Week 7

Student: David Tang

Mentor: Dr. Yue Wang

Project Recap

Motivation

- Standard RL typically assumes a scalar reward
- Many problems have multiple objectives
- Tradeoffs need to be made, but real-world frontiers are noisy

Question

How to effectively navigate trade-offs in a balanced way and determine their stability?

Existing Approaches

- Linear Scalarization, Nash Bargaining Solution, Chebyshev

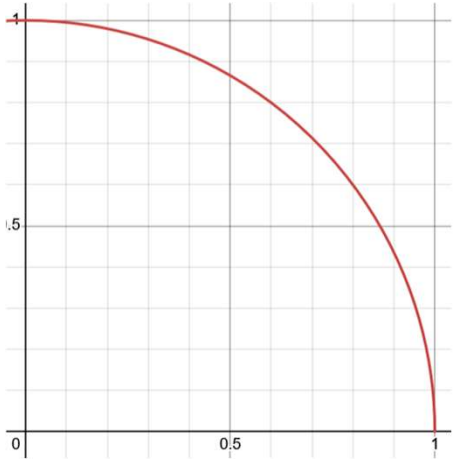

Our Method

- Nash + Chebyshev Convex Combination

Pareto Frontier Evaluation is well explored

How can we evaluate the stability of the evaluation?

Proposed Metric: Sensitivity

Let $\mathcal{D}(\mathbf{y})$ represent a designated multi-objective selection rule

Optimal vector $\boldsymbol{v}^* = \mathcal{D}(\mathbf{y}) = [x^*, y^*]^T$ and $\boldsymbol{v}_E^* = \mathcal{D}(\hat{\mathbf{y}}_E) = [x_E^*, y_E^*]^T$

Sensitivity Curve $S(E) = \|\boldsymbol{v}^* - \boldsymbol{v}_E^*\|_2$ Retained signal energy

$E \in [0, 1]$

Integrated Sensitivity Score $A = \int_0^1 S(E)\, dE$

Intuition: How stable is our selection rule under simplification?

Admissible Simplification Operators

Original empirical frontier $\mathbf{y} = [y_0, y_1, \ldots, y_{N-1}]^T \in \mathbb{R}^N$

Simplified frontier given by operator $\hat{\mathbf{y}}_E = \mathcal{P}_E(\mathbf{y})$

Simplifying the frontier should satisfy the following constraints

- Weak Monotonicity $\Delta y_n \leq 0 \implies \Delta \hat{y}_{E,n} \leq 0$
- Boundary Value Preservation $\hat{y}_{E,0} = y_0 \quad \text{and} \quad \hat{y}_{E,N-1} = y_{N-1}$

Type-II Discrete Cosine Transform as a Simplification Operator

Assumes even symmetry

Frequency coefficients $c_i = \mathrm{DCT}(\mathbf{y})_i$

Main Advantage is Parseval's Identity $\sum_{n=0}^{N-1} |y_n|^2 = \sum_{i=0}^{N-1} |c_i|^2$

We can find the exact cutoff index for the target energy $E(k) = \frac{\sum_{i=k}^{k-1} c_i^2}{\sum_{i=0}^{N-1} c_i^2}$

Simplified Front Construction

1. Truncation determined by cutoff index k $\mathbf{c}_{\text{filt}} = [c_0, c_1, \dots, c_{k-1}, 0, \dots, 0, 0]^T$
2. Reconstruct the front with the Inverse-DCT $\hat{\mathbf{y}}_{\text{recon}} = \text{IDCT}(\mathbf{c}_{\text{filt}})$
3. Enforce the monotonicity
$$\begin{aligned} \tilde{y}_{N-1} &= y_{N-1}, \\ \tilde{y}_n &= \max\left(\hat{y}_{\text{recon},n}, \tilde{y}_{n+1}\right) \quad \forall n \in \{N-2, \dots, 0\}, \end{aligned}$$
4. Rescale + enforce boundary condition $\hat{y}_{E,n} = y_{N-1} + (y_0 - y_{N-1}) \cdot \left(\frac{\tilde{y}_n - \tilde{y}_{N-1}}{\tilde{y}_0 - \tilde{y}_{N-1}} \right)$

