air

AI Redefined

Towards Human-AI Teaming:

Intelligence Ecosystems to Tackle High Stakes Use Cases

Clodéric Mars - VP of Engineering @ AI Redefined

AIs actions don’t align with humans intents because they are not aware of context

AIs don't learn from collaborating with humans

air

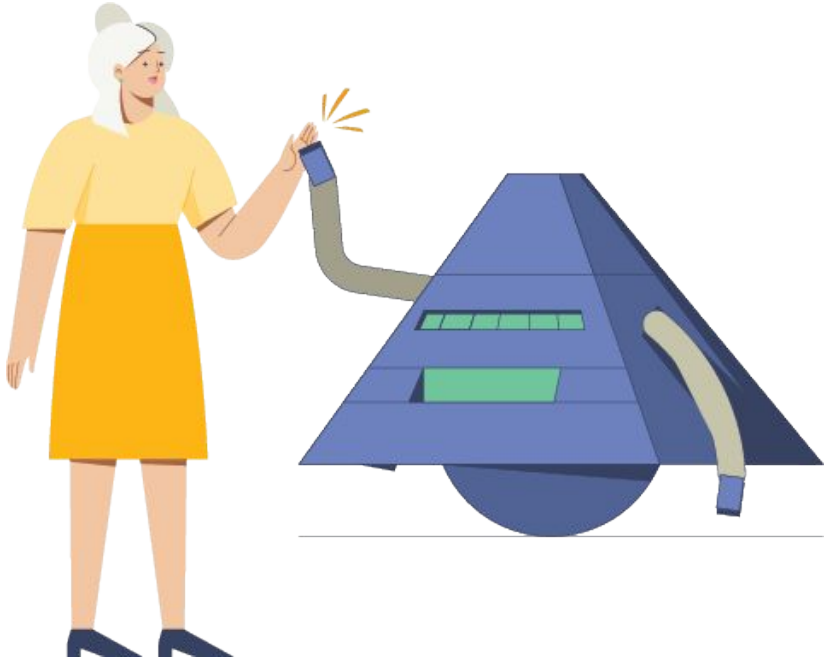

Achieve Human + AI Synergy
Intelligence ecosystems continuous learning through shared experiences

air

Behavior Cloning / Imitation Learning

Humans demonstrating how to achieve a task

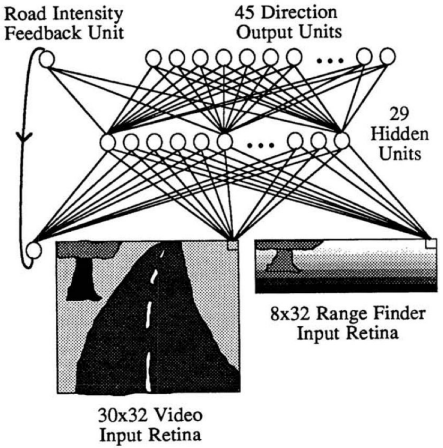

ALVINN: An Autonomous Land Vehicle In A Neural Network
1988 - Pomerlau
https://proceedings.neurips.cc/paper/1988/file/812b4ba287f5ee0bc9d43bbf5bbe87fb-Paper.pdf

- Fast initial training bootstrapping & continuous interactive refinements
- Good alignment
- No additional human skills required
- Bounded performances

air

Reinforcement Learning with Human Feedback

Humans evaluating how a task is completed

Interactive Learning from Policy-Dependent Human Feedback (COACH)
2017 - MacGlashan et al.
https://arxiv.org/abs/1701.06049

