




TeachMyAgent: a Benchmark for Automatic Curriculum Learning in Deep RL

Clément Romac^{1*}, Rémy Portelas^{1*}, Katja Hofmann², Pierre-Yves Oudeyer¹

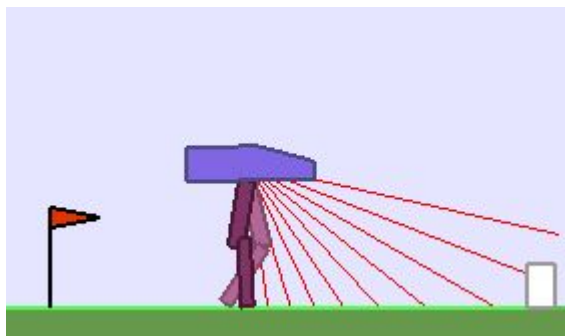
Designing Deep RL agents that can generalise

Single task proposed




Designing Deep RL agents that can generalise

Single task proposed




Multiple tasks using *Procedural Content Generation* (PCG)




Designing Deep RL agents that can generalise

Single task proposed

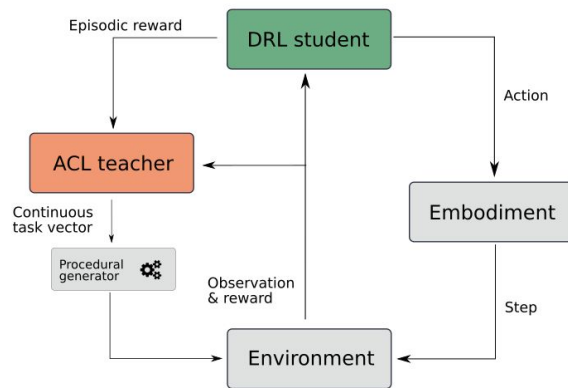


Multiple tasks using *Procedural Content Generation* (PCG)



⇒ How can we design an efficient curriculum ?

Existing ACL methods





Existing ACL methods

ALP-GMM (Portelas et al., 2019)

Comparison:

- Covar-GMM (Moulin-Frier et al., 2014)
- RIAC (Baranes & Oudeyer, 2009)

Environments:



— Maximize Learning Progress


Existing ACL methods

ALP-GMM (Portelas et al., 2019)

Comparison:

- Covar-GMM (Moulin-Frier et al., 2014)
- RIAC (Baranes & Oudeyer, 2009)

Environments:




GoalGAN (Florensa et al., 2018)

Comparison:

- RIAC (Baranes & Oudeyer, 2009)

Environments:




Setter-Solver (Racanière et al., 2020)

Comparison:

- GoalGAN

Environments:



— Maximize Learning Progress

— Generate goals using Neural Networks


Existing ACL methods

ALP-GMM (Portelas et al., 2019)

Comparison:

- Covar-GMM (Moulin-Frier et al., 2014)
- RIAC (Baranes & Oudeyer, 2009)

Environments:




GoalGAN (Florensa et al., 2018)

Comparison:

- RIAC (Baranes & Oudeyer, 2009)

Environments:



Setter-Solver (Racanière et al., 2020)

Comparison:

- GoalGAN

Environments:



SPDL (Klink et al., 2020)

Comparison:

- ALP-GMM
- GoalGAN

Environments:

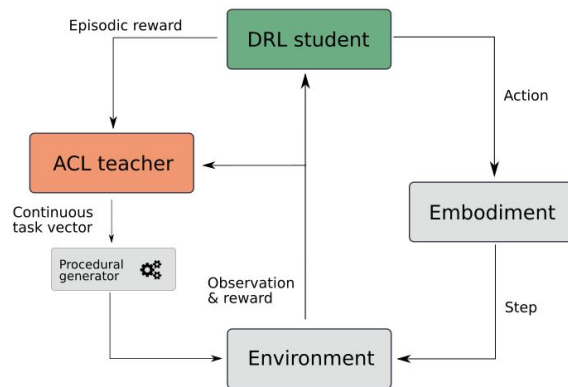


ADR (OpenAI et al., 2019)

Comparisons:

- /

Environments:




