Optimal Sampling for State Estimation of Stochastic Dynamical Systems

Antoine Aspeel

Université Catholique de Louvain

Private PhD Defense - March 4^{th}, 2022

Motivating example: proton therapy

- Cancer treatment method
- Irradiate the tumor while sparing healthy tissue

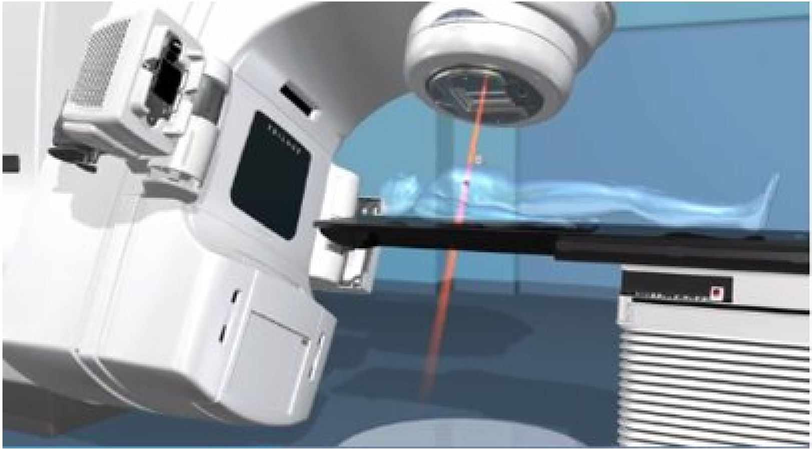

Volumetric Modulated Arc Therapy (VMAT), Varian ®

Motivating example: proton therapy

- Cancer treatment method
- Irradiate the tumor while sparing healthy tissue
- Challenging for mobile tumors (liver or lung)

Motivating example: proton therapy

- Cancer treatment method
- Irradiate the tumor while sparing healthy tissue
- Challenging for mobile tumors (liver or lung)
- Acquire X-ray images of the patient’s chest to locate the tumor

Motivating example: proton therapy

- Cancer treatment method
- Irradiate the tumor while sparing healthy tissue
- Challenging for mobile tumors (liver or lung)
- Acquire X-ray images of the patient's chest to locate the tumor
- X-ray acquisitions are irradiating and must therefore be limited

Motivating example: proton therapy

- Cancer treatment method
- Irradiate the tumor while sparing healthy tissue
- Challenging for mobile tumors (liver or lung)
- Acquire X-ray images of the patient's chest to locate the tumor
- X-ray acquisitions are irradiating and must therefore be limited
- *When* should X-ray images be acquired?

Other motivating examples

Whenever we want to estimate a dynamic quantity and the number of measurements is limited.

- Estimating the position of a drone from its GPS measurements
- Estimating the number of infections based on tests during a pandemic
- Estimating voting intentions from surveys during an election period

Main topic

Considering a stochastic dynamical system over a finite horizon, *when* should we measure to have the best[1] estimate, on average, over the horizon?

[1] In a sense that will be made clear later.

Summary

List of publications

Linear and Gaussian dynamics (Chap. 3 and 6)

Chap. 3 - Optimal intermittent Kalman filter

- A. Aspeel, A. Legay, R. Jungers, and B. Macq. Optimal measurement budget allocation for Kalman prediction over a finite time horizon by genetic algorithms. *EURASIP Journal on Advances in Signal Processing*, 2021.

Chap. 6 - Case study: tumor tracking

- A. Aspeel, D. Dasnoy, R. Jungers, and B. Macq. Optimal intermittent measurements for tumor tracking in X-ray guided radiotherapy. In *Medical Imaging 2019: Image-Guided Procedures, Robotic Interventions, and Modeling*. International Society for Optics and Photonics, 2019.
- A. Aspeel, A. Legay, and B. Macq. Genetic algorithms for optimal intermittent measurements for tumor tracking. In *The International Conference on the Use of Computers in Radiation Therapy*, 2019. (No proceedings).

List of publications

Nonlinear and non-Gaussian dynamics (Chap. 4)

Chap. 4 - Optimal intermittent particle filter

- A. Aspeel, A. Gouverneur, R. Jungers, and B. Macq. Optimal measurement budget allocation for particle filtering. In *27th IEEE International Conference on Image Processing*, 2020.
- A. Aspeel, A. Gouverneur, R. Jungers, and B. Macq. Optimal intermittent particle filter. To appear (under review).
- A. Aspeel, V. François-Lavet, R. Jungers, and B. Macq. Self-triggered measurements for estimation with deep reinforcement learning. To appear.

