Artificial Intelligence in Ultrasound Imaging

A WHITE PAPER

FUNDAMENTALS
APPLICATIONS
OPPORTUNITIES
AND CONSIDERATIONS
IN CLINICAL PRACTICE

Artificial Intelligence in Ultrasound Imaging

Fundamentals, Applications, Clinical Integration

Jim Baun, MEd, RDMS, RVT, FSDMS

What is human intelligence?

The flexible capacity to perceive, learn, remember, reason, and regulate behavior to achieve goals in changing, uncertain environments.

Human intelligence is a flexible set of mental abilities that allows people to perceive the world, focus attention, learn from experience, store and retrieve knowledge, use language, reason through problems, and regulate behavior toward a goal. These processes operate together, often under uncertainty and with incomplete information, so that humans can adapt to new situations rather than simply repeat fixed responses. This psychological framework provides a useful baseline for introducing artificial intelligence, which attempts to reproduce or approximate selected components of these abilities in computational form.

Basic aspects of human intelligence in general psychology that are replicated in AI, especially those pertinent to medical ultrasound practice include:

- **Perception:**
  - *Human process*: taking in information through the senses and organizing it into meaningful patterns
  - *AI replication*: computer vision, speech recognition, signal pattern recognition
- **Attention:**
  - *Human process*: selecting what to focus on and sustaining or shifting focus as needed.
  - *AI replication*: "attention" mechanisms help prioritize information
- **Learning:**
  - *Human process*: changing knowledge or behavior based on experience (*including conditioning and skill learning*).
  - *AI replication*: systems learn statistical patterns from training data
- **Memory:**
  - *Human process*: encoding, storing, and retrieving information (*working, short-term, long-term*).
  - *AI replication*: models store information in parameters and can use short-term context
- **Language:**
  - *Human process*: understanding and producing symbols and speech for communication and thought
  - *AI replication*: text understanding and generation (*LLMs – large language models*), translation, summarization

- **Reasoning:**
  - *Human process*: drawing conclusions, solving problems, and applying logic (*deductive and inductive*)
  - *AI replication*: can perform logical and multi-step problem solving
- **Processing speed:**
  - *Human process*: how efficiently information is handled, especially under time pressure
  - *AI replication*: very fast computation and retrieval once deployed
- **Creativity:**
  - *Human process*: generating novel, useful ideas by combining concepts in new ways
  - *AI replication*: can produce novel combinations (*text/images/music*), though it is pattern-based rather than intention-driven

What is Artificial Intelligence?

The simulation of human intelligence in machines.

Artificial intelligence (*AI*) is the field of computing science that designs systems that can perform tasks often associated with human intelligence - such as recognizing patterns, learning from data, understanding language, and supporting decisions. AI uses algorithms that take an input (*e.g., an ultrasound image or a short text description*) and produce an output (*such as a label, measurement, or suggested next step*). With training and feedback, these systems can improve their performance over time.

What is the Goal of AI?

To enable machines to perform tasks that typically require human intelligence.

The goal of AI is to build computer systems that can take in information, learn patterns from examples, and use that learning to make useful predictions or decisions. In healthcare, the purpose is to support people - not replace them - by helping with tasks that are complex, repetitive, or involve large amounts of data. In practical terms, AI can improve speed and consistency and help handle high volumes of work, while the sonographer and interpreting clinician remain responsible for higher-level judgment, clinical context, and final decisions.

AI Subfields – General

Artificial intelligence is an umbrella term that includes several common subtypes. Each attempts to replicate some aspect of human intellectual processes, not all of which are useful in ultrasound imaging or diagnosis.

- **Machine learning** (*predictive AI*): uses examples (*data*) to learn patterns and then classify something or predict an outcome (*spam filters, risk scores).*
- **Deep learning**: a type of machine learning that uses **neural networks** and is especially good with images, sound, and language (*e.g., image recognition, speech-to-text*).

