Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

On the Importance of Anticipation in Dynamic Optimization

Peter A.N. Bosman

Centre for Mathematics and Computer Science (CWI)
Amsterdam
The Netherlands

Introduction - I

- What makes dynamic optimization hard?
- Problems changes with time.
- Changes may be dramatic.
- Specific problem difficulty: time–dependence.
- Current decisions have future consequences.
- Requires anticipation to solve.

Introduction - II

- Consider vehicle routing
  - Sending vehicle north excludes profitble routing of that vehicle to southern locations in near future.
  - Quality of service influences future demand.
- Consider inventory management
  - Replenishment determines future inventory.
  - Quality of service influences future demand.

Optimization problems

- General definition (unconstrained):

$$\max_{\boldsymbol{x}\in\mathbb{P}} \{\mathfrak{F}(\boldsymbol{x})\} \quad (1)$$

- Dynamic definition (unconstrained):

$$\mathfrak{F}(\boldsymbol{x}) = \int_0^{t^{end}} \mathfrak{F}^{dyn}\left(\boldsymbol{x}_{\boldsymbol{x}}^{dyn}(t)\right) dt \quad (2)$$

Sources of Problem Difficulty in Dynamic Optimization

- Dynamic influences
  1. System influence.
     Solver has no control over it; the way the function changes no matter what.
  2. Control influence.
     Function changes as a result of past choices (i.e. variable settings) made by solver (time–dependence).
- Most EAs specifically designed to handle system influence, i.e. tracking optima.

Sources of Problem Difficulty in Dynamic Optimization

On the Importance of Anticipation in Dynamic Optimization

Peter A.N. Bosman

Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

- Dynamic influences
  1. System influence.
     Solver has no control over it; the way the function changes no matter what.
  2. Control influence.
     Function changes as a result of past choices (i.e. variable settings) made by solver (time–dependence).
- Most EAs specifically designed to handle system influence, i.e. tracking optima.

Pitfall

What if optima themselves depend on past variable settings?
Tracking optima alone will not be enough. . . .

Myopic Approach Falls Victim to Time–Deception

On the Importance of Anticipation in Dynamic Optimization

Peter A.N. Bosman

Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

- The approach:

$$\max_{\boldsymbol{x}(t^{\text{now}})} \left\{ \mathfrak{F}^{\text{dyn}} \left(\boldsymbol{x}(t^{\text{now}}) \right) \right\} \tag{3}$$

- How bad can it be?

Myopic Approach Falls Victim to Time–Deception

Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

- The approach:

$$\max_{\boldsymbol{x}(t^{\text{now}})} \left\{ \mathfrak{F}^{\text{dyn}}\left(\boldsymbol{x}(t^{\text{now}})\right)\right\} \quad (3)$$

- How bad can it be?

- Arbitrarily bad, even assuming smooth system influence.
- Illustration:

$$\max_{\boldsymbol{x}(t)} \left\{ \int_0^{t^{\text{end}}} \varphi(\boldsymbol{x}(t), t)dt \right\} \quad (4)$$

where

$$\varphi(\boldsymbol{x}(t),t) = \begin{cases} -\sum_{i=0}^{l-1} (\boldsymbol{x}(t)_i - t)^2 & \text{if } 0 \leq t < 1 \\ -\sum_{i=0}^{l-1} (\boldsymbol{x}(t)_i - t)^2 + \psi\left(|\boldsymbol{x}(t-1)_i|\right) & \text{otherwise} \end{cases}$$

On the Importance of Anticipation in Dynamic Optimization
Peter A.N. Bosman
Introduction
Online dynamic optimization problems
Myopic Dynamic Optimization
Non–Myopic Dynamic Optimization
Stochasticity
Literature
EA literature
Illustration
Conclusions
Selected references

Myopic Approach Falls Victim to Time–Deception

- Illustration (continued):