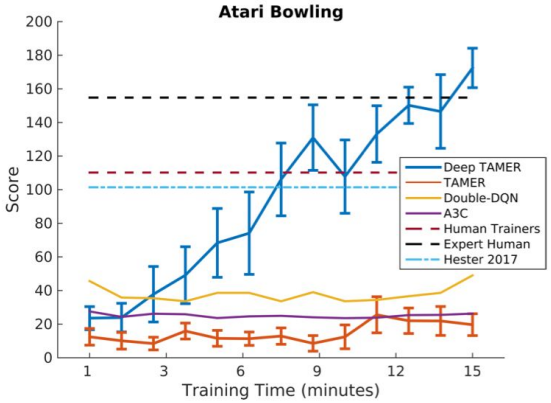

Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces
2018 - Warnell et al.
https://arxiv.org/abs/1709.10163

air

Reinforcement Learning, a very short aside

Discovering instead of reproducing

supervised learning

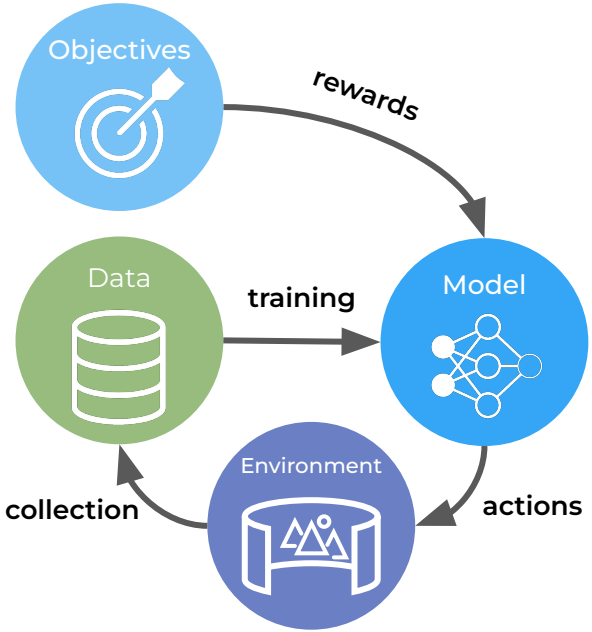

reinforcement learning

air

Reinforcement Learning with Human Feedback

Humans evaluating how a task is completed

Interactive Learning from Policy-Dependent Human Feedback (COACH)
2017 - MacGlashan et al.
https://arxiv.org/abs/1701.06049

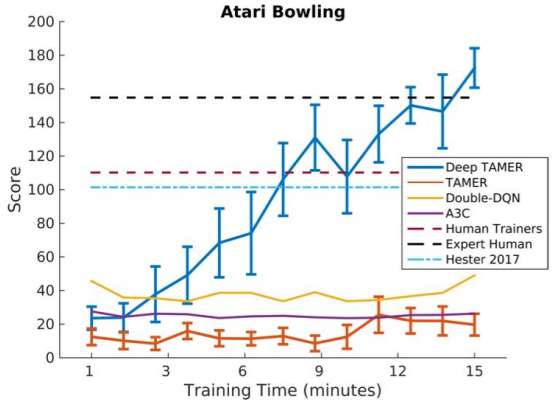

Deep TAMER: Interactive Agent Shaping in High-Dimensional State Spaces
2018 - Warnell et al.
https://arxiv.org/abs/1709.10163

- Superhuman (even optimal) performances
- Indirect alignment
- Some specialized skills required & labor intensive
- Requires the ability to deal with the lag between an action and its evaluation
- Requires safe environment

