Event-Triggered Interactive Gradient Descent for Real-Time Multiobjective Optimization

Pio Ong and Jorge Cortés

Mechanical and Aerospace Engineering
University of California, San Diego
http://carmenere.ucsd.edu/jorge

56th IEEE Conference on Decision and Control:
Event-Triggered and Self-Triggered Control I
Melbourne, Australia
December 12-15, 2017

Flying to Australia

Options	Transit (Hours)	Cost (Dollars)
1	10	1100
2	5	1500
3	2	1400

- How would a robot know which option is the best?

Flying to Australia

Options	Transit (Hours)	Cost (Dollars)
1	10	1100
2	5	1500
3	2	1400

- How would a robot know which option is the best?
- Some options are obviously worse.

Flying to Australia

Options	Transit (Hours)	Cost (Dollars)
1	10	1100
3	2	1400

- How would a robot know which option is the best?
- Some options are obviously worse.
- But, we are left with mathematically ambiguous options (**Pareto Solutions**)

Flying to Australia

Options	Transit (Hours)	Cost (Dollars)	Happiness
1	10	1100	
3	2	1400	

- How would a robot know which option is the best?
- Some options are obviously worse.
- But, we are left with mathematically ambiguous options (**Pareto Solutions**)
- Ask Human!

My Talk in One Slide

- **Motivation:**
  1. Rise of robots that will eventually coexist with human
  2. Robot solve a optimization problem to do something
  3. Robot becomes more complex, can do more than one thing

- **Scenario:**
  Human interacts with robot to help solve multiobjective optimization problem

- **Robot Accommodate Human:**
  1. Human cannot be asked too often
  2. Human needs some time to answer

- **Approach:** Use **Event-Trigger Control** to minimize human interaction.

Outline

Describing Scenario

- Problem Statement and Assumptions
- Our approach: Interactive Gradient Descent

Modeling Humans

- Human needs to rest.
  - Designing Event Trigger
- Adding Human Response Time
  - Limiting design parameter

Wrapping up my Talk

- Simulations
- Conclusions

Problem and Assumptions

Our problem:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}}\ f(x)$$

with $f(x) \in \mathbb{R}^m$, m objective functions

- In general, infinite number of Pareto solutions

Problem and Assumptions

Our problem:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}}\; f(x)$$

with $f(x) \in \mathbb{R}^m$, m objective functions

- In general, infinite number of Pareto solutions
- The human has an implicit cost function, $c : \mathbb{R}^m \to \mathbb{R}$, that ranks them

Problem and Assumptions

Our problem:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \; f(x)$$

with $f(x) \in \mathbb{R}^m$, m objective functions

- In general, infinite number of Pareto solutions
- The human has an implicit cost function, $c : \mathbb{R}^m \to \mathbb{R}$, that ranks them
    1. Implicit because the human cannot express what it is

Problem and Assumptions

Our problem:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}}\ f(x)$$

with $f(x) \in \mathbb{R}^m$, m objective functions

- In general, infinite number of Pareto solutions
- The human has an implicit cost function, $c : \mathbb{R}^m \to \mathbb{R}$, that ranks them
  1. Implicit because the human cannot express what it is
  2. Human can respond to queries; we assume he can give the gradient

Problem and Assumptions

Our problem:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}}\ f(x)$$

with $f(x) \in \mathbb{R}^m$, m objective functions

- In general, infinite number of Pareto solutions
- The human has an implicit cost function, $c : \mathbb{R}^m \to \mathbb{R}$, that ranks them
  1. Implicit because the human cannot express what it is
  2. Human can respond to queries; we assume he can give the gradient

- **Assumptions:** To assure there is a unique solution,
  1. Each objective function is strictly convex.
  2. The implicit function is strictly convex, increasing w.r.t. each objective value.
  3. The implicit function is bounded from below and is radially unbounded.

Restate the problem

- What do we mean by solving a multiobjective optimization problem?

Restate the problem

- What do we mean by solving a multiobjective optimization problem?
  Answer: Find the Pareto solution that the human likes the best.

Restate the problem

- What do we mean by solving a multiobjective optimization problem?
  Answer: Find the Pareto solution that the human likes the best.

- Problem that we will solve:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \; (c \circ f)(x)$$

Restate the problem

- What do we mean by solving a multiobjective optimization problem?
  Answer: Find the Pareto solution that the human likes the best.

- Problem that we will solve:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}}\ (c \circ f)(x)$$

- **Scenario:** human and robot working together to get the best Pareto solution.