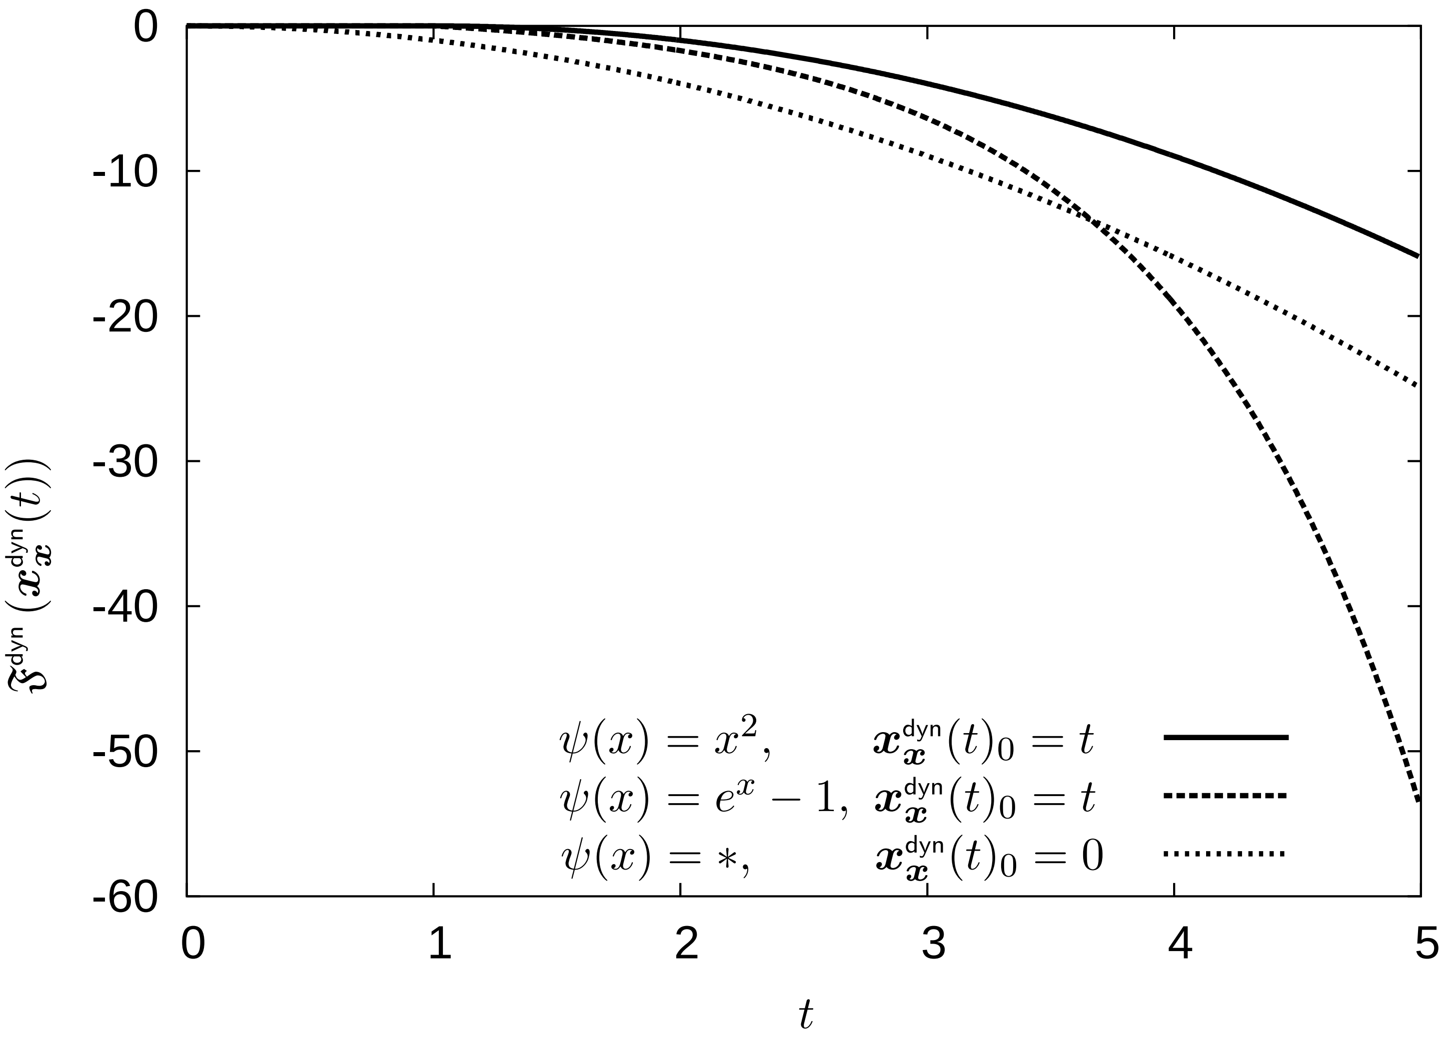

(Theoretical) Non–Myopic Approach Saves the Day

- Observations:
  - Deception because full problem definition not used.
  - Optimization over future decisions mandatory.
- Theoretical approach:

$$\max_{\boldsymbol{x}_{\boldsymbol{x}}^{\text{dyn}}(t)} \left\{ \int_{t^{\text{now}}}^{t^{\text{end}}} \mathfrak{F}^{\text{dyn}}\left(\boldsymbol{x}_{\boldsymbol{x}}^{\text{dyn}}(t)\right) dt \right\} \quad (5)$$

- Observations:
  - Equals problem definition, thus result is optimal.
  - Cannot evaluate the future.
  - Only option: predict the future.

Towards a Practical Non–Myopic Approach

On the Importance of Anticipation in Dynamic Optimization
Peter A.N. Bosman
Introduction
Online dynamic optimization problems
Myopic Dynamic Optimization
Non–Myopic Dynamic Optimization
Stochasticity
Literature
EA literature
Illustration
Conclusions
Selected references

- One approach:
  1. Maintain approximation of $\mathfrak{F}^{\text{dyn}}(\boldsymbol{x}(t))$.
  2. Optimize present and (part of) approximated future:

