Toward Operationalisable Clinical Risk Prediction Models

Sim Mei Choo[1], Sohan Seth[1], Bruce Guthrie[2], Ewen Harrison[2]

[1] School of Informatics, University of Edinburgh; [2] Usher Institute, University of Edinburgh

AI4Bi UKRI AI Centre for Doctoral Training in Biomedical Innovation

Data Science Unit, School of Informatics

The University of Edinburgh

Dynamic **Uncertainty-aware** **Knowledge-informed** **Transparent** **Traceable**

1 WHY THIS MATTERS

- Longitudinal electronic health record (EHR) data capture real-world practice but are sparse, irregular, and shared by care processes.
- Clinical prediction model (CPM) tasks are often simplified into static, binary predictions that ignore competing risks, misfit clinical workflow or fail to support lifecycle use.
- Many CPMs are soley EHR data-driven, while clinical decisions are knowledge-informed.
- Evidence-to-practice gap: critically appraised external evidence needs to become computable, traceable, and updateable within learning health systems (LHS).

2 RESEARCH GAP

Survival modelling	Methodological handling varies across censoring, competing events, missingness, calibration, and UQ.
Uncertainty quantification (UQ)	Survival risk estimates evolve over time, vary by event & horizon. Clinical action cannot be determined from risk estimates alone.
Knowledge-informed ML	Clinical evidence is graded, contextual, and evolving; making it computable and traceable remains difficult.
Neuro-symbolic (NeSy) AI	Clinical NeSy for CPMs remains early-stage, with limited small-scale or classification-based studies; safe rule injection requires weighting, validation, and audit.

