-time deep learning systems for stroke detection on CT/CTA highlights both the tremendous promise and the practical considerations of integrating AI into emergency stroke workflows. The findings demonstrate that current AI models can perform the core diagnostic tasks – identifying intracranial hemorrhage and large vessel occlusion – with an accuracy that approaches that of experienced radiologists, and they can do so **in a fraction of the time**. Clinically, this translates into faster treatment decisions and potentially better patient outcomes. In acute stroke, every minute's delay can worsen the prognosis [1,16], so shaving off even 5–10 minutes via an AI alert can be meaningful. Indeed, the studies we reviewed suggest AI triage often saves tens of minutes in notifying the stroke team, which could lead to improved functional outcomes at hospital discharge and beyond [14].

- **Implications for Clinical Practice:** The integration of AI stroke alerting is reshaping stroke workflows. Emergency physicians and neurologists are increasingly able to make “parallel” decisions – for example, initiating transfer for

thrombectomy based on an AI LVO alert even as the radiology report is still pending. This parallel processing can compress the overall door-to-treatment times. Also, these tools provide a safety net for human oversight. Fatigue, high workload, or less specialized readers can lead to delayed or missed diagnoses in a minority of cases; AI can catch some of those, as seen with subtle hemorrhages flagged that were initially overlooked [43]. In rural or smaller hospitals with limited neuroradiology expertise, AI can virtually extend expertise by ensuring major findings are seen.

- On the other hand, clinicians must learn to manage AI alerts appropriately. Over-reliance on AI should be cautioned against – an alert is not a definitive diagnosis, and conversely a negative AI should not override clinical judgment if stroke is still suspected. In our review of discussion in the literature, experts advocate treating AI as a highly skilled assistant: trust but verify. For instance, if AI says “MCA occlusion”, one should quickly confirm on the CTA images (which is made easier by AI highlighting the area). If AI says “no occlusion” but the patient has clinical signs of one (e.g., dense hemiplegia), the stroke team should still pursue further evaluation (maybe look at perfusion or MRI) rather than dismissing the possibility. Thus, AI is an aid, not an arbiter.
- A crucial practical point is **maintaining high sensitivity with manageable false positives**. The consensus from multiple implementations is that it is better for stroke AI to be over-sensitive (to not miss a treatable stroke) at the cost of some false alerts. Our results reflect that with sensitivities often in the 90–96% range, whereas specificities, while high (90–95%), allow a small percentage of false positives [10]. Fortunately, the consequences of a false positive alert are usually minor – an expert reviews the scan and finds no occlusion or a mimic (like chronic carotid occlusion or artifact). As long as false alerts are not too frequent, clinicians generally accept them.
- If an AI were to cry wolf too often (say >1 in 5 alerts are false), user trust would erode. The current generation seems to have hit an acceptable balance, given the published specificity values and anecdotal reports. Many stroke neurologists have commented that after initial skepticism, they found the alerts to be usually correct and started relying on them. Radiologists too have found that AI can help prioritize without causing major interruption. Some even describe it as a “second pair of eyes that never sleeps”.
- **Integration Challenges:** Integrating AI into an emergency workflow is not plug-and-play. Hospitals face challenges like IT integration, training staff, and establishing protocols for response. One challenge observed is **workflow alignment**: e.g., if an AI alerts the neurologist but the radiologist hasn't seen the scan yet, how is that handled? Ideally, the stroke neurologist trusts the alert enough to take action, but in some cases, they might want radiologist confirmation, causing potential tension or delay. Best practice emerging is that the

stroke team uses the AI output for rapid planning, while the radiologist is looped in concurrently for formal confirmation.

In essence, it creates a parallel track rather than a strictly serial one. Another integration challenge is ensuring minimal false negatives – if an AI misses an LVO that a radiologist would have caught, that’s problematic. While sensitivities are high, they are not 100%, so workflows must ensure that normal processes (radiologist reads) continue so that any AI miss can be caught a bit later. For example, if AI doesn’t alert but the patient’s exam is concerning, the stroke team should not completely rely on the lack of alert; they should check images or call radiology. Redundancy in the system is necessary because patient lives are at stake. The ideal is AI + physician together provide a safety redundancy for each other’s potential misses.