$$\max_{\boldsymbol{x}(t)} \left\{ \int_{t^{\text{now}}}^{\min\{t^{\text{now}}+t^{\text{plen}},t^{\text{end}}\}} \hat{\mathfrak{F}}^{\text{dyn}}_{\boldsymbol{\alpha}}(t, \boldsymbol{x}(t))\, dt \right\} \tag{6}$$

- Alternatively, optimize only current situation, but don't optimize (only) $\mathfrak{F}$, but (also) measure additional information (e.g. flexibility, robustness, sensitivity).

Non–stochastic vs. stochastic - I

- Non–stochastic
  - Single trajectory is always optimal.
- Stochastic
  - Single trajectory is only optimal afterwards.
  - Optimality is scenario–dependent.
  - Need to average somehow over multiple scenarios.
  - Alternatively, take expected–value scenario.
  - Limitation: expected value must be representative.

Non–stochastic vs. stochastic - II

- Example pickup problem:
  - One truck, one load per time-unit.
  - Decide: pick up load or not.
  - If yes, gain 1 - distance traveled.
  - If no, load disappears (no cost, no gain).
- Now consider $x(t), x(t+1)$.
- Time–dependence:
  decision to drive determines new location.
- No sense to plan $x(t+1)$ ahead due to stochasticity.
- Depends on future load location.

Non–stochastic vs. stochastic - III

- Example: load–dropping follows normal distribution, centered at origin.

