Bayesian optimisation for improving accuracy and efficiency of cardio-mechanic parameter estimation

Agnieszka Borowska

School of Mathematics and Statistics, University of Glasgow

09.02.2021

Problem

Joint with: Hao Gao, Alan Lazaru and Dirk Husmeier.

Overall goal: to develop an accurate and efficient (fast) method for parameter estimation in cardiac mechanic models of the left ventricle (LV).

Context: estimation of the heart tissue properties from in-vivo clinical measurements.

- Central problem in biomechanical studies of personalized LV modelling.
- Important properties – they:
  - provide insight into heart function/dysfunction,
  - help to inform on the treatment effectiveness post myocardial infarction (heart attack).

Left ventricle

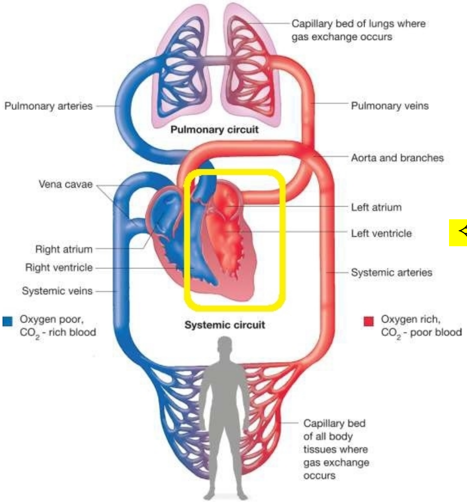

$\Leftarrow$

Left ventricular dysfunction

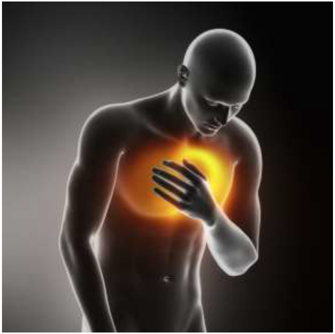

Solution: Bayesian optimisation

Bayesian Optimisation: a method for optimising unknown "black box" objective functions which is

- sequential
- (statistical) model-based
- global

Particularly useful when function evaluations are expensive (Shahriari et al., 2016a).

Outline

Optimisation for cardiac mechanics models

Data: CMR images

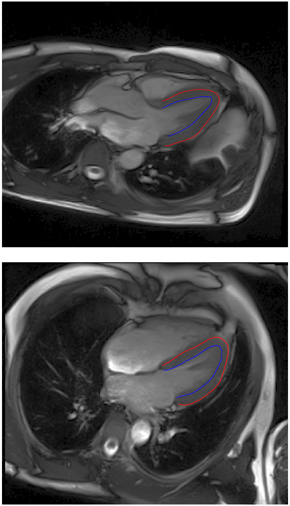

Cardiovascular Magnetic Resonance images

Extracted:

- circumferential strains
- LV cavity volume

(Blue and red lines: LV segmentation)

Cardio-mechanic model

The **myocardium** of the heart can be described by differential equations represented by the Holzapfel and Ogden (2009) law.

Gives a detailed description of the **myocardium response** in diastole.

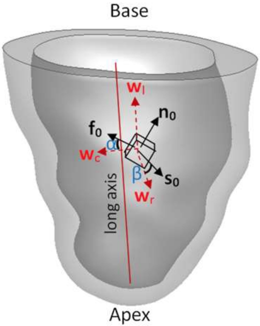

Holzapfel-Ogden law

The strain energy function for the myocardium:

$$\Psi(I_1, I_{4f}, I_{4s}, I_{8fs}) = \frac{a}{2b}\{\exp[b(I_1 - 3)] - 1\} + \sum_{i\in\{f,s\}} \frac{a_i}{2b_i}\{\exp[b_i(I_{4i} - 1)^2] - 1\} + \frac{a_{fs}}{2b_{fs}}[\exp(b_{fs}I_{8fs}^2) - 1],$$

where: I_i, $i \in \{1, 4f, 4s, 8fs\}$ – quantities describing the deformation

$\phi = (a, b, a_f, b_f, a_s, b_s, a_{fs}, b_{fs})^T$ – (unknown) constitutive parameters of interest.

Constitutive parameters

$\phi = (a, b, a_f, b_f, a_s, b_s, a_{fs}, b_{fs})^T$ – (unknown) constitutive parameters of interest.

