

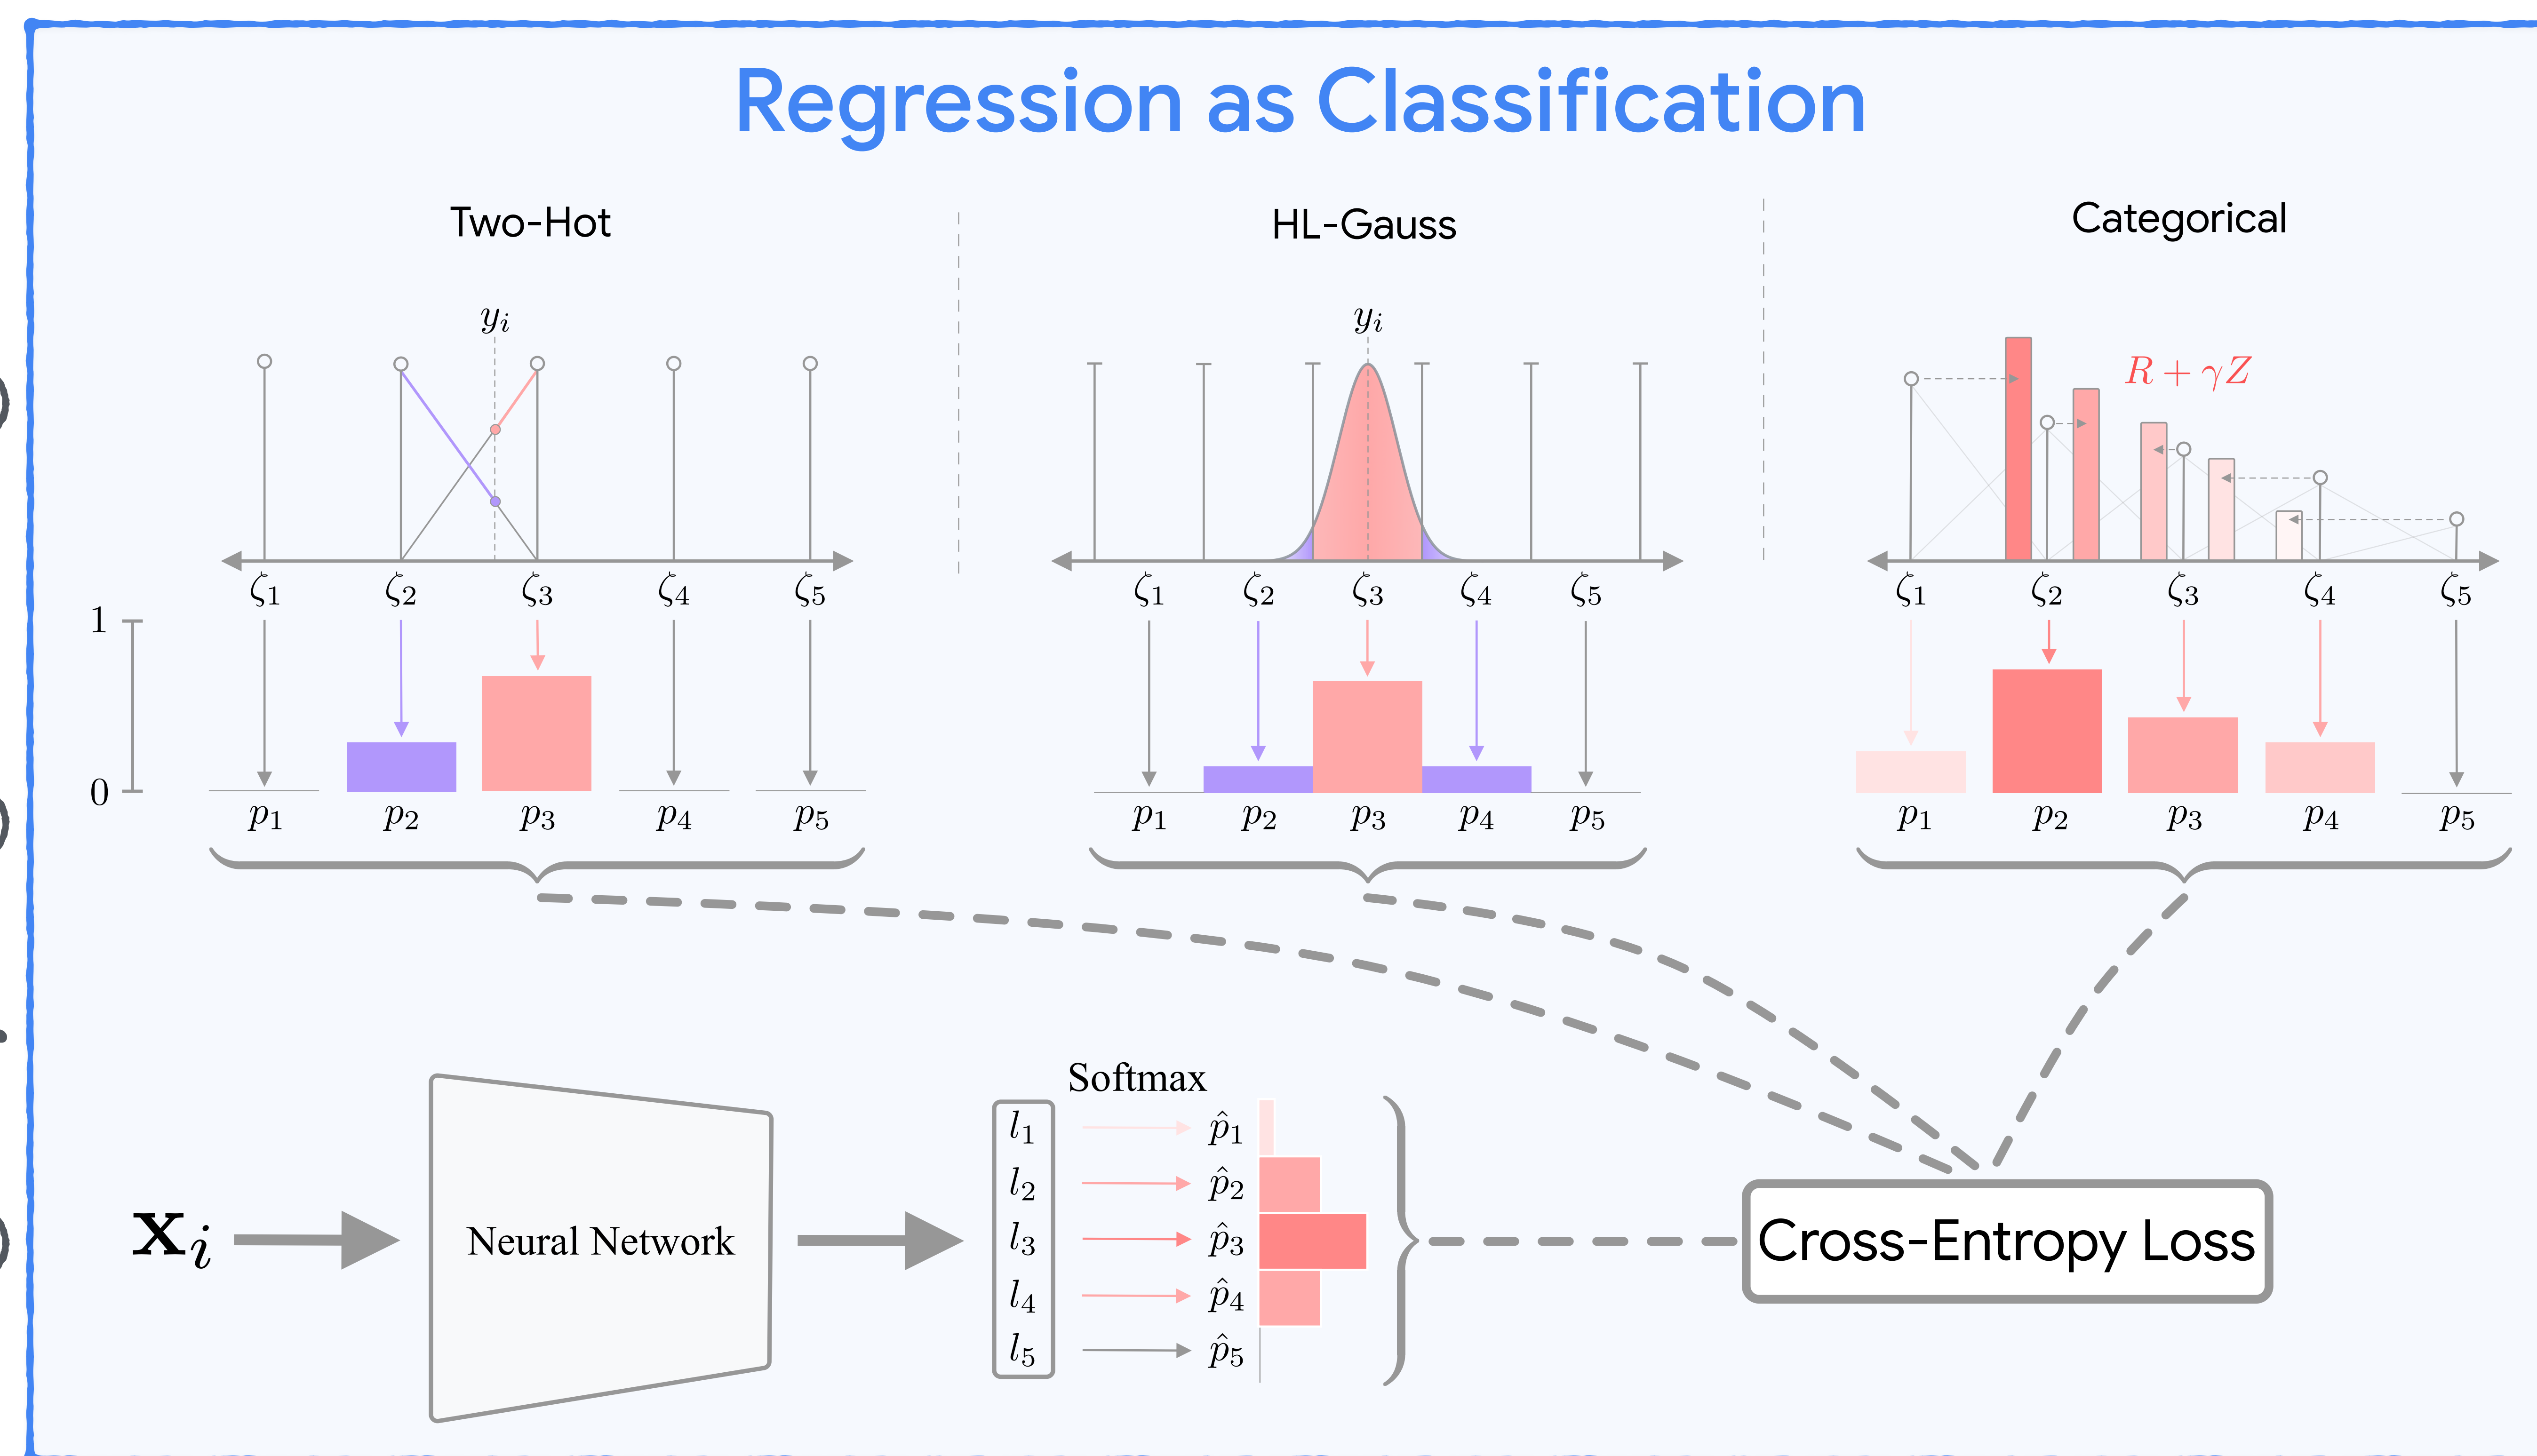
Jesse Farebrother\*, Jordi Orbay, Quan Vuong, Adrien Ali Taïga, Yevgen Chebotar, Ted Xiao, Alex Irpan, Sergey Levine, Pablo Samuel Castro, Aleksandra Faust, Aviral Kumar, Rishabh Agarwal\*

## Why Hasn't RL "Scaled"?

Deep Learning has seen incredible success scaling to large capacity models.

