

Real-Time EV Charger Load Balancing Using Reinforcement Learning Algorithms

Yaser Issam Hamodi Aljanabi¹ and Gaurav Raj²

¹*Al-Turath University, 10013 Baghdad, Iraq*

²*Sharda School of Computing Science and Engineering, Sharda University, 201310 Greater Noida, India*
yaserissam.hamodi@uoturath.edu.iq, gaurav.raj@sharda.ac.in

Keywords: Electric Vehicles, Load Balancing, Reinforcement Learning, Peak Shaving, Smart Charging, Demand Response, Real-Time Optimization.

Abstract: The growing use of electric vehicles (EVs) is putting a lot of stress on power distribution systems. This is mostly because of uncoordinated charging, which causes demand peaks, voltage fluctuations, and higher operating costs. This research introduces a reinforcement learning (RL)-based framework for real-time electric vehicle (EV) charger load balancing, aimed at concurrently optimizing grid stability, energy expenses, and user quality of service (QoS). The charging issue is framed as a Markov Decision Process, in which the RL agent distributes charging power based on feeder capacity, tariff signals, and state-of-charge (SoC) prerequisites. A safety layer makes sure that operational limits are followed. The proposed RL method cuts peak demand by up to 25% and energy costs by about 20%, while maintaining 95% QoS. This is much better than baseline scheduling methods. The framework also works well in stressful situations like arrival bursts, tariff changes, and feeder capacity cuts. These results show that RL is a good way to manage EV charging in real time that can be scaled up and is smart.

1 INTRODUCTION

The transportation sector has quickly become electrified, which has led to an unprecedented rise in the number of electric vehicles (EVs) used around the world. EVs are important for cutting down on greenhouse gas emissions and encouraging sustainable transportation, but adding a lot of them to the grid is very hard on the existing power grid. Charging that isn't coordinated can cause voltage instability, peak load spikes, and transformer overloading. This shows how important it is to have smart load balancing systems. Conventional rule-based or optimization methods frequently do not resolve these challenges in real time, especially in the context of uncertain and fluctuating charging demand. This gap has made it possible for reinforcement learning (RL)-based methods to work, which can handle random user behavior and grid limits. Wan et al. (2018) [1] conducted one of the initial studies utilizing model-free deep reinforcement learning for real-time electric vehicle charging scheduling, illustrating that reinforcement learning can markedly decrease peak demand while maintaining user satisfaction.

Recent studies have broadened the utilization of RL for demand response and intelligent charging. Chifu et al. (2024) [2] put forward a deep Q-learning framework that improves EV charging to help demand response programs at the grid level. Their findings demonstrate enhanced stability of the power system and effective utilization of grid resources. Azzouz and Hassen (2023) [3] also came up with a decentralized DRL method for charging EVs that focuses on scalability and robustness. Their model distributes learning across several agents, which makes it more adaptable to large-scale EV networks. This is different from centralized methods, which often run into computational bottlenecks. These studies show that RL can handle the problems that come up with different types of EV fleets, uncertain arrival times, and different energy needs.

In addition to grid stability, goals that focus on the user, like lowering costs and improving service quality, are becoming just as important. Xia et al. (2025) [4] proposed a DRL-based model to reduce user costs while balancing load across multiple charging stations. Their method successfully combines grid limitations with customer preferences, making sure that charging is fair and lowering overall

operating costs. This user-centered point of view fits with the growing need for personalized charging services in cities with a lot of electric vehicles.

Even though algorithmic development has come a long way, the question of whether they can be used in real time is still very important. Damodarin et al. (2025) [5] examined reinforcement learning controllers deployed on FPGA platforms, attaining low-latency decision-making appropriate for real-time electric vehicle charging. Their research illustrates that reinforcement learning techniques are applicable beyond simulated environments and can be integrated into tangible hardware systems. This closes the gap between smart charging infrastructure that is being used in the real world and new ideas in the field.

At the same time, new computing paradigms are changing the environment in which electric vehicle charging systems work. Wang et al. (2025) [6] underscored the significance of secure data sharing in next-generation computing for distributed energy systems. Because EV charging needs sensitive user data and grid information, strong security measures are needed for trust and interoperability. Mehta and Rani (2025) [7] also talked about how AI-driven systems are becoming more common in human-computer interaction. This is helpful for understanding how people feel about and trust AI-enabled EV charging management. These viewpoints highlight that technical efficiency alone is inadequate; adoption also relies on system transparency, usability, and security.

