Information Theory Maps Hospital Workflow Shifts

Hospitals generate continuous operational data that reveal how care routines form, shift and recover. Information-theory measures provide a practical lens on these patterns without relying on averages alone. Entropy summarises how predictable laboratory ordering is across a day, while surprisal captures how unusual a specific order time is compared with usual practice. When these measures are tracked at the level of individual wards or intensive care units, they show when local routines are stable, when admission activity disperses orders across the clock and when external pressures reshape schedules. They also indicate when clinician-driven actions, such as off-schedule tests, carry hidden signals that can influence predictive models. Together these metrics help teams detect operational change and protect model integrity during deployment.

Consistent Patterns at Unit Level

Ergodicity refers to consistency within a unit over time, where the pattern seen across beds at one moment mirrors the pattern observed across many days. Hourly laboratory order counts were converted to probabilities and then to surprisal and entropy, with short epochs compared against longer baselines using similarity measures. Across medical and surgical wards and intensive care units, short-term distributions closely matched multi-year patterns, indicating approximately ergodic behaviour at unit level. Aggregating across an entire hospital weakened this effect because mixed services and case profiles blur local routines. Unit-level consistency provides a practical baseline for operational monitoring and model development, and divergence between short-term and long-term patterns signals a change in workflow, case mix or documentation.

Entropy Tracks Admission Patterns and Pandemic Impact

Surprisal, defined as the negative logarithm of an order’s probability, is low when ordering follows routine schedules and high when it occurs at unusual hours. Entropy, the weighted average of surprisal, declines after admission as patients move from unscheduled arrival to scheduled care. Heatmaps of lab ordering showed relatively even, higher-surprisal activity within hours of admission, consistent with emergency arrivals and transfers. Over subsequent days, low-surprisal bands emerged at institution-specific release times for routine tests, and overall entropy fell as orders clustered around one or two daily draw windows. This decline occurred across a broad range of order volumes, indicating that entropy contains information about workflow independent of the sheer number of tests.

Institutional electronic health record release rules shaped these patterns. At one site, repeat orders were released 36 hours before the scheduled draw, producing diurnal dispersal across units with unit-specific nuances. At another, routine orders were released four hours before collection, creating visible low-surprisal bands at predictable intervals, while a third released routine day-of tests immediately if due before midnight or at midnight otherwise. These constraints help explain why Emergency Departments, which operate with continuous arrivals, retained high entropy, whereas inpatient wards exhibited lower entropy driven by rounding schedules and routine morning labs.

Must Read: [Guardrailed GenAI Automates Data Extraction from RHC Notes]

The COVID-19 pandemic altered these distributions unevenly. Time-series entropy, calculated on rolling 28-day windows and compared with hospital COVID-19 admissions, changed little at hospitals with low COVID-19 burden. At a site with higher burden, entropy dropped across most units following COVID-19 surges, with only certain medical intensive care units returning to baseline. The Emergency Department remained near maximal entropy throughout. These shifts visualise how surges compressed or reorganised ordering routines, likely reflecting altered workflows, resource allocation or patient profiles.

Surprisal Signals Clinician Behaviour and Model Integrity

Because surprisal measures how atypical an order's timing is, it can proxy clinician judgement embedded in workflows. In retrospective analysis of a medical intensive care cohort used for haemorrhage prediction, a range of cardiorespiratory features and dynamical measures were evaluated alongside the surprisal of haemoglobin reporting time. The haemoglobin surprisal was a statistically significant predictor even when other continuous monitoring parameters were considered, with impact comparable to blood pressure. If a non-routine haemoglobin is ordered based on clinical suspicion of bleeding, its atypical timing contributes information that a model may learn. This illustrates a practical risk to model integrity: label leakage, where clinician actions inadvertently encode outcome signal into inputs. Incorporating surprisal-based features can therefore improve vigilance, helping teams identify when models rely on workflow artefacts rather than physiological change and prompting mitigation before performance degrades during deployment.

Information-theoretic metrics provide interpretable, unit-level signals of stability and change in hospital operations. Ergodicity demonstrates local stationarity suitable for monitoring and modelling, entropy captures predictable transitions from admission to routine care and reveals external shocks, and surprisal bridges clinician behaviour with data-driven prediction. By tracking these measures across units and over time, healthcare teams can detect data drift, understand the operational roots of variation and guard against label leakage in predictive systems, supporting safer deployment and more resilient care pathways.

Source: npj digital medicine

Image Credit: iStock

Published on : Tue, 18 Nov 2025