3 PROPOSED FRAMEWORK

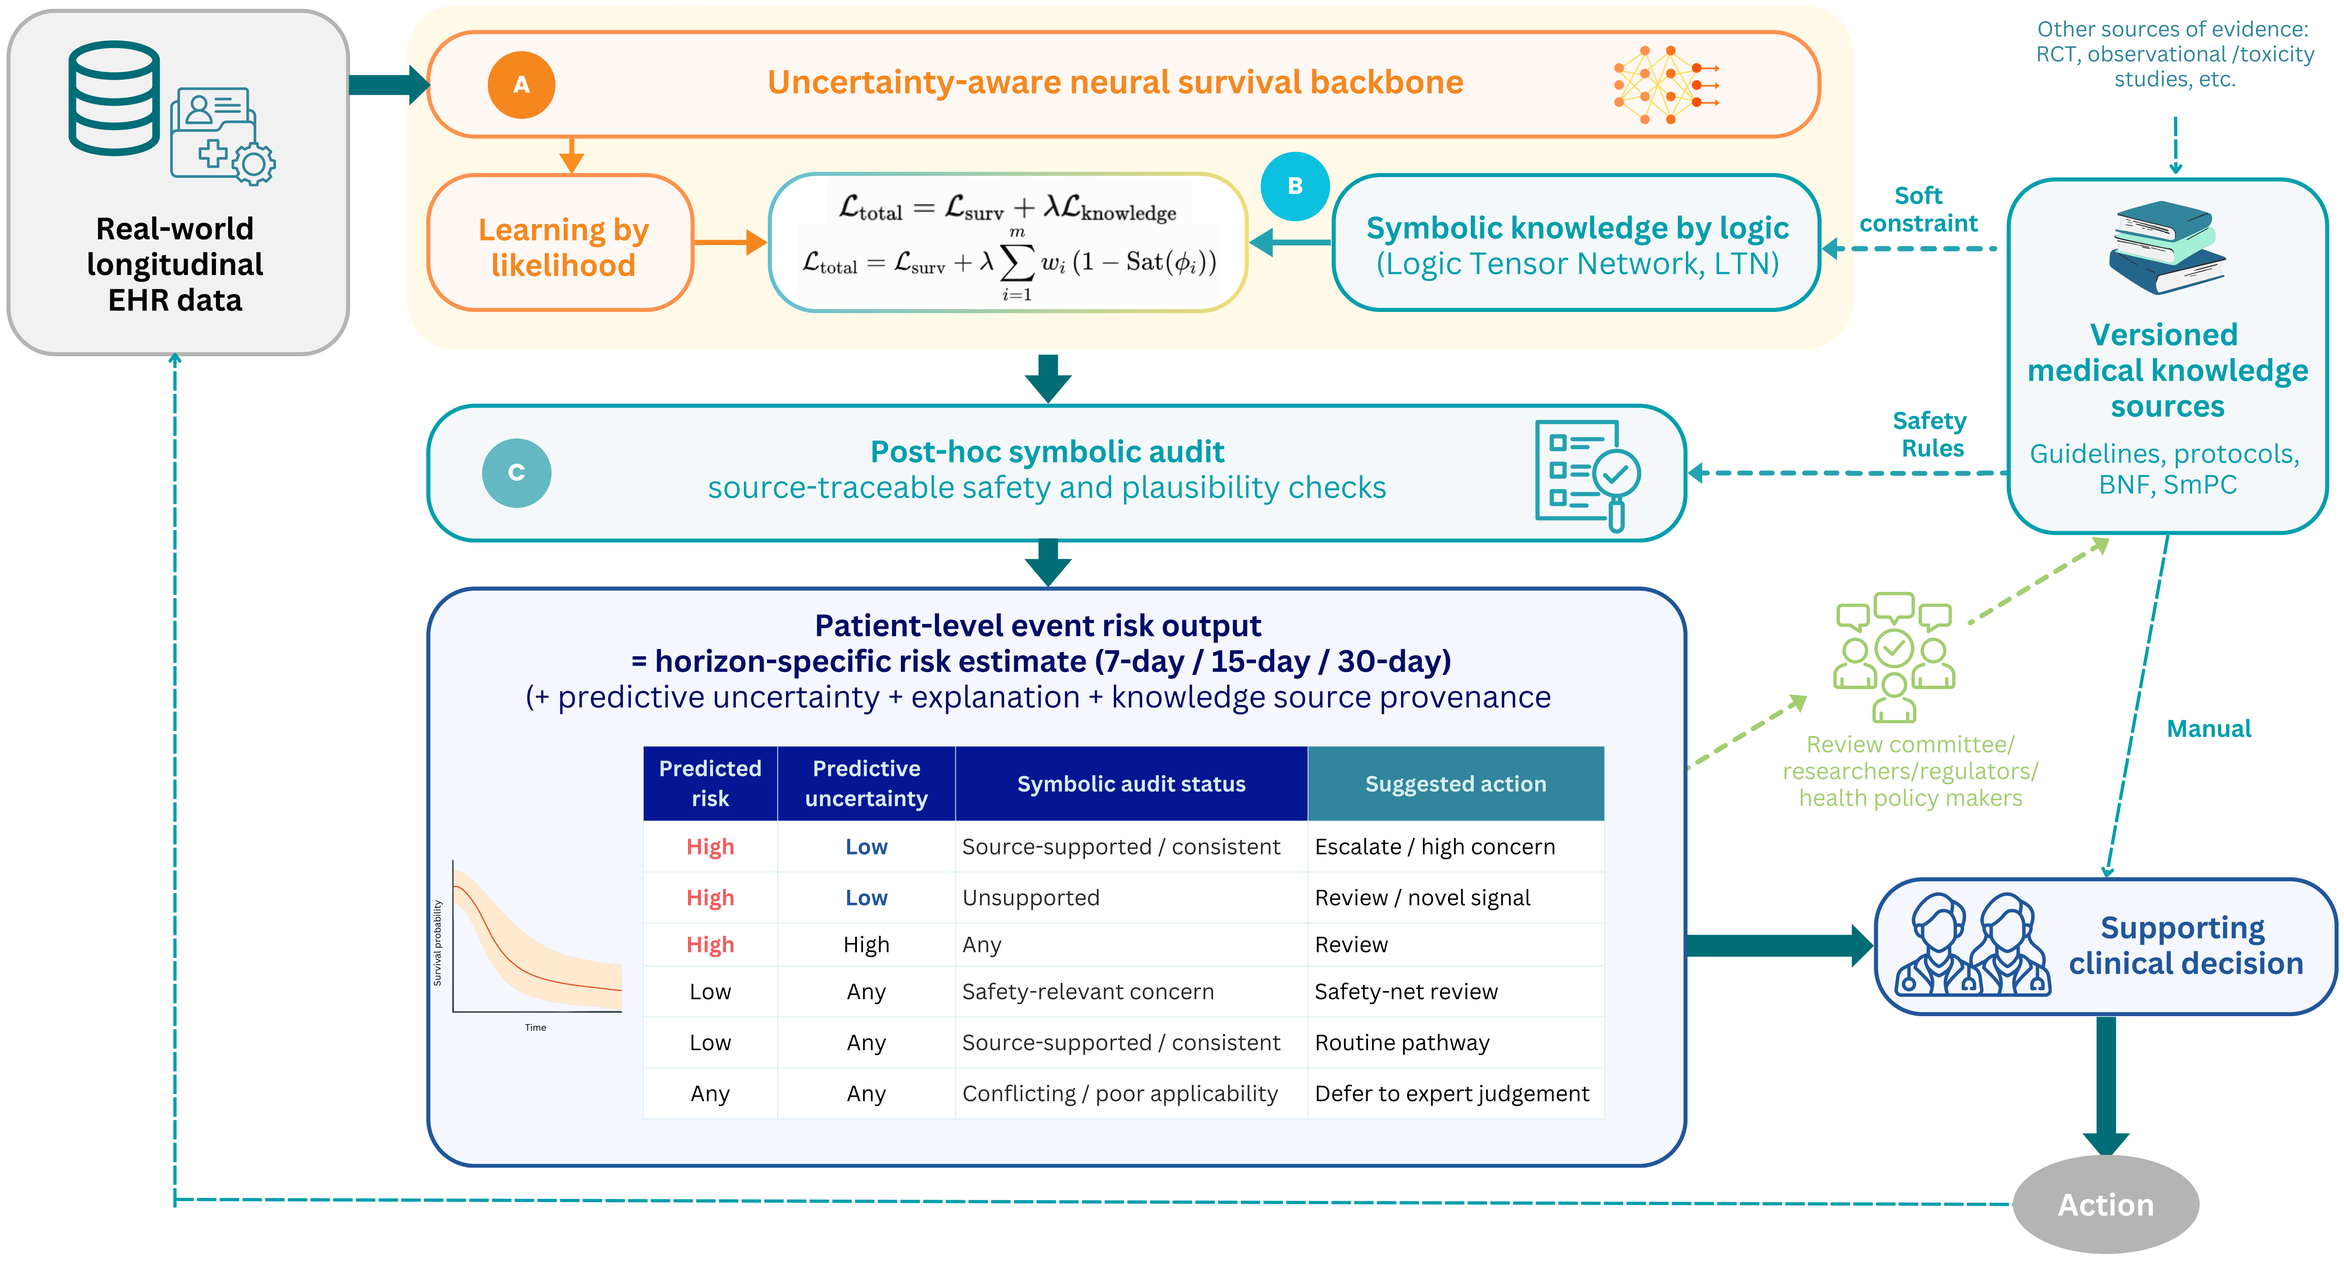

Predicted risk	Predictive uncertainty	Symbolic audit status	Suggested action
High	Low	Source-supported / consistent	Escalate / high concern
High	Low	Unsupported	Review / novel signal
High	High	Any	Review
Low	Any	Safety-relevant concern	Safety-net review
Low	Any	Source-supported / consistent	Routine pathway
Any	Any	Conflicting / poor applicability	Defer to expert judgement

4 PROPOSED METHODS

B Symbolic Integration in Loss Function via Logic Tensor Network (LTN)

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{surv}} + \lambda \mathcal{L}_{\text{knowledge}}$$

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{surv}} + \lambda \sum_{i=1}^{m} w_i \, (1 - \text{Sat}(\phi_i))$$

Where:

$\mathcal{L}_{\text{surv}}$ = competing-risk-aware survival loss
ϕ_i = selected clinical rule
m = number of selected clinical rules
$\text{Sat}(\phi_i) \in [0, 1]$ = fuzzy satisfaction of rule ϕ_i
w_i = optional per-rule weight
λ = global strength of knowledge regularisation

Example:

$$\forall x : HighLactate(x) \wedge LowPlatelets(x) \rightarrow HighMortality(x)$$

5 RESEARCH PLAN

