

Adaptive World Models: Learning Behaviors by Imagination Under Non-Stationarity

Emiliyan Gospodinov¹, Vaisakh Shaj¹, Phillipp Becker¹, Stefan Geyer², Gerhard Neumann¹

1. Karlsruhe Institute of Technology (KIT), Germany

2. University of Stuttgart, Germany


Existing Gap

- Dreamer agents showed SOTA-performance **mainly on single, stationary tasks**:
 - Can Dreamer world models **adapt to changes** and be used to **infer adaptable behaviors**?

Existing Gap

- Dreamer agents showed SOTA-performance **mainly on single, stationary tasks**:
 - Can Dreamer world models **adapt to changes** and be used to **infer adaptable behaviors**?

Observation changes: Color




Dynamics changes: Mass



Objective changes: Skills




Non-Stationary RL Formalisms




POMDP

Non-Stationary RL Formalisms

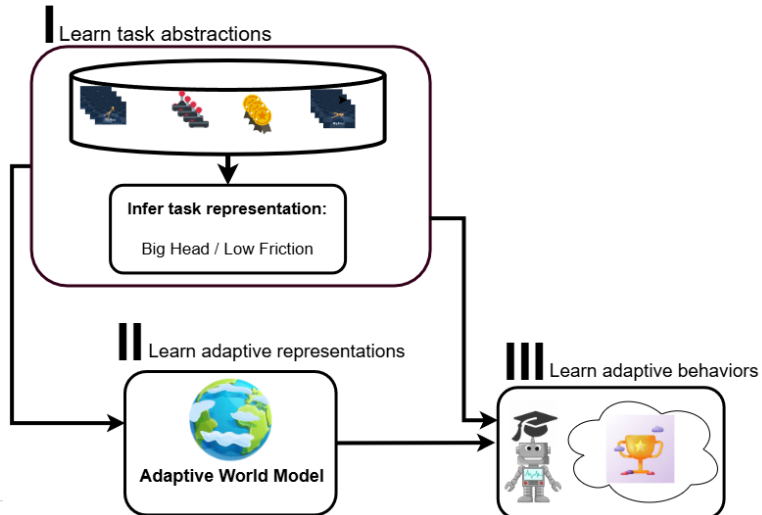


POMDP




HiP-POMDP

Adaptive MBRL Agents

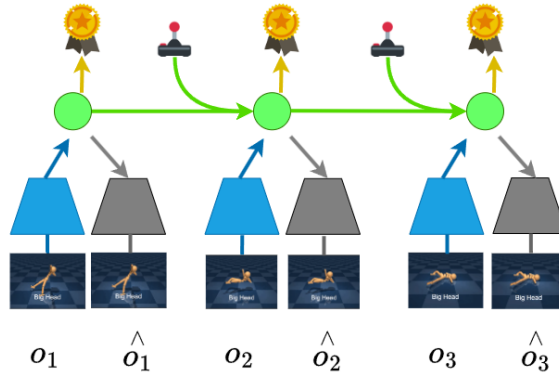


I: Latent Task Inference using Bayesian Context Aggregation




Volpp et al. 2021

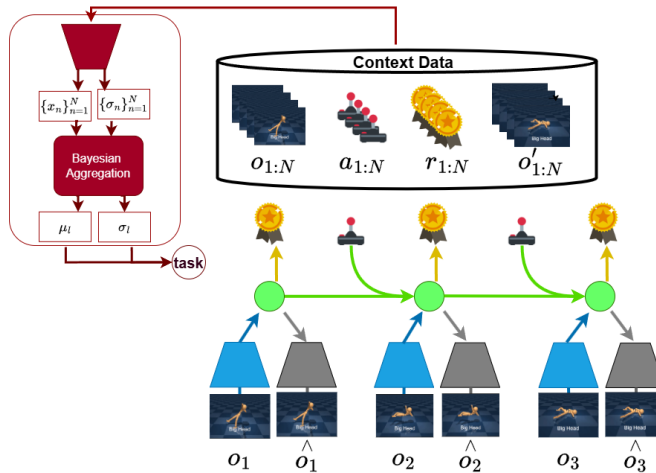
II: Learning Adaptive Representations



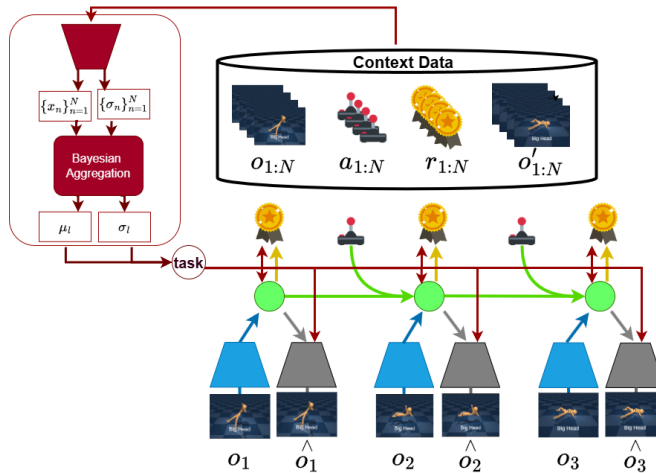
II: Learning Adaptive Representations



