Aggregated Knowledge Learning for Dynamic Vehicle-Task Assignment in Emergency Medical Services

Changle Song David Levinson Emily Moylan
School of Civil Engineering, The University of Sydney, Australia
Corresponding author: changle.song@sydney.edu.au

Extended abstract submitted for presentation at the 12th Triennial Symposium on Transportation Analysis conference (TRISTAN XII)
June 22-27, 2025, Okinawa, Japan

March 31, 2025

Keywords: Aggregated Knowledge Learning; Assignment Problem; Graph Neural Networks; Emergency Medical Services

1 INTRODUCTION

Optimizing vehicle-task assignment is a critical challenge in emergency medical services (EMS), logistics, and ride-sharing, with major implications for operational efficiency, response times, and ultimately, life-saving outcomes (Alonso-Mora *et al.*, 2017). In EMS, efficient assignment of emergency vehicles to tasks—based on real-time data and varying task (referred to as "incidents" in EMS contexts) priorities—is essential for improving patient survival rates and minimizing response times. As demand for EMS continues to rise in densely populated areas, traditional methods like linear programming and heuristic algorithms often struggle to adapt to the complex and dynamic nature of EMS environments (Liu *et al.*, 2024). These methods may lack the scalability and flexibility to handle real-time traffic conditions, varied vehicle capabilities, and task urgency levels, all of which impact response times and service efficacy.

Recent machine learning approaches, particularly those utilizing homogeneous feature spaces, have shown potential for enhancing vehicle-task assignments in EMS. However, they typically fall short in adapting to the heterogeneous features inherent in EMS, such as specific medical equipment requirements, task priority, and dynamic traffic data (Qin *et al.*, 2020, Paul *et al.*, 2024). Graph Neural Networks (GNNs), due to their ability to model relational data, offer promise for EMS applications but remain underutilized in vehicle-task assignment settings (Wu *et al.*, 2020, Dai *et al.*, 2021). In this study, we introduce a framework centered on Aggregated Knowledge Learning (AKL), a novel training method that leverages the outputs from various assignment models, including the Hungarian and Pareto methods, to create a robust GNN-based model for EMS vehicle-task assignments. The AKL approach enhances the model's adaptability to dynamic and heterogeneous EMS scenarios, addressing limitations in traditional and existing machine learning-based methods.

2 METHODOLOGY

In this section, we present our methodology for optimizing vehicle-task assignments in dynamic transportation networks. We begin by formulating the problem, carefully defining all notations and parameters used throughout. We then introduce the Aggregated Knowledge Learning (AKL) framework, which aggregates insights from traditional models to generate high-quality training data. Subsequently, we describe both the proposed learning model, which utilizes Graph Neural Networks (GNNs), and the traditional assignment models used for comparison. Detailed explanations are provided for each component, ensuring clarity and coherence.

2.1 PROBLEM FORMULATION

Vehicle Assignment: $A^{(t)}$ is the assignment matrix at time t. $A^{(t)}_{v,i}$ in Equation 1 indicates the match status whether vehicle v is assigned to task i at time t.

$$A^{(t-1)}_{v,i} = \begin{cases} 1, & \text{if vehicle } v \text{ is assigned to task } i \text{ at time } t, \\ 0, & \text{otherwise.} \end{cases} \tag{1}$$

Efficiently assigning vehicles to tasks in transportation networks is essential for optimizing operational performance metrics such as response time and service quality. This problem is inherently dynamic due to the continuous arrival of new tasks and changes in vehicle availability. Our objective is to develop a model that adapts to these changes and maximizes the overall system benefit.

We denote v as the set of vehicles and V is the vehicle set where $v \in V$. Similarly, i presents the incident and I is the set of incident i. The available vehicles and incidents are different at each time. Here we use T to present the timeline, and $t \in T$ is the discrete time point.

At each time t, we have a set of feasible assignments between vehicles and tasks, denoted by $E^{(t)} \subseteq V \times I$.

For each vehicle-task pair (v, i) at time t, we define $R_{v,i}^{(t)}$ and $S_{v,i}^{(t)}$ are the travel time are survival rate for vehicle v to reach task i at time t respectively.