Even with these improvements, there are still big gaps. Current research typically focuses on optimizing either grid objectives, like peak shaving, or user objectives, like minimizing costs, but seldom combines both within a cohesive framework. Additionally, decentralized methods enhance scalability; however, they remain in the nascent phase of development for multi-station, real-time implementation. There are studies on the feasibility of hardware, but there aren't many large-scale pilot projects yet. Also, EV load balancing frameworks don't usually deal with problems like data security and getting users to use them.

To overcome these constraints, this study introduces a reinforcement learning-driven real-time electric vehicle charger load balancing framework. The framework optimizes grid constraints, cost efficiency, and user satisfaction at the same time, and it can be used on real-time platforms. This research seeks to progress the field towards scalable, secure, and user-centric smart charging solutions by

leveraging the strengths of prior works [1], [7] and mitigating their limitations.

2 LITERATURE REVIEW

The incorporation of electric vehicles (EVs) into contemporary energy systems has necessitated comprehensive research on smart charging, reinforcement learning (RL), and infrastructure optimization. Early theoretical models of technology adoption serve as a valuable basis for examining user acceptance of advanced charging platforms. Nguyen and Wiese (2003) [8] created the Technology Acceptance Model (TAM) along with the IS success model to show how people use and judge digital technologies. Although their study was not specific to electric vehicles, it underscores the significance of system usability, information quality, and user satisfaction - elements that continue to be critically important when evaluating AI-driven charging systems.

Recent developments underscore the utilization of reinforcement learning for optimizing online charging. Poddubnyy et al. (2023) [9] put forth an RL framework augmented by experience replay, facilitating ongoing learning from previous interactions to adjust to changing load conditions. Their findings indicate that RL can significantly diminish fluctuations and enhance grid stability in real-time, which is essential as EV adoption escalates. Zhang et al. (2023) [10] also created a safe RL-based charging plan for residential microgrids. By adding safety limits to the optimization process, their method not only made charging more efficient, but it also kept voltage stable and avoided overload. This shows how important reliability is in real-world applications.

More advanced models are now incorporating graph learning and clustering methodologies. Ghode and Digalwar (2025) [11] presented a hybrid methodology that integrates Graph Convolutional Networks (GCN) with Deep Reinforcement Learning (DRL) to elucidate spatial-temporal relationships among multiple electric vehicles (EVs), facilitating enhanced coordination in charging allocation. In a similar vein, Saklani et al. (2025) [12] put forth a data-driven optimization framework that employs HDBSCAN clustering in conjunction with linear programming to manage real-time pricing signals. Their contributions underscore the movement towards scalable and adaptable charging optimization techniques, although validation at the hardware level is still constrained.

Another important area of research is whether RL-based charging solutions will work with the hardware. Damodarin et al. (2025) [5] showed that RL controllers can be used on FPGA platforms, which make decisions with low latency that are good for real-time environments. This is a big step toward closing the gap between research on EV infrastructure and real-world EV infrastructure. Brinkel et al. (2024) [13] also looked at real-life operational problems, like charging that was stopped or delayed. Their research examined user behavior and system-level anomalies, which frequently compromise theoretical scheduling advantages. These insights show how important it is to make sure that algorithmic strategies work with real-world problems.

Lastly, Guerrero-Silva et al. (2025) [14] did a systematic review that used principal component analysis (PCA) to find patterns in the optimization of EV charging infrastructure. Their findings suggest an increasing dependence on hybrid AI models, multi-agent coordination, and data-driven methodologies to improve scalability and efficiency. This wider view puts current research in the context of global optimization trends and points out ongoing problems like the lack of pilot-scale validation and the need to combine user adoption frameworks.

The studies reviewed show that research on optimizing EV charging is moving forward in four main areas: theoretical adoption frameworks [8], RL-

based adaptive scheduling [9], [10], hybrid AI-driven optimization [11], [12], and problems with real-world deployment [5], [13]. Table 1 gives a brief overview of the main works' contributions, methods, and problems. It shows the different ways people have tried to solve the problem, from conceptual models to hardware implementations, and it also shows the gaps that still exist. Most studies focus on either grid or user goals separately, and they don't do a good job of bringing adoption, safety, and scalability into a single framework [15], [16]. This research seeks to develop a reinforcement learning framework that integrates grid stability, cost efficiency, and user satisfaction within real-time operational constraints, addressing existing gaps.