Experiments (Analytical)

- Tier 1

$$y_n = \frac{1}{2}\left(1 + \cos(\pi x_n)\right)$$

- Tier 2

$$y_n = 1.0 - 0.5x_n - 0.3\tanh\left(10(x_n - 0.5)\right)$$

$$y_n = \frac{y_n - y_{N-1}}{y_0 - y_{N-1}}$$

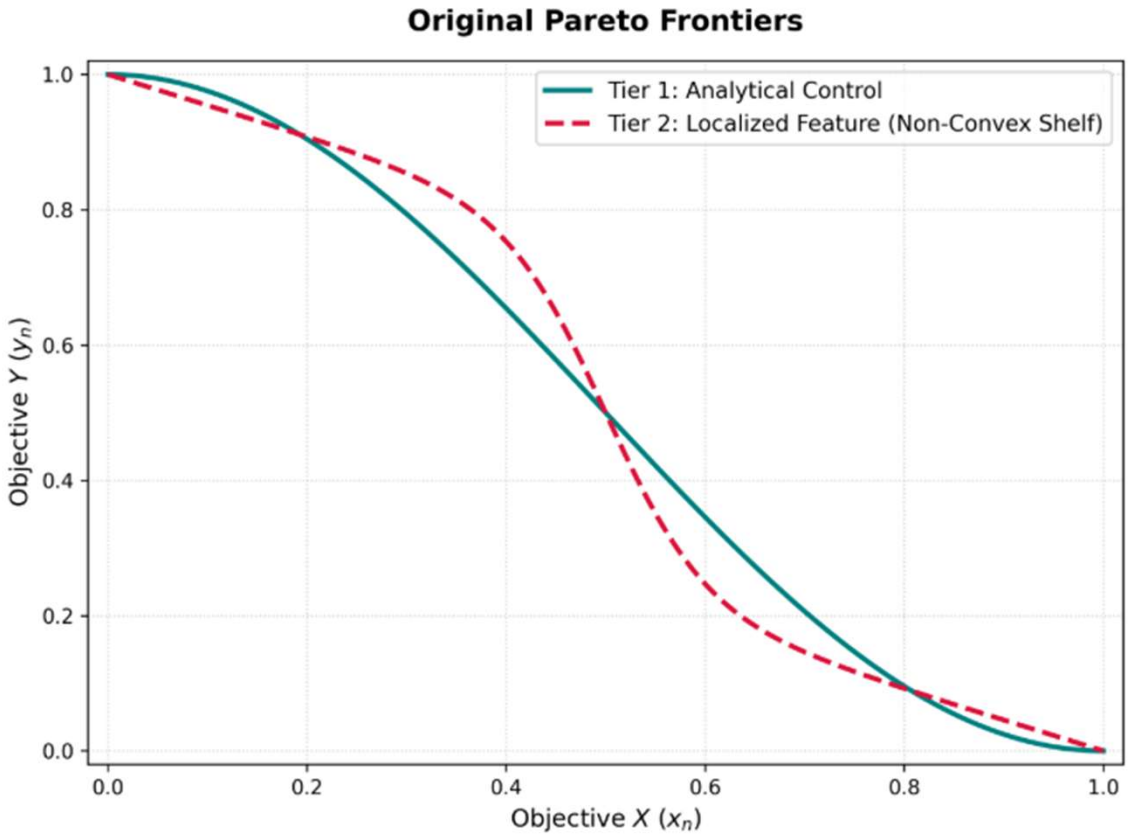

Experiments (Analytical)

