The University of Texas at Austin
Cockrell School of Engineering

TEXAS
The University of Texas at Austin

SIDI Lab

Walker Department of Mechanical Engineering

Cost-Aware Bayesian Agents for Human-AI Teaming

Siyu Chen[1], Dr. Zhenghui Sha[2]

Presenter: Siyu Chen
October 12, 2023

[1] PhD student, Walker Department of Mechanical Engineering, The University of Texas at Austin
[2] Assistant Professor, Walker Department of Mechanical Engineering, The University of Texas at Austin

Outline

Background

- **Unknown design space exploration**: find the best design solution step by step

Background

- **Unknown design space exploration**: find the best design solution step by step $\rightarrow$ sequential decision-making process

Background

- **Unknown design space exploration**: find the best design solution step by step → sequential decision-making process
- How to model this process?

Background

- **Unknown design space exploration**: find the best design solution step by step $\rightarrow$ sequential decision-making process
- How to model this process?
  - Unknown design space: a black-box function.

Background

- **Unknown design space exploration**: find the best design solution step by step → sequential decision-making process
- How to model this process?
  - Unknown design space: a black-box function.

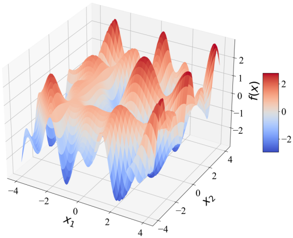

Figure: Black-box objective function

Background

- **Unknown design space exploration**: find the best design solution step by step → sequential decision-making process
- How to model this process?
  - Unknown design space: a black-box function.
  - Unknown design space exploration: find the optimum by sequential sampling in this space.

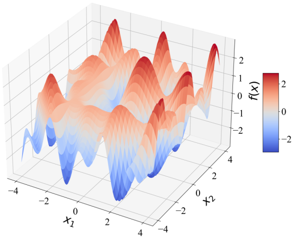

Figure: Black-box objective function

Background

- **Unknown design space exploration**: find the best design solution step by step → sequential decision-making process
- How to model this process?
  - Unknown design space: a black-box function.
  - Unknown design space exploration: find the optimum by sequential sampling in this space.

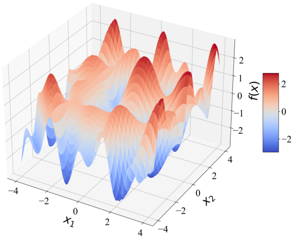

Figure: Black-box objective function

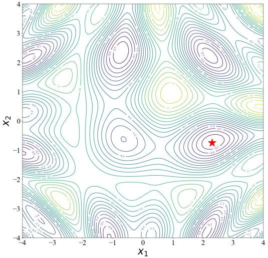

Figure: Contour plot

Motivation

- **Single agent**:

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:
  - mimics the human decision-making process in a design team.

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:
  - mimics the human decision-making process in a design team.
- **Challenges**:

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:
  - mimics the human decision-making process in a design team.
- **Challenges**:
  - In the design team, decisions made by one member could influence others $\rightarrow$ collaboration and communication;

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:
  - mimics the human decision-making process in a design team.
- **Challenges**:
  - In the design team, decisions made by one member could influence others → collaboration and communication;
  - Humans are cost-sensitive: budget, time, and resources.

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:
  - mimics the human decision-making process in a design team.
- **Challenges**:
  - In the design team, decisions made by one member could influence others $\rightarrow$ collaboration and communication;
  - Humans are cost-sensitive: budget, time, and resources.

$\downarrow$

Motivation

- **Single agent**:
  - solve design space exploration using Bayesian Optimization.
    - For *complex* design space, team effort is needed.
- **Multi-agent System (MAS)**:
  - mimics the human decision-making process in a design team.
- **Challenges**:
  - In the design team, decisions made by one member could influence others $\rightarrow$ collaboration and communication;
  - Humans are cost-sensitive: budget, time, and resources.

$\downarrow$

- Develop a cost-aware MAS with collaboration and communication

Objective

- Develop a cost-aware Multi-agent System (MAS) with collaboration and communication based on Bayesian Optimization (BO) to model the sequential decision-making process of a design team in the exploration of complex design spaces.
  - Decision 1: where to sample next;
  - Decision 2: when to stop.

