

READ THIS FIRST

What this file is. A sample project proposal (synopsis) for the IGNOU M.Sc. (Artificial Intelligence and Machine Learning) project course **MCSP-090 Project Work** (12 credits, Semester IV), written by ignouproject.com as study reference. It shows one way to lay out a proposal for one fresh topic. It is not approved by IGNOU and it is not a synopsis you can submit.

What it is built on. The MCSP-090 description in the MSCAIML Programme Guide (School of Computer and Information Sciences, July 2026, version 1). The guide lists eight things a learner should be able to do after the project: identify and define a real-world AI/ML problem; conduct a structured literature survey; collect, preprocess and analyse datasets; apply suitable Machine Learning or Deep Learning algorithms; evaluate model performance using appropriate metrics; interpret results and suggest improvements; prepare a professional technical report; and present and defend the project in the viva voce. Each section below is tied to one of those eight outcomes.

What is not yet official. As of 06 October 2026 IGNOU has not published separate MCSP-090 project guidelines or an MCSP-090 proforma. The guide says they will be provided on the University website. Section headings, page length, spacing and enclosures in this sample are therefore not IGNOU rules. When IGNOU issues the MCSP-090 guidelines, follow them over this sample. Do not use the MSCDSA project (MCSP-079) or the MCA project (MCSP-232) rules as if they were MSCAIML rules.

The [brackets]. Every junction setting, arrival rate, parameter value, date, name and result is left in [brackets]. Those are your own decisions and your own numbers. A result can only come from your own simulation runs.

Why this topic. It uses MCS-089 Reinforcement Learning, a course in the MSCAIML programme. The data is generated by your own simulation, so no personal data, survey or third-party dataset is involved. Choose your own topic with your guide; do not reuse this one.

Academic integrity. IGNOU expects original learner work. Copying this sample, in full or in part, and sending it as your synopsis is not acceptable. The learner remains the author and IGNOU's evaluator decides approval.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

PROJECT PROPOSAL (SYNOPSIS)

MCSP-090: Project Work

Comparing Tabular Q-Learning and a Deep Q-Network for Adaptive Traffic Signal Timing at a Simulated Four-Way Junction [modelled on a junction type in your city]

Submitted by

[Your full name]

Enrolment No.: []

Under the guidance of

[Guide's full name, designation, organisation]

Learner Support Centre: [name and code] Regional Centre: [name and code]

M.Sc. (Artificial Intelligence and Machine Learning) (MSCAIML)

School of Computer and Information Sciences

Indira Gandhi National Open University, New Delhi

[Month Year]

No IGNOU logo is used in this sample. Use the cover layout IGNOU issues for MCSP-090.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

CONTENTS

Section	Title	Guide outcome it evidences
1	Title of the project	1. Identify and define a real-world AI/ML problem
2	Introduction and problem definition	1
3	Objectives	1
4	Literature survey	2. Structured literature survey
5	Data: the simulated environment	3. Collect, preprocess and analyse datasets
6	Methodology	4. Apply suitable ML or DL algorithms
7	Evaluation metrics and comparison plan	5. Evaluate model performance using appropriate metrics
8	Interpretation and improvement plan	6. Interpret results and suggest improvements
9	Tools, platform, hardware and software	4, 7
10	Ethics and responsible AI	6
11	Work plan	7
12	Expected deliverables and report structure	7. Prepare a professional technical report
13	Limitations and scope	1, 6
14	Viva readiness	8. Present and defend the project in the viva voce
15	References	2

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

1. Title of the Project

Comparing Tabular Q-Learning and a Deep Q-Network for Adaptive Traffic Signal Timing at a Simulated Four-Way Junction [modelled on a junction type in your city].

How the title is built. Task (adaptive signal timing), method family (tabular Q-learning versus Deep Q-Network), environment (a simulated four-way junction) and the setting the student will model. Keep the approved title word for word in the final report.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

2. Introduction and Problem Definition

2.1 Context. Many signalised junctions run fixed-time plans, in which each phase receives a set green time regardless of the queue present at that moment. A classical way to set such plans is Webster's method (Webster, 1958), which derives a cycle length from lost time and flow ratios. A fixed plan cannot react when arrivals change within the hour.

2.2 Why reinforcement learning. Signal control is a sequential decision problem: each choice to hold or change a phase affects the queues the controller will face next. Reinforcement learning (RL) learns a policy from interaction with an environment by maximising cumulative reward (Sutton and Barto, 2018), which suits this structure. RL methods for signal control are reviewed by Wei et al. (2019).

2.3 Problem statement. This project will examine how a tabular Q-learning agent and a Deep Q-Network (DQN) agent compare with a fixed-time baseline, in terms of vehicle waiting time, queue length and throughput, when controlling an isolated four-way junction under [number] simulated demand scenarios.

2.4 Why the problem matters. [In two or three sentences, state why junction delay matters in the setting you model. Cite a source you have read for any number you give. Do not invent figures.]