ϕ_0 – **reference parameters** from Wang et al. (2013):

a [kPa]	b	a_f [kPa]	b_f	a_s [kPa]	b_s	a_{fs} [kPa]	b_{fs}
0.236	10.810	20.037	14.154	3.724	5.164	0.411	11.300

→ **Not in vivo!**

→ **Statistical inference**

Statistical inference

Minimise the mismatch between the data and model predictions

- LV volume in diastole
- 24 circumferential strains

from CMR scans or outputs from the forward simulator.

Forward simulator

Solution to the equations associated with the LV model unavailable in a closed form

⇒ Numerical solutions required

⇒ Finite element method

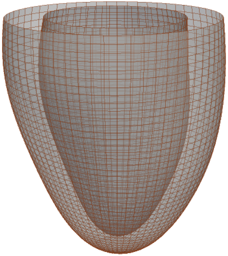

Time consuming: one forward simulation takes ≈15min!

⇒ Standard numerical optimisation or sampling prohibitively expensive.

Existing optimisation algorithm

State-of-the-art optimisation algorithm by Gao et al. (2015, 2018)

- Based on expert domain-specific knowledge
- Multi-step approach
- Each step optimises different subsets of parameters
- Each sub-optimisation using a gradient-based optimisation algorithm

Problem: **slow!**

Bayesian optimisation

BO: an overview

Key ideas

Approximate the costly objective function by a cheaper surrogate function: typically a Gaussian process (GP), see (Rasmussen and Williams, 2006).

Quantify the **exploitation–exploration trade-off** using an acquisition function (to be discussed later).

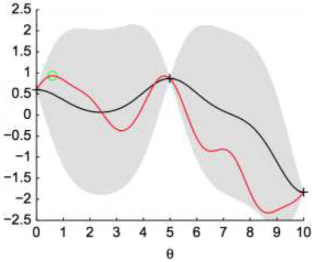

Sequentially update our initial beliefs (prior distribution) about the function of interest after observing the data (likelihood).

Illustration

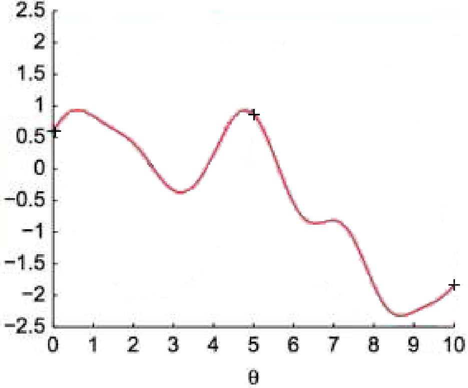

- – Unknown objective function (expensive!)
- + Data points

Here: **likelihood maximisation**

Illustration

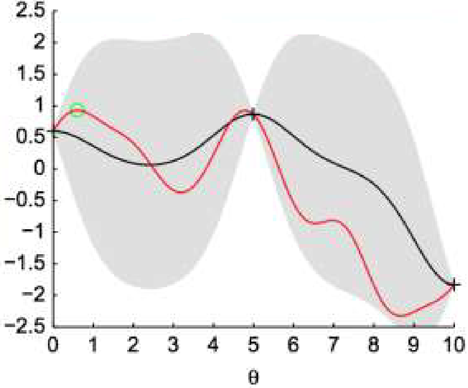

- Surrogate function, typically a GP (cheap!)
- Uncertainty: affects the acquisition function
- × Maximum of acquisition function: exploration–exploitation trade-off

Illustration

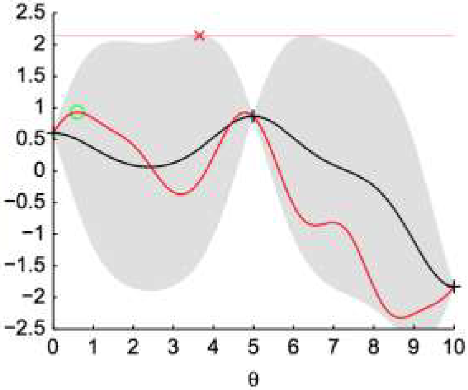

- – Surrogate function, typically a GP (cheap!)
- Uncertainty: affects the acquisition function
- × Maximise the acquisition function: exploration–exploitation trade-off