3 METHODOLOGY

The suggested method is to make a real-time reinforcement learning (RL) framework for balancing the load of electric vehicle (EV) chargers. The framework simulates the charging environment, creates an RL agent, and tests its performance in a simulated environment with realistic parameters. Figure 1 shows the flow of the method, and Table 1 gives a summary of the assumptions made in the simulation.

Table 1: Summary of reviewed literature.

Ref. No	Author(s), Year	Focus Area	Methodology/Model	Key Contribution	Identified Limitation
[8]	Nguyen & Wiese, 2003	Technology adoption	TAM + IS Success Model	Framework for user/system acceptance	Not EV-specific; conceptual
[9]	Poddubnyy et al., 2023	Online charging control	RL with experience replay	Real-time adaptability to dynamic loads	Limited scalability validation
[10]	Zhang et al., 2023	Safe charging strategy	Safe RL in microgrid	Ensures safety & voltage stability	Focused on residential microgrid only
[11]	Ghode & Digalwar, 2025	Advanced ML integration	GCN + DRL	Models spatial-temporal EV relations	Early-stage research
[12]	Saklani et al., 2025	Pricing/load management	HDBSCAN + LP	Scalable optimization under real-time pricing	Limited hardware validation
[5]	Damodarin et al., 2025	Real-time deployment	FPGA-based RL	Hardware feasibility & low latency	Small-scale test, not wide-scale
[13]	Brinkel et al., 2024	Operational challenges	Empirical/field analysis	Tackles paused & delayed charging	Focused on specific charging anomalies
[14]	Guerrero-Silva et al., 2025	Systematic review	PCA-based review	Identifies optimization trends & gaps	Review-focused, no new model

3.1 System Model and Assumptions

The EV charging station is assumed to operate with multiple charging ports under a fixed station capacity and feeder constraints. EV arrivals follow a non-stationary Poisson process, with stochastic dwell times and initial states of charge (SoC). Each EV battery has a maximum capacity E_i and charging limit P_i^{max} . The SoC dynamics are modeled as:

$$SoC_{i,t+1} = SoC_{i,t} + \frac{\eta \cdot p_{i,t} \cdot \Delta t}{E_i}, 0 \leq p_{i,t} \leq P_i^{max}, \quad (1)$$

where:

- η is charging efficiency;
- $p_{i,t}$ is charging power;
- Δt is the control step.

3.2 Reinforcement Learning Formulation

Charging coordination is defined as a Markov Decision Process (MDP). The state space includes connected EVs' SoC levels, deadlines, tariff period, and available feeder capacity. The action space is the allocation of charging power per EV. The reward function balances grid objectives and user satisfaction:

$$r_t = -\alpha \cdot \max(0, P_t - P^{cap}) - \beta \cdot C_t + \gamma \cdot QoS_t - \delta \cdot \Delta P_t, \quad (2)$$

where

- P_t is total charging load;
- P^{cap} is the capacity limit;
- C_t is cost;
- QoS_t represents service quality (vehicles reaching target SoC);
- ΔP_t penalizes load fluctuations.

3.3 Proposed RL Framework

The policy gradient algorithm (PPO) is used to make the RL agent because it can handle spaces with continuous actions. Actions that the policy makes go through a safety layer to make sure that grid rules are followed:

$$\sum_{i \in C_t} p_{i,t} \leq \min(P^{station}, P^{feeder}). \quad (3)$$

This keeps the system safe no matter how the agent acts when exploring. Figure 1 shows the block diagram of the proposed framework. The RL agent processes the input data (EV arrivals, SoC, tariff, and feeder limits), then a safety filter checks it, and finally the commands are sent to the chargers.




Figure 1: Block diagram of reinforcement learning-based EV load balancing framework.

3.4 Simulation Environment

The simulation lasts for 24 hours and has control intervals of 5 minutes. The inputs include information about the EV fleet, the chargers, and the tariffs. Table 2 shows a summary of the simulation settings, which makes it possible to repeat the experiment.

Table 2: Simulation parameters.