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

3. Objectives

The objectives are worded neutrally, without assuming which controller will perform better.

- To formulate traffic signal control at an isolated four-way junction as a Markov Decision Process, defining state, action and reward.
- To implement a fixed-time controller whose cycle length is set by Webster's method, as the baseline.
- To implement a tabular Q-learning agent using a discretised state.
- To implement a Deep Q-Network agent with experience replay and a target network.
- To compare the three controllers on average waiting time, average queue length and throughput across [number] demand scenarios and [number] random seeds.
- To analyse the sensitivity of the learning agents to [learning rate, discount factor, exploration schedule] and to identify possible improvements.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

4. Literature Survey

The survey is organised by theme so that the gap this project addresses is visible. Each row records what the source contributes to this project. Add the studies you read yourself, with full references.

Theme	Source	What it contributes here
Fixed-time baseline	Webster (1958)	Method for setting cycle length from lost time and flow ratios; used to build the baseline controller

Theme	Source	What it contributes here
RL foundations	Watkins and Dayan (1992); Sutton and Barto (2018)	Q-learning update rule; MDP formulation, exploration and evaluation of policies
Deep RL	Mnih et al. (2015)	Deep Q-Network with experience replay and a target network
RL for signal control	Wei et al. (2019)	Survey of transportation approaches and RL literature for traffic signal control; source for state, action and reward choices to compare
Simulation	Lopez et al. (2018)	SUMO, an open-source microscopic traffic simulator that can be controlled from Python
[Your reading]	[Author, year]	[What it adds and how it shapes your design]

4.1 Gap. [After reading, state in two or three sentences what the sources leave open for a single isolated junction of the type you model, for example the comparison of a tabular agent with a deep agent under the same reward and demand settings. Base this only on what you have read.]

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

5. Data: The Simulated Environment

5.1 Source of data. All data will be generated by the student's own simulation runs. No personal data, vehicle registration numbers or images of people will be collected. The environment will be built in [SUMO, version ()], controlled through its Python interface / a self-written grid simulator in Python].

5.2 Junction model.

Element	Setting
Approaches and lanes	[4 approaches; () lanes each]
Phases	[() phases, for example north-south and east-west through movements]
Minimum green, yellow, all-red	[] s, [] s, [] s
Episode length	[] simulated seconds
Vehicle arrivals	[Arrival process and rate per approach, in vehicles per hour, for each scenario]

Element	Setting
Demand scenarios	[Low / medium / high, or a time-varying profile]
Random seeds	[] seeds per scenario

5.3 Preprocessing. Queue lengths and waiting times read from the simulator will be [normalised / discretised into () bins] before being passed to the agents. Logs will be stored per episode in [CSV / Parquet] for analysis with Pandas.

5.4 Exploratory analysis. Before training, the student will plot [arrival counts per approach and queue length over time] under the baseline controller to check that each scenario behaves as intended.

If you calibrate arrival rates from a real count, state who collected it, when, with what permission, and that no personal data was recorded.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

6. Methodology

6.1 MDP formulation.

Component	Definition in this project
State	[Queue length per approach, discretised into () bins for Q-learning and normalised for the DQN] plus the current phase and [time since last change]
Action	Keep the current phase or switch to the next phase, subject to the minimum green and yellow times
Reward	[Negative change in cumulative waiting time between decisions / negative total queue length]; state your choice and why
Discount factor	[]
Decision interval	[] s

6.2 Baseline. A fixed-time plan with cycle length $C_0 = (1.5L + 5) / (1 - Y)$, where L is the total lost time per cycle in seconds and Y is the sum of critical flow ratios (Webster, 1958). Green time will be split in proportion to the critical flow ratios.

6.3 Tabular Q-learning. The action-value table is updated as $Q(s, a) \leftarrow Q(s, a) + \alpha [r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$ (Watkins and Dayan, 1992). Exploration will use an

ϵ -greedy policy with ϵ decayed from [] to [] over [] episodes.

6.4 Deep Q-Network. A fully connected network of [() hidden layers of () units, ReLU activation] will approximate $Q(s, a)$. Training will use an experience replay buffer of [] transitions, mini-batches of [], a target network updated every [] steps (Mnih et al., 2015), the [Adam] optimiser and learning rate [].