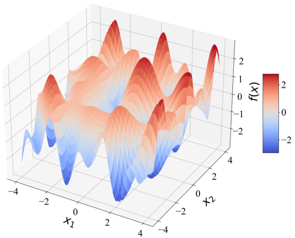

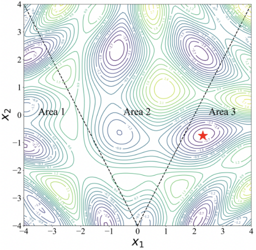

Research Questions (RQs)

- **RQ1**: How can the local-global communication influence the search performance (convergence speed) for the MAS in the varying scenarios considering
  - the complexity of the objective function;
  - the MAS team size?

Research Questions (RQs)

- **RQ1**: How can the local-global communication influence the search performance (convergence speed) for the MAS in the varying scenarios considering
  - the complexity of the objective function;
  - the MAS team size?

- **RQ2**: What impact would the cost-aware stopping criteria have on the search behavior of MAS?

Problem Setups

Problem formulation

Consider a MAS consisting of N agents in a $2D$ design space domain A. The goal of agent is to find the location of the global optimum of a black-box function:

$$\mathbf{x}^* = argmin_{\mathbf{x} \in A} f(\mathbf{x}) \tag{1}$$

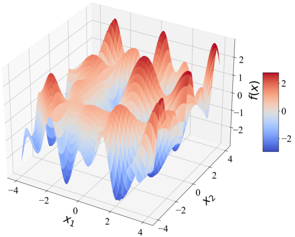

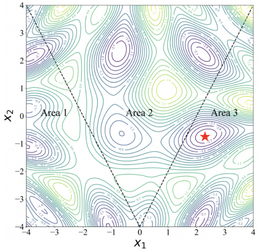

Bayesian Optimization (BO)

Decision 1: where to sample next.

Bayesian Optimization (BO)

Decision 1: where to sample next.

- **Gaussian Process**[a]:
  model the unknown objective function.
- **Acquisition function**:
  determine the next point to sample in the design space.
  - Expected Improvement (EI)
  - Lower Confidence Bound (LCB)[b]

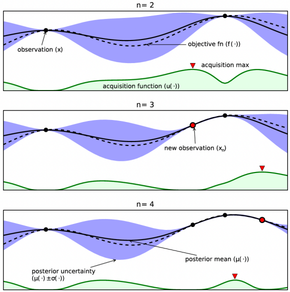

Figure: Illustration of BO procedure over three iterations.

[a] Rasmussen 2003.

[b] Snoek, Larochelle, and Adams 2012.

Cost-aware stopping criterion

Decision 2: when to stop

Cost-aware stopping criterion

$$U = G - K * c, \tag{2}$$

where $G = \sum_{k=0}^{K}(\alpha * PG + \beta * IG)$, IG is Information Gain, PG is Performance Gain, K is the iteration number, c is the cost for each search.

- **Performance Gain (PG)**: the gain already achieved, $PG = f_k^* - f_{k-1}^*$.
- **Information Gain (IG)**: the potential gain can be achieved in the future, value of the acquisition function.
- **Cost-setting strategy**: Different cost for each agent.

