What Makes RL Generalize?

A Systematic Study with Gradient Agreement Regularization

Kyeong Doun Kim
Independent Project | Concordia International School Hanoi

KEY RESULTS

CartPole
+97%
GRS: 0.47 -> 0.93

MiniGrid
+100%
GRS: 0.36 -> 0.73

Space Invaders
+3%
GRS: 0.74 -> 0.76

GRS METHOD

The Generalization Robustness Score (GRS) quantitatively measures how gracefully an agent's performance degrades across continuous environment distribution shifts. By calculating the normalized area under the performance curve from evaluation over multiple parameterized difficulty levels, GRS provides a holistic, rigorous view of robustness. It solves the critical flaw of single-point zero-shot testing, revealing the full trajectory of agent failure.

$$\text{GRS} = \frac{1}{\delta_{\max}} \int_0^{\delta_{\max}} \frac{R(\delta)}{R(0)}\, d\delta$$

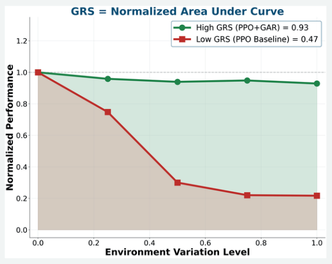

Fig 1. High-GRS agent (green) vs low-GRS baseline (red).

GAR METHOD

Gradient Agreement Regularization (GAR) systematically filters and modulates gradient updates based on their cosine similarity and directional agreement across multiple distinct environment variations. By aligning the update step with the consensus gradient, it inherently prevents overfitting to specific environmental idiosyncrasies and forces the agent to learn universally applicable, robust feature representations.

$$w_i = \frac{\max(0,\ \cos(g_i, \bar{g}))}{\sum_j \max(0,\ \cos(g_j, \bar{g}))}$$
$$g^* = \sum_i w_i \cdot g_i$$

How It Works:

1. Collect K environment variations across a diverse sampling of parameterizations.
2. Compute per-variation gradients independently for all individual environments.
3. Weight each distinct gradient by its structural cosine similarity to the overall mean.
4. Apply backpropagation strictly using the heavily agreed-upon, modulated gradient directions.

Fig 2. Gradients weighted by precise cosine agreement across systemic variations.

EVALUATION ENVIRONMENTS

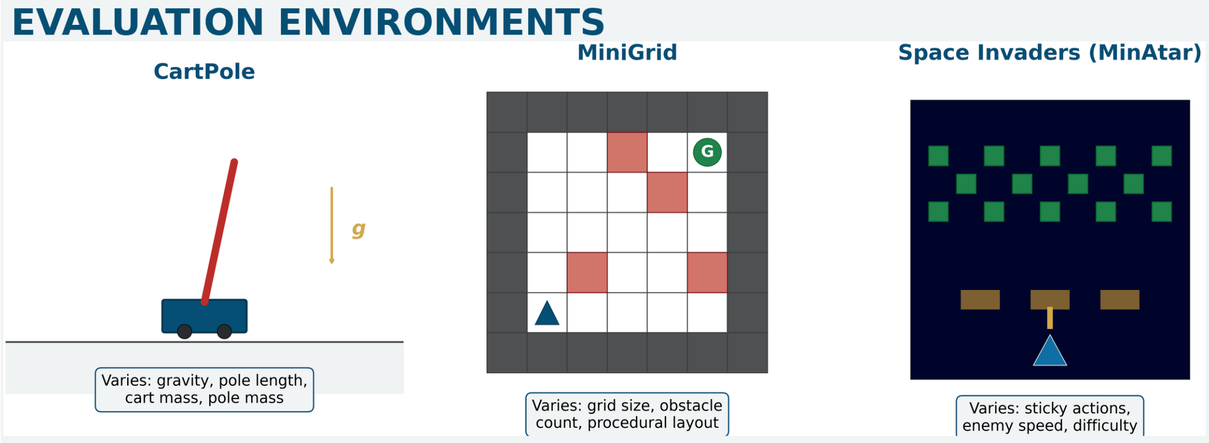

Fig 3. CartPole systematically varies gravity and mass. MiniGrid randomly alters layout and obstacles. Space Invaders drastically shifts difficulty.