6.5 Where the MSCAIML courses support each step.

Step	MSCAIML course and unit (Programme Guide, July 2026)
MDP, policies, Bellman equation	MCS-089 Unit 2
Temporal-difference learning	MCS-089 Unit 6
Deep Q learning	MCS-089 Unit 7; MCS-072 Unit 8
Python implementation	MCS-082; MCS-089 Unit 10 (RL in Python for tic-tac-toe)
Neural network training	MCS-072 Units 1 and 4; MCSL-076 Deep Learning Lab
Hypothesis testing of differences	MCS-066 Unit 10

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

7. Evaluation Metrics and Comparison Plan

Metric	Unit	Why it is used
Average waiting time per vehicle	seconds	Direct measure of delay at the junction
Average queue length	vehicles	Measures congestion over the episode
Throughput	vehicles per hour	Checks that lower delay is not bought by serving fewer vehicles
Mean episode return	reward units	Shows learning progress of each agent
Variation across seeds	standard deviation	Shows how stable each controller is

Each controller will be run on the same scenarios with the same [] seeds. Results will be reported as mean and standard deviation. Differences between controllers will be tested with [a test agreed with the guide, for example Welch's t-test] at the [] significance level. Training and evaluation episodes will be kept separate.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

8. Interpretation and Improvement Plan

Results will be read scenario by scenario, because a controller that helps under one demand level may not help under another. The student will examine learning curves, the actions chosen in [two or three] typical traffic states, and the cases where each agent does worse than the baseline.

Possible improvements to test if time permits: [a different reward definition / a richer state / a policy-based method from MCS-089 Unit 8]. Any improvement will be compared under the same scenarios and seeds as the original agents.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

9. Tools, Platform, Hardware and Software

Item	Planned choice
Language	Python [version]
Libraries	NumPy, Pandas, Matplotlib, [PyTorch / TensorFlow] [versions]
Simulator	[SUMO version () / self-written simulator]
Hardware	[CPU, RAM, GPU if any, or a cloud notebook]
Version control	[Git repository, private until submission]
Reproducibility	Fixed random seeds; requirements file listing exact library versions

MCS-081 Unit 11 introduces NumPy, Pandas, Matplotlib, scikit-learn, PyTorch and TensorFlow for AI programming.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

10. Ethics and Responsible AI

The project uses simulated traffic only and collects no personal data. Results from a simulator do not show that a controller is safe to deploy on a real road; the report will say so. Any real count used for calibration will be described with its source and permission. The discussion will draw on MCS-072 Unit 12 (ethical and societal implications, including transparency and accountability).

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

11. Work Plan

Stage	Weeks	From	To
Literature survey and environment design	[]	[]	[]
Simulator build and baseline	[]	[]	[]
Tabular Q-learning agent	[]	[]	[]
DQN agent	[]	[]	[]
Experiments across scenarios and seeds	[]	[]	[]
Analysis and improvement trials	[]	[]	[]
Report writing and guide review	[]	[]	[]
Viva preparation	[]	[]	[]

Fill in dates once IGNOU publishes the MCSP-090 submission schedule. The guide schedules 4 project counselling sessions (8 hours) in Semester IV.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

12. Expected Deliverables and Report Structure

The project will produce: the simulation environment and agent code; trained model files; experiment logs; and a technical report. The proposed report chapters are problem definition, literature survey, environment and data, methodology, experimental setup, results and evaluation, interpretation and improvements, ethics, and conclusion, followed by references and appendices. No outcome is assumed in

advance.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

13. Limitations and Scope

- A single isolated junction; coordination between neighbouring junctions is outside scope.
- Simulated vehicles follow the simulator's driver model, which may differ from local driving behaviour.
- [Pedestrian phases / emergency vehicles] are [not modelled / modelled as ()].
- Computing time limits the number of training episodes to about [].

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

14. Viva Readiness

The guide's eighth outcome is to present and defend the project during the viva voce. The student will prepare: a short live run of each controller on one scenario; an explanation of the state, action and reward choices; the Q-learning update and the role of the replay buffer and target network; and the reasons behind each metric and statistical test.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.

15. References

Bilgin, E. (2020). *Mastering Reinforcement Learning with Python*. Packt Publishing.

Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.

Indira Gandhi National Open University, School of Computer and Information Sciences (2026). *Programme Guide for Master of Science (M.Sc.) (Artificial Intelligence and Machine Learning) (MSCAIML)*, July 2026, version 1. New Delhi: IGNOU.

Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., and Wießner, E. (2018). Microscopic traffic simulation using SUMO. *21st International Conference on Intelligent Transportation Systems (ITSC)*, IEEE, pp. 2575-2582. doi:10.1109/ITSC.2018.8569938.

Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518, 529-533.

Sutton, R. S., and Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.

Watkins, C. J. C. H., and Dayan, P. (1992). Q-learning. *Machine Learning*, 8, 279-292.

Webster, F. V. (1958). *Traffic Signal Settings*. Road Research Technical Paper No. 39. London: HMSO.

Wei, H., Zheng, G., Gayah, V., and Li, Z. (2019). A survey on traffic signal control methods. arXiv:1904.08117.

[Add every source you actually read and cite.]

What travels with a real proposal (proforma, guide bio-data, signatures) will be stated in the MCSP-090 guidelines when IGNOU issues them. Until then, check with your Regional Centre.

ignouproject.com

Free study reference by ignouproject.com. Not for submission. Reselling or re-uploading this file is not permitted.