

Assessing Risk in Implementing New AI Triage Tools

– How much risk is reasonable in an already risky world?



Centre for Biomedical Ethics
Yong Loo Lin School of Medicine




SHAPES

An NUS Centre for Biomedical Ethics initiative supported by the Singapore Ministry of Health's National Medical Research Council

Coauthors and Contributors

- Michael Dunn – National University of Singapore (NUS)
- Julian Savulescu – NUS
- Jerry Menikoff – NUS
- Sumytra Menon – NUS
- Angela Ballantyne – University of Otago
- Annette Braunack-Mayer – University of Wollongong
- Pavitra Krishnaswamy – A*STAR
- Marcus Ong – Duke NUS
- Liu Nan – Duke NUS
- Tamra Lysaght – University of Sydney
- Mayli Mertens – University of Antwerp

AI Ethics ‘from the bottom up’

Our approach: from overarching frameworks and ‘big picture’ analysis  use-case derived challenges arising in ‘non-ideal’ practical settings

The case leads: ethical analysis springboards out of the specifics of the technology being developed, and the implementation plans of the clinician research teams

Avoids the ‘shoehorning problem’

We aim to shape and refine principled understandings from the ground up.

The “AI Chasm”

- the gap between the development of a robust algorithm and its clinically meaningful application



Assessing Risk in Implementing New AI Triage Tools

- How much risk is reasonable in an already risky world?

Broad Claim: AI risk should be ethically evaluated similarly to other medical interventions and processes appropriate to the level of risk they pose. The assessment of risk needs to be calibrated to levels of risk that are acceptable for translation of new tech or other interventions into clinical contexts, and not necessarily treated as an exceptional type of risk.

Backdrop: Singapore's Patient Acuity Category Scale (PACS)

P1: requires immediate attention

P2: in severe distress and requires critical care

P3: ambulant and presents with mild to moderate symptoms

P4: nonemergency cases more appropriately managed in the primary care setting

- Connects to certain requirements in patient care and resource allocation
- PACS relies heavily on clinical experience
- 70% of doctors see P1 and P2 patients
- Better prioritisation for admission needed

Case Study:


Score for Emergency Risk Prediction (SERP) – Machine Learning Triage Tool for Estimating Mortality After Emergency Admissions

- SERP is not a severity classification index; does not instigate care pathway
- SERP outperformed PACS at predicting mortality in two retrospective studies
- Further evidence is needed to validate real-world predictive capabilities
- An RCT may not be feasible, so how should we implement safely?



Learning Health Systems Framework (LHS) – *a principle-sensitive approach*

1. Respect the rights and dignity of patients
2. Respect clinician judgments
3. Provide optimal clinical care to each patient
4. Avoid imposing nonclinical risks and burdens on patients
5. Address health inequities
6. Conduct continuous learning activities that improve the quality of care and health systems



SERP Implementation Strategy₁

1. Six-month prospective and non-interventional “silent trial”
 - SERP score unknown to triage nurse and not used in practice.
2. If silent trial successful aligns, replace PACS with “PAC+ model 1”
 - SERP score finetunes PACS score

Silent Trial

Benefits

- Prospective data without intervention
- Further evidence base

Risks

- Delaying potential benefits of SERP
- Not accounting for how nurses will use it in practice





Is AI risk exceptional?

Emotional and Environmental Factors

PACS relies heavily on clinical judgment, SERP is unaffected by factors affecting humans

Technology and Automation Bias

Prospective data will help to understand how healthcare professionals will use the tool in practice

Interpretability

SERP is interpretable; points-based algorithm with only five variables to generate raw risk score; nurses more likely to use it correctly

Silent trial is appropriate






PAC + Model 1: *Finetuning of PACS with SERP*

If the SERP score <26 , patient down-triaged by 1 PACS score

If the SERP score ≥ 26 , patient will be up-triaged by 1 PACS score

The final decision to finetune will still be one determined by clinical judgment.



Assessing Risk

- False negatives (low sensitivity) – SERP incorrectly underestimates mortality risks, informing clinical decisions that fail to attenuate risk of death or serious harm correctly, leading to worse patient outcomes
- False positives (low specificity) – SERP incorrectly overestimates mortality risks, leading to the unnecessary use of limited medical resources, put those who are de-prioritised at potential higher risk

Assessing PAC+ Model 1

Benefits

- Retrospective evidence base for finetuning is good
- Addresses the high and growing risks of over-triage and misuse of increasingly limited EM resources

Risks

- Full replacement PACS = major deviation from established clinical care, imposing new risks on patients due to the lack of calibrated workflows to triage score



PAC+ Model 1 approach is supported



References

1. Xie F, Ong MEH, Liew JNMH, et al. Development and Assessment of an Interpretable Machine Learning Triage Tool for Estimating Mortality After Emergency Admissions. *JAMA Netw Open*. 2021;4(8):e2118467. doi:10.1001/jamanetworkopen.2021.18467
2. Ruth R. Faden, Nancy E. Kass, Steven N. Goodman, Peter Pronovost, Sean Tunis, and Tom L. Beauchamp, “An Ethics Framework for a Learning Health Care System: A Departure from Traditional Research Ethics and Clinical Ethics,” *Ethical Oversight of Learning Health Care Systems*, Hastings Center Report Special Report 43, no. 1 (2013): S16-S27. DOI: 10.1002/hast.134
3. Wong, Yvonne. “Integrating the Score for Emergency Risk Prediction (SERP) and Patient Acuity Categorical Scale (PACS) to better predict 30-day mortality among patients presenting to the emergency department (ED)”. MD Scholarly Development Programme requirements for Duke-NUS Medical School. (March 2024)
4. McCradden, Melissa D, James A Anderson, Elizabeth A. Stephenson, Erik Drysdale, Lauren Erdman, Anna Goldenberg, and Randi Zlotnik Shaul. “A Research Ethics Framework for the Clinical Translation of Healthcare Machine Learning.” *The American Journal of Bioethics* 22, no. 5 (January 20, 2022): 1–15. <https://doi.org/10.1080/15265161.2021.2013977>.



Takk

спасибо

Gracias

Gracias

ευχαριστώ

ขอบคุณครับ

Спасибо Takk