The decision variable is $x_{v,i}^{(t)}$, where

- $x_{v,i}^{(t)} \in \{0, 1\}$: indicates if vehicle v is assigned to task i at time t (1 if assigned, 0 otherwise).

Objective Function: Our objective is to maximize the total benefit B over the entire simulation period. Each task's benefit is considered only once, at the time it is last assigned, to prevent double-counting.

Let v_i be the vehicle assigned to task i, and T_i^{assign} be the time when vehicle v_i was last assigned to task i. The total benefit B is:

$$\max_x \quad B = \sum_{i \in I} \left(\beta_1 S_{v_i,i}^{(T_i^{\text{assign}})} - \beta_2 R_{v_i,i}^{(T_i^{\text{assign}})} \right). \tag{2}$$

β_1 is the weight of the survival rate or service benefit and the coefficient β_2 represents the importance of the response time. $\beta_1 >> \beta_2$ in this research.

Assignment Constraints: Each vehicle can be assigned to at most one task at any given time:

$$\sum_{i \in I} x_{v,i}^{(t)} \leq 1, \quad \forall v \in V, \forall t \in T. \tag{3}$$

Each task can be assigned to at most one vehicle at any given time:

$$\sum_{v \in V} x_{v,i}^{(t)} \leq 1, \quad \forall i \in I, \forall t \in T. \tag{4}$$

2.2 AGGREGATED KNOWLEDGE LEARNING FRAMEWORK

The aggregated knowledge learning (AKL) framework generates a high-quality training dataset by aggregating assignment decisions from multiple traditional models. The best-performing model in each simulation, based on total benefit, provides the assignment decisions used for training.

Rationale: Evaluating models based on overall performance over the simulation period accounts for the cascading effects of decisions in a dynamic environment.

AKL Process:

1. **Simulations with Random Seeds:** Run S independent simulations, each with a unique random seed $s \in \{1, 2, \ldots, S\}$.
2. **Application of Traditional Models:** In each simulation s, apply a set of traditional assignment models $\mathcal{M} = \{M_1, M_2, M_3\}$. The models we use in this study include Hungarian algorithm, No Reassignment Policy and Pareto improvement criterion.
3. **Model Evaluation:** For each model $M_j \in \mathcal{M}$, calculate the total benefit $B_{M_j}^{(s)}$ using (2).
4. **Selection of Best Model:** Choose the best-performing model $M^{*(s)}$ in simulation s:
$$M^{*(s)} = \arg \max_{M_j \in \mathcal{M}} B_{M_j}^{(s)}.$$
5. **Aggregation of Assignment Decisions:** Collect assignment decisions $x_{v,i}^{(t,s)}$ from $M^{*(s)}$ at each time t.
6. **Formation of Training Dataset:** Aggregate data D from all simulations $D^{(s)}$ for each s:
$$D = \bigcup_{s=1}^{S} D^{(s)},$$
where
$$D^{(s)} = \left\{ \left(F_{v,i}^{(t,s)}, y_{v,i}^{(t,s)} \right) \,\middle|\, (v, i) \in E^{(t)}, t \in T \right\},$$

$E^{(t)}$ is the simulation experiment running at time t. (v,i) is available vehicle-task pair at the situation $E^{(t)}$. $F_{v,i}^{(t,s)}$ is feature vector for vehicle-task pair (v,i) at time t in simulation s. $y_{v,i}^{(t,s)} = x_{v,i}^{(t,s)}$ denotes assignment decision from $M^{*(s)}$.

7. **Model Training:** Implement supervised learning method to train the core model (GNN in this case study) based on the collected data from step 6.

3 RESULTS AND DISCUSSION

We simulate dynamic EMS operations in Sydney Australia using a road network constructed from OpenStreetMap and travel times obtained from CompassIoT data. Incidents are generated stochastically based on population density, and vehicles are categorized by equipment level (e.g., CCPR, ECPR). Survival functions vary by vehicle-task compatibility and delay. The GNN model, trained using AKL on 200 simulation runs, is benchmarked against the three baseline models.

3.1 PERFORMANCE EVALUATION

Using AKL, the GNN model consistently surpasses traditional methods in maximizing survival rates and minimizing response times in ambulance system. Figure 1 shows the survival rate and response time performance across 10 simulations with 10 vehicles and 30 incidents over 8 hours.