air

AI apprentice: dual control

Interactive human demonstrations to accelerate exploration

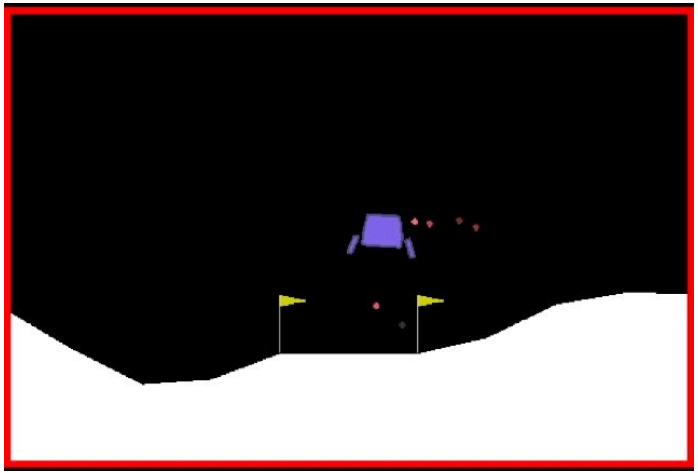

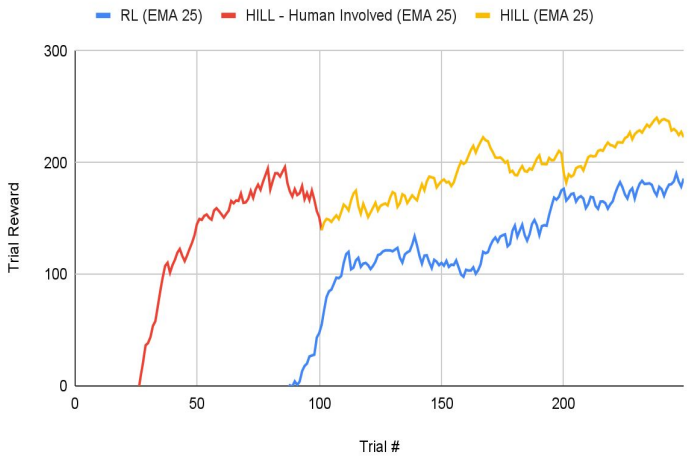

- Superhuman (even optimal) performances
- Indirect alignment
- No additional human skills required
- Requires collaborative UX during training and operation

Powered by

cogment

air

Human-guided summarization

Combining Human-in-the-Loop Learning with language models to improve Human/AI alignment

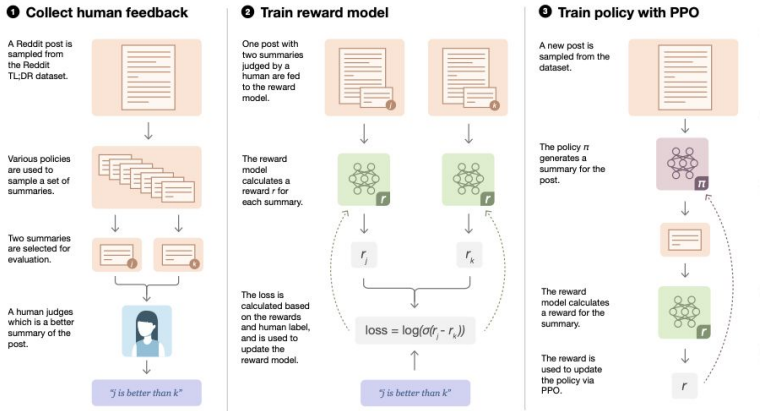

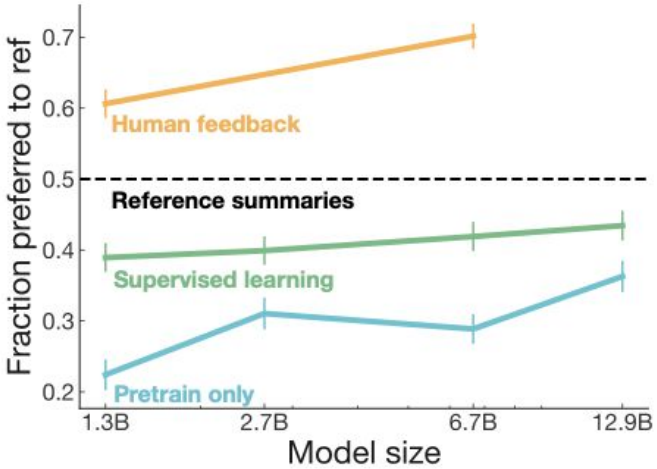

Learning to summarize from human feedback
2020 - Stiennon et al.
https://arxiv.org/abs/2009.01325