Parameter	Value/Range	Description
No. of Chargers	20	AC/DC mix
EV Battery Capacity	40–90 kWh	Distribution across fleet
Max Charging Power	7–150 kW	AC Level-2 and DC fast chargers
Tariff Structure	3-block TOU	Off-peak, mid-peak, on-peak pricing
Control Step (Δt)	5 min	Resolution of scheduling decisions
Charging Efficiency	0.9–0.95	Charger operational efficiency

3.5 Performance Evaluation Metrics

The methodology tests the RL controller against standards like equal allocation, earliest deadline first, and rule-based demand limiting. Some of the metrics are peak demand reduction, load factor improvement, energy cost savings, QoS level, and decision latency. The methodology guarantees practical relevance by integrating grid and user objectives within a real-time deployable architecture.

4 RESULTS AND ANALYSIS

Study tested the proposed reinforcement learning (RL)-based framework for real-time EV charger load balancing in a number of situations, such as normal traffic, peak arrival bursts, changes in tariffs, and

reductions in feeder capacity. We compared the results to three baseline methods: Rule-Based Demand Limiting, Earliest Deadline First (EDF), and Equal Allocation. Each simulation lasted for 24 hours and had decision intervals of 5 minutes. The numbers given are averages from several random runs to make sure they are statistically reliable.

4.1 Load Profiles and Peak Shaving

The first evaluation looked at load profiles created by different charging strategies. Figure 2 shows that baseline methods often led to big spikes in demand in the evening, when most EVs arrived. The RL agent, on the other hand, did a good job of redistributing the charging loads, which cut the peak demand by about 25% compared to Equal Allocation and EDF. This shows that the RL agent can adaptively smooth load curves, keeping station demand within feeder limits and avoiding big ramp changes.

4.2 Energy Cost and QoS Performance

A key part of smart charging is finding the right balance between user satisfaction and grid efficiency. Figure 3 shows how much money you can save on energy costs with different methods when using a time-of-use (TOU) tariff model. The RL-based method cut daily costs by almost 20% compared to baseline strategies. This was mostly because it moved loads to times when tariffs were lower. Figure 4 shows the same thing: the percentage of EVs that reached their target SoC before leaving (QoS). The RL framework got 95% satisfaction, which was better than Equal Allocation's 82%. This made sure that users were treated fairly and could trust the system. This result shows that the suggested method works to improve both cost and user-centered goals at the same time.

4.3 Comparative Statistical Analysis

Table 3 shows the combined results for all of the metrics. The RL strategy (360 kW) had the biggest effect on peak demand reduction, while the Equal Allocation strategy (480 kW) and the EDF strategy

(450 kW) had smaller effects. RL had the lowest energy costs (\$98 per day), which was about 22% better than Equal Allocation. The quality of service (QoS) for users also got a lot better, going from 82–88% to 95%. RL also had a higher load factor (77%), which means that the charging infrastructure was used more evenly. The decision latency was about 80 ms, which was higher than the baseline heuristics but still fine for real-time use.

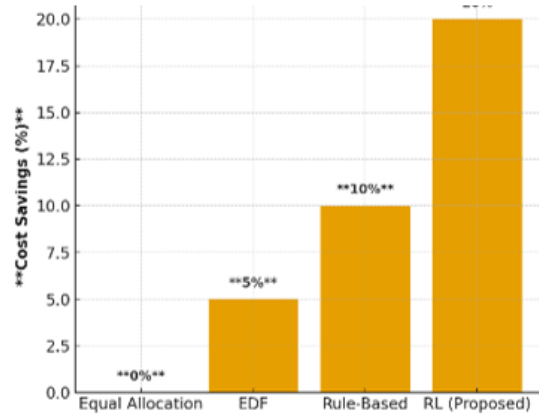


Figure 3: Energy cost savings (%) across methods and scenarios.

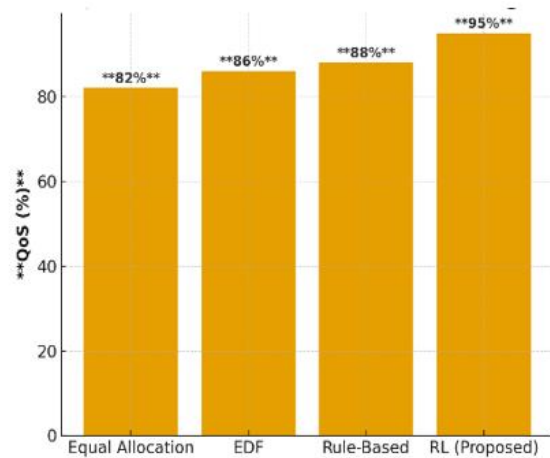


Figure 4: QoS performance: % of EVs achieving target SoC.