— Maximize Learning Progress

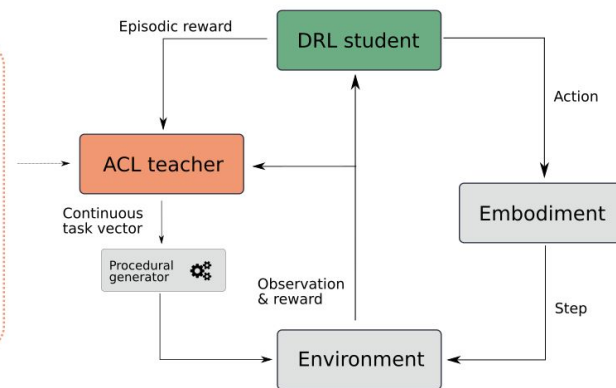

— Generate goals using Neural Networks

— Shift a task distribution according to student's capabilities


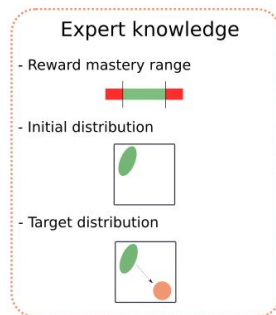
Diverse expert knowledge used



Diverse expert knowledge used



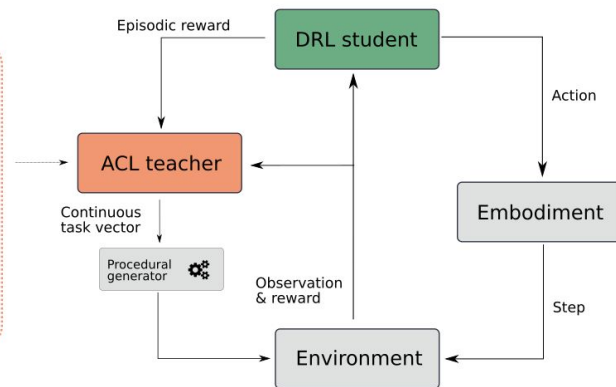
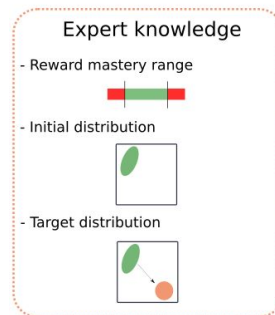
Diverse expert knowledge used



ALGORITHM	\mathcal{D}_{init}	\mathcal{D}_{target}	REWARD MASTERY RANGE
ADR	REQ.		REQ.
ALP-GMM	OPT.		
COVAR-GMM	OPT.		
GOAL-GAN	OPT.		REQ.
RIAC			
SPDL	REQ.	REQ.	
SETTER-SOLVER		OPT.	REQ.

REQ.: Required
OPT.: Optional

Diverse expert knowledge used



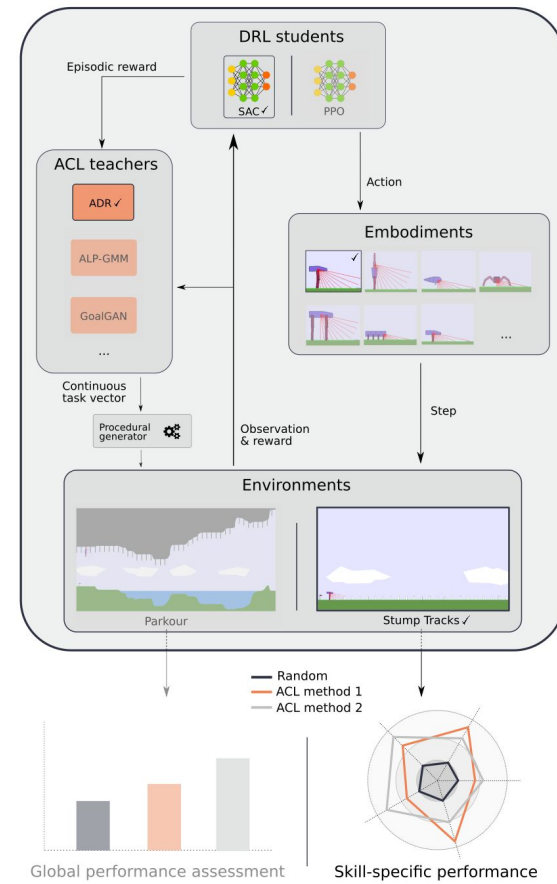
ALGORITHM	\mathcal{D}_{init}	\mathcal{D}_{target}	REWARD MASTERY RANGE
ADR	REQ.		REQ.
ALP-GMM	OPT.		
COVAR-GMM	OPT.		
GOAL-GAN	OPT.		REQ.
RIAC			
SPDL	REQ.	REQ.	
SETTER-SOLVER		OPT.	REQ.

REQ.: Required
OPT.: Optional

⇒ No existing systematic comparison

TeachMyAgent

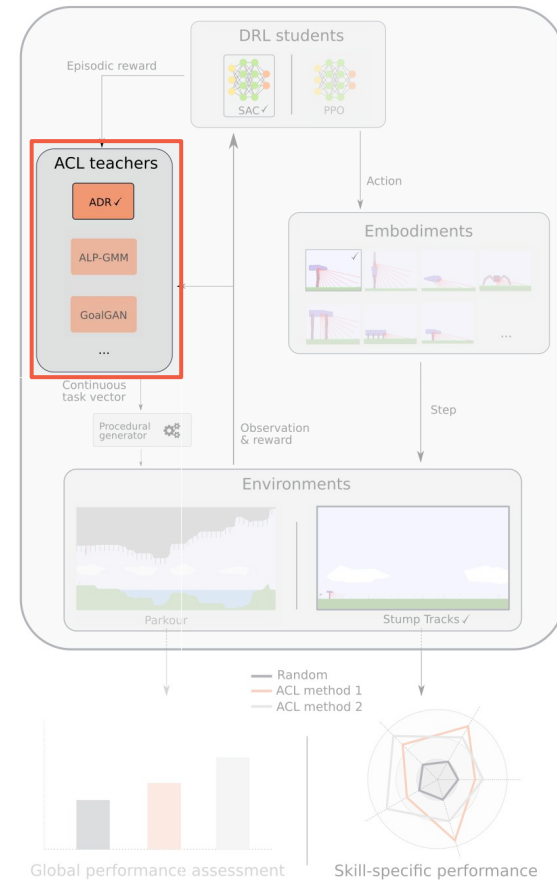
ACL methods | DRL students | PCG Environments | Embodiments | Two experiment types