Chap. 5 - Beyond mean squared error: Optimal sampling for robust estimation

- A. Aspeel, K. Rutledge, R. Jungers, B. Macq, and N. Ozay. Optimal control for linear networked control systems with information transmission constraints. In *The 60th IEEE International Conference on Decision and Control*, 2021.

List of publications

Other publications

- M. Fanuel, A. Aspeel, J.-C. Delvenne, and J. Suykens. Positive semi-definite embedding for dimensionality reduction and out-of-sample extensions. In *SIAM Journal on Mathematics of Data Science*, 2022.
- D. Dasnoy, A. Aspeel, K. Souris, and B. Macq. Locally tuned deformation fields combination for 2D cine-MRI-based driving of 3D motion models. In *Physica Medica*, 2022.
- A. Aspeel, J.-C. Delvenne, M. Fanuel, and M. Schaub. Ellipsoidal embedding of graphs. To appear.[2]
- T. Duterme, A. Aspeel, and A. Germain. Poverty measurement and Adaptive preferences: a synthesis. To appear.

[2] The order of the authors is not determined, alphabetical here.

Summary

Linear and Gaussian dynamics

Linear system dynamics:

$$x(t+1) = Ax(t) + w(t)$$

$$y_{\sigma(t)}(t) = \begin{cases} Cx(t) + v(t) & \text{if } \sigma(t) = 1 \\ \emptyset & \text{if } \sigma(t) = 0 \end{cases}$$

Normally distributed unknowns:

$$w(t) \sim \mathcal{N}(0, Q)$$

$$v(t) \sim \mathcal{N}(0, R)$$

$$x(0) \sim \mathcal{N}(\bar{x}_0, \bar{P}_0)$$

Linear and Gaussian dynamics

Linear system dynamics:

$$x(t+1) = Ax(t) + w(t)$$

$$y_{\sigma(t)}(t) = \begin{cases} Cx(t) + v(t) & \text{if } \sigma(t) = 1 \\ \emptyset & \text{if } \sigma(t) = 0 \end{cases}$$

Normally distributed unknowns:

$$w(t) \sim \mathcal{N}(0, Q)$$
$$v(t) \sim \mathcal{N}(0, R)$$
$$x(0) \sim \mathcal{N}(\bar{x}_0, \bar{P}_0)$$

Mean squared error estimate:

$$\hat{x}_{\sigma(0:t)}(t) = \mathbb{E}[x(t) | y_{\sigma(0:t)}(0:t)]$$

Optimization program:

$$\min_{\sigma(0:T)} \sum_{t=0}^{T} \underbrace{\mathbb{E}[\|x(t) - \hat{x}_{\sigma(0:t)}(t)\|^2]}_{\text{Expected MSE}}$$

$$\text{such that } \underbrace{\sum_{t=0}^{T} \sigma(t) \leq N}_{\text{Budget constraint}}$$

Linear and Gaussian dynamics

Remarks

- Fits the Kalman filtering formalism
- Recursive “closed form” formula for the objective function
- Reduces to a deterministic combinatorial optimization program

Linear and Gaussian dynamics

Remarks

- Fits the Kalman filtering formalism
- Recursive “closed form” formula for the objective function
- Reduces to a deterministic combinatorial optimization program

Contributions

- A genetic algorithm implementing a count preserving crossover is well suited to solve it

Optimization algorithms

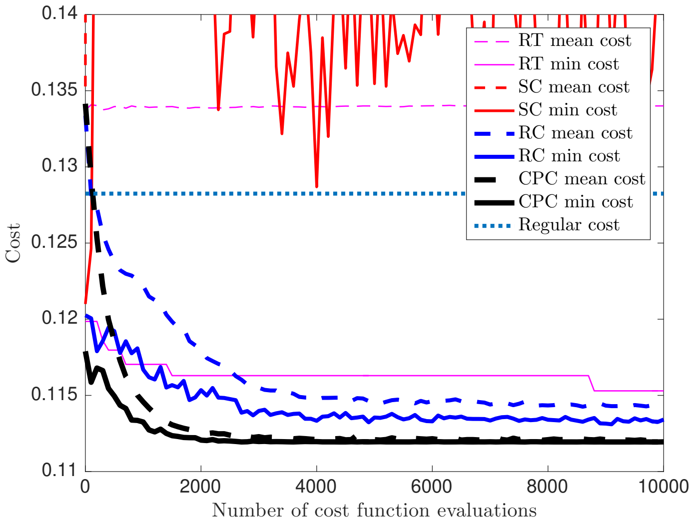

Genetic algorithm with count preserving crossover (CPC) outperforms other variants