Cost-aware Multi-agent Bayesian Optimization

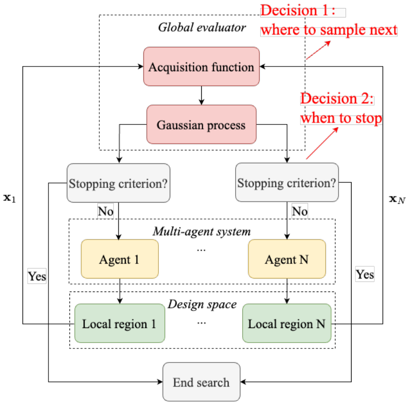

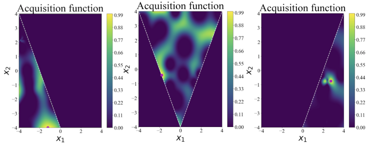

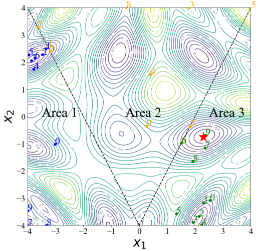

Cost-aware Multi-agent Bayesian Optimization

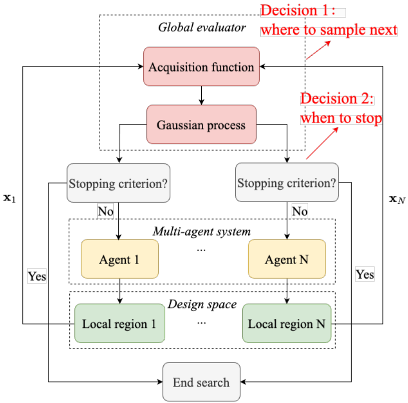

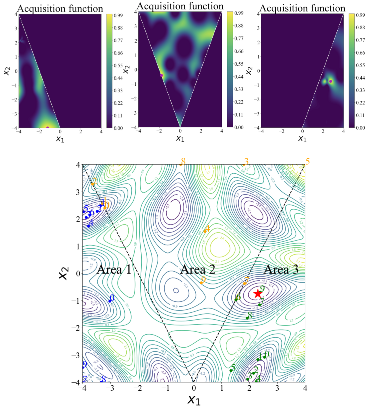

- Division of design space

Cost-aware Multi-agent Bayesian Optimization

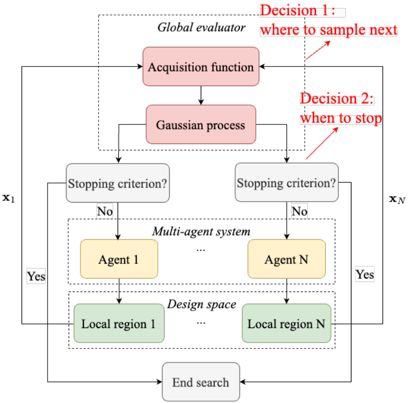

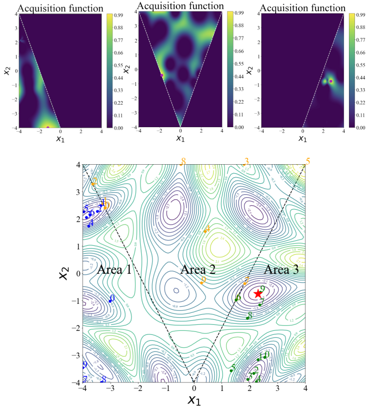

- Division of design space
- Local-global communication

Cost-aware Multi-agent Bayesian Optimization

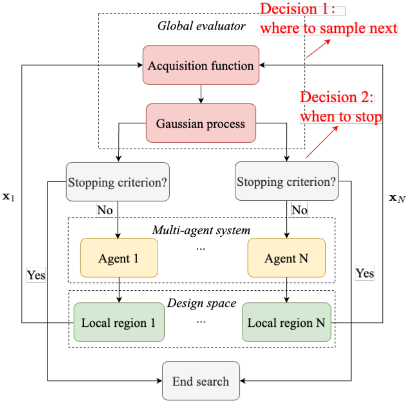

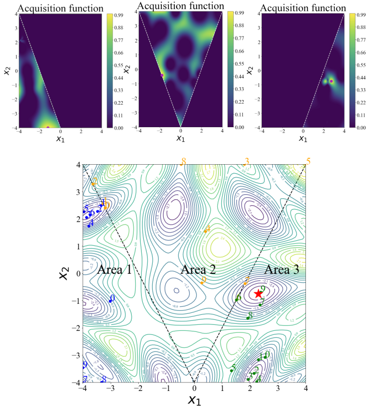

- Division of design space
- Local-global communication
- Decision-making

Cost-aware Multi-agent Bayesian Optimization

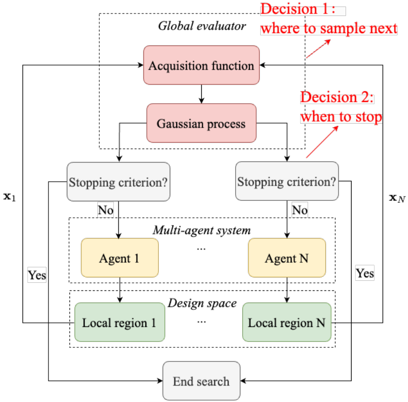

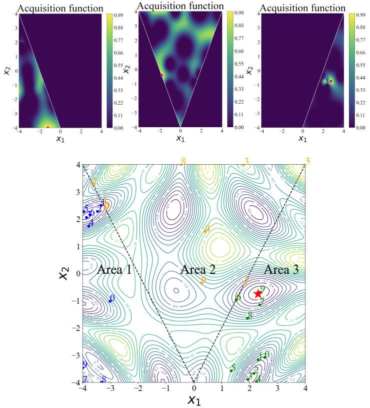

- Division of design space
- Local-global communication
- Decision-making $\rightarrow$ Decision 1: where to sample next;

Cost-aware Multi-agent Bayesian Optimization

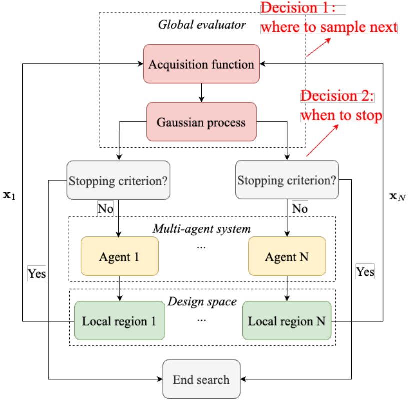

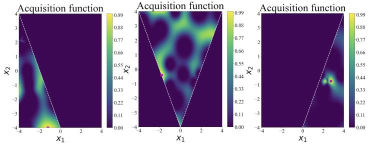

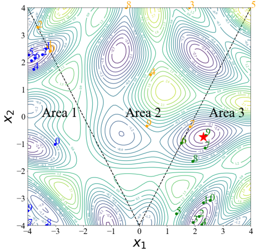

- Division of design space
- Local-global communication
- Decision-making $\rightarrow$ Decision 1: where to sample next; Decision 2: whether to stop

Experimental setup

- **RQ1**: How can local-global communication influence convergence speed?
  - complexity of the objective function
  - MAS team size

Experimental setup

- **RQ1**: How can local-global communication influence convergence speed?
  - complexity of the objective function
  - MAS team size

Method comparison

- **Method 1**: the MABO process without a global evaluator;
- **Method 2**: the proposed MABO with a global evaluator enabled.

Experimental setup

- **RQ1**: How can local-global communication influence convergence speed?
  - complexity of the objective function
  - MAS team size

Method comparison

- **Method 1**: the MABO process without a global evaluator;
- **Method 2**: the proposed MABO with a global evaluator enabled.

Scenarios without stopping criterion

- MABO of the Cosines function with a MAS of three agents;
- MABO of the Eggholder function with a MAS of three agents;
- MABO of the Eggholder function with a MAS of five agents.

Experimental setup

Scenarios without stopping criterion

- MABO of the Cosines function with a MAS of three agents
- MABO of the Eggholder function with a MAS of three agents
- MABO of the Eggholder function with a MAS of five agents