TeachMyAgent

ACL methods | DRL students | PCG Environments | Embodiments | Two experiment types

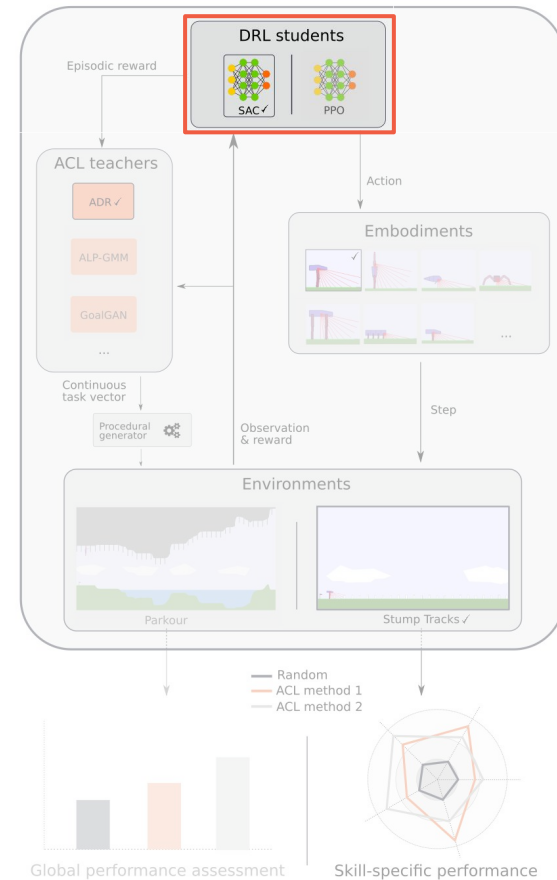
- Multiple **SOTA ACL algorithms**, in a **common API**
 - Taken from open-source code or **re-implemented**



TeachMyAgent

ACL methods | **DRL students** | PCG Environments | Embodiments | Two experiment types

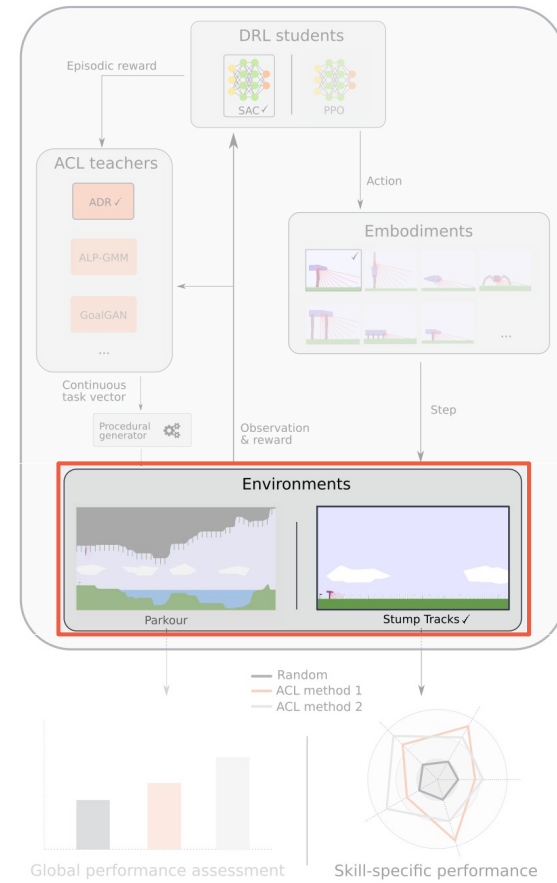
- Multiple **SOTA ACL algorithms**, in a **common API**
 - Taken from open-source code or **re-implemented**
- Two **SOTA Deep RL students**



TeachMyAgent

ACL methods | DRL students | **PCG Environments** | Embodiments | Two experiment types

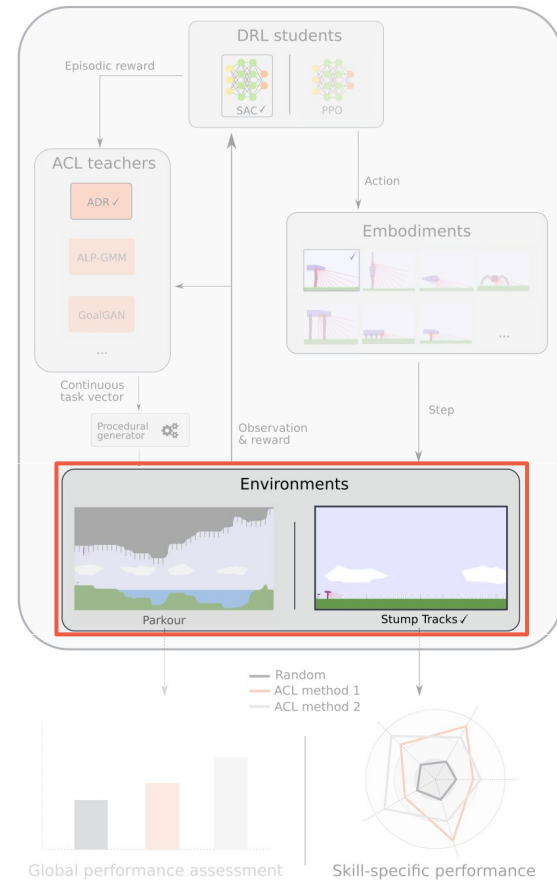

- Multiple **SOTA ACL algorithms**, in a **common API**
 - Taken from open-source code or **re-implemented**
- Two **SOTA Deep RL students**
- **PCG** environments