Linear and Gaussian dynamics

Remarks

- Fits the Kalman filtering formalism
- Recursive “closed form” formula for the objective function
- Reduces to a deterministic combinatorial optimization program

Contributions

- A genetic algorithm implementing a count preserving crossover is well suited to solve it
- It has been shown on a lung tumor patient that the optimal intermittent Kalman filter outperforms the regular Kalman filter (Chap. 6)

Summary

Nonlinear and non-Gaussian dynamics

Nonlinear system dynamics:

$$x(t+1) = f(x(t), w(t))$$

$$y_{\sigma(t)}(t) = \begin{cases} g(x(t), v(t)) & \text{if } \sigma(t) = 1 \\ \emptyset & \text{if } \sigma(t) = 0 \end{cases}$$

Arbitrarily distributed unknowns:

$$w(t) \sim \mathcal{W}$$
$$v(t) \sim \mathcal{V}$$
$$x(0) \sim \mathcal{X}_0$$

Nonlinear and non-Gaussian dynamics

Nonlinear system dynamics:

$$x(t+1) = f(x(t), w(t))$$
$$y_{\sigma(t)}(t) = \begin{cases} g(x(t), v(t)) & \text{if } \sigma(t) = 1 \\ \emptyset & \text{if } \sigma(t) = 0 \end{cases}$$

Arbitrarily distributed unknowns:

$$w(t) \sim \mathcal{W}$$
$$v(t) \sim \mathcal{V}$$
$$x(0) \sim \mathcal{X}_0$$

Mean squared error estimate:

$$\hat{x}_{\sigma(0:t)}(t) = \mathbb{E}[x(t) | y_{\sigma(0:t)}(0:t)]$$

Optimization program (time-triggered):

$$\min_{\sigma(0:T)} \sum_{t=0}^{T} \underbrace{\mathbb{E}[\|x(t) - \hat{x}_{\sigma(0:t)}(t)\|^2]}_{\text{Expected MSE}}$$

$$\text{such that } \underbrace{\sum_{t=0}^{T} \sigma(t) \leq N}_{\text{Budget constraint}}$$

Nonlinear and non-Gaussian dynamics

Remarks

- Fits the particle filtering formalism
- No “closed form” formula for the objective function

Monte Carlo approximation

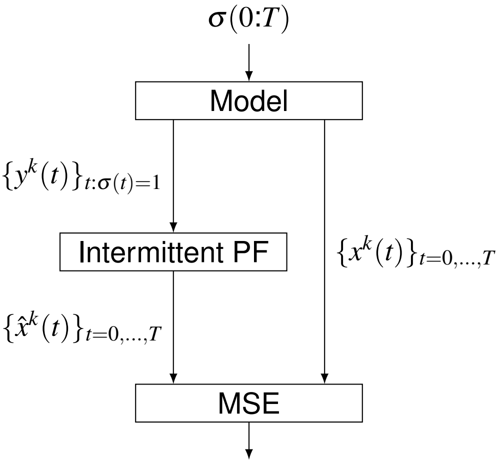

$\mathrm{MSE}^k = \sum_{t=0}^{T} \|x^k(t) - \hat{x}^k(t)\|^2$

Objective function approximation:

$$\mathbb{E}\left[\sum_{t=0}^{T} \|x(t) - \hat{x}(t)\|^2\right] \approx \frac{1}{K} \sum_{k=1}^{K} \mathrm{MSE}^k$$

Nonlinear and non-Gaussian dynamics

Remarks

- Fits the particle filtering formalism
- No "closed form" formula for the objective function

Contributions

- Objective function approximated numerically by a Monte-Carlo method
- A genetic algorithm implementing a count preserving crossover is well suited to solve it
- Also developed an Ex post version: recomputes the next measurement times after each new measurement