- Objective functions

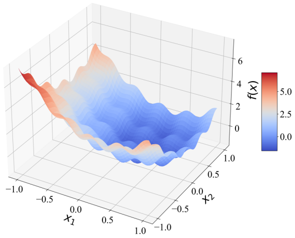

Figure: Cosines function

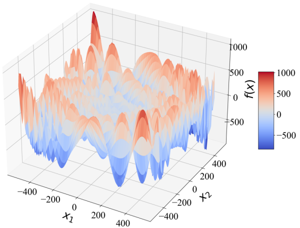

Figure: Eggholder function

Experimental results

Scenarios without stopping criterion

- MABO of the Cosines function with a MAS of three agents

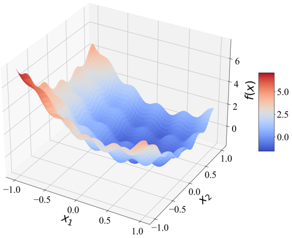

Figure: Cosines function

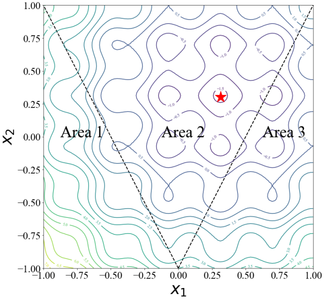

Figure: Space division

Experimental results

Scenarios without stopping criterion

- MABO of the Cosines function with a MAS of three agents

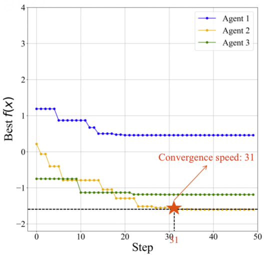

Figure: Convergence speed, Method 1

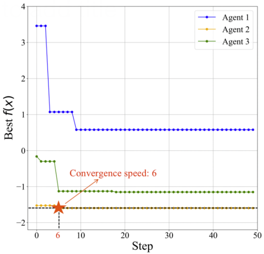

Figure: Convergence speed, Method 2

- Observations:
  – Faster convergence speed to the global optimum
  – Faster convergence speed to local optimum for each agent

Experimental results

Scenarios without stopping criterion

- MABO of the Eggholder function with a MAS of three agents

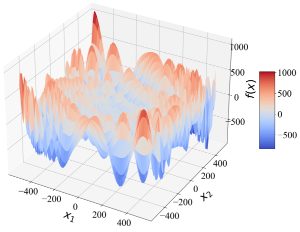

Figure: Eggholder function

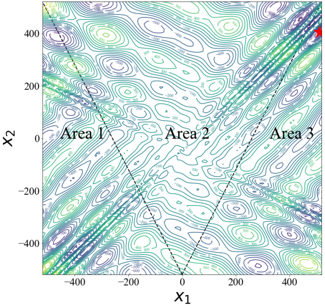

Figure: Space division

Experimental results

Scenarios without stopping criterion

- MABO of the Eggholder function with a MAS of three agents

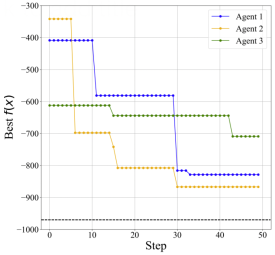

Figure: Convergence speed, Method 1

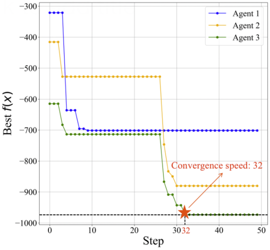

Figure: Convergence speed, Method 2

- Observations:
  - Faster convergence speed to the global optimum
  - Faster convergence speed to local optimum for each agent

Experimental results

Scenarios without stopping criterion

- MABO of the Cosines function with a MAS of three agents
- MABO of the Eggholder function with a MAS of three agents

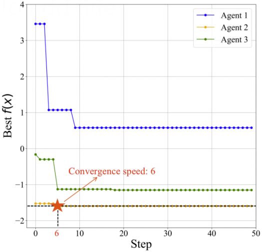

Figure: Convergence speed, Method 2

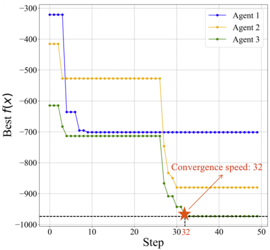

Figure: Convergence speed, Method 2

- Observation:
  - Slower convergence speed when the complexity increases

Experimental results

Scenarios without stopping criterion

- MABO of the Eggholder function with a MAS of five agents

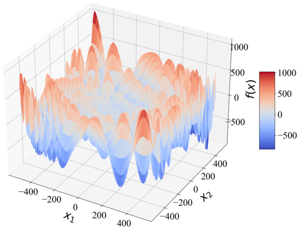

Figure: Eggholder function

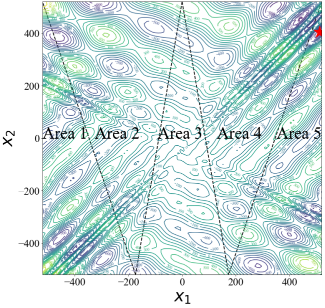

Figure: Space division

Experimental results

Scenarios without stopping criterion

- MABO of the Eggholder function with a MAS of five agents

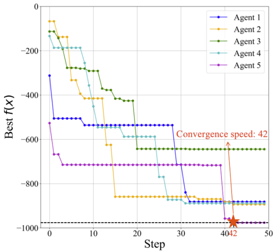

Figure: Convergence speed, Method 1

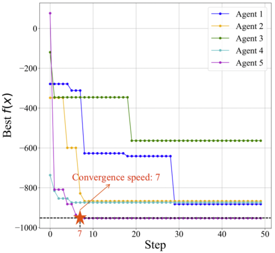

Figure: Convergence speed, Method 2

- Observations:
  - Faster convergence speed to the global optimum
  - Faster convergence speed to local optimum except for Agent 3