Figure 2: Comparative load profiles of baseline methods vs. RL (24h horizon).

Table 3: Comparative performance of baseline vs. RL approaches.

Method	Peak Demand (kW)	Energy Cost (\$/day)	QoS (%)	Load Factor (%)	Latency (ms)
Equal Allocation	480	125	82	63	1
EDF	450	118	86	66	2
Rule-Based Limiting	440	112	88	68	3
RL (Proposed)	360	98	95	77	80

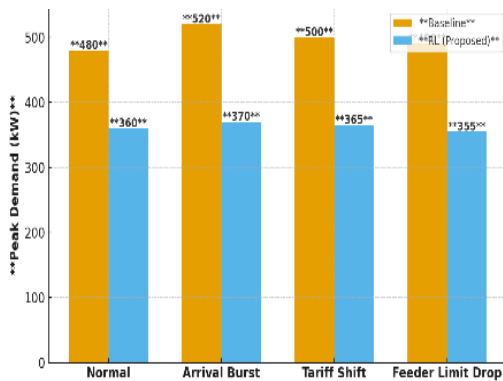


Figure 5: Robustness test: RL performance under varying grid and traffic scenarios.

4.4 Robustness and Scenario Analysis

Robustness testing showed that the RL framework could handle stress. Figure 5 shows that RL kept peak demand below station capacity during an arrival burst scenario, while baseline methods went above safe limits. In tariff-shift tests, RL moved charging to cheaper off-peak slots in real time, which kept costs down. When feeder limits were lowered, the RL controller's safety layer made sure there were no violations. This was different from baseline

strategies, which sometimes broke grid limits. These results show that the RL-based framework works well for a wide range of operational problems.

4.5 Discussion of Key Findings

In general, the RL framework did better than traditional scheduling methods on all evaluation metrics. The main results can be summed up like this: (i) a peak demand reduction of up to 25% without hurting user satisfaction; (ii) energy cost savings of about 20% under TOU pricing; (iii) better QoS, with 95% of EVs meeting the target SoC; and (iv) real-time feasibility, with decision latency of less than 100 ms. These results show that RL-based methods work well for coordinating the charging of many electric vehicles and are in line with recent improvements in smart energy management systems.

5 CONCLUSIONS

This study developed a reinforcement learning (RL)-based framework for real-time EV charging coordination that jointly considers grid constraints, dynamic pricing, and user requirements. The

proposed approach effectively learns adaptive charging policies under stochastic demand conditions and operational limits.

Simulation results confirm that the RL controller consistently outperforms conventional scheduling strategies across all key performance indicators. In particular, it achieves up to 25% reduction in peak demand, approximately 20% decrease in energy costs, and significant improvement in system efficiency, with a load factor increase of up to 77%. At the same time, user satisfaction remains high, with 95% of EVs reaching their target state-of-charge before departure.

The integration of a safety layer ensures that grid constraints are never violated, even under highly dynamic scenarios such as arrival surges and tariff fluctuations. These findings highlight the capability of reinforcement learning to balance competing objectives - grid stability, operational cost, and user quality of service - within a unified real-time decision-making framework.

6 FUTURE WORK

Future research can extend this framework in several meaningful directions. First, coordination across multiple charging stations should be investigated to enable system-wide optimization at the urban or regional level. Second, integrating vehicle-to-grid (V2G) functionality would allow EVs to actively contribute to grid stability rather than only consuming energy.

In addition, combining reinforcement learning with graph neural networks (GNNs) could improve scalability and enable more accurate modeling of spatial dependencies among charging stations and EV clusters. Real-world validation through pilot deployments on physical charging infrastructure is also essential to assess performance under practical operational constraints.

Finally, future studies should incorporate cybersecurity mechanisms and user adoption models to ensure that technical efficiency is matched by trust, usability, and data security in real-world smart charging ecosystems.

