DOPE CARS: Deciding Optimally How to Efficiently Charge Automotives with Renewable Energy Sources

Yash Chandramouli, Arec Jamgochian, Emily Jewell
Aeronautics/Astronautics Department, Stanford University, Stanford, CA, USA
{yashc3, arec, ejewell}@stanford.edu

Abstract

Predicted trends show a growth in the number of electric vehicles and a rise in the use of renewable energy in the coming decades. While the environmental and societal benefits of these technologies will be large, they also bring with them numerous challenges. In particular, the temporal volatility of renewable energy means that managing and timing the flow of energy to systems with high energy consumption, such as electric vehicles, will be of great importance. It is therefore necessary to create a framework for decision-making that will allow future systems to optimize charging schedules. We formulate the problem of charging electric vehicles with a time-dependent energy source as a Markov Decision Process (MDP), with states defining the presence of cars, their individual levels of charge, and the level of available renewable energy. Using Value Iteration over this state space and a 7-hour timespan, we were able to generate optimal policies that balanced the energy toll on the electric grid with the final charge levels of each vehicle. We hope that future work can build upon this foundation to incorporate more aspects of the real-world problem and design policies that will be able to tackle this major issue of the next few decades.

Background and Motivation

As transportation evolves for the 21st century, it must meet the needs of modern society by being efficient for personal travel while simultaneously utilizing environmentally conscious and sustainable methods. Electrical vehicles (EV) are one solution that perfectly satisfies these requirements. EVs have already been successfully introduced globally and are projected to account for up to 40% of auto sales by 2030 [1]. This large market penetration is due to EV's abilities to reduce global dependence on oil, an environmentally conscious public (over 50% of American households support the idea of using an EV in their household), and likely continued government incentives to purchase EVs [2]. Ultimately, the positive environmental impact of EVs can be most directly seen by eliminating private transportation's 95% energy dependence on oil, which in turn accounts for 50% of the world's oil dependence [3].

Hybrid electric vehicles (HEV), plug-in hybrid electric vehicles (PHEV) and battery electric vehicles (BEV) are examples of EVs that can revolutionize personal transportation. Not only for their reliance on energy from the grid over combustion, but also for their improved efficiency rates. A traditional internal combustion engine vehicles (ICEV) efficiency is approximately 15% while an EV can have up to 70% efficiency [3].

To further maximize EV's positive environmental impacts, the use of Vehicle-to-Grid (V2G) and Smart Charging have been suggested [4, 5]. V2G allows for energy from the grid to be stored in an EV when they are parked and plugged in, transforming the EV to a critical fuel storage device for the electric grid. Since most EV are plugged in overnight, V2G can allow for a greater use of renewable energies for powering the grid, such as wind and solar wherein the current bottleneck for using more renewable energy sources is the inability to store such energy efficiently. Denmark has extensively researched this technology for implementation with wind energy, with promising results [5]. However, scholars agree that Smart Charging will bring the necessary changes to EVs and the grid so that EV can become a sustainable future technology [5]. Smart Charging simply requires some form of feedback from the EV so that it only is pulling energy from the grid during off-peak times, and only to the extent that the grid can sustainably handle [1, 4]. Therefore, since both of these ideas rely on a knowledgeable car charging energy distribution cycle, a problem wrapped in uncertainties, it becomes the basis for this research. This project model's the uncertainty of ideal energy distribution for overnight charging of electric vehicles so as to maximize the positive environmental impact of an EV over an ICEV.

In order to solve this decision-making problem, it is useful to model the scenario as a Markov Decision Process (MDP). An MDP is a type of sequential decision process whereby an agent can choose from a set of actions at each timestep after observing their state at the current time. It is fundamentally based on the Markov assumption, which is the assumption that the next state only depends on the current state and the action chosen by the agent. There are a variety of solution methods for MDPs; however, the one that shall be implemented here is known as value iteration. This is a method that approximates the optimal value function iteratively and then extracts the optimal policy from the determined value function [6].

Methodology

Problem Definition

Since deep market penetration of EVs will impose substantial current demands on an already fragile electricity grid, an optimal policy is sought to schedule smooth charging of electric vehicles overnight and minimize the need for non-renewable electricity sources. This is done by structuring the problem as an MDP and performing Value Iteration using the POMDPs.jl Julia library [7].

We scale our MDP to the level of a suburban street of homes. In this case, we model n homes on the electric grid, each with an EV charging port capable of charging one car. It is assumed that each car follows a unique driving route during the day and arrives back to the home for charging at night at a variable time and with variable amount of current charge c.