Experimental results

Scenarios without stopping criterion

- MABO of the Eggholder function with a MAS of five agents

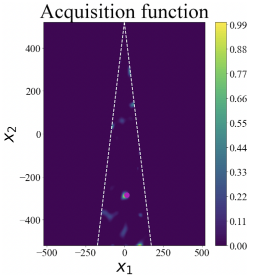

Figure: Acquisition function from Agent 3

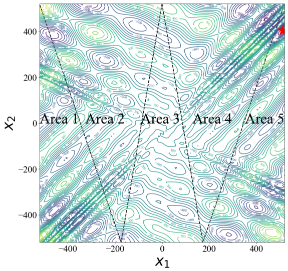

Figure: Space division

- Observations:
  - Faster convergence speed to the global optimum
  - Faster convergence speed to local optimum except for Agent 3

Experimental results

Scenarios without stopping criterion

- MABO of the Eggholder function with a MAS of three agents
- MABO of the Eggholder function with a MAS of five agents

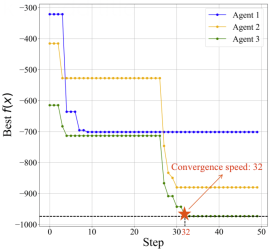

Figure: Convergence speed, Method 2

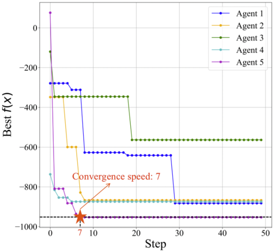

Figure: Convergence speed, Method 2

- Observation:
  - Faster convergence speed to the global optimum when the MAS team size increases

Experimental setup

- **RQ2**: What impact would cost-aware stopping criteria have on the search behavior of the MAS?

Experimental setup

- **RQ2**: What impact would cost-aware stopping criteria have on the search behavior of the MAS?

Cost-aware stopping criterion

$$U = G - K * C, \tag{3}$$

where $G = \sum_{k=0}^{K}(\alpha * PG + \beta * IG)$, IG is Information Gain, PG is Performance Gain, K is the iteration number, C is the cost for each search.

Experimental setup

- **RQ2**: What impact would cost-aware stopping criteria have on the search behavior of the MAS?

Cost-aware stopping criterion

$$U = G - K * C, \tag{3}$$

where $G = \sum_{k=0}^{K}(\alpha * PG + \beta * IG)$, IG is Information Gain, PG is Performance Gain, K is the iteration number, C is the cost for each search.

Scenarios with cost-aware stopping criterion

- MABO of the Cosines function with a MAS of three agents;
- MABO of the Eggholder function with a MAS of three agents.

Experimental results

Scenarios with cost-aware stopping criterion

- MABO of the Cosines function with a MAS of three agents

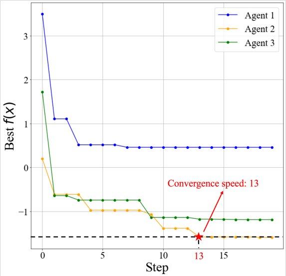

Figure: Convergence speed without stopping criterion

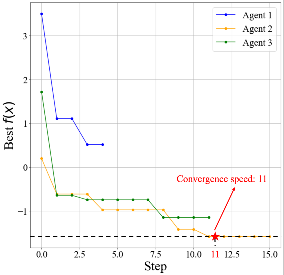

Figure: Convergence speed with cost-aware stopping criterion

- Observation: Agent stopping early does not have a great impact on the convergence in a simple objective function.