TeachMyAgent

ACL methods | DRL students | **PCG Environments** | Embodiments | Two experiment types


Stump Tracks (Simple 2D task space)



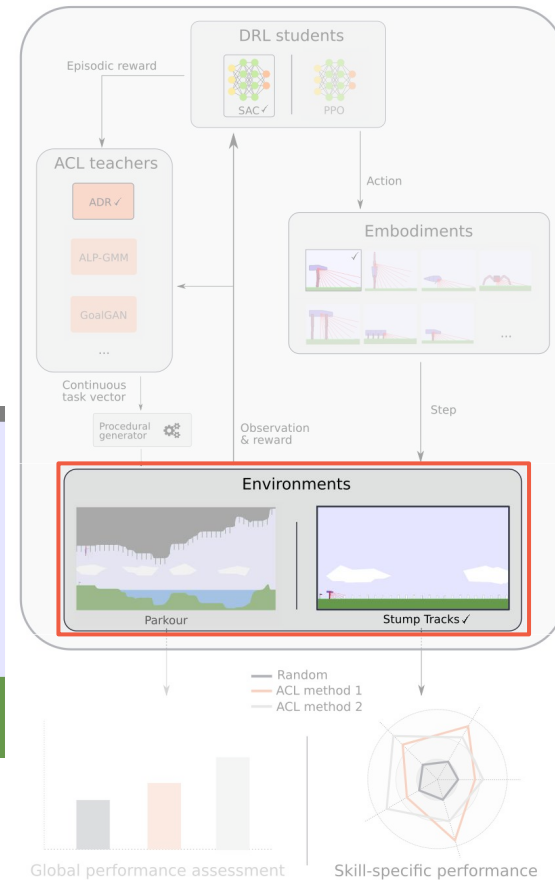

TeachMyAgent

ACL methods | DRL students | **PCG Environments** | Embodiments | Two experiment types

Stump Tracks
(Simple 2D task space)



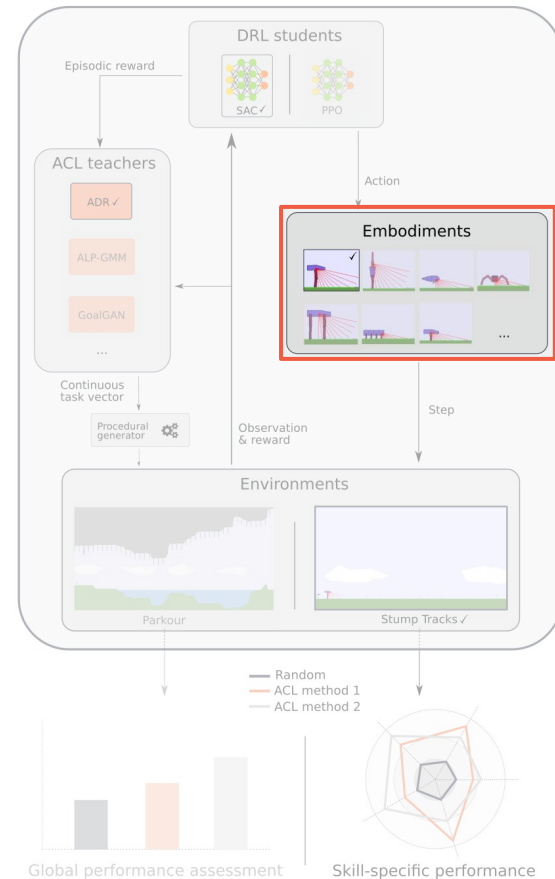
Parkour
(Complex 6D task space)



TeachMyAgent

ACL methods | DRL students | PCG Environments | **Embodiments** | Two experiment types