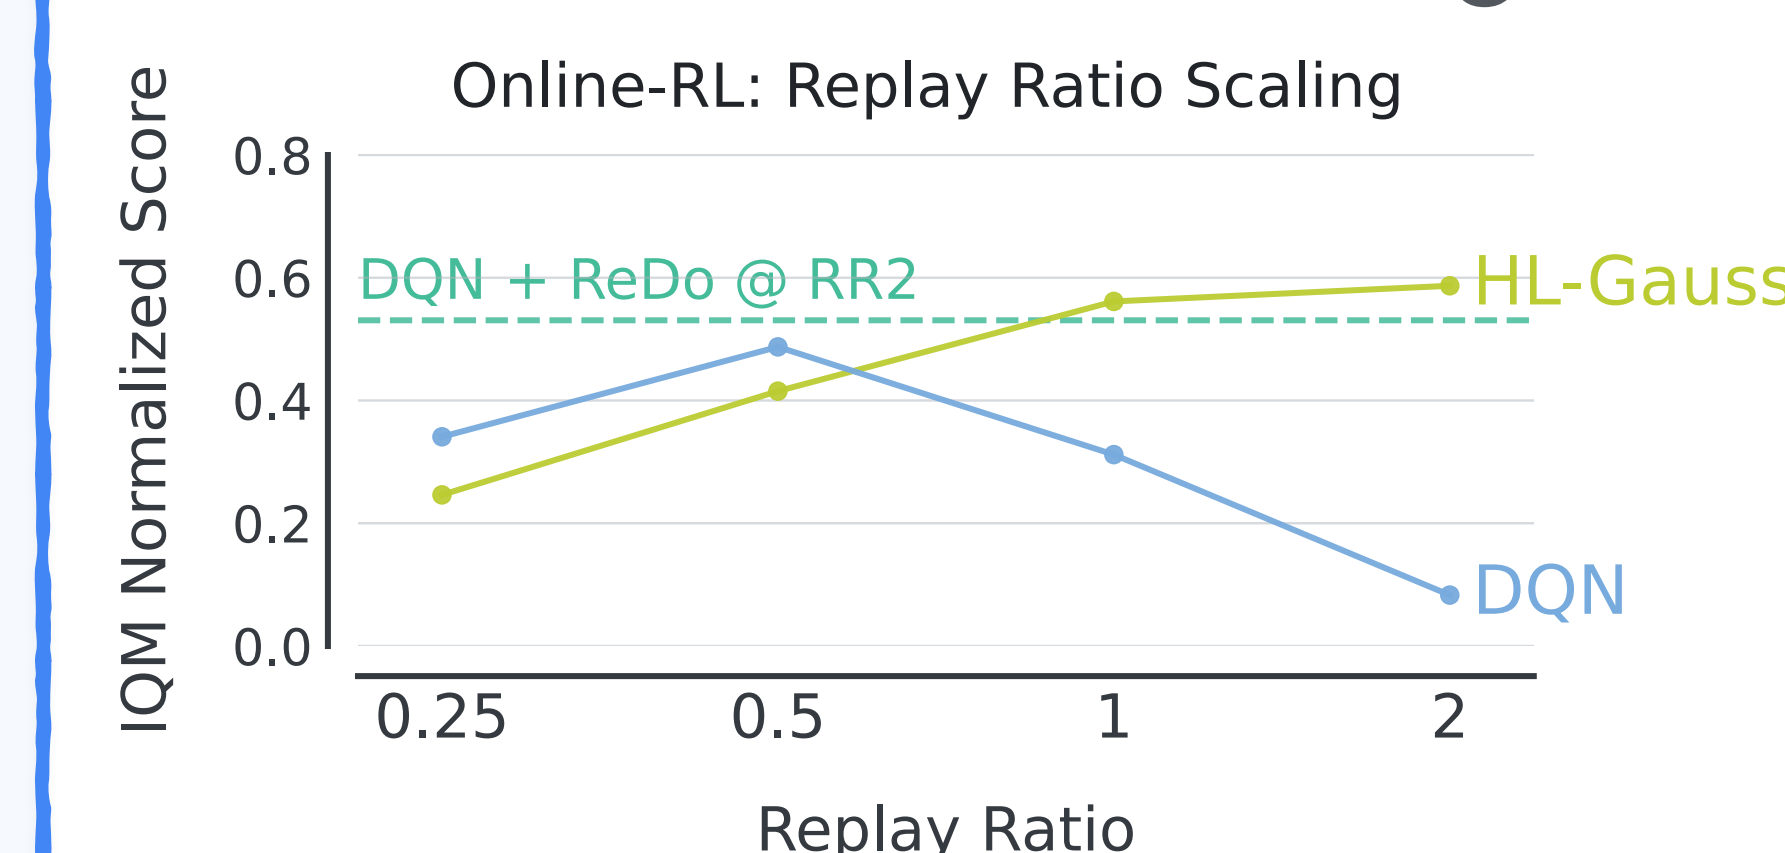
But these results are all on **classification** tasks... too bad RL uses regression 🤖

