Skip the Prompts. Encode the Logic Directly!

Multimodal Bayesian Network for Robust Assessment of Casualty in Autonomous Triage

Szymon Rusiecki[1,2], Cecilia G Morales[2], Kimberly Elenberg[2], Leonard Weiss[3], Artur Dubrawski[2]
1 - AGH University of Krakow, 2 - Carnegie Mellon University, 3 - University of Pittsburgh

NEURAL INFORMATION PROCESSING SYSTEMS

The DARPA Triage Challenge (DTC)

Context: Mass Casualty Incidents (MCIs) overwhelm human responders. The DTC seeks to develop autonomous systems capable of providing situational awareness and triage in the first minutes of an event.

Problem: Existing autonomous perception models are brittle. In the chaotic, noisy environment, standalone computer vision often fails to make reliable predictions.

Our solution: We propose a Bayesian network (**BN**) that integrates fragmented sensor data into a coherent clinical assessment to build a reasoning layer that survives data loss.

Methodology: Expert-Driven Design

Knowledge Elicitation: Instead of training on datasets, which are scarce for MCIs, we constructed the network structure and Conditional Probability Tables (CPTs) via a structured process:

- **Expert Elicitation**: Iterative sessions with emergency medicine specialists.
- **Medical Logic**: Encoding rules like “Amputations often imply Severe Hemorrhage”.

System Architecture & Hardware

Platform: Field-deployable robotic unit (ROS2 based).

Sensor Suite: High-resolution RGB camera, thermal camera, LiDAR, radar, microphone.

Inference Engine: SMILE (Structural Modeling, Inference, and Learning Engine) operating in real-time (<1ms/update).

Integration: Each physiological estimator operates as an independent ROS2 node publishing to the central Bayesian Network.