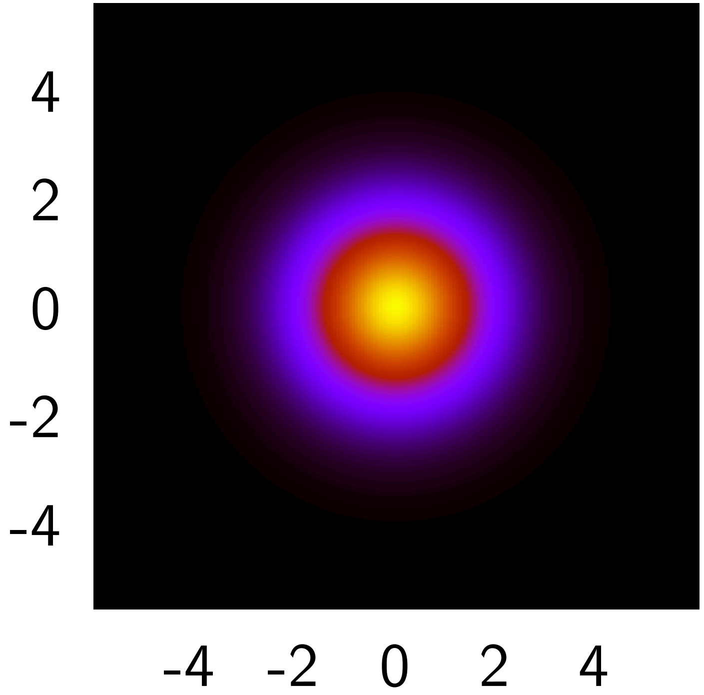

- Expected value is origin.
- Picks up loads within $\frac{1}{2}$ of origin and truck.
- Good strategy.

Introduction
Online dynamic optimization problems
Myopic Dynamic Optimization
Non–Myopic Dynamic Optimization
Stochasticity
Literature
EA literature
Illustration
Conclusions
Selected references

Non–stochastic vs. stochastic - IV

- Another example: load–dropping follows 4 normal distributions, centered at (2,2), (-2,2), (-2,-2), (2,-2).

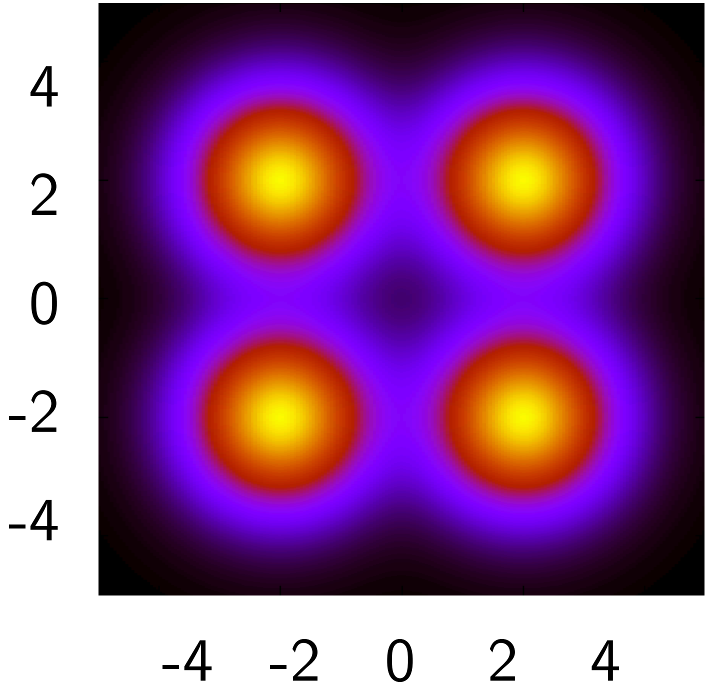

- Expected value is origin again.
- Leads to same strategy.
- This time: bad strategy.

Non–EA literature - I

- Chang, Givan & Chong (2000): expectation method.
- Optimize future trajectory once for each alternative.
- Future trajectory starts with that alternative.
- Once for expected value or repeat for scenarios.
- Choose decision with average best result.

Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

Non–EA literature - II

- Bent & Van Hentenryck (2004): Consensus.
- Faster than expectation method.
- Can lead to inferior results.
- Remove loop over each alternative.
- Solve expected value or each scenario only once.
- Choose decision with average best result.

Non–EA literature - III

- Bent & Van Hentenryck (2004): Regret.
- Faster than expectation method.
- Closer to results of expectation method.
- Requires approximation of regret.
- Regret: what if suboptimal alternative was chosen.
- Solve expected value or each scenario only once.
- Then loop over all alternatives and compute regrets.
- Choose decision with average best result.