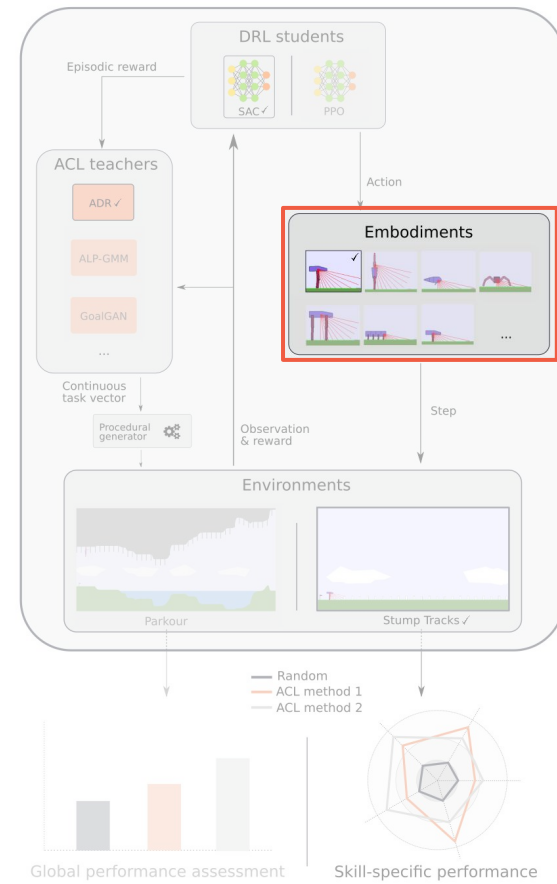
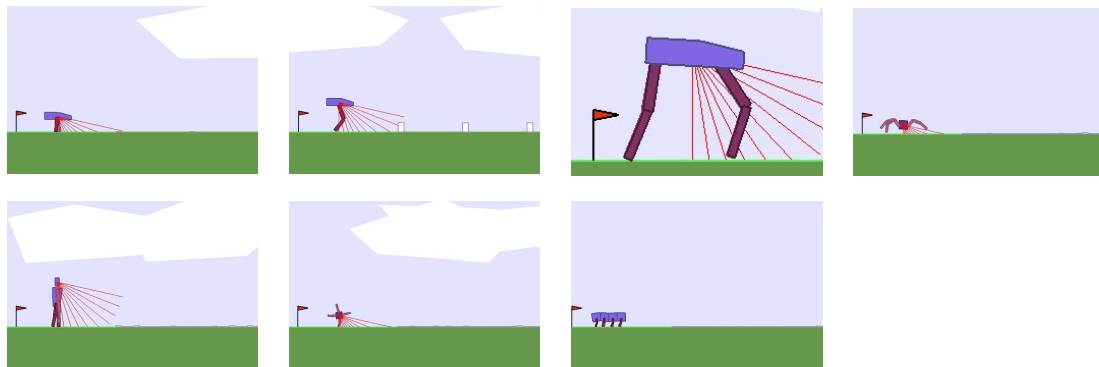
- Multiple **SOTA ACL algorithms**, in a **common API**
 - Taken from open-source code or **re-implemented**
- Two **SOTA Deep RL students**
- **PCG** environments
 - Stump Tracks (simple 2D task space)
 - Parkour (complex 6D task space using a CPPN)
- Multiple **embodiments**



TeachMyAgent

ACL methods | DRL students | PCG Environments | **Embodiments** | Two experiment types


Walkers



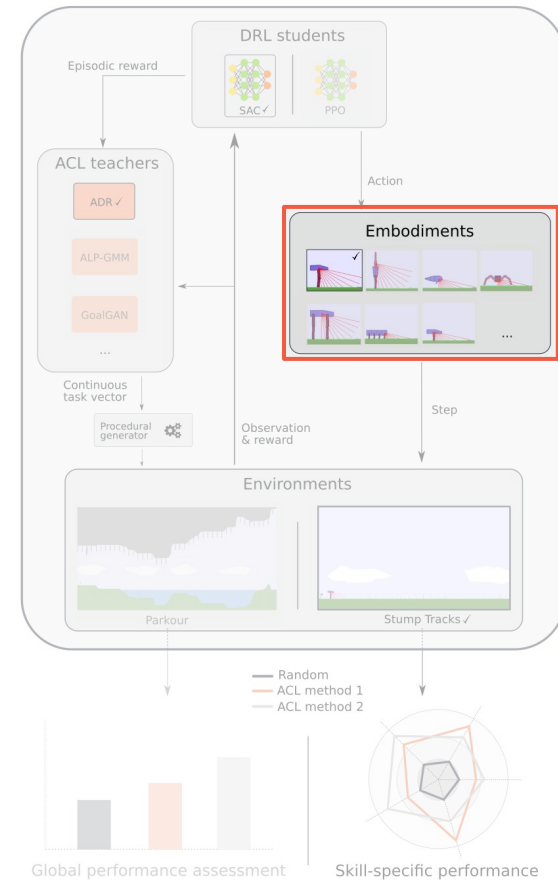

TeachMyAgent

ACL methods | DRL students | PCG Environments | **Embodiments** | Two experiment types

Walkers




Climbers




TeachMyAgent

ACL methods | DRL students | PCG Environments | **Embodiments** | Two experiment types