- **Generative AI**: creates **new content** - text, images, audio, or code - based on a prompt and learned patterns (*e.g., chatbots, summarizers, image generators*).
- **Rule-based AI** (*expert systems*): uses hand-written **if–then** rules to make predictable decisions (*e.g., simple protocol or triage logic*).
- **Recommendation systems**: a type of machine learning that **suggests or ranks** items based on user behavior and similarity (*e.g., streaming or shopping suggestions*).
- **Reinforcement learning**: learns which actions work best by using **rewards and penalties** (e.g., game-playing agents, some robotics and control systems).

AI Subfields – Use in Ultrasound

In ultrasound systems and ultrasound-based diagnosis, the AI tools used most often are **machine learning (*ML*)** and **deep learning (*DL*)**. In **machine learning** the computer learns from many examples to help make a decision or give a number - such as labeling a finding as normal vs abnormal or estimating a measurement or risk. Many ML methods work best when people first tell the computer *what to look at* by providing specific measurements or descriptors (*e.g., a lesion size, a fetal biometric measurements, texture patterns, or Doppler values*). **Deep learning** is a type of machine learning that can learn these "what to look for" patterns directly from the ultrasound data itself, such as still images or short video clips, which can help when images vary from patient to patient or when scanning technique differs. **Generative AI** is usually used in a supporting role - helping with words and workflow tasks like summarizing findings, assisting with report wording, or explaining decision support - rather than creating the ultrasound image or making the main diagnosis by itself.

Machine Learning (ML)

Often used alongside DL for workflow logic, quality scoring, and traditional predictive tasks. Some specific uses in real-time ultrasound imaging include:

Image quality assessment ("is this clip usable?")

ML models score whether an image or cine loop is *good enough* for measurement, interpretation, or saving to the record. Typical checks include:

- **Focus and sharpness**: is the region of interest crisp, or is it blurred by motion/out-of-plane scanning?
- **Depth and zoom appropriateness**: is the target anatomy too small, too deep, or cropped?
- **Gain / dynamic range**: is the image washed out (*too bright*) or too dark, masking anatomy?
- **Shadowing / dropout**: are key structures hidden by rib shadow, bowel gas, or poor contact?
- **View completeness**: is the required plane captured (*e.g., full organ in view, correct landmarks present*)?

Concrete examples:

- **Obstetric**: “This is not an acceptable fetal head plane for BPD/HC - midline/falx not clearly seen.”
- **Cardiac/POCUS**: “This parasternal long-axis clip is low quality due to motion and incomplete LV visualization.”
- **Abdomen**: “Gallbladder image flagged because posterior wall is not fully visualized (*dropout*).”

Triage / prioritization and workflow productivity (varies by implementation)

Triage is the process of quickly sorting patients or cases by urgency, so the most critical problems are handled first. In imaging, triage can also mean prioritizing studies or findings (*for example, flagging an exam that may show an acute bleed or a critical vascular problem*) so it is reviewed sooner. ML can help route exams, reduce repetitive steps, and keep the exam moving by detecting patterns that suggest urgency or by organizing study content. Typical functions include:

- **Auto-labeling and sorting:** recognize common views (*e.g., kidney long vs transverse*) and label/sort clips automatically.
- **Key-image suggestion:** identify the “best” frames within a cine loop to save.
- **Protocol guidance:** prompt next steps (*“acquire additional view”*) when required images are missing.
- **Abnormality flags for attention:** highlight clips that may contain concerning findings, so they are reviewed first.
- **Worklist prioritization (*department level*):** in some workflows, studies with potential critical findings may be pushed higher in the read queue (*implementation-dependent and policy-controlled*).

Concrete examples:

- **Obstetrics:** automatically identify and label standard anatomy views (*e.g., head, abdomen, femur*) and alert if a required view is missing.
- **Vascular:** detect that a Doppler sample volume is off-center, or angle correction is incorrect and prompt a correction before saving.
- **ED/POCUS**: flag a FAST exam clip as “possible free fluid” so it is reviewed immediately, while still requiring clinician confirmation.

