RB Weekly AI Brief - Issue 18 - 05.08.2026

Covering the week of 05.08.2026 · Issue 18 of the RB Weekly AI Brief

Recurring themes: Regulatory & HTA Signals (3 of last 4 issues) · Regulation & Policy (3 of last 4 issues) · Healthcare & Life Sciences (3 of last 4 issues) · Models & Research (3 of last 4 issues)

AI News Roundup

Regulatory & HTA Signals

FDA Draft AI Guidance for Drug Development Remains Pending Finalisation at Mid-2026

[ONGOING] — A mid-2026 analysis by Drug Discovery News confirms that the FDA's January 2025 draft guidance on AI for regulatory decision-making — covering nonclinical, clinical, post-marketing and manufacturing phases — remains in draft form, with finalisation still pending. The joint FDA-EMA ten-point guiding principles published in January 2026 provide a non-binding framework but stop short of binding requirements, and no dedicated, enforceable global standard for AI in drug development yet exists. The EMA's own reflection paper is similarly non-binding, leaving sponsors to interpret overlapping, non-converged frameworks across jurisdictions.

So what? *Until the FDA finalises its guidance and the FDA-EMA principles harden into binding requirements, market access and HEOR teams must build AI validation packages to the most stringent plausible standard — the draft seven-step credibility framework — to ensure regulatory submissions are not vulnerable to challenge.*

Drug Discovery News

Regulation & Policy

MHRA National Commission AI Recommendations Imminent This Summer

[ONGOING] — The MHRA's National Commission into the Regulation of AI in Healthcare — an independent advisory body bringing together global AI leaders, clinicians, and regulators — is expected to publish its regulatory recommendations in summer 2026, following publication in June 2026 of two evidence reports drawing on 760 responses from patients, clinicians, industry, and system partners. The Commission's remit is to advise the MHRA on a new regulatory framework for AI as a medical device, with the MHRA committed to acting on its recommendations. The MHRA's draft Medical Devices (Amendment) Regulations 2026, published in May, already introduce Predetermined Change Control Plans (PCCPs) for AI software devices but have left AI-specific provisions pending the Commission's findings.

So what? *The Commission's imminent recommendations will directly shape MHRA's AI-as-a-medical-device framework — including post-market surveillance, clinical validation standards, and market access pathways — making it critical for pharma and healthtech companies to monitor the output and engage now with the MHRA's parallel draft regulation consultation to influence the final rules.*

GOV.UK (MHRA)

Healthcare & Life Sciences

FDA Accepts First AI-Driven Digital Liver Model Under ISTAND Programme

[IN CASE YOU MISSED IT] — The FDA's Centre for Drug Evaluation and Research accepted the first Letter of Intent for a purely in silico, AI-driven Drug Development Tool under its ISTAND qualification programme — Absentia Labs' Digital Liver Model for predicting drug-induced liver injury (DILI) in small molecule candidates. The tool uses AI to compare new chemical structures against reference drugs with known DILI risk, aiming to improve predictive toxicology and reduce reliance on animal testing. DILI is one of the leading causes of clinical trial termination and IND-stage drug attrition, and current preclinical models fail to reliably capture human liver response variability.

So what? *For HEOR and pharma regulatory teams, FDA's formal acceptance of an AI-based DILI prediction tool into a qualification pathway signals that AI-generated safety evidence may eventually be submitted as part of a weight-of-evidence nonclinical package, creating both opportunity and obligation to understand how regulators and HTA bodies will weigh such model-derived data in benefit–risk assessments.*

FDA

Models & Research

Anthropic Discovers Hidden 'Global Workspace' Inside Claude's Neural Network

[IN CASE YOU MISSED IT] — Anthropic published research revealing that Claude models have spontaneously developed an internal structure called 'J-space' — a privileged mental workspace that holds concepts the model is reasoning about but has not yet verbalised, identified using a new technique called the Jacobian lens. The findings, inspired by neuroscience's Global Workspace Theory, show J-space supports deliberate multi-step reasoning and flexible reuse across tasks, while most of the model's processing bypasses it entirely. Anthropic released open-source code and partnered with Neuronpedia for an interactive demo, and invited independent commentary from leading neuroscientists and AI interpretability researchers.