Not so fast... maybe casting **regression as classification** enables scaling for RL?

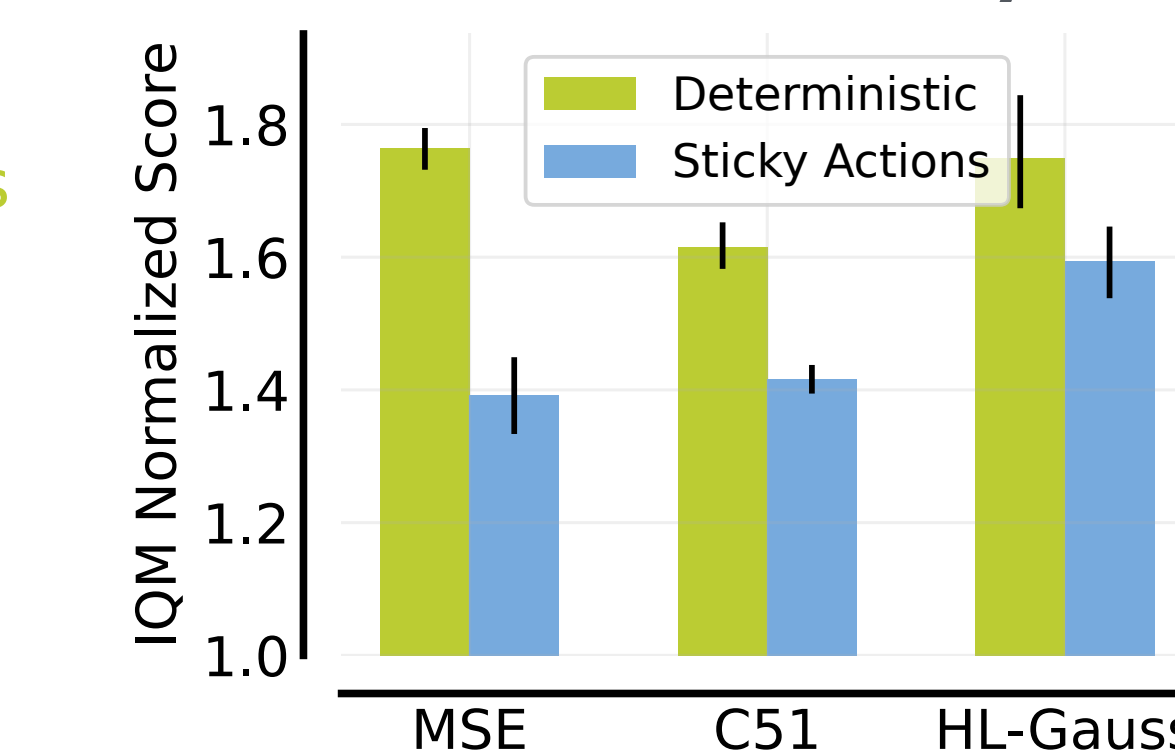