Deep Learning (DL)

Deep learning (*DL*) is a subset of machine learning that uses neural networks and is especially strong at recognizing patterns in images and video. This makes it a good fit for ultrasound, because most exams are created from real-time images and cine loops. In many ultrasound applications, DL is paired with smaller “classic” machine learning methods that handle supporting automation and workflow logic. Common DL-enabled features include:

A **neural network** is a computer model made of many small, connected processing units (“*neurons”*) that learns by adjusting connection strengths so it can turn inputs (*like pixels or numbers*) into useful outputs (*like a label, a measurement, or a prediction*).

Scan guidance / view recognition (identify the correct plane, coach the operator)

DL can recognize whether the probe is showing the *right view* and provide real-time feedback to help the operator adjust. Guidance may include prompts for **probe angle, tilt, rotation, or position** and may signal when a plane is acceptable. Specific tasks include:

- **Detect the anatomy and exam type**
  - Recognizes the organ/region (heart, liver, kidney, thyroid, fetus, vessel).
  - Confirms the intended exam preset/workflow is being followed.
- **View (*plane*) recognition**
  - Identifies standard views, such as:
  - **Echo:** apical 4-chamber, apical 2-chamber, parasternal long-axis, parasternal short-axis.
  - **OB:** correct biometry planes (*head/HC plane, abdomen/AC plane, femur/FL plane*).
  - **Vascular:** longitudinal vs transverse vessel views.
  - Flags “non-standard” or off-axis planes when a measurement view is required.
- **Probe-motion coaching (*real-time guidance*)**
  - Provides prompts such as **slide, rock/heel-toe, tilt, rotate, or fan**, guiding the probe to the correct angle and position.
  - Helps center the target structure in the field of view.
- **Quality scoring / accept–reject gating**
  - Rates whether the clip/frame is “good enough” for automation (measurement or reporting).
  - Rejects frames with motion blur, poor contact, shadowing over key anatomy, or incomplete visualization.
- **Landmark detection**
  - Locates key points (e.g., valve annulus, chamber apex, fetal skull edges, vessel walls).
  - Uses these landmarks to stabilize view recognition and prepare for measurement.

Concrete examples:

- **Obstetrics:** “Fetal head plane not correct—rotate slightly; midline and cavum septi pellucidi not visible.”
- **Echo:** Identify apical 4-chamber vs apical 2-chamber and prompt small adjustments to reduce foreshortening.
- **POCUS:** Confirm IVC long-axis view and coach to center the vessel and optimize depth.

Segmentation

Segmentation is an AI method that automatically "traces" a structure on an ultrasound image - for example, the border of the left ventricle, a fetus's head, or a mass. The software separates the structure from the background, so the system can measure it quickly and consistently without the operator drawing the outline by hand. Common segmentation tasks include:

- **Boundary (*contour*) tracing:** draws the outline of the target (*automatic "auto-trace"*).
- **Multi-structure separation:** segments more than one structure at once (*e.g., LV cavity + myocardium; fetal skull + brain region*).
- **Frame-to-frame tracking**: keeps the outline stable across a cine loop (important for beating heart or fetal motion).
- **Area/volume calculation inputs:** produces masks/contours used to compute area (*cm²*), volume (*mL*), and derived values (*e.g., EF*).
- **Measurement readiness / confidence:** outputs a quality or confidence score so unreliable frames can be rejected.
- **Region-of-interest (*ROI*) placement:** automatically places an ROI for follow-on tools (*e.g., strain analysis, lesion sizing, Doppler sampling location in some workflows*).
- **Change tracking (follow-up):** supports consistent re-measurement by outlining the same type of structure in repeat exams.

Concrete examples:

- **Heart (*echocardiography*):** borders of the left ventricle and right ventricle, and sometimes the atria - used to calculate chamber size, volumes, and ejection fraction (*EF*).
- **Obstetrics:** fetal head, abdomen, and long bones (*such as the femur*) used for biometry (*dating and growth assessment*).
- **Vascular:** the inside of a vessel (*lumen*) and sometimes plaque—used to estimate diameter, area, or narrowing (*stenosis*).
- **Abdomen and small parts:** outlines of organs or focal findings (*e.g., thyroid nodules, liver/kidney lesions*) - used for size and volume measurements and to improve follow-up consistency.