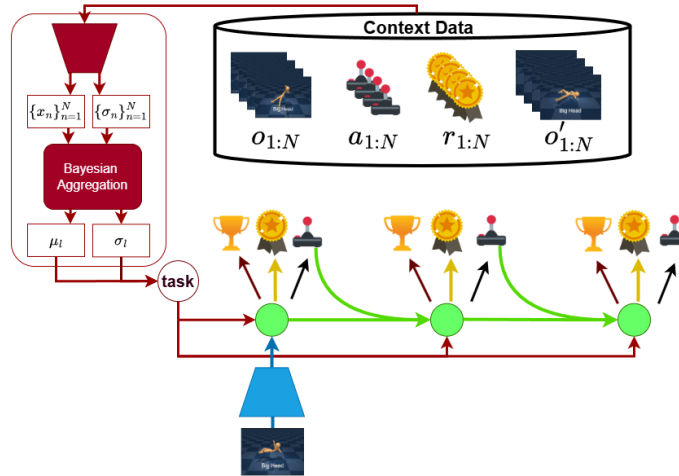
II: Learning Adaptive Representations



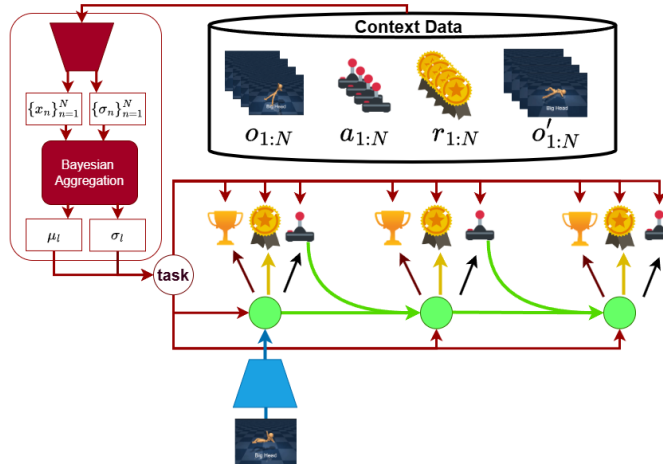
II: Learning Adaptive Representations



III: Learning Adaptive Behaviors



III: Learning Adaptive Behaviors




Modified Gymnasium + DMC Benchmarks

Gymnasium

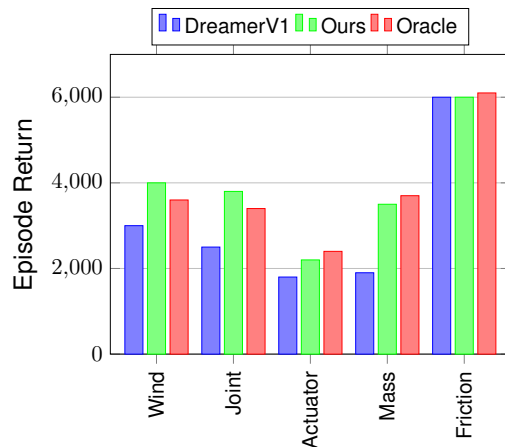
DMC Multi-task

Can agents handle changing dynamics?

Inter-Episodic

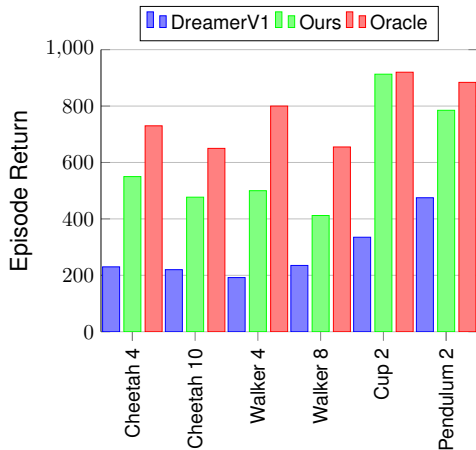


Intra-Episodic

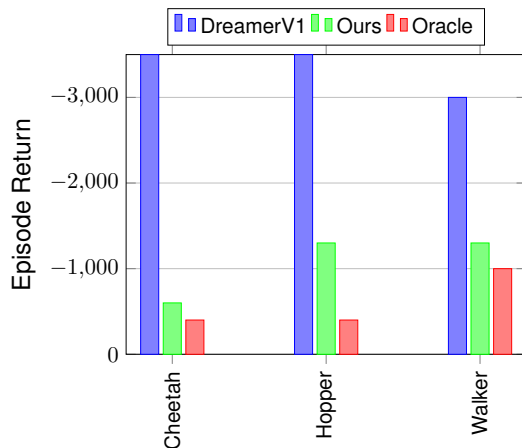


Can agents handle objective changes?

Multi-task




Target Velocity




Latent space visualizations


Task Inference Dreamer
Latent Task Space




Task Inference Dreamer
Latent Task Space




Task Inference Dreamer
Latent State Space




Task Inference Dreamer
Latent State Space




Vanilla Dreamer
Latent State Space




Vanilla Dreamer
Latent State Space



Oracle Dreamer
Latent State Space



Oracle Dreamer
Latent State Space



Conclusion

- **Bayesian context aggregation** as a **general** task inference **approach**.
- **POMDP** formalism **can handle** dynamical changes, **not** objective changes.
- **HiP-POMDP** results in **more structured latent spaces**, **improving** the **performance**.