Evidential Deep Learning and Recurrent World Models for Hypersonic Reentry Aerothermodynamics

Maneth Perera[1] and Poojak Patel[2]

[1]University of Illinois Urbana-Champaign, [2]Purdue University

Background

- Hypersonic reentry vehicles traverse extreme aerothermodynamic conditions across multiple flow regimes, from continuum flow at lower altitudes to free-molecular flow at the edge of the atmosphere.
- High-fidelity tools such as the Free Open Source Tool for Re-entry of Asteroids and Debris (FOSTRAD) are too computationally expensive for real-time guidance or Monte Carlo uncertainty propagation.
- Neural surrogates offer a fast alternative, but a surrogate that is confidently wrong in an unseen flow regime is more dangerous than one that correctly expresses its ignorance.

Methods

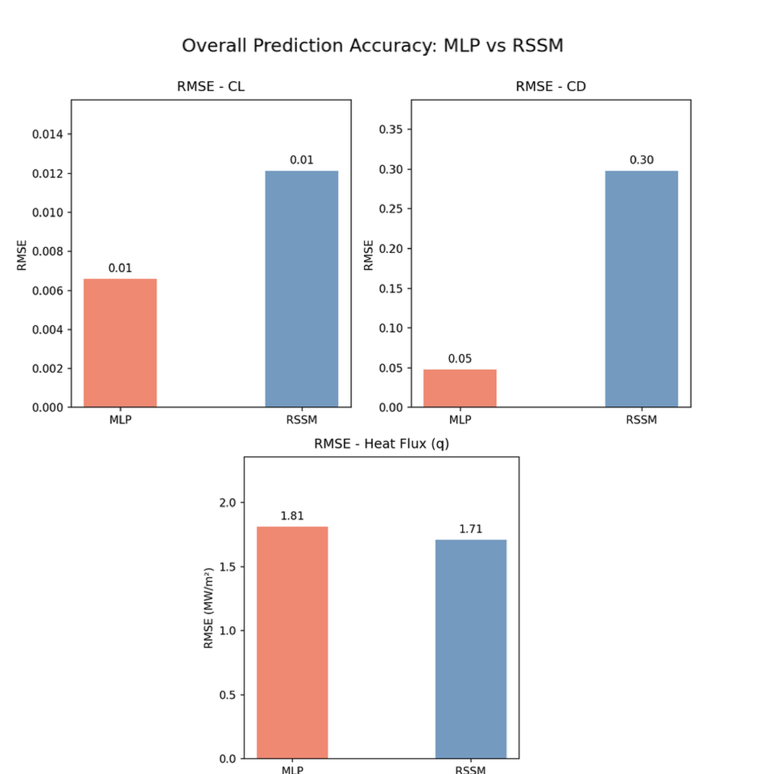

Figure 1:Comparison of prediction accuracies for lift coefficient, drag coefficient, and stagnant heat flux.

- **Dataset:** 340-point Latin Hypercube Sample from FOSTRAD, spanning altitude (30–120 km), velocity (3000–10000 m/s), AoA (0–20°), and wall temperature (300–1500 K), with 40 extra points targeting the free-molecular regime
- Four heat-flux correlations (Schaaf–Chambre, Kemp–Rose–Detra, Fay–Riddell, Van Driest), trajectory-level aerothermal targets come from an RBF interpolant over the full LHS sweep.
- **Evidential MLP:** 4-layer network with a Normal-Inverse-Gamma (NIG) head, trained with the evidential loss of Amini et al. (2020) for closed-form aleatoric/epistemic uncertainty.
- **Simplified RSSM:** GRU deterministic path with a stochastic latent state, trained on 300 synthetic 3DOF trajectories.

Important Results

- **RSSM (boundary-aware, not altitude-fixed):** Epistemic uncertainty peaks near the training range's edge at 28.6 km, not at a fixed physical altitude (tracks distance from the training boundary, not a hard-coded regime cutoff).
- **MLP (regime-conditioned, but not pooled):** MLP epistemic uncertainty on heat flux correlates stronger with Knudsen number within the continuum regime ($\rho = -0.409$) than when pooled across all regimes ($\rho = -0.120$), showing that uncertainty tracks local physics difficulty.