Generative AI (GenAI)

Generative AI (*GenAI*) refers to models that can **create new content** or **rewrite existing content** - most commonly text, images, video, or audio. In diagnostic ultrasound, GenAI is usually used as a **support tool** to assist with communication and workflow (*for example, drafting report text or summarizing findings*) rather than independently "reading" and interpreting the ultrasound image. Common uses include:

Report drafting / wording support (turn findings into clear language)

GenAI can create a first-pass report from structured inputs (*measurements, selected findings, normal templates*) and then the clinician edits/approves.

Concrete examples:

- **Convert measurements** into a clean sentence: "Gallbladder without stones; wall thickness normal; no pericholecystic fluid."
- **Obstetrics**: Draft a standard "normal anatomy survey" paragraph once required items are checked off.
- **Echo**: Generate a coherent impression section using ejection fraction (*EF*), chamber sizes, valve summaries.

Summarization of the exam (short, readable takeaways)

GenAI can help produce brief, plain-language summaries for the chart, patient handoff, or the ordering provider - especially when an exam includes many images, measurements, and technical details.

Concrete examples:

- **Protocol summary:** condenses what was performed (*exam type, required views obtained, key measurements completed*) into a short checklist-style recap.
- **Findings summary:** compresses the main observations into a few high-level categories (e*.g., normal vs abnormal patterns, presence/absence flags*) without detailed narrative.
- **Impression draft:** generates a brief closing statement that organizes the overall takeaway and can be expanded by the clinician into a final report.

Decision-support explanations (plain-language rationale, not a diagnosis)

When another algorithm flags a concern, GenAI can explain what that flag means in plain language and outline common next steps, while keeping the clinician in control of interpretation and decisions.

Concrete examples:

- **Vascular:** An AI tool flags a spectral Doppler waveform as potentially abnormal. GenAI could explain that the waveform shows reduced phasicity or an unexpected velocity pattern and suggest that the sonographer review probe position, Doppler angle, and the possibility of upstream or downstream disease before drawing a clinical conclusion.
- **Echocardiography:** An algorithm identifies a possible reduction in left ventricular systolic function. GenAI could explain that the concern was triggered by chamber motion or an estimated ejection fraction below the expected range and recommend review of image quality, endocardial border tracing, and additional views.
- **Obstetrics:** An automated system flags fetal growth measurements as potentially discordant with gestational age. GenAI could summarize which measurements contributed to the flag and suggest confirming measurement technique, reviewing dating information, and assessing other biometric parameters.

- **POCUS / Emergency ultrasound:** An AI system flags a FAST examination for possible free fluid. GenAI could explain where the suspected fluid was detected, identify the relevant view, and remind the operator that the finding should be confirmed with additional views and clinical correlation.

Protocol and scanning guidance "chat" help (training and just-in-time coaching)

GenAI can answer questions during learning or scanning - like a built-in tutor - using the department's protocol language and providing just-in-time reminders about steps, views, and documentation.

Concrete examples:

- "What views are required for a RUQ ultrasound?"
- "How should the Doppler angle be aligned for accurate velocity measurement?"
- "What should be documented for a DVT study?"

Workflow automation with text and labels (organize and clean up the study)

GenAI can help apply consistent labels, suggest standard phrases, and promote uniform terminology across reports, which improves organization and reduces rework.

Concrete examples:

- Suggest standardized labels for saved images/clips (*e.g., "Right kidney longitudinal"*).
- Convert free-text notes into structured fields (*indication, technique, limitations*).
- Translate or simplify report language for patient-facing summaries (*when enabled*).

Patient communication aids (education and after-visit summaries)

In some settings, GenAI can create a patient-friendly explanation of what the exam evaluated and what the results mean, using clear language and minimal medical jargon.