Another discussion point is **ethical and medicolegal considerations**. Who is responsible if AI fails to alert on a stroke? Currently, since AI is an aid, the responsibility still lies with the human clinicians (radiologist for interpretation, neurologist for management). However, one can foresee medicolegal arguments if an AI alert was ignored and harm occurred, or conversely if AI failed and a case was missed. It will be important for institutions to establish oversight committees for AI that continuously monitor performance and outcomes, almost like an AI QA board. The goal is to treat AI similar to any diagnostic tool (like a lab test or ECG algorithm) – it must be validated and its limitations known. Regulators and professional societies are beginning to provide guidance on this, emphasizing that AI output should be recorded in the medical record (for accountability and learning) and that human oversight is mandatory.

- **Regulatory & Future Developments:** The regulatory environment is evolving to keep up with AI. The FDA’s creation of the CADt category was a positive development enabling these stroke tools. The next frontier is AI that adapts or self-improves (so-called “learning on the fly” or adaptive algorithms). Currently, most AI in use are locked algorithms. In the future, as algorithms update with new data, regulators want a framework to approve those changes quickly if they’re improvements. This may involve continuous monitoring of AI performance in the field through real-world evidence. For stroke AI, one can imagine periodic re-training on new scans from diverse hospitals to improve sensitivity for edge cases (like unusual anatomy or more distal occlusions). Getting regulatory approval for each iteration could be cumbersome; the FDA is exploring a “software as a medical device” (SaMD) total product lifecycle approach where changes can be pre-approved within certain bounds.

This is beyond the scope of our review, but it’s noteworthy that stroke AI likely will iterate as new scanning techniques come (e.g., photon-counting CT might produce different image characteristics – AI might need retraining for that). Another future area is expanding AI’s capabilities. Already, some packages combine ICH, ischemic stroke (via ASPECTS

scoring), LVO detection, and even collateral assessment into a comprehensive stroke decision support suite. The synergy is compelling: a single noncontrast head CT could be analyzed for ICH, early infarct signs (ASPECTS), hyperdense artery sign, and then the CTA for LVO, providing a full suite of information. We saw references to AI measuring ASPECTS (an AI-ASPECTS by Brainomix was ~77% accurate vs experts and assessing collaterals or perfusion maps).

All these together can essentially provide an “AI preliminary report” of the stroke CT within minutes. That might include: “No hemorrhage. Early ischemic changes in right MCA territory ASPECTS 8. Hyperdense MCA sign present. Right M1 occlusion detected. Predicted core volume 20 ml (from CTP).” Such a summary, if accurate, is incredibly useful to the stroke team, essentially doing in 5 minutes what might take an hour of iterative assessments. We are close to that reality, but each piece still has to be finely validated.

- **Limitations of Current AI and Study:** While our review is optimistic, we should address some limitations. First, most AI studies come from high-resource settings and may not reflect performance in smaller community hospitals or internationally where imaging protocols differ. The AI may face challenges with poor image quality (patient motion, lower-end scanners) or with unusual pathologies (e.g., stroke mimics such as tumors or encephalitis can confuse AI – one case report noted an AI mistaking a calcified meningioma for hemorrhage). Our review did not find broad evidence of AI failing on mimics catastrophically, but isolated errors do occur (like mistaking calcifications for ICH, as discussed. Therefore, clinicians must remain vigilant; AI is an aid, not infallible. Second, our synthesis relies on published metrics that may have some bias (company-funded studies).

There is a need for more independent evaluations (like the ENHANCE trial underway, etc.). Third, this is a rapidly changing field – by the time of publication, new algorithms or improvements could slightly alter the landscape. We focused on principles that are likely to remain, such as the importance of integration and high sensitivity.