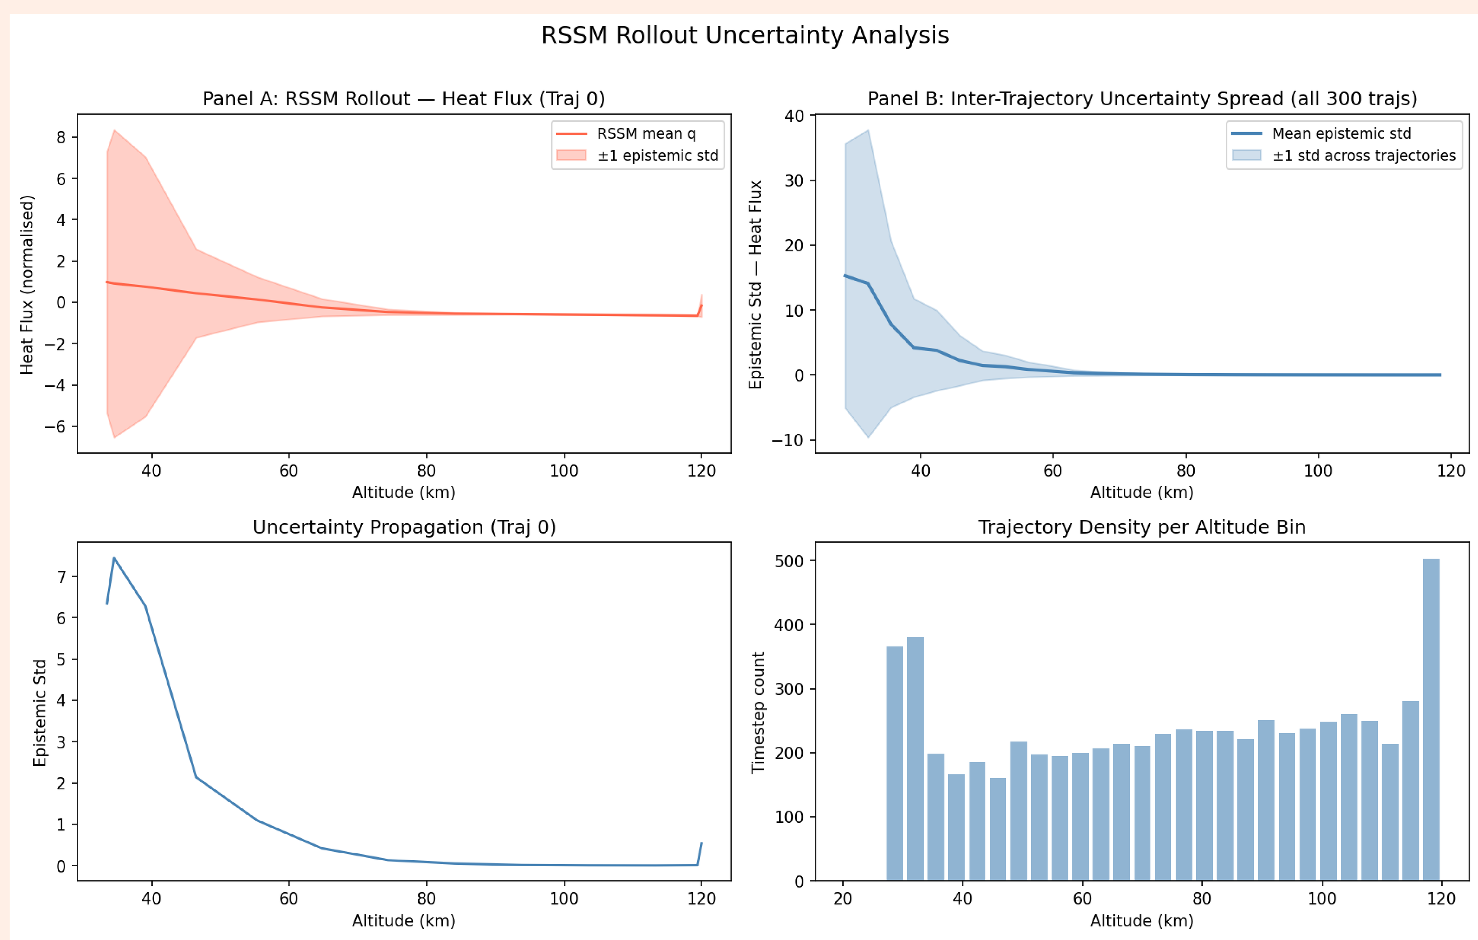

Figure 2:RSSM epistemic uncertainty vs. altitude during descent, peaking near the training boundary (28.6 km) rather than at a fixed interior altitude.