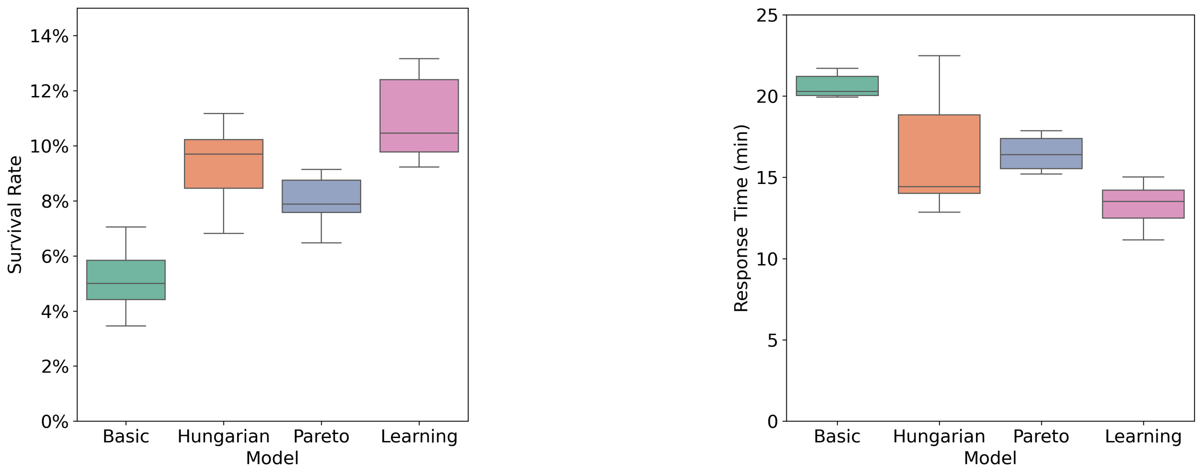

(a) *Survival Rate* (b) *Response Time*

Figure 1 – *Performance Comparison with 30 Incidents and 10 Vehicles*

In Figure 1a, the Learning model (GNN with AKL training approach) achieves the highest median survival rate of **11.5%**, significantly outperforming the Hungarian model (**9.7%**), Pareto model (**7.9%**), and Basic model (**5.3%**). The lower quartile of the Learning model also surpasses the median survival rates of the Hungarian and Pareto models, demonstrating superior performance even in less favorable scenarios.

Figure 1b presents the response time results. The Learning model achieves the lowest median response time of **13.3 minutes**, markedly outperforming the Basic/no re-assignment model (**20.4 minutes**), Hungarian model (**16.2 minutes**), and Pareto model (**16.4 minutes**). The Learning model also exhibits a tighter interquartile range compared to the Hungarian model, indicating more consistent and reliable performance across different scenarios.

3.2 SCALABILITY WITH INCIDENT VOLUME

AKL's scalability was tested across scenarios with 30, 40, and 50 incidents, as shown in Figure 2. The GNN model, leveraging AKL, maintained the highest survival rates and shortest response times, demonstrating its robustness and adaptability under varying load conditions.

In Figure 2a, the Learning model maintains the highest median survival rates across all incident volumes: **11.5%** for 30 incidents, **9.4%** for 40 incidents, and **7.1%** for 50 incidents. These values significantly outperform the Hungarian model (**9.7%**, **8.1%**, **5.6%**), Pareto model (**7.9%**, **7.7%**, **6.8%**), and Basic model (**5.3%**, **4.8%**, **4.3%**) at corresponding incident levels.

Similarly, in terms of response time (Figure 2b), the Learning model achieves the lowest median response times: **13.3 minutes** for 30 incidents, **14.7 minutes** for 40 incidents, and **23.8 minutes** for 50 incidents. In contrast, the Basic model records the longest response times, increasing from **20.4 minutes** to **29.2 minutes** as incidents increase. The Hungarian model shows a significant rise in response time at 50 incidents (**42.4 minutes**), indicating scalability issues under higher demand.

