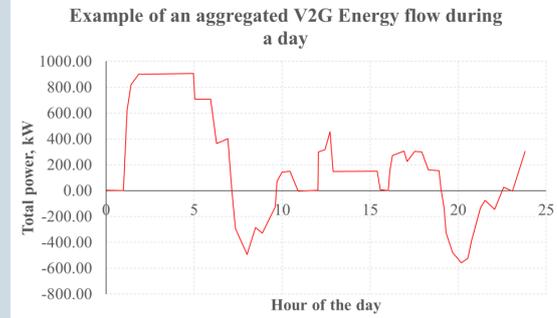


Simulation demonstrates up to a 20-30% improvement compared to conventional unmanaged charging strategies.

The aggregate model reduces peak demand stress by preventing simultaneous high-power charging events and minimizes charging costs by leveraging off-peak tariffs and vehicle-to-grid (V2G) capabilities.



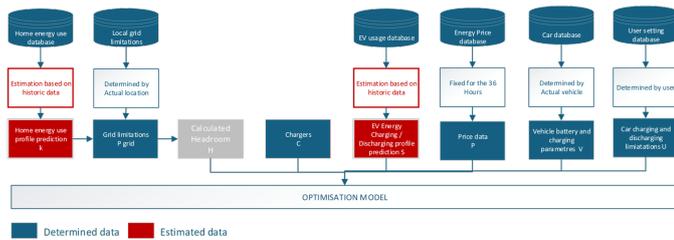
Introduction

The rapid electrification of transportation necessitates the development of intelligent energy management solutions to ensure the efficient operation of electric vehicle (EV) fleets while maintaining grid stability.

Traditional energy balancing methods primarily target large-scale power systems, leaving small-scale energy storage aggregators, such as distributed EV charging hubs, without effective optimisation frameworks.

Methodology

Our model rests on a unified notation that pulls together every signal the aggregator can observe:



Household base load.

5 000 smart-meter profiles are split into fixed-price and market-rate customers. For each tariff we fit a generalised additive model $\mu_g(t)$ (hour, weekday, season) plus a price-elasticity term $\beta(p_t - \bar{p})$ for market users. Drawing Gaussian noise with hour-dependent variance and scaling by household counts N_g yields the stochastic head-room series h_t

Vehicle trips.

15-min charger logs from 10 real EVs are segmented into energy sessions, converted to daily driving loads $E_{i,d}$ and distilled into weekday-model Gamma laws

$$\Gamma(k_{m,w}, \theta_{m,w})$$

Methodology (cont.)

Monte-Carlo sampling those laws—rescaled by each synthetic car's efficiency—produces a full-year trip set S_i and a closed-form next-day forecast

$$\tilde{E}_{i,d+1} = k_{m(i),w_{d+1}} \theta_{m(i),w_{d+1}}$$

Car fleet aggregation

The 278-car fleet composition follows Latvian 2024 market shares; 1 MW of aggregate charger power is guaranteed.

Place	Model	Share of EVs	Battery Size E_i^{cap} , kWh	Energy consumption η_i , kWh/km	Share in model s_m	Calculated vehicles in model n_m
1	TESLA MODEL 3	9.2%	57.5	0.139	23%	64
2	TESLA MODEL Y	5.1%	57.5	0.164	13%	36
3	NISSAN LEAF 40KWH	4.9%	95	0.165	12%	34
4	BMW I13	4.0%	37.9	0.161	10%	28
5	TESLA MODEL S	3.8%	95	0.165	10%	28
6	HYUNDAI KONA	3.3%	48	0.164	8%	23
7	NISSAN LEAF 62KWH	3.1%	39	0.166	8%	23
8	VW ID.3 PRO 150 KW	2.1%	58	0.166	5%	14
9	AUDI E-TRON 55	2.0%	86.5	0.237	5%	14
10	VW ID.4 PRO 150 KW	1.9%	77	0.177	5%	14

Optimisation layer

For every horizon $T=24T=24T=24$ h we solve.

$$\min_{P_{i,t}^{ch}} \sum_{t \in T} P_t \left(\sum_i P_{i,t}^{ch} \right) s.t. \begin{cases} \sum_i P_{i,t}^{ch} \leq h_t & \forall t \\ E_i^{SoC}(T) \geq E_{e,d+1} & \forall i \\ 0 \leq P_{i,t}^{ch} \leq P_{j,i}^{max} & \forall i, t \end{cases}$$

where E_i^{SoC} is updated via battery balance

$$E_i(t+1) = E_i(t) + \eta_{chg} P_{i,t}^{ch} \Delta t$$

Results

Data-driven statistical modules (GAM for households, Gamma sampling for mobility) feed deterministic and stochastic demand forecasts into a convex cost-minimisation.

Because every block speaks the same symbol language, the pipeline scales from ten measured cars to hundreds of synthetic ones while preserving calendar effects, price response and physical battery limits—delivering a ready-to-deploy scheduler for small-scale EV-storage aggregators.