



# Visual Reinforcement Learning with Imagined Goals

Ashvin Nair\*, **Vitchyr Pong\***, Murtaza Dalal, Shikhar Bahl, Steven Lin, Sergey Levine

# Autonomous Learning





# Autonomous Learning



# Autonomous Learning

**Reinforcement Learning?**



# Autonomous Learning: Proof of Concept



# Reinforcement Learning with Imagined Goals (RIG)

# Reinforcement Learning with Imagined Goals (RIG)

Autonomous goal setting

# Reinforcement Learning with Imagined Goals (RIG)

Autonomous goal setting

Vision input




# Reinforcement Learning with Imagined Goals (RIG)




Autonomous goal setting

Vision input


Real-world results



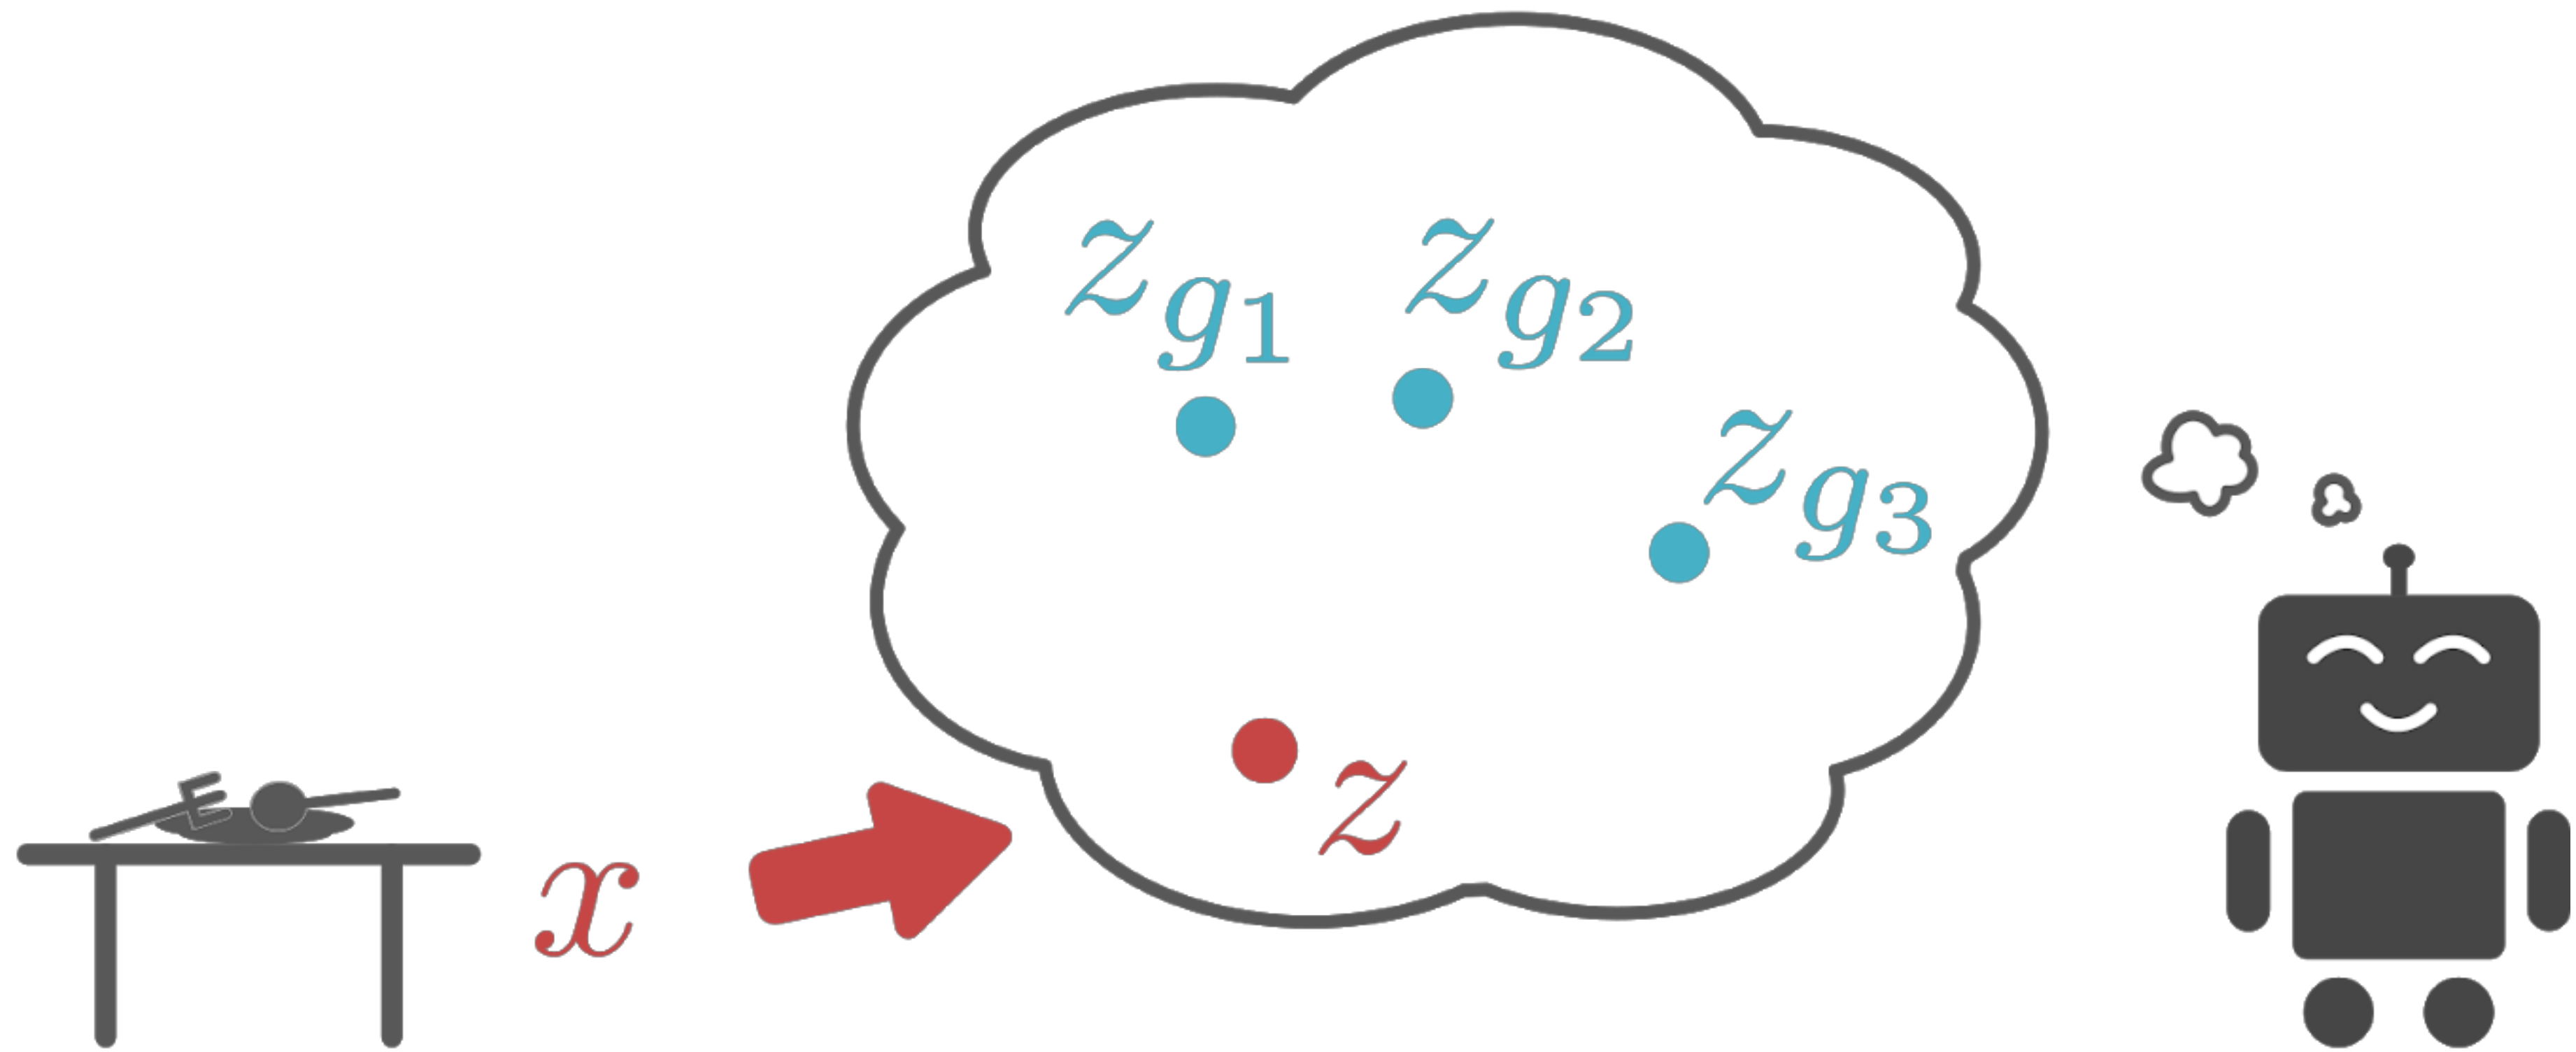