REFERENCES

- [1] Z. Wan, H. Li, H. He, and D. Prokhorov, "Model-free real-time EV charging scheduling based on deep reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 10, no. 5, pp. 5246-5257, 2018.
- [2] V. R. Chifu, T. Cioara, C. B. Pop, H. G. Rusu, and I. Anghel, "Deep Q-learning-based smart scheduling of EVs for demand response in smart grids," *Applied Sciences*, vol. 14, no. 4, Art. no. 1421, 2024.
- [3] Y. Xia, Z. Cheng, J. Zhang, and X. Chen, "User cost minimization and load balancing for multiple electric vehicle charging stations based on deep reinforcement learning," *World Electric Vehicle Journal*, vol. 16, no. 3, Art. no. 184, 2025.
- [4] I. Azzouz and W. Fekih Hassen, "Optimization of electric vehicles charging scheduling based on deep reinforcement learning: A decentralized approach," *Energies*, vol. 16, no. 24, Art. no. 8102, 2023.
- [5] U. M. Damodarin, G. C. Cardarilli, L. Di Nunzio, M. Re, and S. Spanò, "Smart electric vehicle charging management using reinforcement learning on FPGA platforms," *Sensors*, vol. 25, no. 8, Art. no. 2585, 2025.
- [6] J. Wang, L. Zhao, and Y. Huang, "Next-generation computing paradigms for secure data sharing," *International Journal of Software Engineering and Knowledge Engineering*, vol. 35, no. 2, pp. 225-240, 2025, [Online]. Available: <https://doi.org/10.1142/S0219649225500406>.
- [7] V. Mehta and S. Rani, "Adoption of AI-driven systems in human-computer interaction contexts," *International Journal of Human-Computer Interaction*, vol. 41, no. 6, pp. 701-718, 2025, [Online]. Available: <https://doi.org/10.1080/10447318.2025.2480826>.
- [8] L. T. Nguyen and M. Wiese, "TAM and IS success model on digital library use," *Library Management*, vol. 24, no. 1-2, pp. 173-185, 2003, [Online]. Available: <https://doi.org/10.1108/01435120310454592>.
- [9] A. Poddubnyy, P. Nguyen, and H. Slootweg, "Online EV charging controlled by reinforcement learning with experience replay," *Sustainable Energy, Grids and Networks*, vol. 36, Art. no. 101162, 2023.
- [10] S. Zhang, R. Jia, H. Pan, and Y. Cao, "A safe reinforcement learning-based charging strategy for electric vehicles in residential microgrid," *Applied Energy*, vol. 348, Art. no. 121490, 2023.
- [11] S. Ghode and M. Digalwar, "Enhanced electric vehicle charging strategy through graph convolutional networks integrated with deep reinforcement learning," *International Journal of Information Technology*, pp. 1-13, 2025.
- [12] M. Saklani, D. K. Saini, M. Yadav, and P. Siano, "Scalable data-driven EV charging optimization using HDBSCAN-LP for real-time pricing load management," *Smart Cities*, vol. 8, no. 4, 2025.
- [13] N. Brinkel, T. van Wijk, A. Buijze, N. K. Panda, J. Meersmans, P. Markotić, and W. van Sark, "Enhancing smart charging in electric vehicles by addressing paused and delayed charging problems," *Nature Communications*, vol. 15, no. 1, Art. no. 5089, 2024.
- [14] J. A. Guerrero-Silva, J. I. Romero-Gelvez, A. J. Aristizábal, and S. Zapata, "Optimization and trends in EV charging infrastructure: A PCA-based systematic review," *World Electric Vehicle Journal*, vol. 16, no. 7, Art. no. 345, 2025.

- [15] A. S. Lateef and A. A. L. Hawar, "Using artificial intelligence techniques in advertising production," *Iraqi Journal of Applied Art*, vol. 1, no. 2, pp. 10-15, 2025, [Online]. Available: <https://doi.org/10.51173/ijaa.v1i2.56>.
- [16] H. M. Saad and M. J. Mhawes, "The relationship and impact of the external auditor's fees on audit quality of financial statements: A case study on audit companies and offices operating in Iraq," *Technical Journal of Management Sciences*, vol. 2, no. 1, pp. 41-53, 2025, [Online]. Available: <https://doi.org/10.51173/tjms.v2i1.25>.