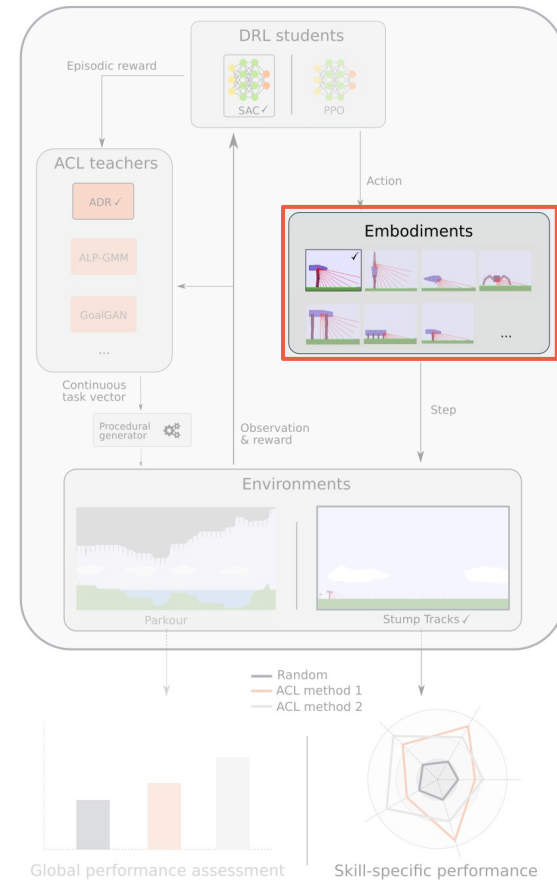

Walkers



Climbers



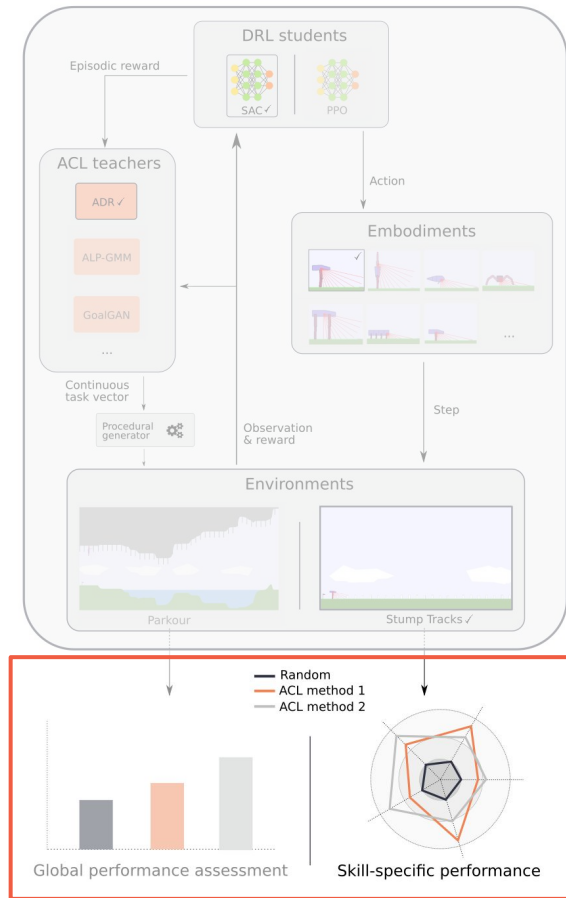
Swimmers



TeachMyAgent

ACL methods | DRL students | PCG Environments | Embodiments | **Two experiment types**


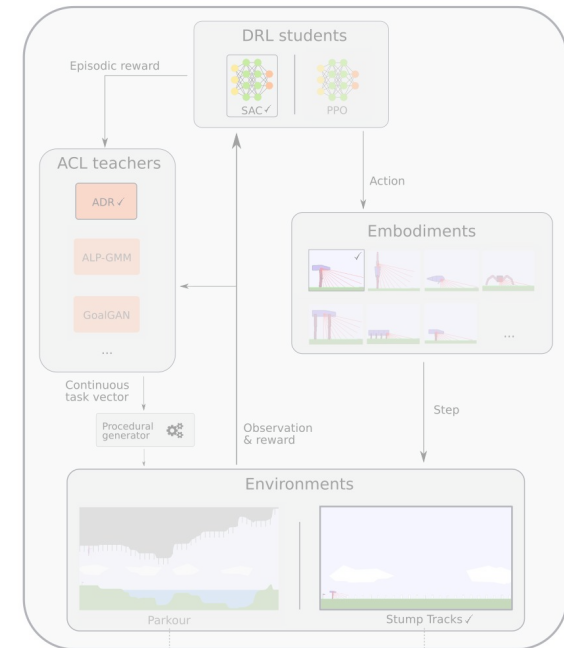
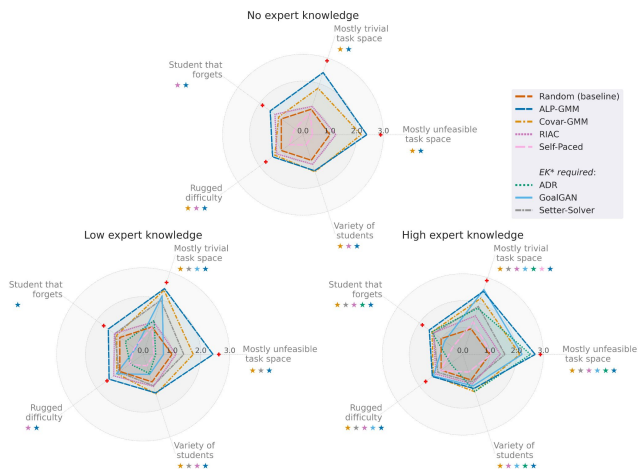
- Multiple **SOTA ACL algorithms**, in a **common API**
 - Taken from open-source code or **re-implemented**
- Two **SOTA Deep RL students**
- **PCG** environments
 - Stump Tracks (simple 2D task space)
 - Parkour (complex 6D task space using a CPPN)
- Multiple **embodiments**
 - Walkers / Climbers / Swimmers
- Two types of experiment



TeachMyAgent

ACL methods | DRL students | PCG Environments | Embodiments | **Two experiment types**