Restate the problem

- What do we mean by solving a multiobjective optimization problem?
  Answer: Find the Pareto solution that the human likes the best.

- Problem that we will solve:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \ (c \circ f)(x)$$

- **Scenario:** human and robot working together to get the best Pareto solution.
  - Single objective optimization.

Restate the problem

- What do we mean by solving a multiobjective optimization problem?
  Answer: Find the Pareto solution that the human likes the best.

- Problem that we will solve:

$$\underset{x \in \mathbb{R}^n}{\text{minimize}}\ (c \circ f)(x)$$

- **Scenario:** human and robot working together to get the best Pareto solution.
  - Single objective optimization.
  - **No objective function avaliable**. Only gradient value available upon requests.

Interactive Gradient Descent

Let's try gradient descent!

$$\dot{x}(t) = -(\nabla(c \circ f)(x(t)))^T$$

Role: What's the human and robot role in this optimization?

Interactive Gradient Descent

Let's try gradient descent!

$$\dot{x}(t) = -(\nabla(c \circ f)(x(t)))^T$$

Role: What's the human and robot role in this optimization?

$$\dot{x}(t) = -(\nabla c(f(x(t)))J_f(x(t)))^T$$

Interactive Gradient Descent

Let's try gradient descent!

$$\dot{x}(t) = -(\nabla(c \circ f)(x(t)))^T$$

Role: What's the human and robot role in this optimization?

$$\dot{x}(t) = -(\underbrace{\nabla c(f(x(t)))}_{\text{human}} \underbrace{J_f(x(t))}_{\text{robot}})^T$$

Interactive Gradient Descent

Let's try gradient descent!

$$\dot{x}(t) = -(\nabla(c \circ f)(x(t)))^T$$

Role: What's the human and robot role in this optimization?

$$\dot{x}(t) = -(\underbrace{\nabla c(f(x(t)))}_{\text{human}}\underbrace{J_f(x(t))}_{\text{robot}})^T$$

Humans cannot update the value continuously!

Event-Triggered Interactive Gradient Descent

Preferably, only ask for human help only when it really needs to.

$$\dot{x}(t) = -(\underbrace{\nabla c(f(x(t_k)))}_{\text{human}}\underbrace{J_f(x(t))}_{\text{robot}})^T$$

with t_k to be determined by the robot iteratively.

Event-Triggered Interactive Gradient Descent

Preferably, only ask for human help only when it really needs to.

$$\dot{x}(t) = -(\underbrace{\nabla c(f(x(t_k)))}_{\text{human}} \underbrace{J_f(x(t))}_{\text{robot}})^T$$

with t_k to be determined by the robot iteratively.

When to ask human?

Event-Triggered Interactive Gradient Descent

Preferably, only ask for human help only when it really needs to.

$$\dot{x}(t) = -(\underbrace{\nabla c(f(x(t_k)))}_{\text{human}} \underbrace{J_f(x(t))}_{\text{robot}})^T$$

with t_k to be determined by the robot iteratively.

When to ask human?

Our proposition: robot monitors

- $\|x(t) - x(t_k)\|$
- $\sigma \frac{\|\dot{x}(t)\|}{L_c \|J_f(x(t))\|}$ with $\sigma \in (0, 1)$

When these two things are equal, ask human.

Design Guarantees

Properly stated

$$t_{k+1} = \min_t \left\{ t \geq t_k \mid \|x(t) - x(t_k)\| = \sigma \frac{\|\nabla c(f(x(t_k)))J_f(x(t))\|}{L_c \|J_f(x(t))\|} \right\}$$

Design Guarantees

Properly stated

$$t_{k+1} = \min_{t} \left\{ t \geq t_k \mid \|x(t) - x(t_k)\| = \sigma \frac{\|\nabla c(f(x(t_k)))J_f(x(t))\|}{L_c \|J_f(x(t))\|} \right\}$$

Guarantees:

1. **Global asymptotic stability:**
   Lyapunov Function: $V(x) = c \circ f(x) - c \circ f(x^*)$

$$\implies \frac{d}{dt} V(x(t)) \leq -\frac{1-\sigma}{(1+\sigma)^2} \|\nabla c(f(x(t)))J_f(x(t))\|^2 < 0$$

$\implies$ **Asymptotic Stability**

Moreover, if $c \circ f$ is strongly convex with a parameter μ, the optimizer is **exponentially stable** with the following bound,

$$V(x(t)) \leq V(x_0) e^{-\frac{2\mu(1-\sigma)}{(1+\sigma)^2} t}$$

Design Guarantee - Continued

$$t_{k+1} = \min_t \left\{ t \geq t_k \mid \|x(t) - x(t_k)\| = \sigma \frac{\|\nabla c(f(x(t_k)))J_f(x(t))\|}{L_c \|J_f(x)\|} \right\}.$$