## Why Classification benefits RL?

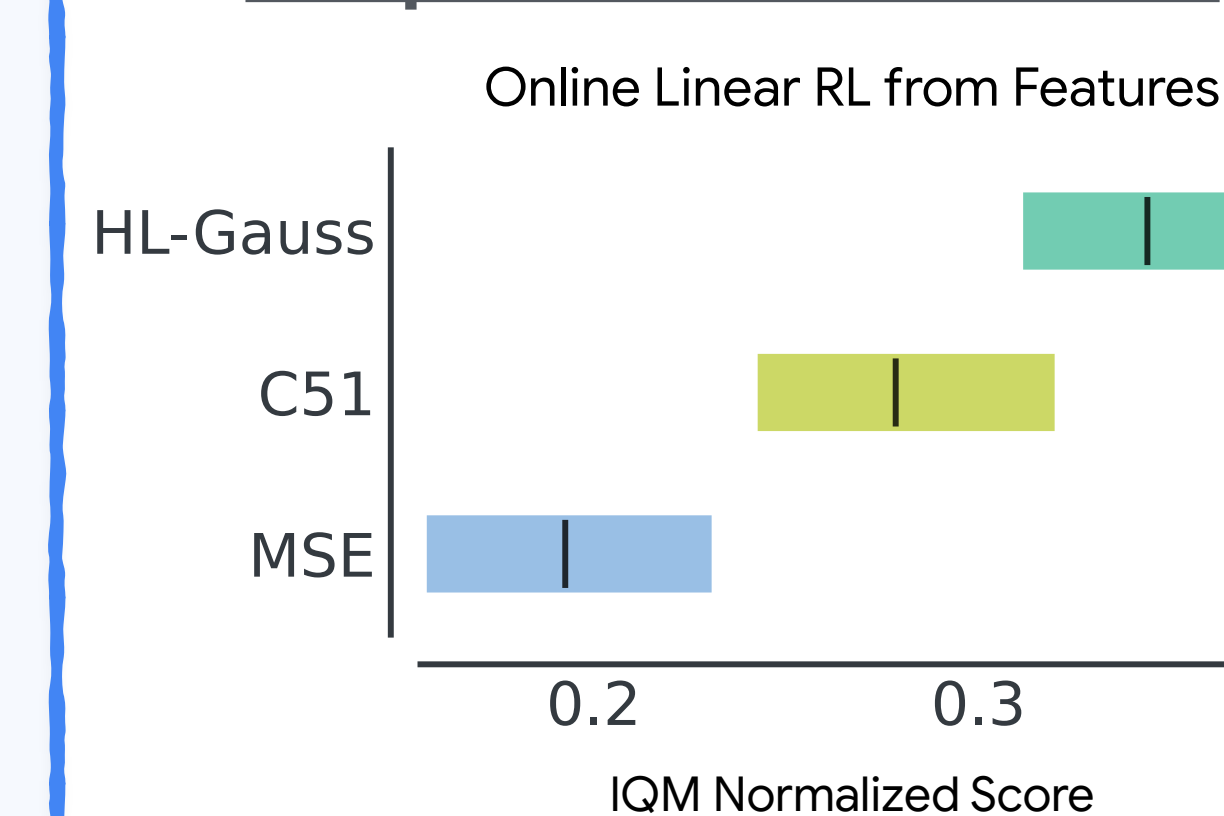
### Reduced Overfitting



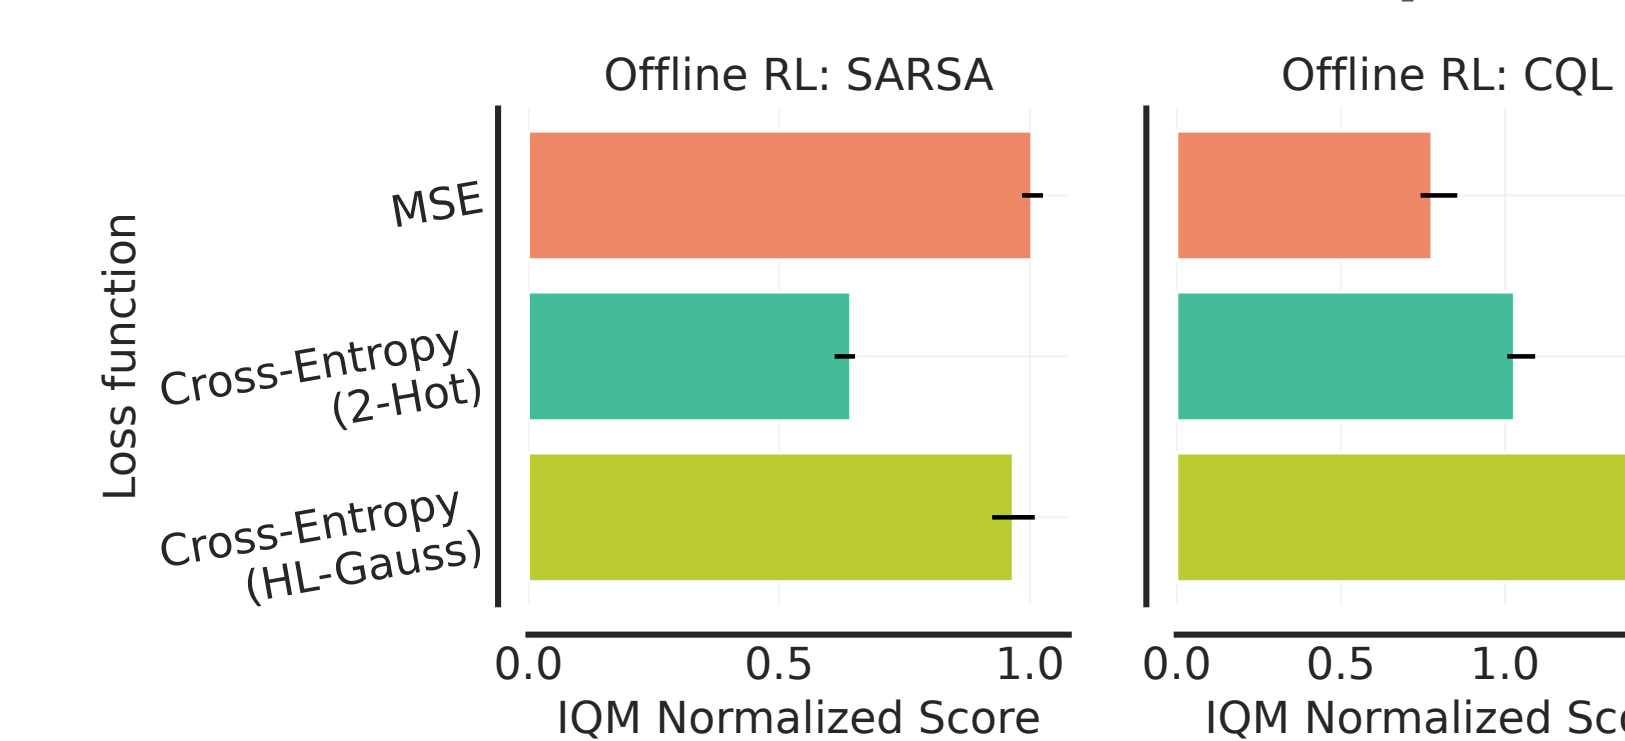
### Robustness to Stochasticity



### Improved Representations



### Robustness to Policy Non-Stationarity



## Classification Enables Scaling Capacity in Deep RL