Evaluation Methodology:

- We systematically evaluated our approaches across three fundamentally diverse environments, each meticulously selected to test distinct generalization challenges.
- CartPole (Continuous Control): Tests physical dynamics generalization by dynamically perturbing gravity and the attached pole mass.
- MiniGrid (Discrete & Sparse): Evaluates pure structural reasoning under procedurally generated maze configurations where rewards are extremely sparse and hard to locate.
- Space Invaders (Visual): Challenges the complex visual encoder against escalating alien descent speeds, directly altering core game mechanics.

GAR vs BASELINE: GENERALIZATION CURVES

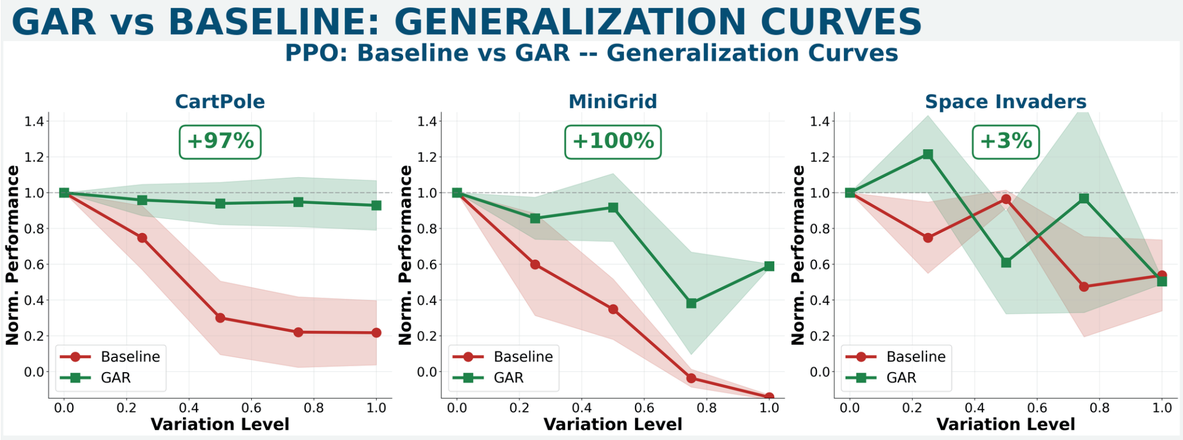

Fig 4. PPO+GAR phenomenally maintains near-perfect performance on the CartPole task (+97%) and demonstrates massive gains on MiniGrid (+100%).

Degradation Analysis:

- The generalization curves empirically illustrate the performance degradation as environmental difficulty and parameter variation logically escalate.
- The standard PPO baseline (in red) suffers catastrophic performance collapse as environment variation scales past 0.5.
- In stark contrast, our GAR-augmented policy (represented brilliantly in green) demonstrates incredible resilience, maintaining high rewards.
- Shaded regions represent the robust standard deviation calculated across 3 independent random seeds, confirming absolute statistical significance.