Experimental results

Scenarios with cost-aware stopping criterion

- MABO of the Eggholder function with a MAS of three agents

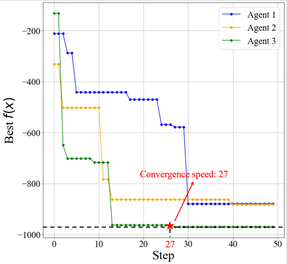

Figure: Convergence speed without stopping criterion

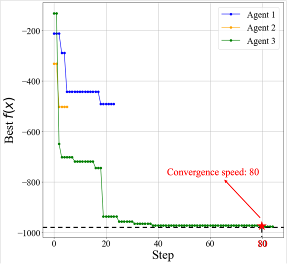

Figure: Convergence speed with cost-aware stopping criterion

- Observation: Agent stopping early could influences the convergence in a complex objective function.

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

- **MAS team size:**

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

- **MAS team size:**
- **Complexity of objective function**

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

- **MAS team size:**

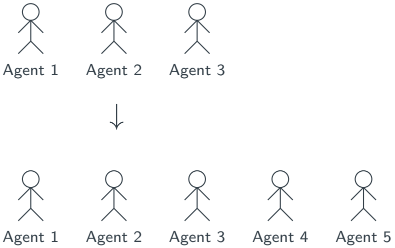

- **Complexity of objective function**

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

- **MAS team size:**

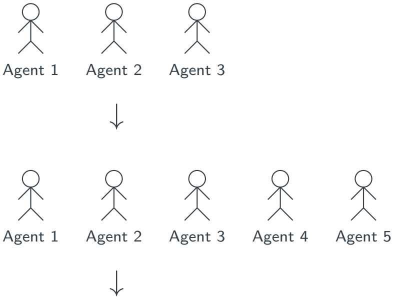

Faster convergence speed

- **Complexity of objective function**

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

- **MAS team size:**

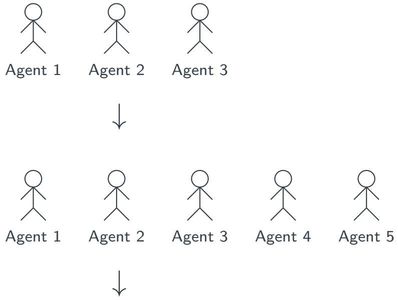

Faster convergence speed

- **Complexity of objective function**

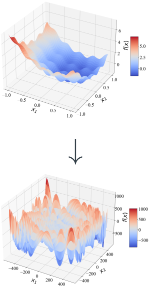

Conclusion

Allowing global-local communication significantly improves convergence speed to the global optimum, but not necessarily to the local optimum for every agent.

- **MAS team size:**

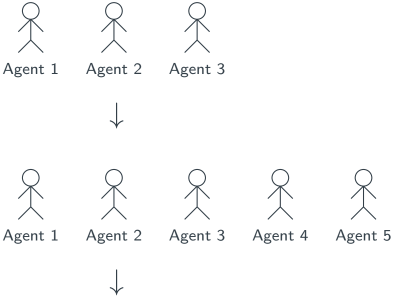

Faster convergence speed

- **Complexity of objective function**

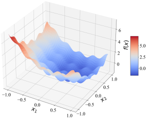

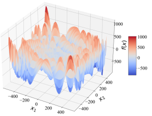

Slower convergence speed

Conclusion

Agent stops early would have a great impact on convergence in a complex objective function but not a simple objective function.

Conclusion

Agent stops early would have a great impact on convergence in a complex objective function but not a simple objective function.

$\downarrow$

For a design team, communication mechanisms and incentive structures for solution search shall be designed and tailored according to the complexity of the problem to be solved.

Conclusion

Agent stops early would have a great impact on convergence in a complex objective function but not a simple objective function.

↓

For a design team, communication mechanisms and incentive structures for solution search shall be designed and tailored according to the complexity of the problem to be solved.