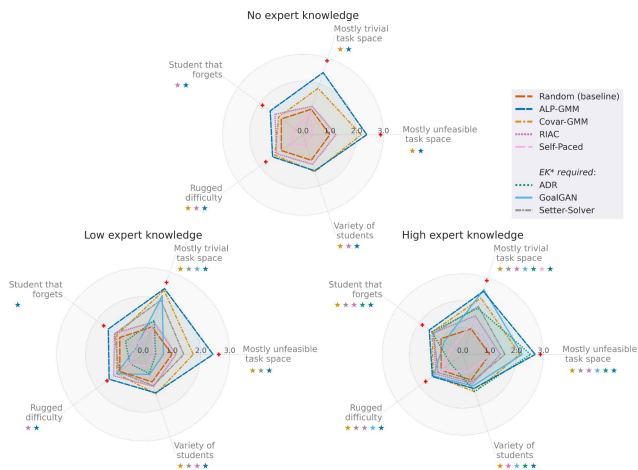
Skill-specific (using Stump Tracks)



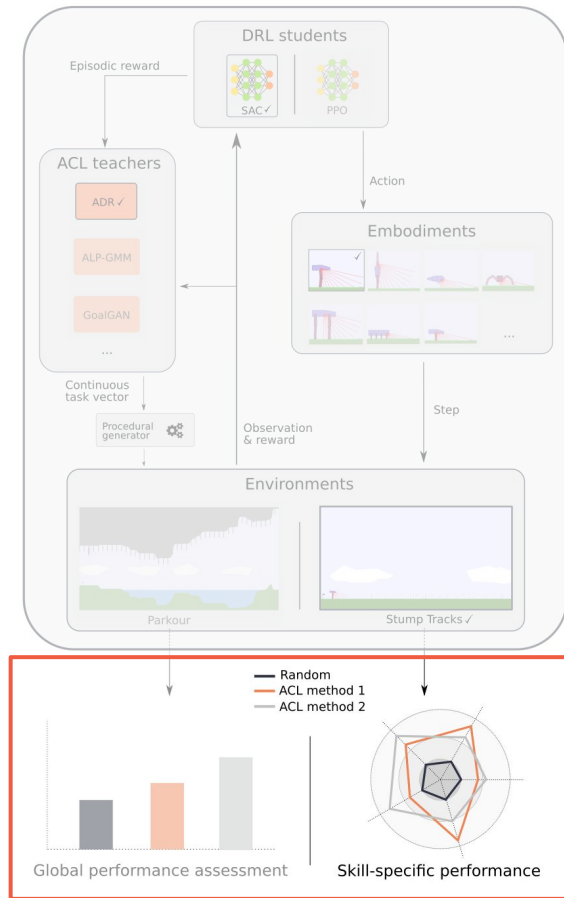

TeachMyAgent

ACL methods | DRL students | PCG Environments | Embodiments | **Two experiment types**

Skill-specific (using Stump Tracks)



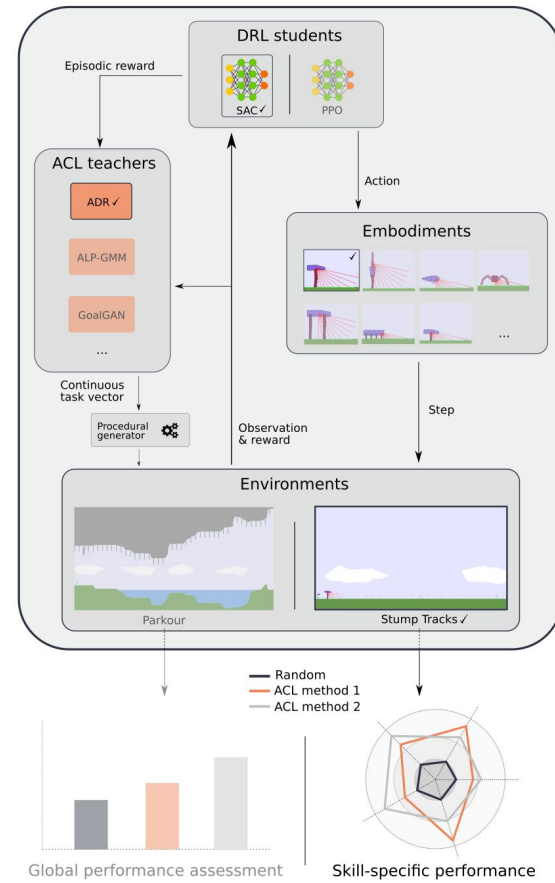
Global performance (using Parkour)



TeachMyAgent

ACL methods | DRL students | PCG Environments | Embodiments | Two experiment types

- Multiple **SOTA ACL algorithms**, in a **common API**
 - Taken from open-source code or **re-implemented**
- Two **SOTA Deep RL students**
- **PCG** environments
 - Stump Tracks (simple 2D task space)
 - Parkour (complex 6D task space using a CPPN)
- Multiple **embodiments**
 - Walkers / Climbers / Swimmers
- Two types of experiment
 - **Skill-specific**
 - **Global performance**