Arm

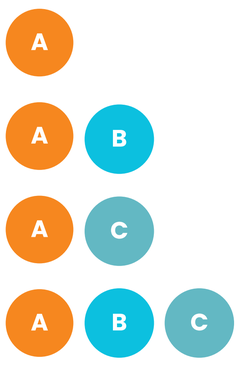

Evaluation metrics

- Cause-specific C-index / time-dependent AUC
- Competing-risk negative log-likelihood
- Cause-specific calibration curves
- Cumulative Incidence Function (CIF)-based Brier score / integrated Brier score (IBS)
- Uncertainty–error alignment by event and horizon
- Rule satisfaction / violation
- λ-sensitivity: calibration vs rule satisfaction
- Case review of discordant predictions

Abbreviations:
AUC: area inder the curve
BNF: British National Formulary;
C-index: concordance index
CPM: clinical prediction model;
EHR: electronic health record;
LHS: learning health system;
LNN: Logic Neural Network;
LTN: Logic Tensor Network;
RCT: randomised clinical trial;
SmPC: Summary of Product Characteristics;
UQ: uncertainty quantification.

Key references:
Guise et al., 2018, doi: 10.1007/s11606-018-4633-1; Badreddine et al., 2022, doi: 10.1016/j.artint.2021.103649; Sirocchi et al., 2024, doi: 10.1186/s12911-024-02582-4;. Ries and Seth, 2025, doi:10.48550/arXiv.2512.18129; Mondal et al., 2025, doi: 10.3390/app152111806; De Santis et al., 2026, doi:10.1016/j.engappai.2026.114920.