Figure: Simple design

Figure: Complex design

Future work

- Decision 3: where not to sample

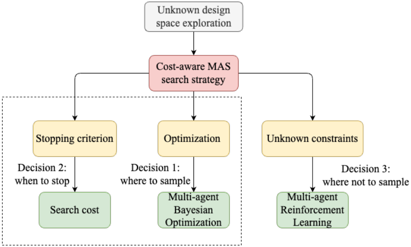

Future work

- Decision 3: where not to sample

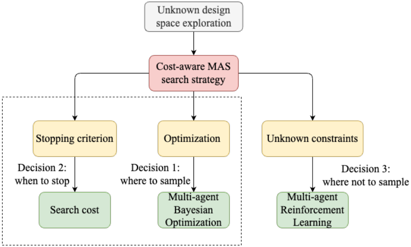

- Topology of communication

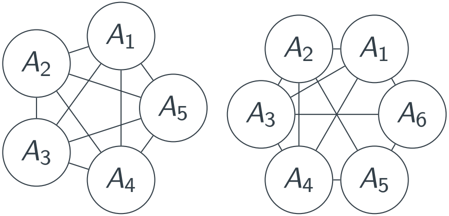

Future work

- Decision 3: where not to sample

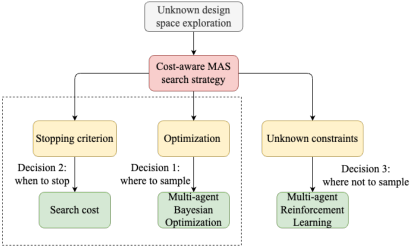

- Topology of communication

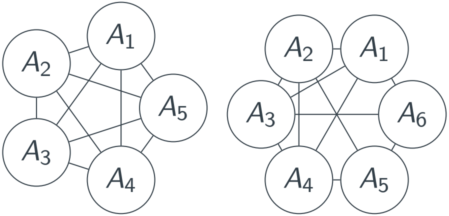

- Allocate appropriate amount of initial funds or budget to take care of the risk-averse attitude of human designers and enhance team resilience.

Future work

- Decision 3: where not to sample

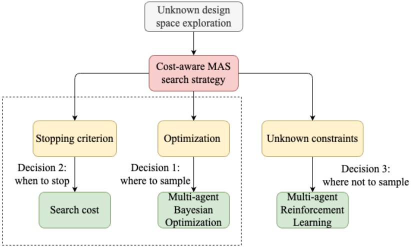

- Topology of communication

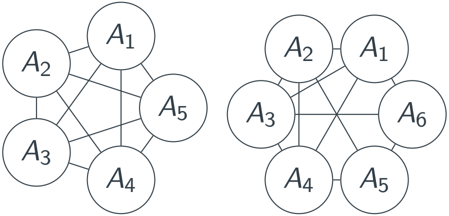

- Allocate appropriate amount of initial funds or budget to take care of the risk-averse attitude of human designers and enhance team resilience.
- The impact of MAS team size on the cost-setting strategy.

Future work

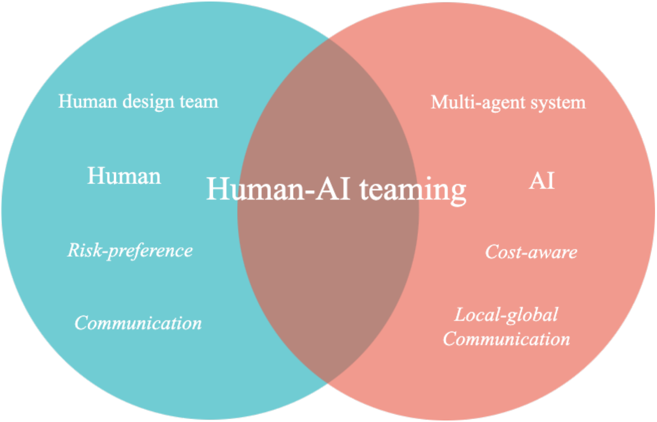
Human design team
Human
Risk-preference
Communication
Human-AI teaming
Multi-agent system
AI
Cost-aware
Local-global
Communication

Published paper

- S. Chen, A. E. Bayrak[3], Z. Sha, "Multi-Agent Bayesian Optimization for Unknown Design Space Exploration", ASME 2023 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference, Boston, MA, Aug. 20-23, 2023.

[3] Assistant Professor, School of Systems and Enterprises, Stevens Institute of Technology

Thank you!

We gratefully acknowledge the supports from

For more information, please contact:
Siyu Chen at siyu.chen@utexas.edu

The University of Texas at Austin
Cockrell School of Engineering