Concrete examples:

- "This ultrasound checked the gallbladder for stones and signs of inflammation…"
- "The ultrasound of the pregnancy measured growth and checked major structures…"

General Applications in Medical Imaging

AI in medical imaging is best understood as a set of tools that either (1) improve how images are acquired/formed, (2) extract quantitative information from images, or (3) support clinical and operational decisions around imaging. In practice, most regulated/marketed AI systems today are still "narrow" (*task-specific*) rather than general intelligence, and many are deployed as software functions inside imaging workflows. The core categories of AI use in medical imaging include:

- **Image formation and enhancement (*improving the image*)**
  - *Examples*: reconstruction/denoising, artifact reduction, faster acquisition, resolution enhancement.
  - *Typical impact*: improved image quality or reduced dose/time without changing the scanner hardware.

- **Detection, classification, and triage (*finding abnormalities*)**
  - *Examples*: flagging suspected findings for prioritization, abnormality detection, incidental finding detection support.
  - *Typical impact*: faster routing of urgent cases and decision support for readers (*not "autonomous diagnosis" in most real-world deployments*).

- **Segmentation and quantitative imaging (*measuring anatomy & physiology*)**
  - *Examples*: automated contouring, measurements, volumes, ejection fraction, strain, biometry, and Doppler indices.
  - *Typical impact*: improved measurement consistency, reduced repetitive manual tasks.
- **Workflow automation and reporting (*efficiency & consistency*)**
  - *Examples*: auto-labeling, structured report support, text summarization, consistency checks, QA checks (*wrong patient/side, missing views*).
  - *Typical impact*: fewer clicks, fewer omissions, and faster turnaround times.
- **Prognosis, risk prediction, and personalization (*downstream analytics*)**
  - *Examples*: combining imaging features with clinical data to estimate risk, recommend follow-up intervals, or predict treatment response.
  - *Typical impact*: shifts imaging from "what is present" toward "what is likely next," which can improve planning but requires careful validation and ongoing monitoring.
- **Foundation models and generative AI (*cross-cutting, emerging*)**
  - *Examples*: multimodal assistants, report drafting, image/text retrieval, generalized segmentation or representation learning.
  - *Typical impact*: broader workflow support and consistency, but with higher governance needs (*hallucination risk, privacy, traceability*).

Ultrasound Applications - Summary

- **Real-time scan guidance (*probe coaching*)**
  - AI provides on-screen prompts (*tilt/rotate/slide*) to help obtain a diagnostic view, especially for POCUS and for less-experienced operators.
- **View recognition and exam standardization**
  - AI identifies standard planes/views (*e.g., common echo views*) and may assist with consistent placement of Doppler/sample regions of interest to reduce operator variability.
- **Automatic image quality assessment**
  - AI scores whether an acquired view is “good enough” (*focus, centering, anatomy inclusion*) and prompts re-acquisition when needed.
- **Automated segmentation of anatomy**
  - AI outlines structures (*LV endocardium, fetal head/abdomen, thyroid nodules, etc*.) to support consistent downstream measurements and quantification.
- **Automated measurements and quantification**
  - AI performs repetitive measurements (*e.g., echo chamber quantification, fetal biometry*) and exports values directly into the report.
- **Doppler automation**
  - AI helps identify the correct Doppler site/view, traces spectral envelopes, and calculates indices (*e.g., fetal-placental Doppler measurements*).
- **Detection/classification of pathology patterns**
  - AI flags suspicious findings or classifies disease patterns (*e.g., thyroid nodule risk features, echo pattern-recognition tasks*), depending on the specific model and clinical use case.
- **Task-specific screening/triage tools**
  - AI supports rapid identification of studies that may need priority review (*more common in echo and OB workflows*).
- **Automated reporting and structured documentation**
  - AI can pre-populate report sections with recognized views, measurements, and standardized labels, reducing manual clicks and transcription.
- **Blind-sweep and low-resource workflows**
  - AI can extract clinically useful outputs from simplified acquisition approaches (*e.g., blind sweeps*) expanding access where expert scanning is limited.
- **Image enhancement (*signal and image-domain*)**
  - AI-based denoising/speckle suppression, contrast improvement, and artifact-aware enhancement can improve readability *(implementation varies by vendor and device*).