EA literature - I

- Non–EA literature important and interesting, but. . .
  - Many re–optimizations required.
  - What if # alternative decisions is large?
  - What if decisions are real–valued?
  - Only tackles time–dependence partly.
    - Influence on future decisions (i.e. where the truck drives) is tackled.
    - Influence on future model/simulation/real–world (i.e. customer satisfaction) is not tackled.
    - Optimization and simulation need to be intertwined.
    - Makes algorithmic design even harder.
    - Quickly need to return to enumerative search.
- Also, for learning, diverse population can help.

EA literature - II

- Branke and Mattfeld (2000): flexibility.
- Considers online scheduling.
- Fitness is not just tardiness.
- Also includes idle–time penalty.
- Focuses on early use of capacity.
- Warrants flexibility for future decisions.

EA literature - III

- Van Hemert and La Poutré (2004): implicit anticipation.
- Considers online vehicle routing.
- Allows a solution to insert anticipated moves.
- Anticipated move: move to location without load.
- Self-adaptation of valuation of anticipated moves.
- No explicit anticipation of loads.

Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

EA literature - IV

- Bosman (2005), Bosman and La Poutré (2006, 2007): explicit anticipation.
- Considers new benchmark and online vehicle routing.
- Performs explicit anticipation.
- Predicts future situations.
- Optimizes future decisions.
- Prediction quality influences solution quality.
- For EAs (use of adaptivity characteristic):

	Expected value	**Scenarios**
Decision list	✓	✗
Strategy	!	✓

Conclusions

Introduction

Online dynamic optimization problems

Myopic Dynamic Optimization

Non–Myopic Dynamic Optimization

Stochasticity

Literature

EA literature

Illustration

Conclusions

Selected references

- Consequences of current decisions are important.
- Not only track optima, but also require anticipation.
- Relatively novel in dynamic optimization.
- Possible to obtain better results.
- Much room for new results exists.

Selected references

- H. Chang, R. Givan and E. Chong (2000). Online Scheduling Via Sampling. In Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems (AIPS), 62–71.
- R. Bent and P. Van Hentenryck (2004). Regrets Only! Online Stochastic Optimization Under Time Constraints. In Proceedings of the Nineteenth National Conference on Artificial Intelligence (AAAI), 501–506.
- R. Bent and P. Van Hentenryck (2004). The Value of Consensus in Online Stochastic Scheduling. In Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS), 219–226.
- J. Branke and D. Mattfeld (2000). Anticipation in dynamic optimization: The scheduling case. In M. Schoenauer *et al.*, editors, Parallel Problem Solving from Nature — PPSN VI, pages 253–262, Springer–Verlag, Berlin.
- J. I. van Hemert and J. A. La Poutré (2004). Dynamic routing problems with fruitful regions: models and evolutionary computation. In X. Yao *et al.*, editors, Parallel Problem Solving from Nature — PPSN VIII, pages 692–701, Springer–Verlag, Berlin.
- P. A. N. Bosman (2005). Learning, Anticipation and Time–Deception in Evolutionary Online Dynamic Optimization. In S. Yang and J. Branke, organisers, Proceedings of the Evolutionary Algorithms for Dynamic Optimization Problems EvoDOP Workshop at the Genetic and Evolutionary Computation Conference – GECCO–2005, pages 39–47, ACM Press, New York, New York.
- P. A. N. Bosman and H. La Poutré (2006). Computationally Intelligent Online Dynamic Vehicle Routing by Explicit Load Prediction in an Evolutionary Algorithm. In T.P. Runarsson *et al.*, editors, Parallel Problem Solving from Nature – PPSN IX, pages 312–321, Springer–Verlag, Berlin.
- P. A. N. Bosman and H. La Poutré (2007). Learning and Anticipation in Online Dynamic Optimization with Evolutionary Algorithms: The Stochastic Case. In Proceedings of the Genetic and Evolutionary Computation Conference — GECCO–2007, ACM Press, New York, New York.