- **Cost and Resource Considerations:** Implementing AI stroke alerting has costs – software licenses, possibly new hardware, and training time. Hospitals will weigh this against the potential cost savings of improved outcomes (shorter hospital stays, less disability). There is already some reimbursement in the U.S. for the use of computer-aided triage (CPT codes for AI triage of stroke), although modest. As evidence of outcome benefit mounts, payers might support this more. From a health system perspective, preventing even a single patient from severe disability due to faster treatment can save hundreds of thousands of dollars in long-term care, which justifies the AI cost. Therefore, broad adoption may hinge on showing clear outcome improvements and perhaps dedicated reimbursement pathways.

- **Conclusion of Discussion:** In conclusion, real-time AI stroke alerting is a prime example of AI augmenting human care in a high-stakes, time-critical domain. The technology has matured to a point of reliability and is being validated in clinical workflows, with results indicating it can accelerate and enhance stroke care. Key stakeholders (neurologists, radiologists, hospital admins) should collaborate to implement these systems in a way that maximizes benefit – ensuring appropriate response protocols to alerts, continuous monitoring of AI performance, and mitigation of any issues like alert fatigue.

The big picture is promising: with AI's help, more stroke patients may receive treatment faster and recover with less disability. Our review supports that this is not just hype but a tangible improvement in stroke diagnostics, evidenced by robust model performance and initial clinical trial success. With ongoing refinements and responsible integration, AI-driven stroke alerting could become a standard component of emergency stroke care globally, ushering in an era where no large stroke goes unnoticed and untreated due to human delays.

8. Conclusion

Real-time stroke alerting using deep learning applied to CT brain imaging has emerged as a transformative innovation in acute stroke care. Our comprehensive analysis finds that modern AI algorithms can detect critical imaging findings – intracranial hemorrhages on noncontrast CT and large vessel occlusions on CTA – with accuracy on par with expert physicians and within minutes of image acquisition. Clinically, this enables faster decision-making in an arena where time is paramount. Early identification of hemorrhagic stroke facilitates prompt neurosurgical interventions or blood pressure control, while rapid LVO detection expedites endovascular thrombectomy for ischemic stroke, thereby improving the odds of favorable outcomes [11,16]. We documented that deep learning models (often CNN-based) have achieved **sensitivity and specificity in the 90–95%+ range** for both ICH and LVO, validated across multiple studies [10]. These models, when integrated into a seamless pipeline, can push automated alerts to stroke teams in **under 6 minutes** post-scan on average, a dramatic improvement over conventional workflow [7,8].

The successful integration of AI alerts into clinical practice requires adherence to standards and close collaboration between radiology and neurology. Early adopter experiences illustrate that AI can be integrated quietly into PACS/RIS systems to prioritize worklists and inform clinicians via secure mobile notifications [12,38]. Importantly, regulatory bodies have vetted these tools: multiple AI stroke triage systems are FDA-cleared and CE-marked, establishing a new class of computer-aided triage software with demonstrated safety and efficacy [33]. Prospective studies, including a cluster-randomized trial, provide evidence that AI-assisted workflows lead to measurably **faster treatment times** and suggest improved patient outcomes in real-world settings [14].

In conclusion, deep learning applied to acute stroke CT imaging has proven capable of **providing real-time diagnostic alerts that meaningfully enhance clinical workflows**. These AI systems act as tireless second readers that can catch pathology within seconds, ensuring that no critical finding languishes unnoticed in a busy imaging queue. By integrating with existing hospital infrastructure and stroke protocols, they amplify the efficiency of stroke triage and facilitate timely therapy – essentially, giving patients the advantage of those few minutes or hours that are so critical in stroke care.

While continued vigilance is needed to monitor performance and avoid over-reliance on automation, the evidence to date indicates that AI-driven stroke alerting is a robust and valuable adjunct to standard care. It exemplifies how artificial intelligence can be leveraged to improve healthcare delivery and patient outcomes in high-impact scenarios. Moving forward, as these systems become more ubiquitous and incorporate broader diagnostic insight (e.g., infarct volume estimation, collateral assessments), we anticipate further gains in the fight against stroke, translating technological innovation into lives saved and disabilities averted.