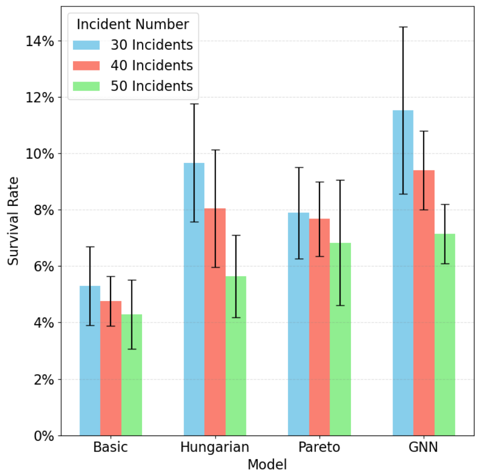

(a) *Survival Rate Across Incident Volumes*

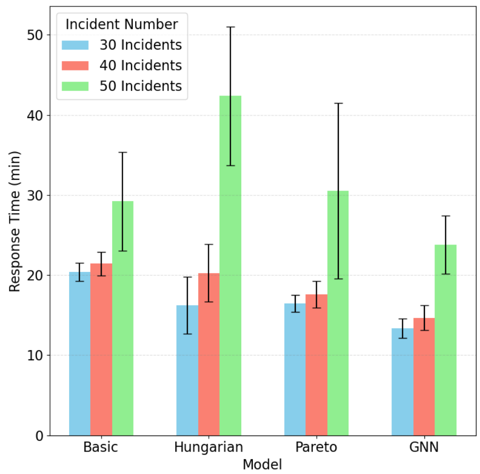

(b) *Response Time Across Incident Volumes*

Figure 2 – *Scalability of AKL Model with Increasing Incidents*

3.3 PERFORMANCE IMPLICATIONS

The strong inverse correlation between survival rates and response times highlights the effectiveness of the AKL-based model. Figures 2a and 2b show that as incident volume increases, the AKL model consistently delivers superior survival rates and reduced response times, proving its scalability and robustness.

In comparison, the Basic model's low survival rates and high response times reveal its limitations in dynamic scenarios. While the Hungarian model achieves moderate survival rates, its scalability is limited under high-demand conditions, and the Pareto model offers a balance but does not reach the performance levels of AKL. The Learning model's adaptability under varying loads suggests that it is particularly well-suited to the unpredictable demands of emergency medical services.

4 CONCLUSION

This study presents a novel framework that combines the structural strengths of GNNs with the flexibility of Aggregated Knowledge Learning (AKL) for dynamic vehicle-task assignment. By learning from the best-performing traditional model in each simulation, AKL enables the GNN to capture diverse, context-specific policies, resulting in robust and adaptive performance.

AKL's modular design allows future integration of more advanced models, including reinforcement learning or larger neural architectures, without altering the training process. It also exhibits knowledge distillation properties, enabling smaller, efficient learners to approximate complex policies. Compared to reinforcement learning, this supervised approach significantly reduces training time while maintaining strong decision quality in dynamic, heterogeneous environments. The framework is scalable and generalizable to various domains such as EMS, ride-hailing, and disaster response.

References

Alonso-Mora, Javier, Samaranayake, Samitha, Wallar, Alex, Frazzoli, Emilio, & Rus, Daniela. 2017. On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment. *Proceedings of the National Academy of Sciences*, **114**(3), 462–467.

Dai, Genan, Hu, Xiaoyang, Ge, Youming, Ning, Zhiqing, & Liu, Yubao. 2021. Attention based simplified deep residual network for citywide crowd flows prediction. *Frontiers of Computer Science*, **15**, 1–12.

Liu, Shuangxi, Lin, Zehuai, Huang, Wei, & Yan, Binbin. 2024. Current development and future prospects of multi-target assignment problem: A bibliometric analysis review. *Defence Technology*.

Paul, Aditya, Levin, Michael W, Waller, S Travis, & Rey, David. 2024. Data-driven Optimization for Drone Delivery Service Planning with Online Demand. *arXiv preprint arXiv:2404.02442*.

Qin, Zhiwei, Tang, Xiaocheng, Jiao, Yan, Zhang, Fan, Xu, Zhe, Zhu, Hongtu, & Ye, Jieping. 2020. Ride-hailing order dispatching at didi via reinforcement learning. *INFORMS Journal on Applied Analytics*, **50**(5), 272–286.

Wu, Zonghan, Pan, Shirui, Chen, Fengwen, Long, Guodong, Zhang, Chengqi, & Philip, S Yu. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, **32**(1), 4–24.