



Labor-AI Lab

Personalizing Obstetric Care Through AI

www.labor-ai.org

labor.ai.research@gmail.com

Reinforcement Learning for Clinical Recommendations During Labor

Gabriella Gross^{1,2}, Simcha Yagel¹, Yishai Sompolinsky¹, Michal Lipschuetz^{1,3}

¹Obstetrics & Gynecology Division, Hadassah Medical Center, Faculty of Medicine of the Hebrew University of Jerusalem, Israel,

²The Rachel and Selim Benin School of Computer Science and Engineering, The Hebrew University of Jerusalem, Israel,

³Henrietta Szold Hadassah Hebrew University School of Nursing in the Faculty of Medicine, Jerusalem, Israel



Introduction

During labor, clinicians repeatedly choose whether to titrate pitocin, perform artificial rupture of membranes (AROM), or continue observation as maternal, fetal, and labor conditions evolve.

Aim: To develop a reinforcement learning policy for intrapartum management from historical labor data, and to compare its recommendations with those of clinicians.

Data

Retrospective study of a stratified cohort of 8,803 labor events, represented by 20-minute decision steps.

Patient-level 80/20 split: 7,036 training and 1,767 held-out test labors.

Methods

Reinforcement Learning Framework for Intrapartum Decision-Making



Labor was modeled as a Markov decision process with 31 data-driven clinical states and 3 action classes (AROM, pitocin dose, wait).

From past labors we estimated how each action changes the next clinical state and the expected reward. Favorable outcomes received a reward and adverse events received a punishment. The policy picks the best valid action in each state by maximizing expected reward.

Results

Policy-clinician agreement was 95.2% overall and 74.7% for non-wait decisions. The AI favored slightly more conservative pitocin use.





Figure 2. *estimated expected reward of each action by clinical state. Darker = higher expected reward.*

(a) Repeated held-out evaluation of policy-iteration strategy.



(b) AI advantage over chart actions and passive care




Figure 3. *Estimated expected reward over 100 patient holdouts (higher is better).*

Conclusion

The AI reproduced clinical decisions with high agreement. When it differed, it usually recommended a lower pitocin dosage. Under the modeled reward, the AI policy scored higher than clinician actions and wait-only care. Whether this improves safety in practice requires large prospective trials.