- Fine tune offline trained model
- Hybrid technique: demonstration + feedback
- Leverage human expertise on language task
- Improved alignment

air

Man-Machine Teaming: Multiple AIs and Humans collaborating to complete a task

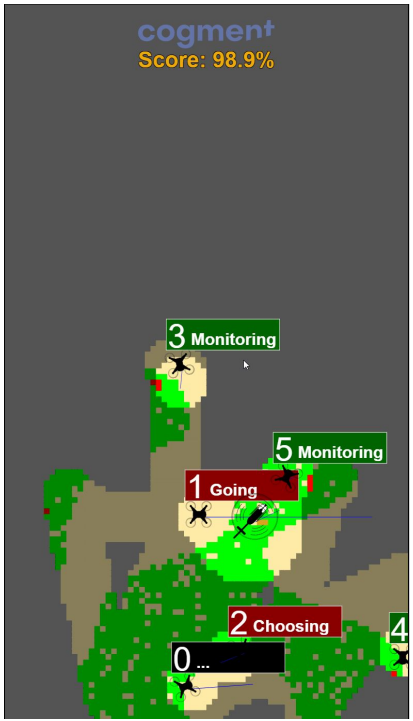

Training AI agents to coordinate together and with humans and vice versa

Powered by cogment

- Implicitly train AIs from their interaction with humans
- Train Humans alongside AIs
- Requires bootstrapped AIs
- Enables human supervision

air

Intelligence Ecosystem: Heterogeneous actors collaborating

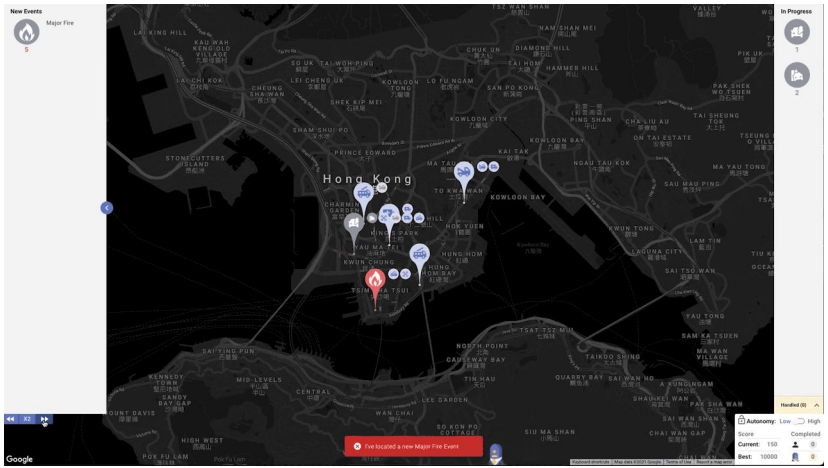

Training and operating complex topologies of roles and tasks within a common environment

- Faster deployment and iterations
- Enable supervision & learning by keeping humans in the loop
- Compliance & accountability where it matters