Estimation performance

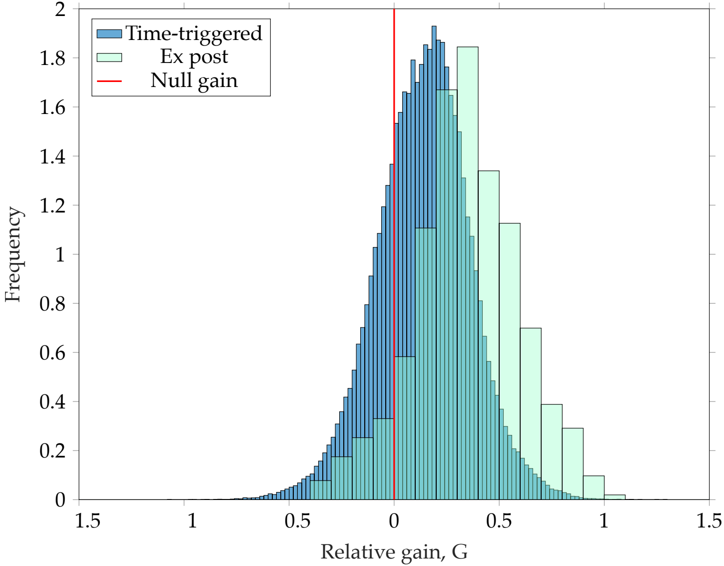

Histogram of the gain

$$G := \log_{10}\left(\frac{\text{MSE}_{\text{REG}}}{\text{MSE}}\right).$$

Summary

Linear dynamics under bounded disturbances

Linear system dynamics:

$$x(t+1) = Ax(t) + w(t)$$

$$y_{\sigma(t)}(t) = \begin{cases} Cx(t) + v(t) & \text{if } \sigma(t) = 1 \\ \emptyset & \text{if } \sigma(t) = 0 \end{cases}$$

Unknowns in convex polyhedra:

$$w(t) \in \mathcal{W}$$

$$v(t) \in \mathcal{V}$$

$$x(0) \in \mathcal{X}_0$$

Linear dynamics under bounded disturbances

Linear system dynamics:

$$x(t+1) = Ax(t) + w(t)$$

$$y_{\sigma(t)}(t) = \begin{cases} Cx(t) + v(t) & \text{if } \sigma(t) = 1 \\ \emptyset & \text{if } \sigma(t) = 0 \end{cases}$$

Unknowns in convex polyhedra:

$$w(t) \in \mathcal{W}$$
$$v(t) \in \mathcal{V}$$
$$x(0) \in \mathcal{X}_0$$

Luenberger-like estimate:

$$\hat{x}(t+1) = A\hat{x}(t) - \gamma(t)$$
$$\gamma(t) = f_t + \sum_{\tau \le t \text{ s.t. } \sigma(\tau)=1} F_{(t,\tau)}\left(y_{\sigma(\tau)}(\tau) - C\hat{x}(\tau)\right)$$

Optimization program:

$$\min_{\sigma(0:T), f_t, F_{(t,\tau)}} \max_{w(t), v(t), x(0)} \sum_{t=0}^{T} \|x(t) - \hat{x}(t)\|_\infty$$

$$\text{such that } \underbrace{\sum_{t=0}^{T} \sigma(t) \le N}_{\text{Budget constraint}}$$

Networked control systems

Co-design the controller, the measurement times and the control times.

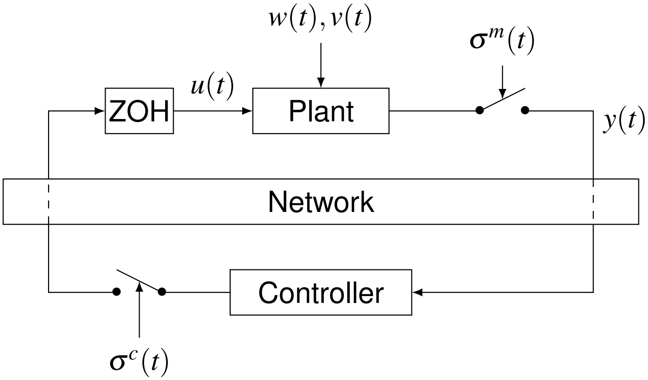

Networked control systems

Co-design the controller, the measurement times and the control times.

