Transformer-Based Patient Response Modeling Across All Perioperative Phases for Cardiac Surgery-Associated Acute Kidney Injury Management*

Yao Ji[1], Yan Li[2], Guanghui Lan[1], Zehua Dong[3], Shuoling Li[3], Ilker Guven[4], Xiaoyu Chen[3], Lihui Bai[4], Jiapeng Huang[4]

[1]Georgia Institute of Technology • [2]Texas A&M University • [3]University at Buffalo, SUNY • [4]University of Louisville

*This work is partially supported by Air Force Office of Scientific Research grant FA9550-22-1-0447 and American Heart Association grant 23CSA1052735.

1 Background & Motivation

- **Acute Kidney Injury (AKI)** is a common complication after cardiac surgery, affecting up to 30% of patients with significant morbidity and mortality.
- AKI risk evolves over time across perioperative phases; actions now can prevent or worsen downstream injury.
- Clinicians repeatedly choose interventions as new labs arrive — a naturally sequential decision-making problem.
- **Goal:** Learn policies that reduce AKI risk while incorporating clinical constraints and expert knowledge.

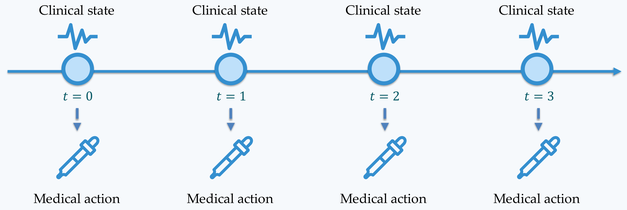

Fig. 1: AKI management as sequential decision-making across perioperative phases.

We treat AKI management as a sequence of decisions where the intervention shapes future

2 Dataset

- Scale: 108,258 total samples from the STS Adult Cardiac Surgery Database.
- States: 340 state features (199 pre-op + 23 intra-op + 96 post-op + 22 discharge).
- Actions: 122 action features (32 pre-op + 54 intra-op + 19 post-op + 17 discharge).

Table 1: AKI Staging (KDIGO) — 66,335 stageable patients 23.1% AKI incidence

Stage	Count	%	KDIGO Criteria
0	50,837	76.9	No AKI
1	11,121	16.8	Increase ≥ 0.3 or 1.5–1.9x baseline
2	1,518	2.3	2.0–2.9x baseline
3	2,618	4.0	≥ 3.0x baseline or peak ≥ 4.0

Table 2: Per-Task Loss Design and Evaluation Metrics

Task Type	Preprocess	Loss
Classification	Learned embedding per class	Cross-Entropy (CE)
Ordinal	Same as classification	CE with class weights
Abnormality	Same as classification	Focal Loss
Continuous	log1p + min-max normalization	Weighted MSE

Missing values are explicitly excluded from loss and gradient updates.

3 Methods: RL Framework

- **State s_t**: Patient status (demographics, vitals, labs, urine output, creatinine).
- **Action a_t**: Clinical interventions (fluids, medications, blood transfusion).
- **Transition kernel $s_{t+1} \sim p(\cdot\,|s_t, a_t)$**: learned via transformer.
- **Policy $\pi(a|s)$:** Optimized to minimize expected AKI risk.

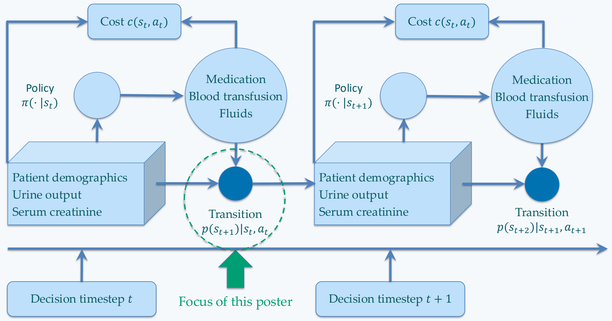

Fig. 2: RL formulation for perioperative AKI management.

4 Model Architecture

- Decoder-only transformer for kernel estimation $p(s_{t+1}|s_t, a_t)$
- Generative simulator for policy evaluation via Monte Carlo rollouts.
- Policy learning: behavior cloning initialization + actor-critic with regularization.