2. **Autonomous operation:** Robot has all the information to calculate above

Design Guarantee - Continued

$$t_{k+1} = \min_t \left\{ t \geq t_k \mid \|x(t) - x(t_k)\| = \sigma \frac{\|\nabla c(f(x(t_k)))J_f(x(t))\|}{L_c \|J_f(x)\|} \right\}.$$

2. **Autonomous operation:** Robot has all the information to calculate above
3. **No Zeno behavior:** There exists a uniform lower bound for the inter-event times $\tau_\sigma \leq t_{k+1} - t_k$ for all $k \in \mathbb{N} \cup \{0\}$ where τ_σ is a constant given by

$$\tau_\sigma = \frac{1}{\beta} \ln(1 + \beta \frac{\sigma}{L_c J_{\max}})$$

Design Guarantee - Continued

$$t_{k+1} = \min_{t} \left\{ t \geq t_k \mid \|x(t) - x(t_k)\| = \sigma \frac{\|\nabla c(f(x(t_k)))J_f(x(t))\|}{L_c \|J_f(x)\|} \right\}.$$

2. **Autonomous operation:** Robot has all the information to calculate above
3. **No Zeno behavior:** There exists a uniform lower bound for the inter-event times $\tau_\sigma \leq t_{k+1} - t_k$ for all $k \in \mathbb{N} \cup \{0\}$ where τ_σ is a constant given by

$$\tau_\sigma = \frac{1}{\beta} \ln(1 + \beta \frac{\sigma}{L_c J_{\max}})$$

Delay in Human

- **Previous model** human responds instantaneously
- **Better model** Human requires some time to work, has some response time.

Delay in Human

- **Previous model** human responds instantaneously
- **Better model** Human requires some time to work, has some response time.
- The more accurate gradient descent is

$$\dot{x}(t) = \underbrace{\nabla c(f(x(t_k)))}_{\text{human}} \underbrace{J_f(x(t))}_{\text{robot}}, \quad t \in [t_k + D_k, t_{k+1} + D_{k+1})$$

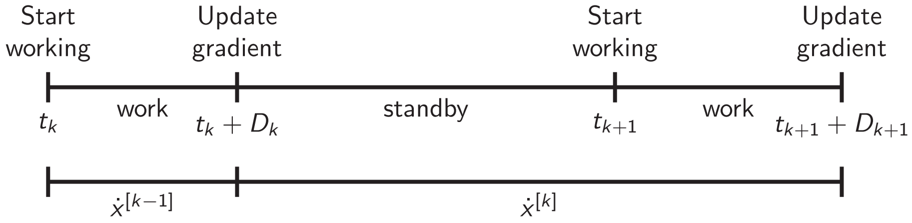

Trigger Design - Delay Case

Assuming there is a maximum delay D, we propose a similar trigger design: robot monitors

- $\|x(t) - x(t_k)\|$
- $\sigma' \frac{\|\dot{x}(t)\|}{L_c J_{\max}}$ but σ' not $(0, 1)$

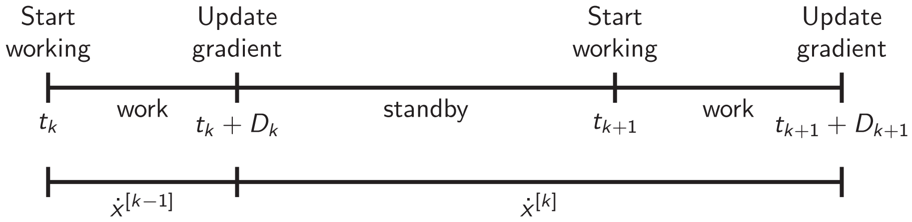

Trigger Design - Delay Case

Assuming there is a maximum delay D, we propose a similar trigger design: robot monitors

- $\|x(t) - x(t_k)\|$
- $\sigma' \frac{\|\dot{x}(t)\|}{L_c J_{\max}}$ but σ' not $(0, 1)$