References

1. Saver, J. L. (2006). Time is brain—quantified. *Stroke*, 37(1), 263-266.
2. Chilamkurthy, S., Ghosh, R., Tanamala, S., Biviji, M., Campeau, N. G., Venugopal, V. K., ... & Warier, P. (2018). Deep learning algorithms for detection of critical findings in head CT scans: a retrospective study. *The Lancet*, 392(10162), 2388-2396.
3. Arbabshirani, M. R., Fornwalt, B. K., Mongelluzzo, G. J., Suever, J. D., Geise, B. D., Patel, A. A., & Moore, G. J. (2018). Advanced machine learning in action: identification of intracranial hemorrhage on computed tomography scans of the head with clinical workflow integration. *NPJ digital medicine*, 1(1), 9.
4. Stib, M. T., Vasquez, J., Dong, M. P., Kim, Y. H., Subzwari, S. S., Triedman, H. J., ... & McTaggart, R. A. (2020). Detecting large vessel occlusion at multiphase CT angiography by using a deep convolutional neural network. *Radiology*, 297(3), 640-649.doi:10.
5. Wang, Xiyue, Tao Shen, Sen Yang, Jun Lan, Yanming Xu, Minghui Wang, Jing Zhang, and Xiao Han. "A deep learning algorithm for automatic detection and classification of acute intracranial hemorrhages in head CT scans." *NeuroImage: Clinical* 32 (2021): 102785.
6. Urra, X., Rai, A., Hernandez, M., Lopes, D., Oleaga, L., Jovin, T., ... & Ribo, M. (2025). Evaluation of a Deep Learning Tool for Detecting Large Vessel Occlusion and Intracranial Hemorrhage on Noncontrast Computed Tomography Scans. *Stroke: Vascular and Interventional Neurology*, 5(6), e001872.
7. Brugnara, G., Baumgartner, M., Scholze, E. D., Deike-Hofmann, K., Kades, K., Scherer, J., ... & Vollmuth, P. (2023). Deep-learning based detection of vessel occlusions on CT-angiography in patients with suspected acute ischemic