Bayesian Network Influence Graph

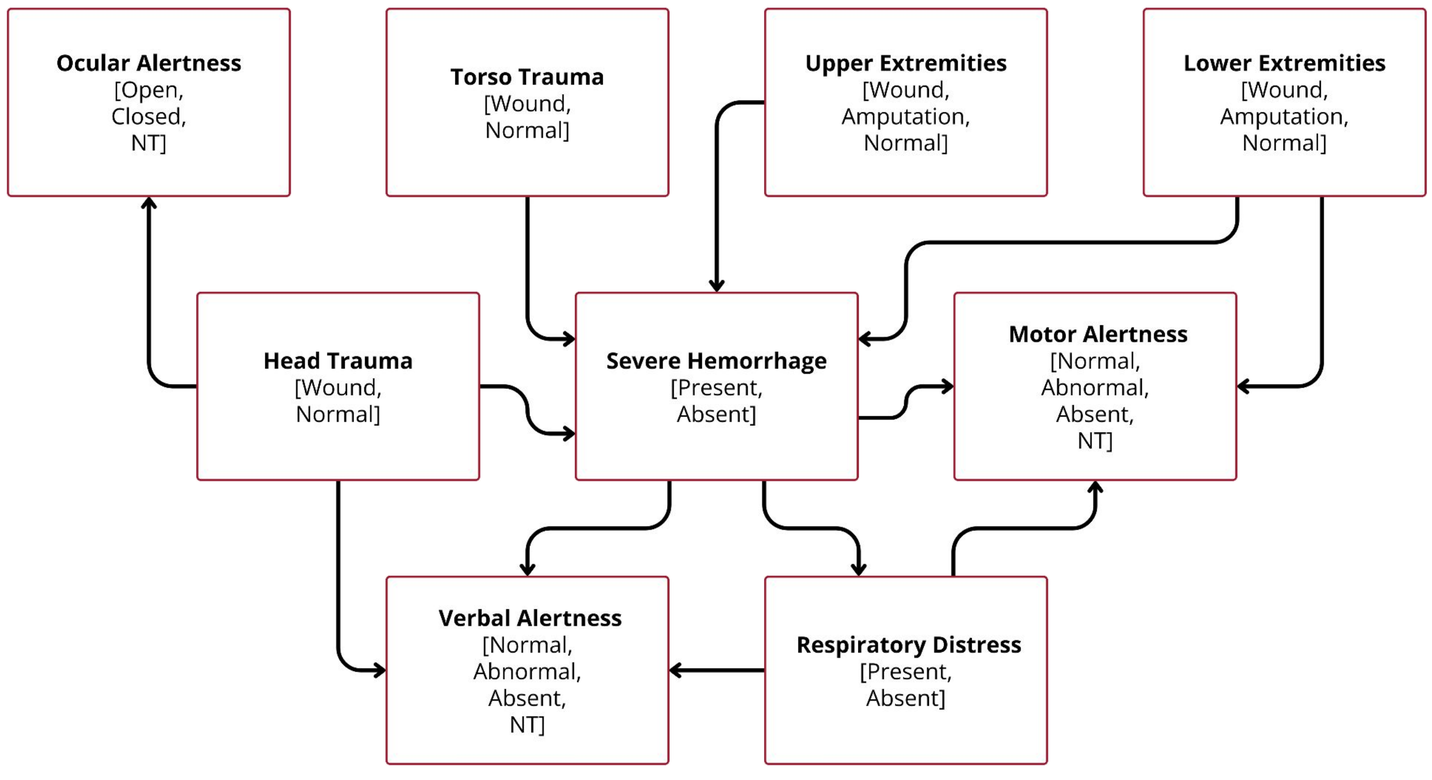

Influence Graph established in Expert Elicitation process. Every node represents CPT whose weights were set manually.

Results

	Correct Assignments		Assignment Attempts	
Vital	**Robot**	**Robot + BN**	**Robot**	**Robot + BN**
Severe Hemorrhage	6	12	12	19
Respiratory Distress	5	16	6	19
Head Trauma	0	15	0	19
Torso Trauma	0	11	0	19
Lower Ext. Trauma	8	11	14	19
Upper Ext. Trauma	3	8	14	19
Motor Alertness	0	8	0	19
Verbal Alertness	3	7	8	19
Ocular Alertness	0	8	1	19
Total	25	96	55	171
Reliability	–	–	0.31	**0.95**
Performance	14%	**53%**	–	–
Accuracy	46%	**56%**	–	–

Comparison of Correct Assignments and Assignment Attempts with and without BN across all casualties in the 1st year of the DTC

Conclusions

We already have tools that allow us to easily create resource-less baselines for decision-making systems, with virtually no integration into the existing architecture. In our case, this enabled us to **quadruple** the total number of points gained without loss in accuracy of the system.

CMU's DARPA Triage Challenge TEAM Chiron

DTC DARPA TRIAGE CHALLENGE

References

O. Adebayo et al. (2023). "Exploring the effectiveness of artificial intelligence, machine learning and deep learning in trauma triage: A systematic review and meta-analysis." In: Digital Health: 9:20552076231205736.

J. J. Wadden (2022). "Defining the undefinable: the black box problem in healthcare artificial intelligence." In: Journal of Medical Ethics, 48(10):764-768.

B. J. Gabbe et al. (2022). "Review of the requirements for effective mass casualty preparedness for trauma systems. A disaster waiting to happen?" In: British Journal of Anaesthesia, 128(2):e158–e167.

S. Sadeghi et al. (2006). "A bayesian model for triage decision support" In: International Journal of Medical Informatics, 75(5):403–411, 2006.

M. J. Druzdzel (1999). "SMILE: Structural modeling, inference, and learning engine and GeNIe: a development environment for graphical decision-theoretic models.". In: AAAI/IAAI, pages 902–903.

J. Pearl et al. (1988). “Probabilistic reasoning in intelligent systems: networks of plausible inference”. In: Synthese-Dordrecht 104(1).