σ' must satisfy

$$\sigma' \leq \frac{L_c J_{\max}}{\beta}(e^{\beta(\tau_\sigma - D)} - 1)$$

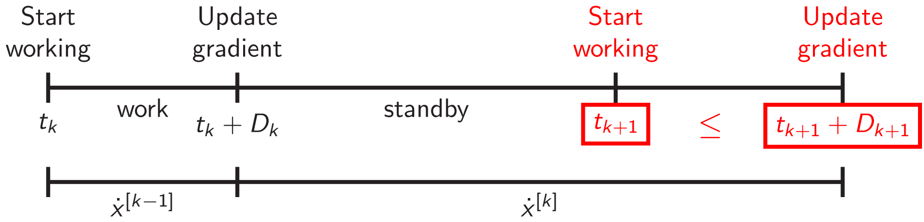

Trigger Design - Delay Case

Assuming there is a maximum delay D, we propose a similar trigger design: robot monitors

- $\|x(t) - x(t_k)\|$
- $\sigma' \frac{\|\dot{x}(t)\|}{L_c J_{\max}}$ but σ' not $(0, 1)$

σ' must satisfy.

$$\frac{L_c J_{\max}}{\beta} \left(\frac{1+\sigma}{1-\sigma} \right)^2 (e^{\beta D} - 1) < \sigma' \leq \frac{L_c J_{\max}}{\beta} (e^{\beta(\tau_\sigma - D)} - 1)$$

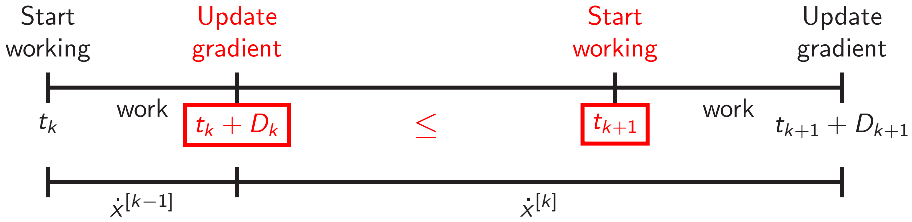

Design Guarantees - Delay Case

$$t_{k+1} = \min_t \left\{ t \geq t_k \mid \|x(t) - x(t_k)\| = \sigma' \frac{\|\nabla c(f(x(t_k)))J_f(x(t))\|}{L_c J_{\max}} \right\}.$$

Guarantees:

1. **Global asymptotic stability:** Same
2. **Autonomous Operation:** Same
3. **No Zeno Behavior:** The uniform lower bound to the interevent times for the delay case is

$$\tau_{\sigma'} = \frac{1}{\beta} \ln \left(\frac{1 + \beta \frac{\sigma'}{L_c J_{\max}}}{1 + \left(\frac{1+\sigma}{1-\sigma}\right)^2 (e^{\beta D} - 1)} \right) + D$$

Simulations

A robot trying to get close to two objects

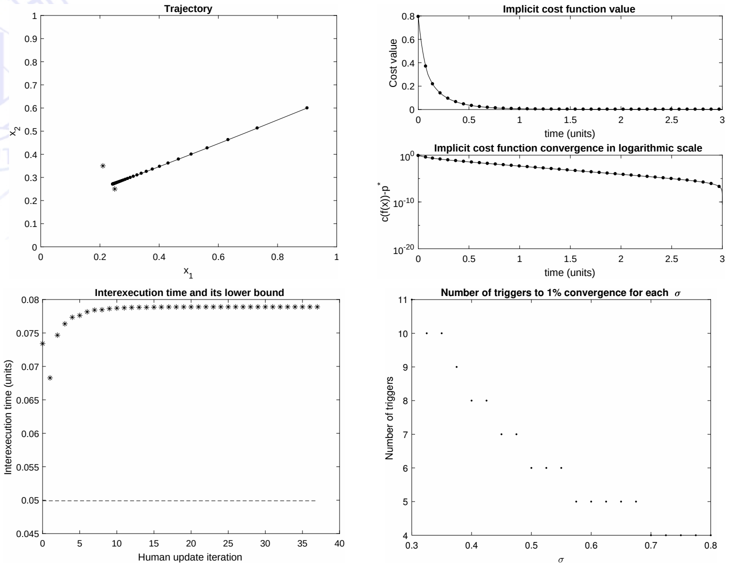

Conclusions

Event-triggered design for human-robot interaction

- Human works as supervisor, robot works as extension of human capability
- Bound on inter-event time guarantees human has time to do other things.
- Provably correct: achieves multiobjective optimization task

Future Work

- Richer models for human engagement
  1. Rest time and Response time
  2. Human inputs with errors
- Scenarios where human needs to rest for longer than interevent time
- Online Learning of human model using human responses

Thank You!

- **My advisor**: Jorge Cortés

- **Our Funding Source:** NSF award CNS-1329619

Question?

Questions and feedbacks are welcome!