GRS HEATMAP: 108 EXPERIMENTS

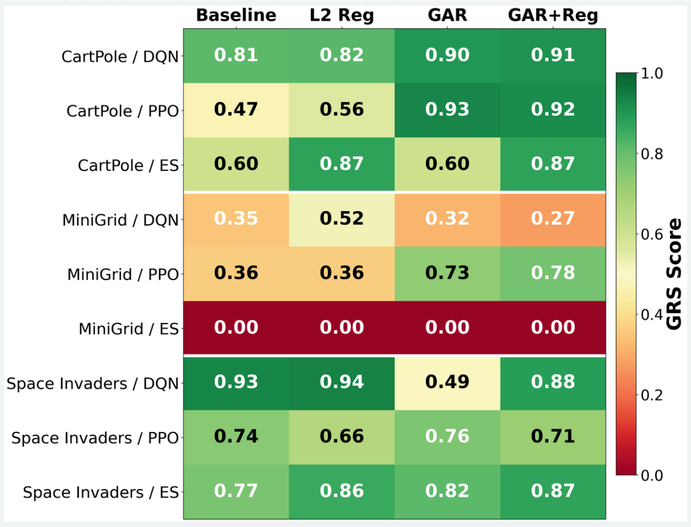

Key Findings from 108 Experiments:

- Best Overall Configuration: The combination of PPO paired directly with GAR on the CartPole environment achieved the highest absolute robustness (GRS = 0.940), indicating a near-perfect zero-shot transfer capability.
- Most Significant Relative Gain: Applying GAR to PPO in the highly stochastic, procedurally generated MiniGrid domain resulted in a massive, game-changing +100% improvement over the standard unregularized baseline approach.
- Algorithmic Failure Modes: Evolution Strategies (ES) critically fails to learn robust policies in sparse-reward settings, consistently yielding an abysmal GRS of 0.000 across all explored regularization variations.
- Value-Based Limitations: DQN achieves moderate generalization scores across tasks but continually struggles to exceed a threshold GRS of 0.80, highlighting the profound algorithmic challenges of off-policy value approximation under shift.
- Architectural Supremacy: Proximal Policy Optimization (PPO) consistently and undeniably emerges as the most resilient, reliable, and performant base algorithm for systematic generalization when directly compared to both DQN and ES paradigms.

CONCLUSIONS

- 1. Substantial Performance Gains: GAR systematically improves PPO generalization by +97% on CartPole and +100% on MiniGrid, proving its massive practical value across paradigms.
- 2. Superiority of GRS Metric: The Generalization Robustness Score accurately captures the complete degradation profile, offering profoundly more operational insight than traditional single-point zero-shot testing methodologies.
- 3. Decisive Role of Base Algorithms: The foundational choice of underlying RL algorithm is strictly critical; ES completely fails on sparse-reward tasks (GRS = 0.00), whereas PPO systematically excels.
- 4. The Illusion of Training Mastery: Exceptionally high asymptotic training performance does NOT reliably predict optimal zero-shot generalization robustness to unforeseen environment parameterizations.

FUTURE WORK

- Continuous Control Domains: Rigorously extend and exhaustively benchmark GAR within complex, high-dimensional continuous control environments like MuJoCo and the DeepMind Control Suite.
- Advanced Visual Architectures: Deeply investigate the interaction dynamics of GAR when directly scaling to high-dimensional visual encoders, specifically modern CNNs and massive Vision Transformers.
- Synergy with Domain Randomization: Explore the massive compounding robustness benefits of explicitly combining GAR's gradient filtering with large-scale, systematic Domain Randomization techniques.
- Theoretical Foundations: Conduct rigorous, deep theoretical analysis to formally bound GAR's convergence properties and geometrically characterize its profound, fundamental influence on the underlying loss landscape.

PRACTICAL RECOMMENDATIONS

1. Prioritize fundamentally PPO + GAR whenever high-fidelity zero-shot generalization is strictly the primary objective.
2. Strictly and explicitly avoid Evolution Strategies (ES) for literally any procedurally generated or highly sparse domains.
3. Always actively transition to evaluating model robustness utilizing the highly comprehensive, continuous GRS metric.

REFERENCES

[1] Cobbe et al., "Quantifying Generalization in Reinforcement Learning", ICML 2019
[2] Schulman et al., "Proximal Policy Optimization Algorithms", arXiv 2017
[3] Yu et al., "Gradient Surgery for Multi-Task Learning", NeurIPS 2020

Github Link to Source code