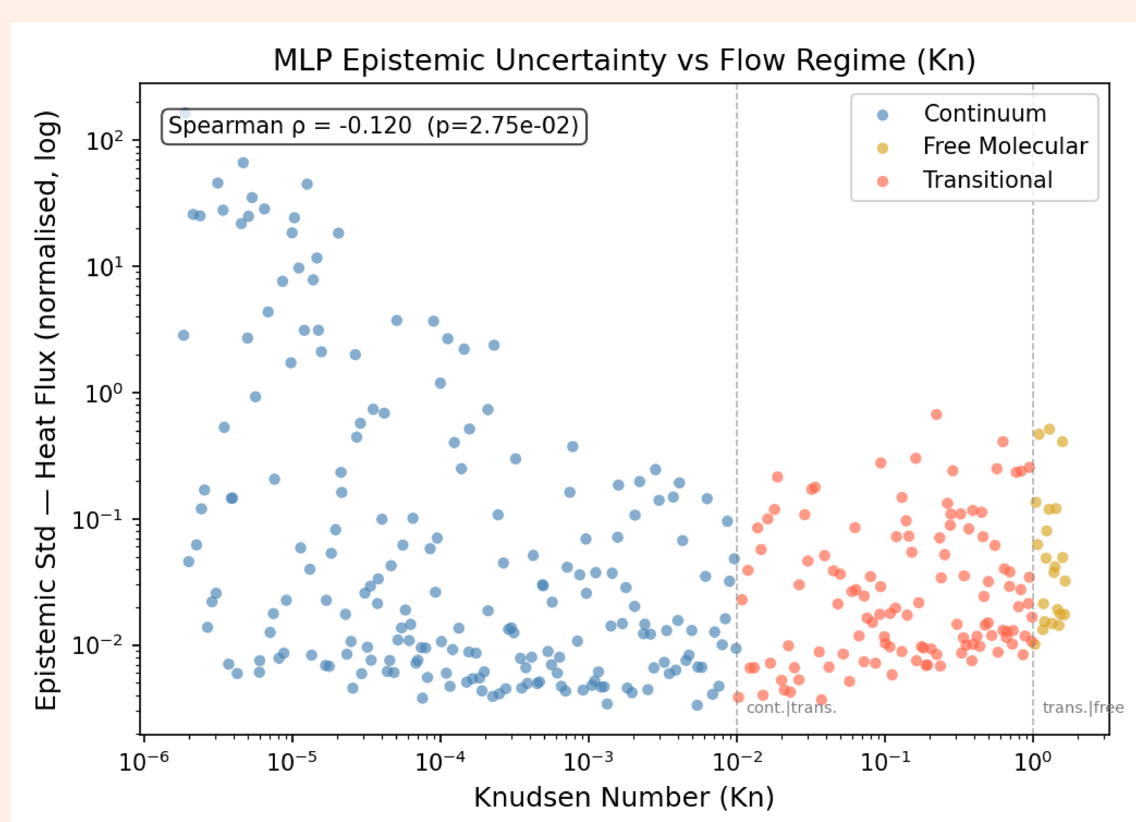

Figure 3:MLP epistemic uncertainty vs. Knudsen number (Kn), colored by flow regime. Correlation is strongest within the continuum regime ($\rho = -0.409$).

- **OOD Evaluation:** Three FOSTRAD-evaluated extrapolation sets (high-altitude, low-altitude, shifted-velocity), each 30 points beyond the training envelope.
- **Active Learning:** MLP epistemic uncertainty on heat flux selects new FOSTRAD query points.
- **Seed Robustness:** Evidential MLP is retrained across 15 seeds to quantify sensitivity of headline RMSE numbers.

Results

The MLP is consistent throughout, while the fixed KL weight leaves the RSSM overconfident below 0.7 confidence.

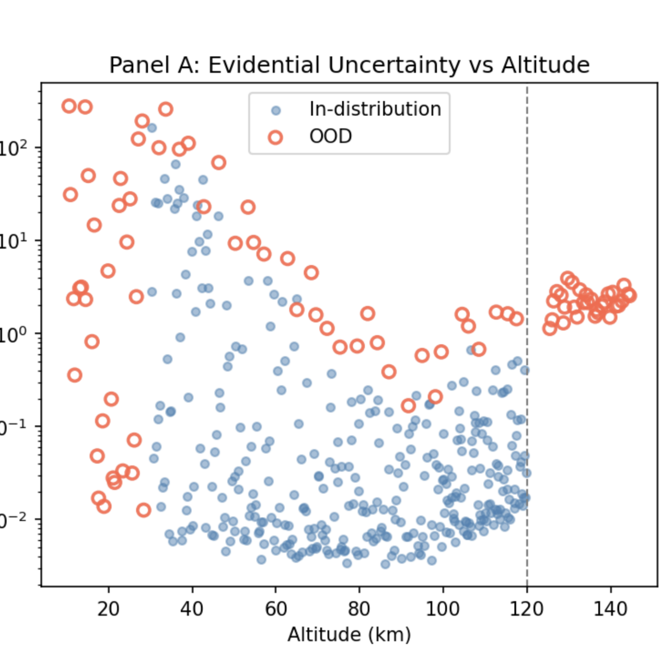

Figure 4:OOD points are on average 1.14× more uncertain, with 2.80× separation given by k-NN input distance.