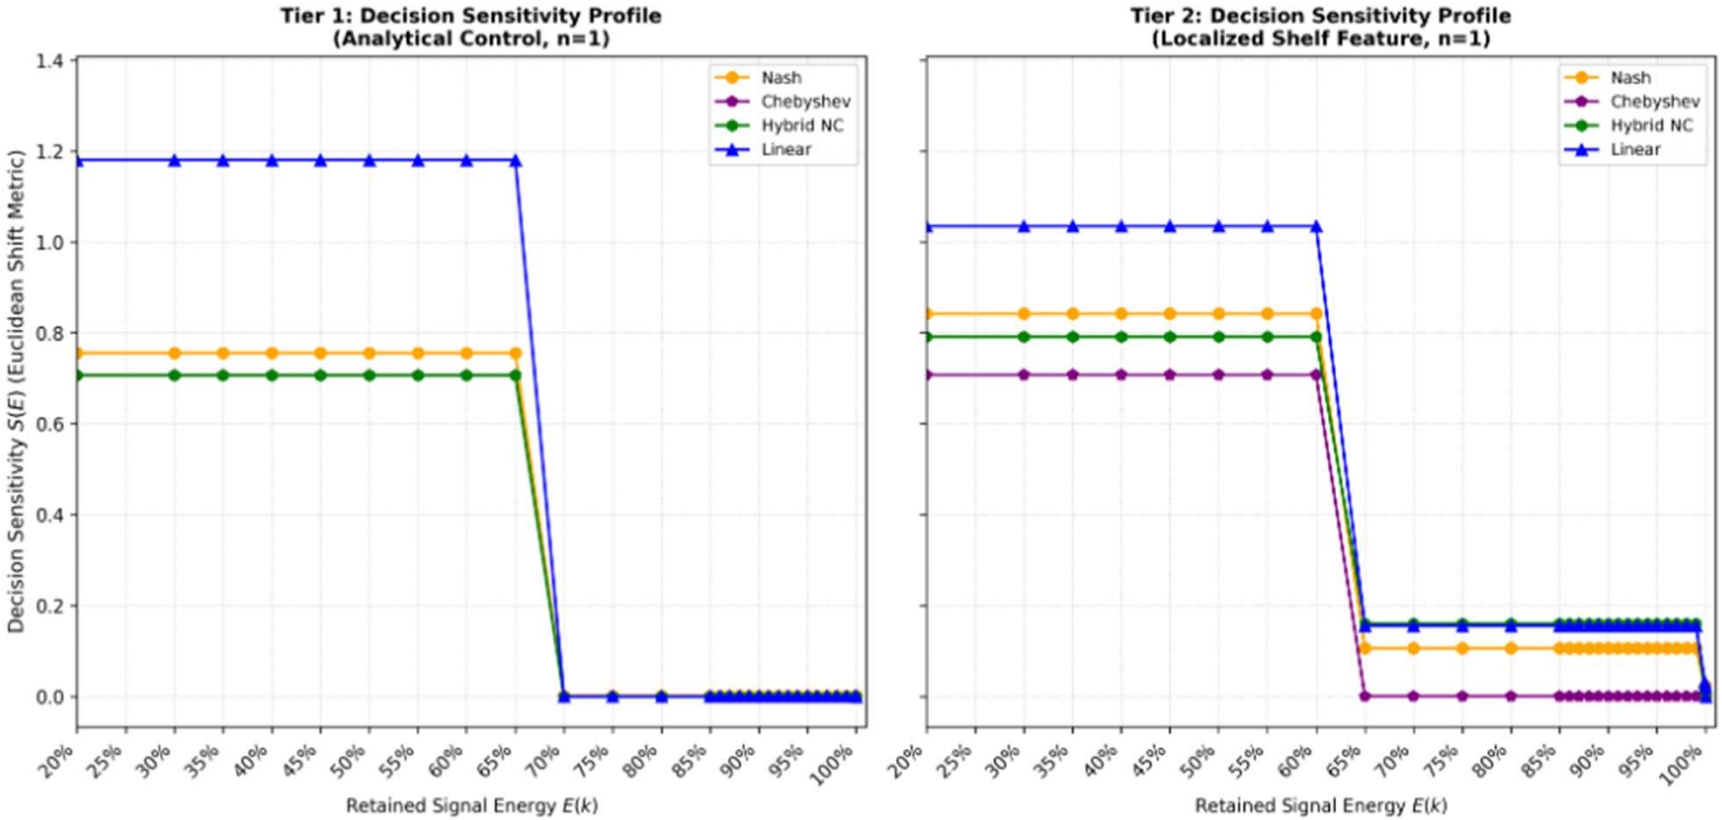

Experiments (Stochastic)