Illustration

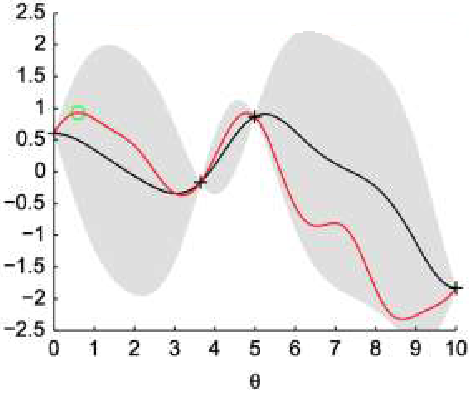

+ Query at the previous maximum ×
  ⇒ uncertainty gets reduced

× Find a new maximum of acquisition function

Illustration

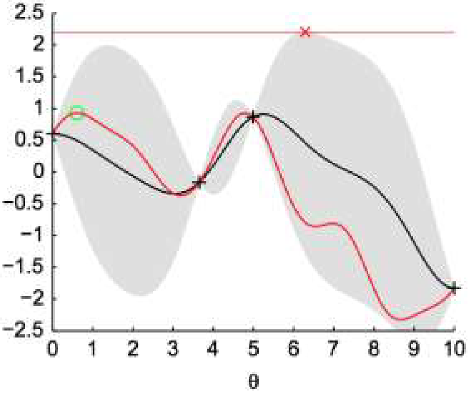

+ Query at the previous maximum ×
  ⇒ uncertainty gets reduced

× Find a new maximum of acquisition function

Illustration

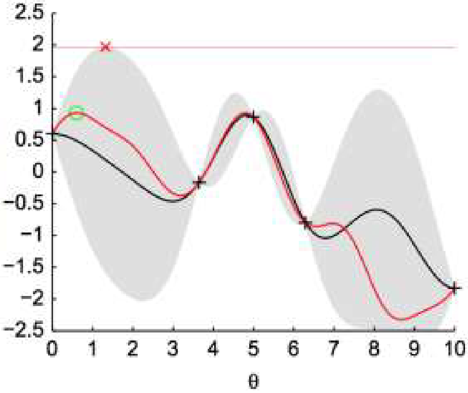

Iterate:

- + Evaluate the objective at the current maximum × (expensive!)
- – Update the surrogate model (cheap!)
- × Find a new maximum of the acquisition function (cheap!)

Illustration

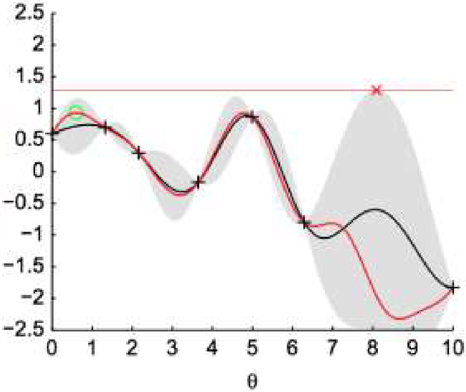

Iterate:

- + Evaluate the objective at the current maximum × (expensive!)
- – Update the surrogate model (cheap!)
- × Find a new maximum of the acquisition function (cheap!)

Illustration

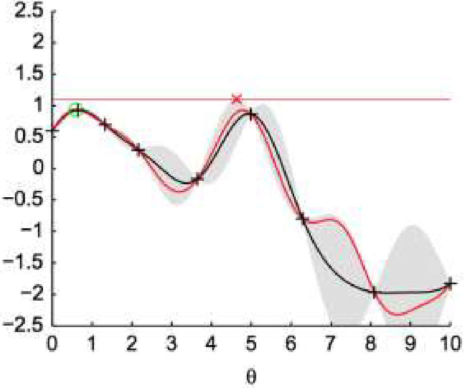

Iterate:

- $+$ Evaluate the objective at the current maximum $\times$ (expensive!)
- – Update the surrogate model (cheap!)
- $\times$ Find a new maximum of the acquisition function (cheap!)

Illustration

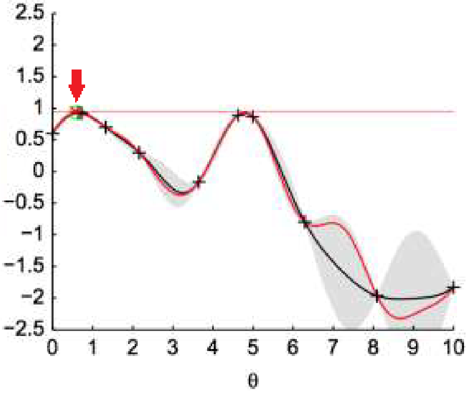

Continue until:
global maximum ⇓

Acquisition functions

- Dictate where to query next – where to carry out the expensive evaluation step.
- Cheap: optimised instead of the true objective function (so-called "inner optimisation").
- Correspond to the exploration-exploitation trade-off.
- Several different AFs have been proposed:
  - probability of improvement
  - expected improvement
  - upper confidence bound
  - entropy search
  - portfolios acquisition functions
  - ...

