

Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron Courville, Marc G. Bellemare

# **Statistical Uncertainty in Deep RL Exacerbated by Handful of Runs**



## **Report Interval Estimates rather than Point Estimates.**



# Deep RL at the Edge of the Statistical Precipice

# 2019 2020 2015 2005 2017 Year

## How to Reliably Evaluate Performance in Deep RL?

variable and should be treated as such.

- A. Use interval estimates as opposed to point estimates for accounting for uncertainty in results.
- B. More is more: Performance profiles for qualitative summarization
- D. Provide individual runs for thorough statistical comparisons.





# **Outstanding Paper Award**

- Performance estimates based on a finite number of runs is a random
  - C. Better aggregate performance measures from robust statistics such
    - as interquartile mean (IQM) and probability of improvement.