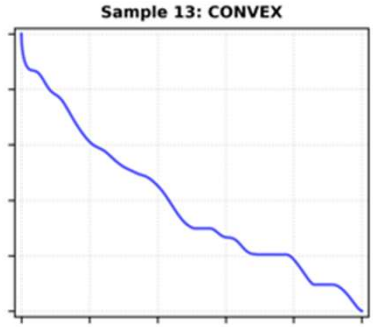

- Tier 3

Convex, Concave, S-Shape, Linear, Knee, Plateau

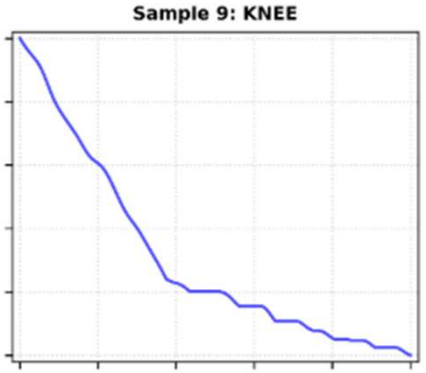

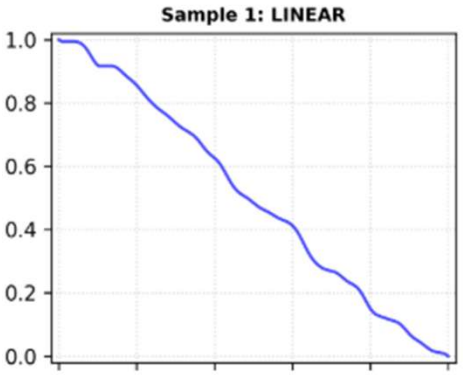

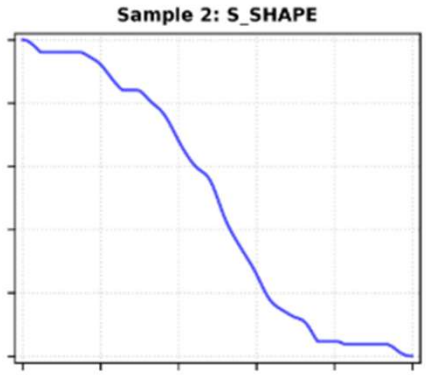

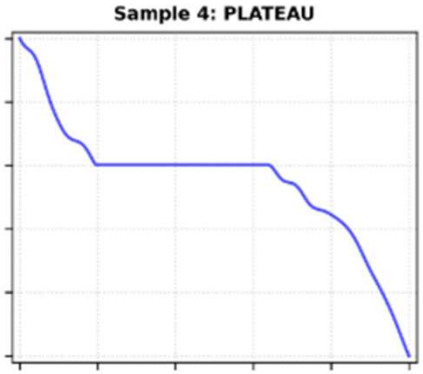

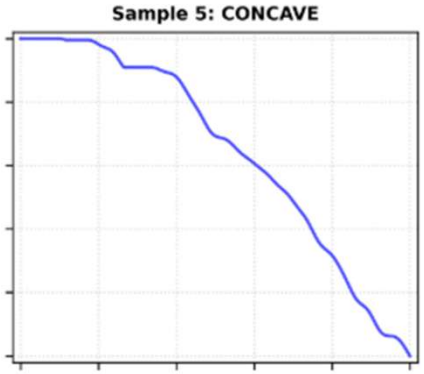

Experiments (Stochastic)