The RSSM's temporal smoothing gives it the lowest continuum-regime RMSE, but both MLPs outperform it by a wide margin in the transitional regime.

$$\mathcal{L} = \mathcal{L}_{\text{NIG}} + \beta_{\text{KL}} \cdot D_{\text{KL}}[q(z_t|h_t, o_t) \,\|\, p(z_t|h_t)]$$

Discussion

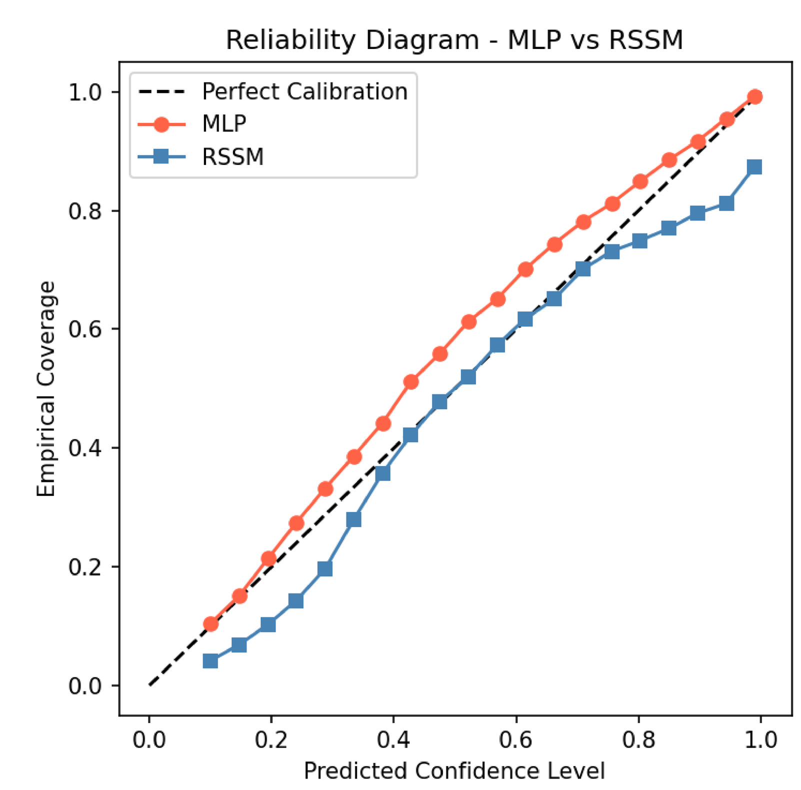

Figure 5:Reliability diagrams for the MLP (orange) and RSSM (blue). The MLP over-estimates and the RSSM underestimates.

- Heat-flux RMSE varies substantially across 15 seeds (CV = 45.1%). RSSM reversal and active-learning gain (6–9%) fall below this noise floor.
- Transitional-regime reversals in accuracy and uncertainty-vs-Kn correlation are connected but unexplained, plausibly tracing to FOSTRAD's bridging formula.
- Distance-aware heads (SNGPs, deterministic UQ, Prior Networks) are natural candidates for closing the OOD gap.
- Multi-vehicle generalization and onboard posterior state estimation remain open extensions.

Deployment & Active Learning

- FOSTRAD takes 7.28 s/call. MLP and RSSM inference take 0.056 ms and 0.082 ms, speedups of 130,000× and 88,000×.
- Uncertainty-guided learning beats the baseline (−6.2%) and random selection (−3.1%).
- This gain doesn't transfer to C_L/C_D. Their uncertainty correlates with each other ($r = 0.954$) but not with heat flux ($r \approx -0.10$).

References

- Amini et al. (2020) - Deep Evidential Regression, NeurIPS
- Hafner et al. (2019) - Dream to Control (Dreamer), ICLR
- Lips & Fritsche (2005) - FOSTRAD, Acta Astronautica
- Lakshminarayanan et al. (2017) - Deep Ensembles, NeurIPS
- McKay et al. (1979) - Latin Hypercube Sampling, Technometrics

Contact Info

mapere@illinois.edu
pate3495@purdue.edu