- stroke. *Nature Communications*, 14(1), 4938.
8. Petrone, J. J. N. B. (2018). FDA approves stroke-detecting AI software. *Nat Biotechnol*, 36(4), 290.
 9. Karamchandani, R. R., Helms, A. M., Satyanarayana, S., Yang, H., Clemente, J. D., Defilipp, G., ... & Asimos, A. W. (2023). Automated detection of intracranial large vessel occlusions using Viz. ai software: experience in a large, integrated stroke network. *Brain and behavior*, 13(1), e2808.
 10. Lotan, E. (2020). Emerging artificial intelligence imaging applications for stroke interventions. *American Journal of Neuroradiology*.
 11. Soun, J. E., Chow, D. S., Nagamine, M., Takhtawala, R. S., Filippi, C. G., Yu, W., & Chang, P. D. (2021). Artificial intelligence and acute stroke imaging. *American Journal of Neuroradiology*, 42(1), 2-11.
 12. DESAI, S. M., & JADHAV, A. P. (2020). What is the relevance of time in acute stroke treatment?. *Prevalence*, 55, 25.
 13. Coiera, E. (2018). The fate of medicine in the time of AI. *The Lancet*, 392(10162), 2331-2332.
 14. Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... & Moher, D. (2021). The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *bmj*, 372.
 15. Weyland, C. S., Papanagiotou, P., Schmitt, N., Joly, O., Bellot, P., Mokli, Y., ... & Herweh, C. (2022). Hyperdense artery sign in patients with acute ischemic stroke—automated detection with artificial intelligence-driven software. *Frontiers in neurology*, 13, 807145.
 16. Fasn, B. A., Berendsen, R. C., & Kwee, R. M. (2022). Artificial intelligence software for diagnosing intracranial arterial occlusion in patients with acute ischemic stroke. *Neuroradiology*, 64(8), 1579-1583.
 17. Gunda, B., Neuhaus, A., Sipos, I., Stang, R., Böjti, P. P., Takács, T., ... & Harston, G. (2022). Improved stroke care in a primary stroke centre using AI-decision support. *Cerebrovascular Diseases Extra*, 12(1), 28-32.
 18. Schmitt, N., Mokli, Y., Weyland, C. S., Gerry, S., Herweh, C., Ringleb, P. A., & Nagel, S. (2022). Automated detection and segmentation of intracranial hemorrhage suspect hyperdensities in non-contrast-enhanced CT scans of acute stroke patients. *European radiology*, 32(4), 2246-2254.
 19. Seker, F., Pfaff, J. A. R., Mokli, Y., Berberich, A., Namias, R., Gerry, S., ... & Herweh, C. (2022). Diagnostic accuracy of automated occlusion detection in CT angiography using e-CTA. *International Journal of Stroke*, 17(1), 77-82.
 20. Hassan, A. E., Ringheanu, V. M., Rabah, R. R., Preston, L., Tekle, W. G., & Qureshi, A. I. (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.
 21. Martinez-Gutierrez, J. C., Kim, Y., Salazar-Marioni, S., Tariq, M. B., Abdelkhaleq, R., Niktabe, A., ... & Sheth, S. A. (2023). Automated large vessel occlusion detection software and thrombectomy treatment times: a cluster randomized clinical trial. *JAMA neurology*, 80(11), 1182-1190.
 22. Delora, A., Hadjialiakbari, C., Percenti, E., Torres, J., Alderazi, Y. J., Ezzeldin, R., ... & Ezzeldin, M. (2024). Viz LVO versus Rapid LVO in detection of large vessel occlusion on CT angiography for acute stroke. *Journal of NeuroInterventional Surgery*, 16(6), 599-602.
 23. Figurelle, M. E., Meyer, D. M., Perrinez, E. S., Paulson, D., Pannell, J. S., Santiago-Dieppa, D. R., ... & Meyer, B. C. (2023). Viz. ai implementation of stroke augmented intelligence and communications platform to improve indicators and outcomes for a comprehensive stroke center and network. *American Journal of Neuroradiology*, 44(1), 47-53.
 24. Mallon, D. H., Taylor, E. J., Vittay, O. I., Sheeka, A., Doig, D., & Lobotesis, K. (2022). Comparison of automated ASPECTS, large vessel occlusion detection and CTP analysis provided by Brainomix and RapidAI in patients with suspected ischaemic stroke. *Journal of Stroke and Cerebrovascular Diseases*, 31(10), 106702.
 25. Alwood, B. T., Meyer, D. M., Ionita, C., Snyder, K. V., Santos, R., Perrotta, L., ... & Bolar, D. S. (2024). Multicenter comparison using two AI stroke CT perfusion software packages for determining thrombectomy eligibility. *Journal of Stroke and Cerebrovascular Diseases*, 33(7), 107750.
 26. Scavasine, V. C., Stoliar, G. A., de Almeida Teixeira, B. C., Zétola, V. D. H. F., & Lange, M. C. (2024). Automated evaluation of collateral circulation for outcome prediction in acute ischemic stroke. *Journal of Stroke and Cerebrovascular Diseases*, 33(4), 107584.
 27. Mallon, D., Fallon, M., Blana, E., McNamara, C., Menon, A., Ip, C. L., ... & Doig, D. (2024). Real-world evaluation of Brainomix e-Stroke software. *Stroke and Vascular Neurology*, 9(5).
 28. Vacek, A., Mair, G., White, P., Bath, P. M., Muir, K. W., Salman, R. A. S., ... & Wardlaw, J. M. (2024). Evaluating artificial intelligence software for delineating hemorrhage extent on CT brain imaging in stroke: AI delineation of ICH on CT. *Journal of Stroke and Cerebrovascular Diseases*, 33(1), 107512.
 29. Chan, N., Sibtain, N., Booth, T., de Souza, P., Bibby, S., Mah, Y. H., ... & U-King-Im, J. M. (2023). Machine-learning algorithm in acute stroke: real-world experience. *Clinical Radiology*, 78(2), e45-e51.
 30. Mohapatra, S., Lee, T. H., Sahoo, P. K., & Wu, C. Y. (2023). Localization of early infarction on non-contrast CT images in acute ischemic stroke with deep learning approach. *Scientific reports*, 13(1), 19442.
 31. Soun, J. E., Zolyan, A., McLouth, J., Elstrott, S., Nagamine, M., Liang, C., ... & Chow, D. S. (2023). Impact of an automated large vessel occlusion detection tool on clinical workflow and patient outcomes. *Frontiers in neurology*, 14, 1179250.
 32. Lee, S. J., Park, G., Kim, D., Jung, S., Song, S., Hong, J. M., ... & Lee, J. S. (2024). Clinical evaluation of a deep-learning model for automatic scoring of the Alberta stroke program early CT score on non-contrast CT. *Journal of NeuroInterventional Surgery*, 16(1), 61-66.
 33. Chen, W., Wu, J., Wei, R., Wu, S., Xia, C., Wang, D., ... & Zhang, X. (2022). Improving the diagnosis of acute ischemic stroke on non-contrast CT using deep learning: a multicenter