Shows consistency with known behavior of the different aggregation methods

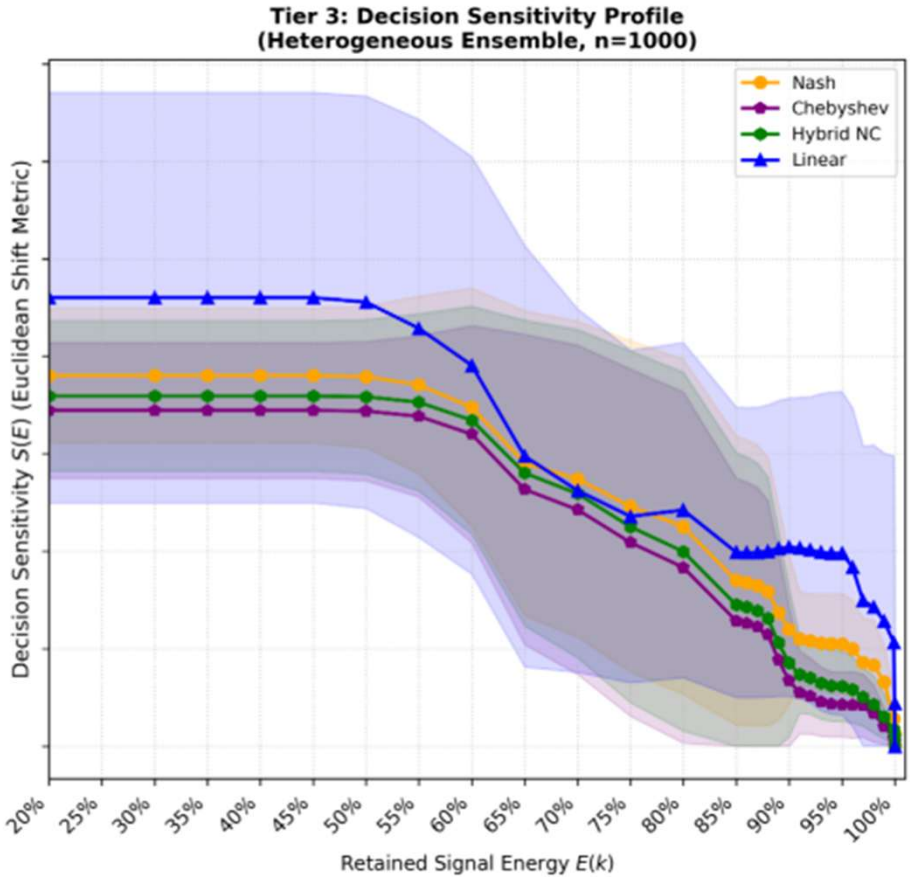

Experiments Table Summary

Table 1: Integrated Sensitivity Score Summary

Decision Metric	**Tier 1** ($n = 1$) Mean ($\pm$ Std)	**Tier 2** ($n = 1$) Mean ($\pm$ Std)	**Tier 3 Ensemble** ($n = 1000$) Mean ($\pm$ Std)
MaxMin	0.3360 ($\pm$0.0000)	0.3011 ($\pm$0.0000)	0.4052 ($\pm$0.1297)
Chebyshev	0.3360 ($\pm$0.0000)	0.3011 ($\pm$0.0000)	0.4052 ($\pm$0.1297)
Hybrid NMM	0.3360 ($\pm$0.0000)	0.3224 ($\pm$0.0000)	0.4186 ($\pm$0.1433)
Hybrid NC	0.3360 ($\pm$0.0000)	0.3959 ($\pm$0.0000)	0.4298 ($\pm$0.1407)
Nash	0.3596 ($\pm$0.0000)	0.3973 ($\pm$0.0000)	0.4653 ($\pm$0.1292)
Linear	0.5611 ($\pm$0.0000)	0.4977 ($\pm$0.0000)	0.5495 ($\pm$0.1134)

Next Steps

- Apply sensitivity metric to the MORL environments
- Hyperparameter tuning and ablations
- Clean up code base
- Statistical Analysis
- Continue paper writing