air

Applications to other verticals

Utilities

Education

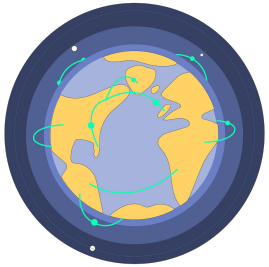

Logistics
Manufacturing

Transportation

Health

air

Intelligence Ecosystem: Benefits

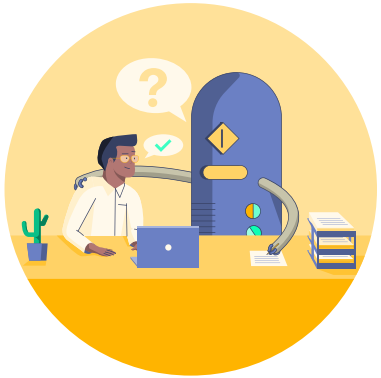

Leverage human expertise: training, supervision, collaboration

Deploy faster & continuously get better

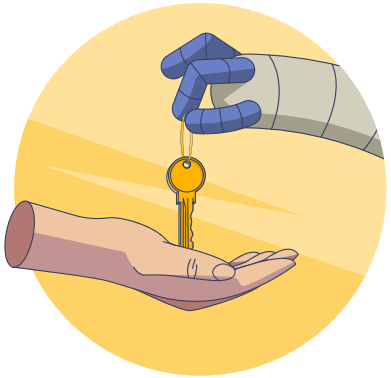

Trust, compliance & accountability

air

Intelligence Ecosystem: Requirements

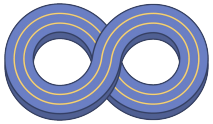

Continuous learning from building to operation

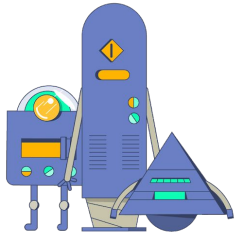

Multiple AI techniques

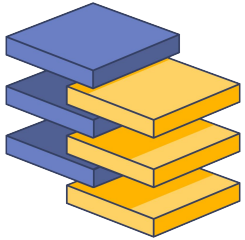

Tech agnostic

air

Cogment: Build, train, and operate AI agents in simulated or real environments shared with humans

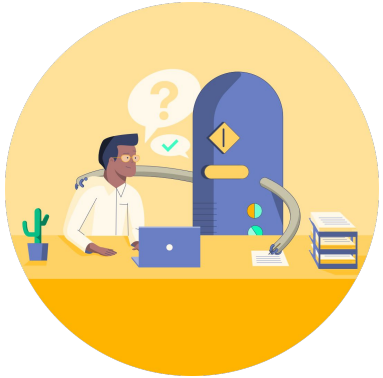

Continuously train AIs & Humans together

Less data required
Real time adaptation
Faster training
Fostering trust

Operate intelligence ecosystems

Best of human & AI capabilities
Human supervision when it matters
Hybrid AI: compliance and high performance
Modular approach: reduce compute usage & facilitate validation

Iterate smoothly from sim to real

Safe and simple design and training in simulation
Progressive deployment to real environment
Real environments, digital twins, numerical simulations, etc.

Available open-source & with further information at https://cogment.ai

air

Lessons learned designing AI-enabled products

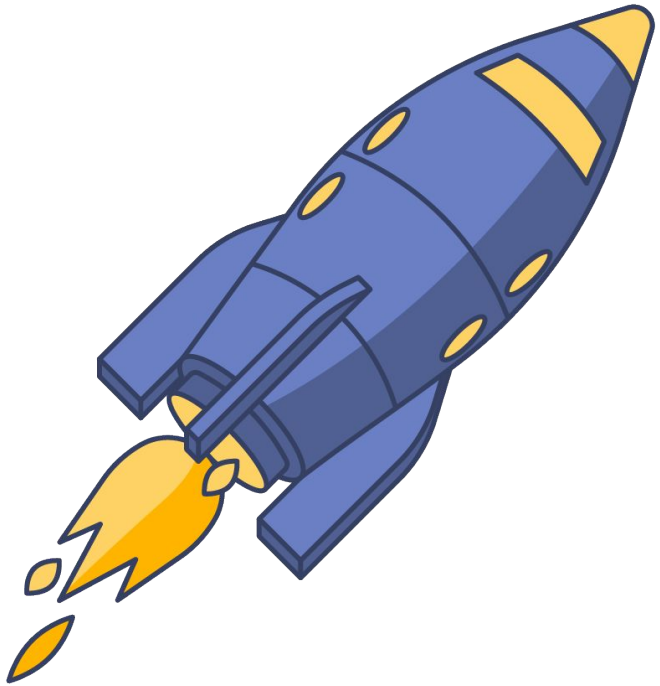

- Take into account the Human/AI relationship
- Consider AI apprenticeship approach
- Think in terms of intelligence ecosystem

air

AI Redefined: Humans and AI elevating each other

P.S. We are hiring!

Clodéric Mars

cloderic@ai-r.com
@cloderic

ai-r.com
cogment.ai
github.com/cogment