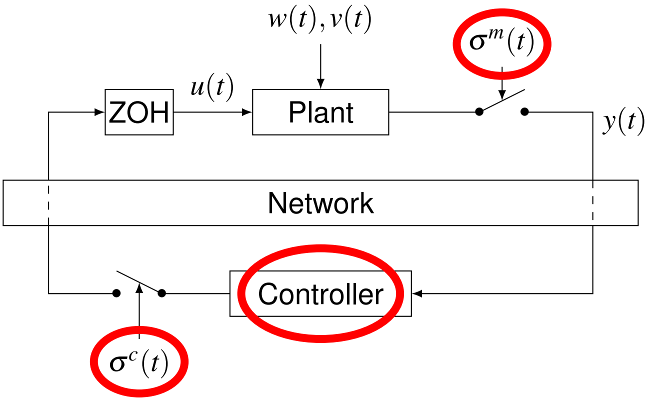

Linear dynamics under bounded disturbances

Remark

- Without the binary variables, gives a linear program using Youla-parameterization and polytope containment methods

Linear dynamics under bounded disturbances

Remark

- Without the binary variables, gives a linear program using Youla-parameterization and polytope containment methods

Contributions

- Technical result: some linear constraints remain linear after Youla's nonlinear change of variable
- Reduces to a mixed integer linear program

Example

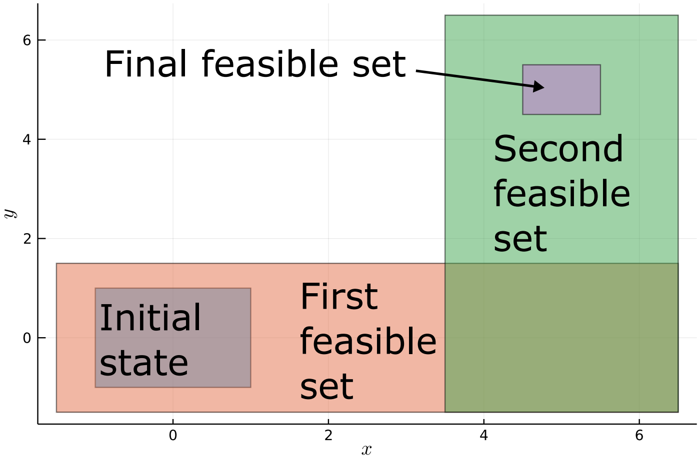

Drone model

- 20 time steps
- 3 measurements
- 4 control actions

Example

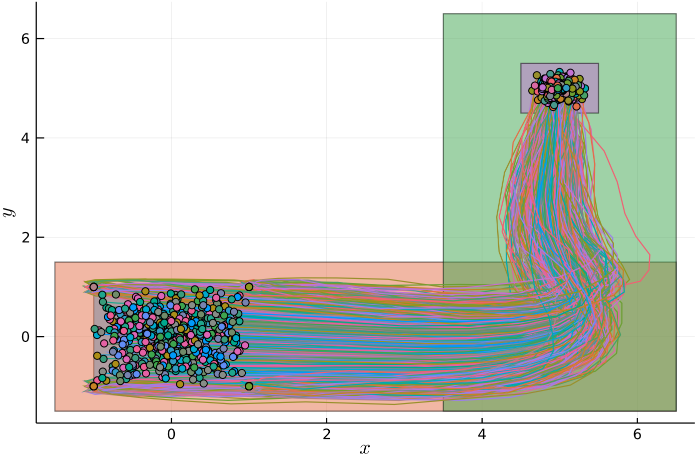

Drone model

- 20 time steps
- 3 measurements
- 4 control actions

Summary

Conclusion

Take home message

Optimize measurement times for state estimation of stochastic dynamical systems in three cases:

1. Minimizing the expected MSE for linear and Gaussian systems
2. Minimizing the expected MSE for nonlinear and non-Gaussian systems
3. Minimizing the worst case estimation error for linear systems subject to bounded disturbances

Potential to improve the treatment of mobile tumors by proton therapy

Conclusion

Take home message

Optimize measurement times for state estimation of stochastic dynamical systems in three cases:

1. Minimizing the expected MSE for linear and Gaussian systems
2. Minimizing the expected MSE for nonlinear and non-Gaussian systems
3. Minimizing the worst case estimation error for linear systems subject to bounded disturbances

Potential to improve the treatment of mobile tumors by proton therapy

Further work

- Extend to networked control systems: co-design the activation times of sensors and actuators and the transmission times between sensors and actuators

Thank you for listening!

Questions?