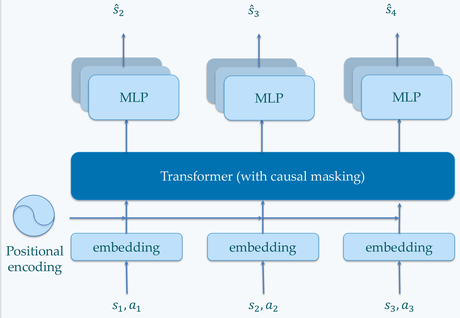

Fig. 3: Transformer-based kernel estimator:

32.2M parameters

132 prediction heads

0.5 Head dropout

0.3 Transformer dropout

300 epochs

AdamW with peak LR 1.5×10^4, weight decay 0.01, linear warmup (5%) + cosine decay to 1% of peak LR

5 Results: Patient Response (Transition Kernel)

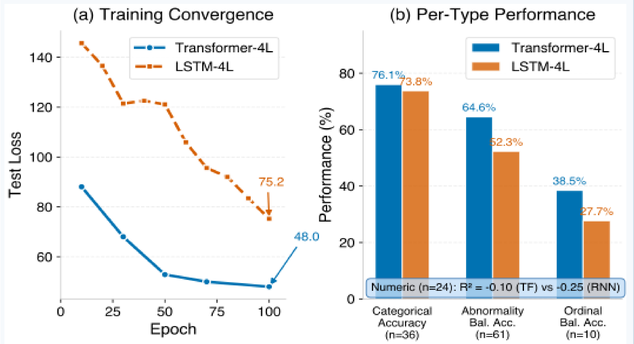

Fig. 4: Transformer vs. LSTM kernel performance. (a) The Transformer converges to 36% lower test loss than the LSTM, reaching its minimum by epoch 50 while the LSTM is still improving at epoch 100. (b) The Transformer outperforms the LSTM across all discrete feature types, with the largest gap on rare abnormality detection (64.6% vs 52.3% balanced accuracy).

6 Results: AKI Prediction

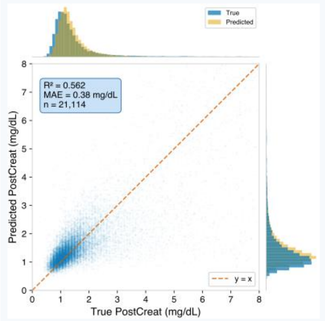

Fig. 5: The kernel predicts post-operative creatinine with $R^2 = 0.562$ in mg/dL (MAE ≈ 0.3mg/dL), accurately tracking the clinically critical high-creatinine tail.

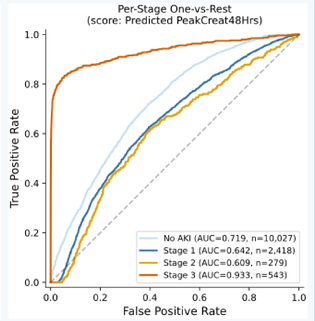

Fig. 6: KDIGO stage classification achieves 93.3% AUROC for Stage 3 AKI, with 95%+ within ±1 stage accuracy across all severity levels.

7 Results: Complication Prediction

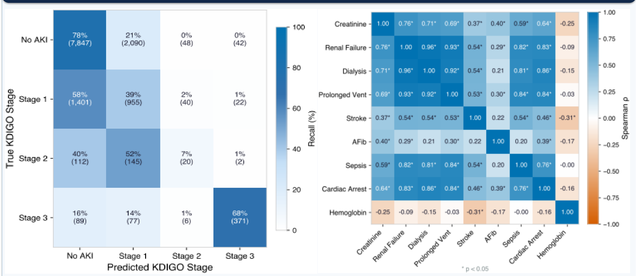

Fig. 7: KDIGO stage classification derived from predicted creatinine shows strong diagonal dominance, with most misclassifications within ±1 stage.

Fig. 8: The kernel learns clinically correct action–outcome relationships: higher intraop hemoglobin is renoprotective while longer cross-clamp and blood product use increase creatinine.

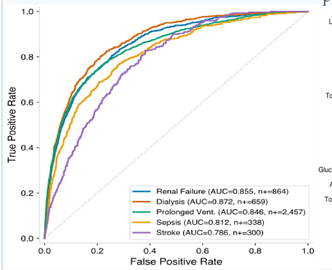

Fig. 9: Complication ROC. The kernel simultaneously predicts five major postoperative complications with AUROCs 0.79–0.87, with dialysis (0.872) and renal failure (0.855)

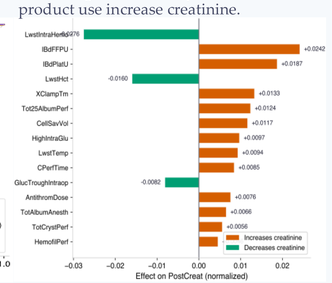

Fig. 10: Actions worsen creatinine also increase renal failure ($\rho = 0.76$), dialysis ($\rho = 0.71$), and cardiac arrest ($\rho = 0.64$), validating creatinine as a safe RL optimization target.

8 Discussion & Future Work

High concordance ($\rho = 0.76$ Creat ↔RenFail) validates creatinine as an RL optimization target, but a composite reward incorporating mortality, cardiac arrest, and sepsis penalties would provide stronger safety guarantees against reward hacking.

Next steps: train an offline RL policy using the frozen kernel, evaluate via off-policy estimation against observed clinician actions, and validate on held-out institutions before any clinical pilot.

#SCA2026

SOCIETY OF CARDIOVASCULAR ANESTHESIOLOGISTS
Knowledge • Care • Investigation