-
- study. *Insights into Imaging*, 13(1), 184.
34. Reidler, P., Stueckelschweiger, L., Pühr-Westerheide, D., Feil, K., Kellert, L., Dimitriadis, K., ... & Kunz, W. G. (2021). Performance of automated attenuation measurements at identifying large vessel occlusion stroke on CT angiography. *Clinical neuroradiology*, 31(3), 763-772.
 35. Yahav-Dovrat, A., Saban, M., Merhav, G., Lankri, I., Abergel, E., Eran, A., ... & Sivan-Hoffmann, R. (2021). Evaluation of artificial intelligence-powered identification of large-vessel occlusions in a comprehensive stroke center. *American Journal of Neuroradiology*, 42(2), 247-254.
 36. Schlossman, J., Ro, D., Salehi, S., Chow, D., Yu, W., Chang, P. D., & Soun, J. E. (2022). Head-to-head comparison of commercial artificial intelligence solutions for detection of large vessel occlusion at a comprehensive stroke center. *Frontiers in Neurology*, 13, 1026609.
 37. Hoelter, P., Muehlen, I., Goelitz, P., Beuscher, V., Schwab, S., & Doerfler, A. (2020). Automated ASPECT scoring in acute ischemic stroke: comparison of three software tools. *Neuroradiology*, 62(10), 1231-1238.
 38. Sawicki, M., Safranow, K., Wiska, L., Pasek, I., Gajdział, A., Gruszeński, M., & Poncyłjusz, W. (2021). Diagnostic value of artificial intelligence-based software in detection of large vessel occlusion in acute ischemic stroke. *Applied Sciences*, 11(21), 10017.
 39. Slater, L. A., Ravintharan, N., Goergen, S., Chandra, R., Asadi, H., Maingard, J., ... & Phan, T. G. (2024). RapidAI compared with human readers of acute stroke imaging for detection of intracranial vessel occlusion. *Stroke: Vascular and Interventional Neurology*, 4(2), e001145.
 40. Bushnaq, S., Hassan, A. E., Delora, A., Kerro, A., Datta, A., Ezzeldin, R., ... & Ezzeldin, M. (2024). A comparison of CT perfusion output of RapidAI and Viz. ai software in the evaluation of acute ischemic stroke. *American Journal of Neuroradiology*, 45(7), 863-870.
 41. Pisani, L., Haussen, D. C., Mohammaden, M., Perry da Camara, C., Jillella, D. V., Rodrigues, G. M., ... & Nogueira, R. G. (2023). Comparison of CT perfusion software packages for thrombectomy eligibility. *Annals of neurology*, 94(5), 848-855.
 42. Shahrouki, P., Kihira, S., Tavakkol, E., Qiao, J. X., Vagal, A., Khatri, P., ... & Nael, K. (2024). Automated assessment of ischemic core on non-contrast computed tomography: a multicenter comparative analysis with CT perfusion. *Journal of NeuroInterventional Surgery*, 16(12), 1288-1293.
 43. Royal College of Radiologists. (2021). **Integrating artificial intelligence with the radiology reporting workflows (RIS and PACS)**. RCR Guidance.

Copyright: ©2026 Netanel Stern. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.