# RIG: Autonomous Learning




# RIG: Autonomous Learning



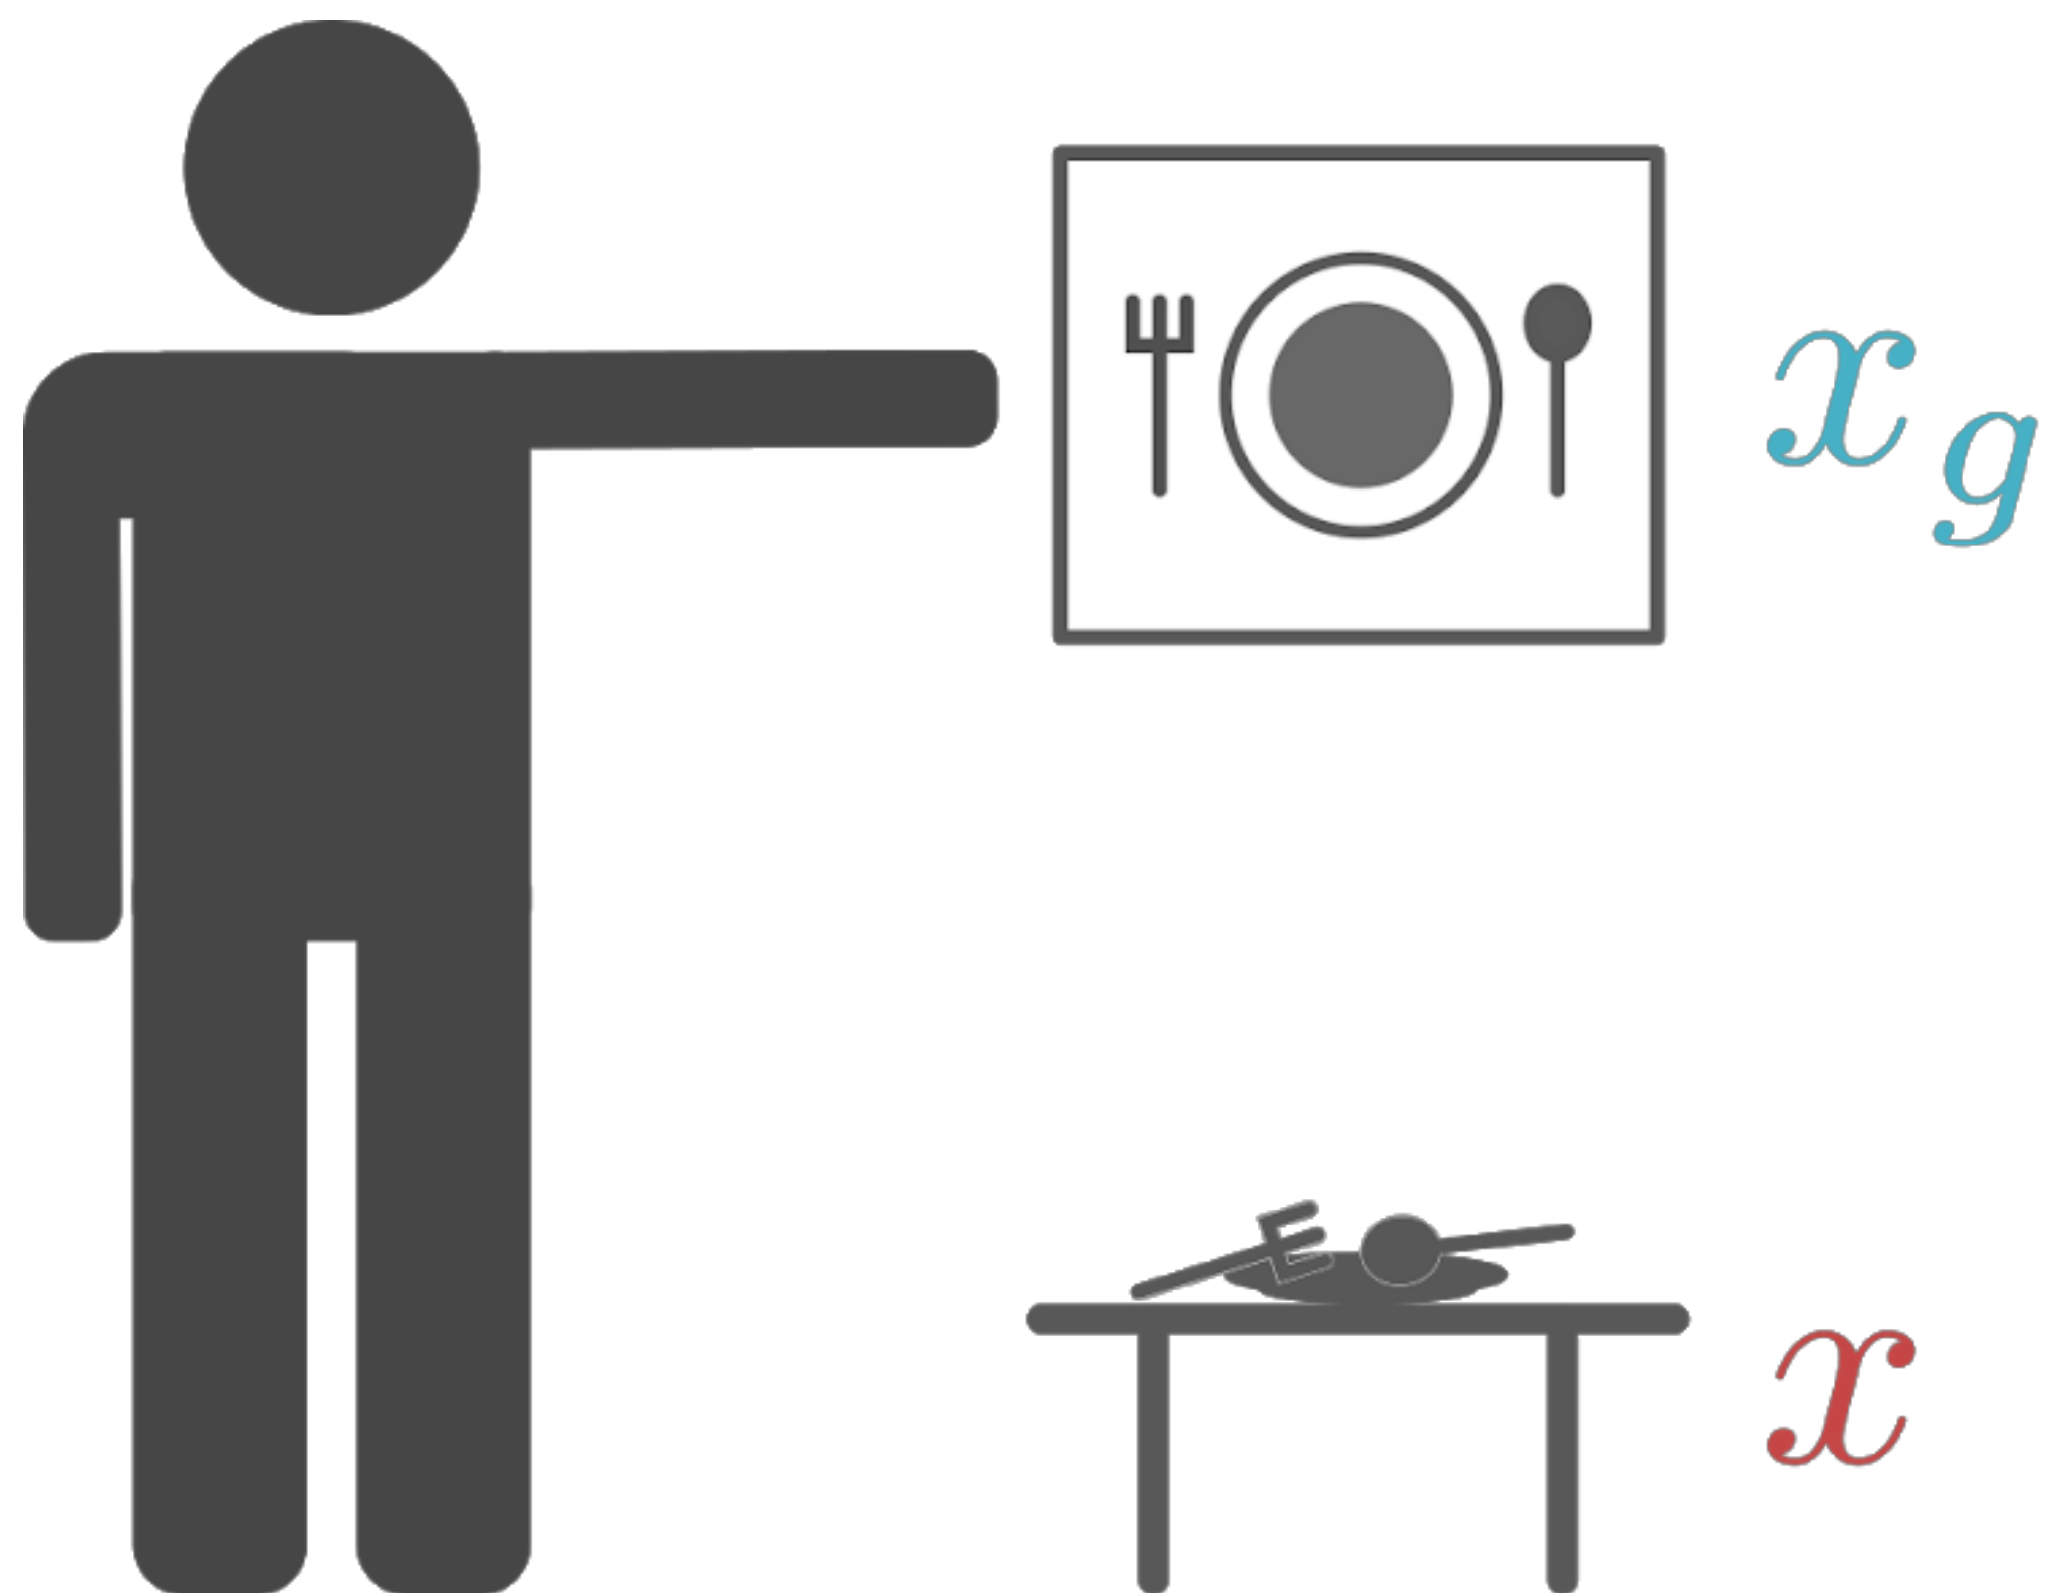
# RIG: Autonomous Learning



# RIG: Autonomous Learning

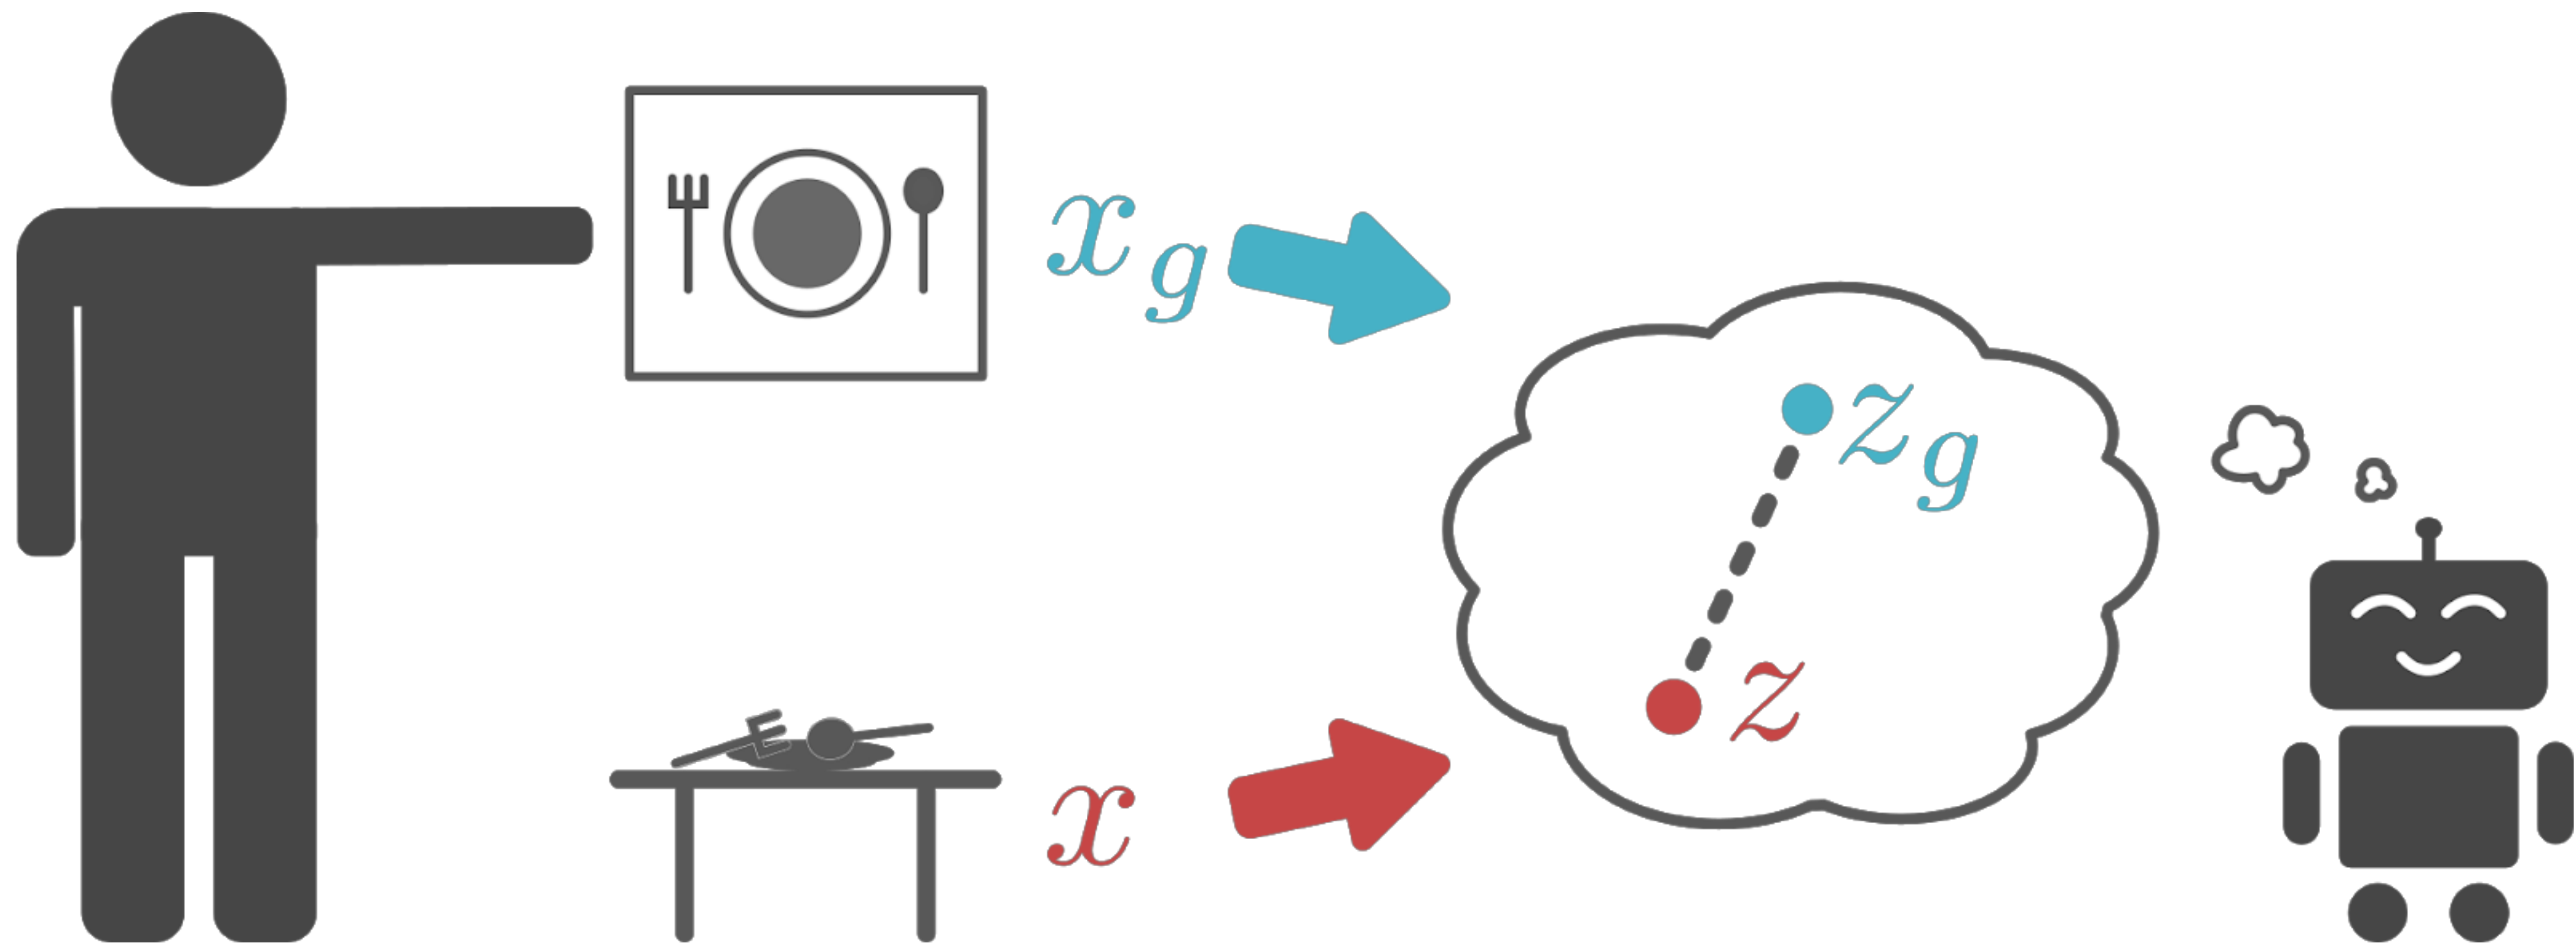


# RIG: Test Time









# RIG: Test Time







# Simulation: Autonomous Exploration













# Simulation: Autonomous Exploration



# Simulation: Test Time




# Simulation: Other Environments




State

$\mathcal{X}$




Human-specified Goal

$\mathcal{X}_g$



# Real-world Experiments

- No human supervision
- Directly from vision
- 0.5-2.5 Hours




State

$\mathcal{X}$



Human-specified Goal

$\mathcal{X}_g$



# Reinforcement Learning with Imagined Goals: Conclusion

Environments: <https://github.com/vitchyr/multiworld/>

Algorithm: <https://github.com/vitchyr/rllkit>

Website: <https://sites.google.com/site/visualrlwithimaginedgoals/>

Poster: #141