All vehicles draw energy from the same grid, which can have a renewable energy level up to R. Charging a vehicle will increment its charge level up towards C, the maximum charge level. However charging too many cars at once will reduce the energy level. An ideal policy for this model would discern which cars to charge at every time step, with the goal that by the end of the night (when $t = T$), all cars would be fully or nearly fully charged while keeping the grid's renewable energy mixture level high.

States

The states are defined by a four component tuple. First, we include a boolean vector of length n describing whether there is a car in each port. Second, we include a vector providing the charge level for the vehicle at that indexed location. Each car can take on discrete charge levels in the range of [0,C]. If there is no car present, the charge level reported is zero. Third, to model car arrival probabilities and renewable level changes as a function of time, we include the current time step t up to the maximum simulated value T. Last, we include the current renewable energy level, which can take discrete values in [0,R]. The state space therefore has $|S| = 2^n(C+1)^nT(R+1)$ possible states.

The initial state is defined with all cars being away, each car having zero charge (though this may change once they arrive), time t of 1, and a maximum renewable energy level.

Actions

The action to take at each time step is whether to charge each car. This is represented as a boolean array of length n. The total number of possible actions is therefore $|A| = 2^n$.

Transitions

The transition model mandates that time increments by one, or $t_t = t_{t-1} + 1$. Additionally, if a vehicle is present, it will remain until the end of the simulation. If a vehicle is not present, the probability of it arriving at the next time step is

$$P_{n,t} = \frac{1}{1 + \exp(-20\frac{t-1}{T})}$$

When a car first arrives, it's initial charge may take any value between 0 and C-1 with equal probability. If a car was present and the action to charge that car is taken, the level is incremented by one, up to a maximum charge potential C. The renewable energy level is updated based on the number of actions taken and a function modeling how the renewable energy source levels change naturally throughout the day. The update equation follows:

$$r_{t+1} = r_t + naturalAddition(t) - \frac{\sum_n(actions^1)}{n}$$

which will be rounded to the nearest discrete value in [0, R]. With each update, we decrement our renewable level if we charge more than half of all cars. The naturalAddition(t) function allows us to add a level of complexity by modeling how a renewable energy changes over time due to external causes. In our default case, we assume there is no naturalAddition, but we also explore how our policy changes when we allow naturalAddition(t) to add a level per time-step for the first half of the simulation.

Rewards

The reward function is modeled using the current renewable energy level *r*, charge amounts in each car c, and time *t*, by

$$R_t(s, a) = \lambda r - 1(t = T) \sum_{i=1:n} \exp((C - c_i)/C)$$

The first part of the equation gives reward at each time-step based on the current renewable energy level r. This inherently penalizes for dropping the renewable energy level. The second part of the equation gives penalties based on the level of charge in each car at the terminal state. Remaining charge is exponentiated to more heavily penalize cars with less charge. The relative weighting of the two components is dictated by λ. In our finite-horizon MDP, future rewards are not discounted.

Results and Discussion

We tested our implemented solver with a 14,336-state MDP, where *n* is 3 cars, *C* is a maximum charge level of 3, T is 7 timesteps, and R is a maximum renewable energy level of 3. Using the POMDPs.jl package Discrete Value Iteration solver, this we can solve for the exactly optimal policy in 30 seconds. Below we show plots of the same simulation run with policies solved with λ= 0.01 vs λ= 50.0. The color of each block represents the charge of that car at that time step, with black indicating a car is yet to arrive.