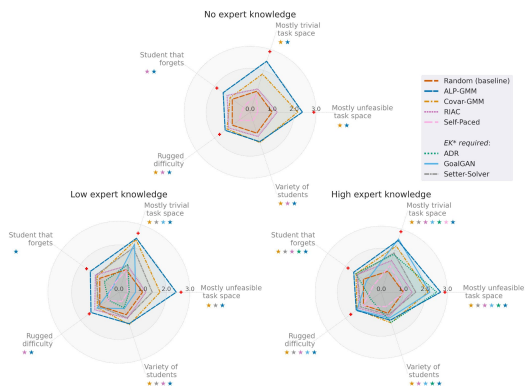
TeachMyAgent

Large-scale experiments

TeachMyAgent

Large-scale experiments

+ Benchmark comparison


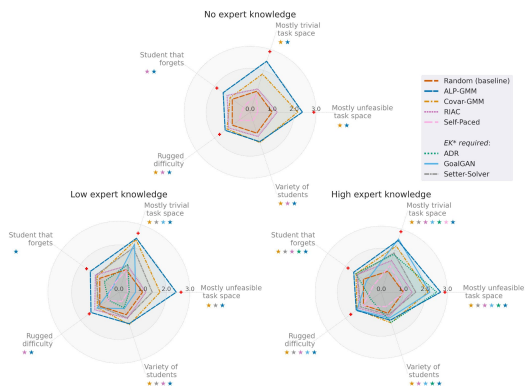


TeachMyAgent

Large-scale experiments

+ Benchmark comparison

+ Baselines release



TeachMyAgent

Large-scale experiments

Open Source release

+ Benchmark comparison

+ Baselines release




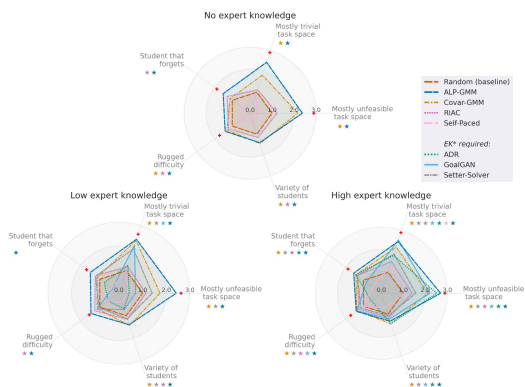
TeachMyAgent

Large-scale experiments

Open Source release

+ Benchmark comparison

+ Baselines release



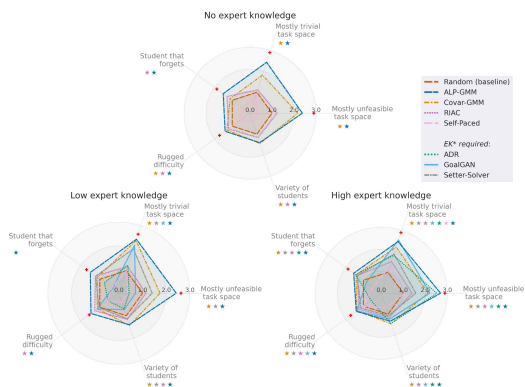
+ Our (well-documented) full platform

TeachMyAgent

Large-scale experiments

+ Benchmark comparison

+ Baselines release



Open Source release

+ Our (well-documented) full platform


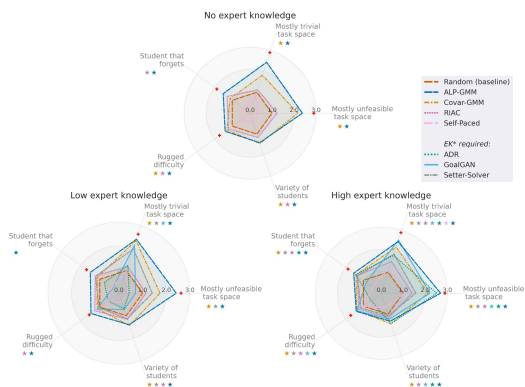
+ Simple APIs to extend it

TeachMyAgent

Large-scale experiments

+ Benchmark comparison

+ Baselines release




Open Source release

+ Our (well-documented) full platform


+ Simple APIs to extend it

+ Jupyter notebooks to reproduce figures


Discussion




Robust walker



Realistic fish




Jumping chimp




Realistic creepers


Discussion




Robust walker



Realistic fish




Jumping chimp




Realistic creepers

TeachMyAgent: a platform for systematic study of ACL methods !


Discussion




Robust walker



Realistic fish



Jumping chimp




Realistic creepers


TeachMyAgent: a platform for systematic study of ACL methods !

Version 1.0 ⇒ More teachers / environments could be added


Discussion




Robust walker



Realistic fish



Jumping chimp



Realistic creepers


TeachMyAgent: a platform for systematic study of ACL methods !

Version 1.0 ⇒ More **teachers** / **environments** could be added


Website: <http://developmentalsystems.org/TeachMyAgent/>

Code: <https://github.com/flowersteam/TeachMyAgent>


Discussion




Robust walker



Realistic fish



Jumping chimp



Realistic creepers

TeachMyAgent: a platform for systematic study of ACL methods !

Version 1.0 ⇒ More **teachers** / **environments** could be added

Website: <http://developmentalsystems.org/TeachMyAgent/>

Code: <https://github.com/flowersteam/TeachMyAgent>

Thank you !