So what? *For pharma and HEOR professionals deploying AI for clinical evidence synthesis or regulatory writing, this interpretability breakthrough means it may soon be possible to audit what a model was 'thinking' before it generated an output — a capability that could become a regulatory expectation for AI tools used in high-stakes decisions such as HTA dossier preparation or pharmacovigilance.*

Anthropic

DeepSeek Releases Official V4-Flash-0731 with Major Agentic Gains

DeepSeek released the official production version of DeepSeek-V4-Flash (checkpoint 0731) on July 31, 2026, pushing it into public beta via its API changelog with no formal launch event. The 284-billion-parameter Mixture-of-Experts model — activating only 13 billion parameters per token — was retrained entirely via an improved post-training pipeline, with no architectural changes, yet outperformed the V4-Pro preview across nine agent benchmarks and added native Responses API and Codex support. The MIT-licensed open weights were published on Hugging Face, and the preview had already ranked as the most-used model on OpenRouter for seven consecutive weeks.

So what? *For pharma and market access teams evaluating AI tools for evidence generation or literature review, DeepSeek V4-Flash's combination of frontier agentic performance, open weights under an MIT licence, and aggressive pricing lowers the barrier to deploying capable in-house AI pipelines — but also intensifies the need for governance frameworks that address model provenance, safety testing, and regulatory acceptability of outputs.*

DeepSeek

Academic Paper Summaries

Selected from PubMed · Published within the last 12 months · New selections each week

Domain Paper — HEOR / Health Economics / Market Access

Artificial Intelligence in Health Technology Assessment Submissions: A Targeted Review of Global Policy and Practice.

Gaind N, Badin M, Jacob J, et al. · Value in health regional issues · 2026

#HTA · #MarketAccess · #Regulation

A targeted review screened 1,309 abstracts and identified guidance documents from nine HTA agencies worldwide, finding that NICE (England) and CDA-AMC (Canada) are the only bodies with a clear position statement and implementation strategy for AI use in HTA submissions — the review's own stated conclusion, not an oversimplification. AI is primarily applied to systematic literature reviews, study screening, data extraction, evidence synthesis, and economic modelling across these agencies. For pharma and HEOR teams submitting dossiers globally, this means AI-assisted evidence cannot be assumed acceptable beyond these two agencies without explicit methodological documentation and human oversight.

PMID: 42334393	PubMed →	DOI →

AI Research Paper 1

Explainable AI for Healthcare.

Li Y, Sun Q, Akman A, et al. · Studies in health technology and informatics · 2025

#ClinicalAI · #Regulation · #Diagnostics

This paper reviews the field of explainable AI in healthcare, examining how different techniques can make complex AI models more transparent and understandable to clinicians and patients across a range of data types including medical images, clinical notes, and patient monitoring data. The core argument is that without explainability, AI tools risk being rejected or misused in high-stakes clinical settings. For executives, this underscores that regulatory acceptance and clinical adoption of AI tools will increasingly depend on built-in interpretability, not just predictive performance.

PMID: 41728714	PubMed →	DOI →

AI Research Paper 2

Multimodal Large Language Models in Medical Imaging: Current State and Future Directions.

Nam Y, Kim DY, Kyung S, et al. · Korean journal of radiology · 2025

#ClinicalAI · #Diagnostics · #Regulation

This review examines the current state of multimodal large language models in radiology, where AI systems can simultaneously process medical images and clinical text to generate report drafts, answer diagnostic questions, and support clinical decision-making. While performance is promising, significant barriers remain including risk of AI-generated errors, limited access to large curated medical datasets, and lack of transparency. For healthcare executives, this highlights both the near-term commercial opportunity in AI-assisted radiology and the regulatory and safety hurdles that must be addressed before broad deployment.

PMID: 41015856	PubMed →	DOI →

This brief is produced using an AI-assisted workflow with expert human editorial review. Stories are curated and sources verified before publication. Readers are encouraged to consult source links directly.