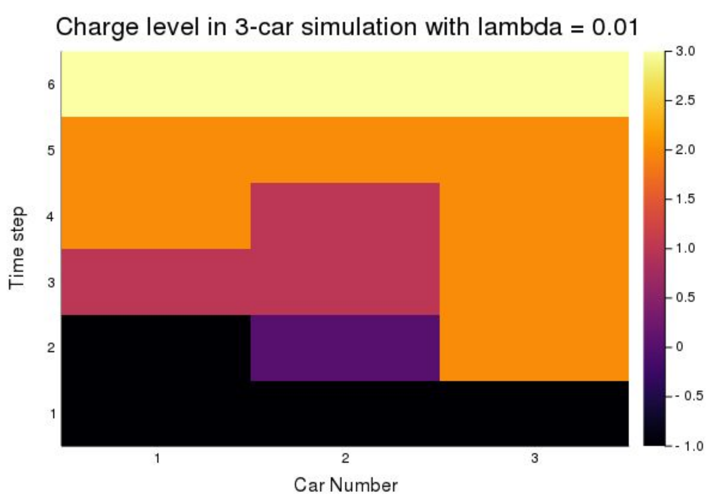

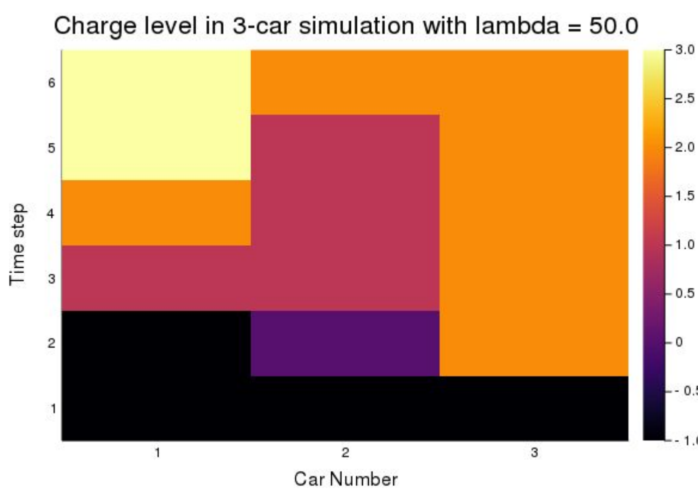

As we can see, the λ= 0.01 policy cares much more that every car is charged by the end of the simulation. It doesn't mind taking the action to charge all cars at the final timestep, even though doing so decrements the renewable energy level (not shown). It does however try to incrementally charge the cars for most of the time duration in order to keep the renewable energy level high.

The λ= 50.0 policy on the other hand focuses more on keeping the renewable energy level high throughout the simulation. It therefore charges only one car at a time, even if in doing so, the car will not be fully charged at the end of the simulation. It does make the decision, however, to have one fully charged and two cars mostly charged rather than two cars fully charged and one car partially charged. This is because our reward function exponentially penalizes lack of charge, meaning that one car having a small amount of charge is much more costly than having two cars with a moderate amount of charge. This mirrors what we would wish to see in a real world implementation of such a tool, since making the decision to not charge one person's car and fully charge another's could lead to ethical issues and consumer distrust of the system if there is no reasoning to determine which car is charged less.

Next, we observe what happens when we model additional renewable energy being added to the system as it is collected from available sunlight . With the simulation shown on the left, we model the naturalAddition(t) function to add a level of renewable energy for each of the first T/2 time steps. This would correspond to starting the simulation when there is daylight still available. One may notice that even though we use λ= 50.0, as we did in the top right model, our policy still informs us to take the greedy action, and charge every car it can during timestep 3, since there would be no reduction in renewable energy level in the first few timesteps. As the number of timesteps in "daylight" (i.e. with renewables added each timestep) increases, one would expect the decision to become more greedy. A run with all 6 timesteps in sunlight yielded a similarly greedy result, indicating this hypothesis to be true. This implies that in a real world implementation, one could potentially include the time of charging as a state variable as well.

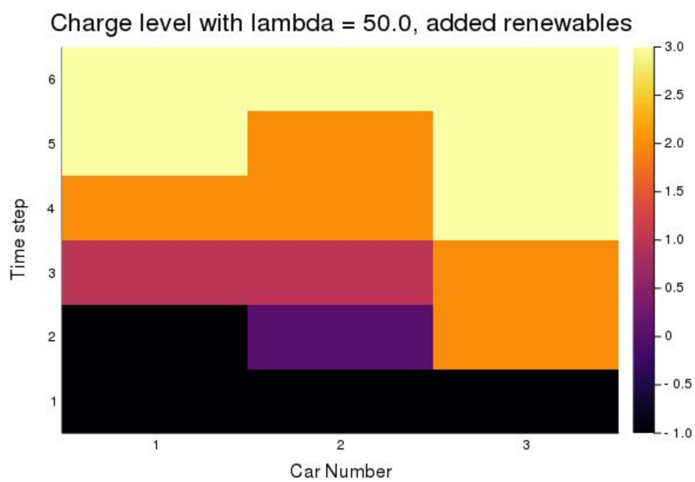

One point of note is that the discretization of the state space currently has a large effect on the predicted policy. Currently there will not be a reduction in the energy level until more than 2 cars are charged simultaneously. This means that early on the tradeoffs for charging only one or two cars are not as severe as they would be in a state space where the renewable energy level was closer to a continuous variable. See Future Work below for discussions on how to increase this state space in future additions to this work.

Conclusions and Future Work

Conclusions

Overall, we were able to create a model of the electric vehicle charging process and compute a non-trivial, optimal policy with respect to the number of cars that need to be charged, the length of time available for charging, and different renewable energy level dynamics. Our model can be tuned between prioritizing the final level of charge in each car and the renewable energy level throughout charging using a free parameter λ. When we added external sources of energy, as expected, the model chose to greedily charge the cars when renewable energy levels were high, and become more conservative as these levels dropped. While the model required various assumptions, primarily regarding the discretization of the state space, the groundwork now exists for future work to incorporate more aspects of this large-scale and complex real world problem.

Limitations and Future Work

While Value Iteration converges quickly on an optimal policy when we have three cars, we note that our approach is not very scalable. Every time we add a car, we multiply the size of our action-space by 2 and our state-space by $2(C+1)$. Since the algorithmic complexity of each iteration of Value Iteration is $O(|A|\ |S|^2)$, every time we add a car, we multiply our runtime by $8(C+1)^2$ per iteration. To scale this problem, we will therefore have to change our approach.

There are a variety of options available to further simplify this problem. One approach would be to utilize a local or global approximation to find an approximate solution for a larger state space. This would potentially allow for a larger set of timesteps and a more "continuous-esque" state space but sacrifices the guarantee of global optimality. Another option would be to use an online method. This would only consider the states that are reachable from the current state and would therefore limit the computational power and storage required for computation by again trading off the certainty of optimality.

Lastly, another approach would be to take advantage of the inherent structure within the problem and implement a form of hierarchical planning. One could abstract the policy for a single car and use that division of policies to compute the optimal policy over a set of N cars. In scaling this method up further, one could further abstract the policy for a "street" of cars and scale it up to a whole neighborhood. Without some form of hierarchical planning, the state space for such a problem would otherwise be intractably large.

Future work, in addition to implementing the above suggestions to further scale up the model, could also utilize other methodologies to capture more of the real-world dynamics and constraints inherent in the problem. For example, cars do not all arrive at the same time, so a future model could utilize a structure more similar to a POMDP to determine an optimal policy despite uncertainty as to each car's arrival time. This could also implement a form of machine learning to determine approximate arrival and departure times for each car over time.

Another option for future work could be to implement a more complete model of available power. The current model relies on estimates of energy grid availability. However, discretizing the energy state space further would allow for greater exploration of the effects of the time-dependency of renewable energy on the computed policy.

Contributions

Each group member contributed approximately evenly. Emily wrote a lot of the final report, as well as the state/action indexing, reward, and simulation functions for the MDP. Arec wrote a lot of the MDP structure, solver, and visualization code, and part of the final report. Yash wrote the MDP transition function, as well as a part of the final report.

The code is made available at https://github.com/jamgochiana/evMDP

References

[1] Mwasilu F., Justo, J. J., Kim, E.-K., Do, T. D., and Jung, J.-W., "Electric vehicles and smart grid interaction: A review on vehicle to grid and renewable energy sources integration," *Renewable and Sustainable Energy Reviews*, vol. 34, pp. 501–516, Jun. 2014.

[2] Hadley, S. W. and Tsvetkova, A. A., "Potential Impacts of Plug-in Hybrid Electric Vehicles on Regional Power Generation," *The Electricity Journal*, vol. 22, no. 10, pp. 56–68, Dec. 2009.

[3] Andersen, P. H., Mathews, J. A., and Rask, M, "Integrating private transport into renewable energy policy: The strategy of creating intelligent recharging grids for electric vehicles," *Energy Policy*, vol. 37, no. 7, pp. 2481–2486, Jul. 2009.

[4] Lopes, J. A. P., Soares, F. J., and Almeida, P. M. R., "Integration of Electric Vehicles in the Electric Power System," *Proceedings of the IEEE*, vol. 99, no. 1, pp. 168–183, Jan. 2011.

[5] Richardson, D.B., "Electric vehicles and the electric grid: A review of modeling approaches, Impacts, and renewable energy integration," *Renewable and Sustainable Energy Reviews*, vol. 19, pp. 247–254, Mar. 2013.

[6] Kochenderfer, M. J. (2015). Decision Making Under Uncertainty: Theory and Application. MIT Press. pp. 77-82

[7] Egorov, M., Sunberg, Z. N., Balaban, E., Wheeler, T. A., Gupta, J. K., & Kochenderfer, M. J. (2017). POMDPs. jl: A framework for sequential decision making under uncertainty. Journal of Machine Learning Research, 18(26), 1-5.