Expected improvement

$$\mathrm{EI}(\mathbf{x}) = \mathbb{E}_{p(y|\mathbf{x},\mathcal{D})}[\min(f^* - f(\mathbf{x}), 0)],$$

$\mathcal{D}$ – the set of inputs and outputs recorded so far,
f^* – incumbent value i.e. the lowest value of f found so far.

With a GP surrogate $\mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}))$ EI can be expressed as (Jones et al., 1998; Shahriari et al., 2016b)

$$\mathrm{EI}(\mathbf{x}) = (f^* - \mu(\mathbf{x}))\Phi(z) + \sqrt{k(\mathbf{x})}\phi(z),$$

$z = (f^* - \mu(\mathbf{x}))/\sqrt{k(\mathbf{x})}$,
Φ and ϕ are the CDF and PDF of the standard normal distribution, respectively.

BO: our extensions

Extensions

1. Ex-vivo data extended objective function
2. Unknown constraints
3. Partial error surrogates

Re 1: standard objective function

- Standard objective function for **minimisation**:

  mismatch between the simulated values (depending on the constitutive parameter ϕ and LV geometry $\mathcal{H}$) and the measurements:

$$f(\phi, \mathcal{H}) = \underbrace{\frac{(V(\phi, \mathcal{H}) - V^*)^2}{V^*}}_{\text{LV volume}} + \sum_{i=1}^{24} \underbrace{(\varepsilon_i(\phi, \mathcal{H}) - \varepsilon_i^*)^2}_{i\text{th circumferential strain}} .$$

- Measurements only for physiologically typical, low LV pressures (8 mmHg).
- Need for **accurate predictions** also for high LV pressures (30 mmHg): those may reveal LV stiffness with impaired relaxation (which characterises diastolic heart failure).

Re 1: Klotz et al. (2006) curve

But: high pressure volume measurements unavailable in vivo.

We propose to **predict them** using the empirical law found by Klotz et al. (2006) based on **ex vivo data**:

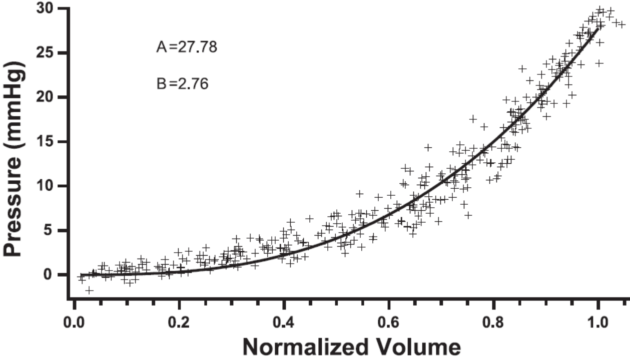

Re 1: high-pressure volume predictions

Normalised end-diastolic volume:

$$\tilde{V}^* = \frac{V^* - V_0}{V_{30} - V_0},$$

where V^* – the measured unnormalised volume at P^*, V_0 – the zero-pressure volume (load-free volume).

Predicted high-pressure end-diastolic volume:

$$\hat{V}_{30}^{Kl} = V_0 + \frac{V^* - V_0}{\tilde{V}^*} = V_0 + \frac{V^* - V_0}{\left(\frac{P^*}{A}\right)^{1/B}}.$$

Re 1: extended objective function

$$f_{\text{Klotz}}(\phi, \mathcal{H}) = \left(\frac{V(\phi, \mathcal{H}) - V^*}{V^*} \right)^2 + \sum_{i=1}^{24} \left(\varepsilon_i(\phi, \mathcal{H}) - \varepsilon_i^* \right)^2 + \left(\frac{V_{30}(\phi, \mathcal{H}) - \hat{V}_{30}^{Kl}}{\hat{V}_{30}^{Kl}} \right)^2,$$

Re 2: unknown constraints

Why? Simulator crashing or failing to terminate for some ϕ.

Solution: weighting the AF, e.g. expected improvement (EI), by the probability of the constraint being satisfied (Snoek, 2013; Gelbart et al., 2014)

$$\mathrm{EI}_{con}(\phi, \mathcal{H}) = \mathrm{EI}(\phi, \mathcal{H})\mathbb{P}(\phi \in \mathcal{C}|\mathcal{H}),$$

where $\mathbb{P}(\phi \in \mathcal{C}|\mathcal{H})$ is the probability of ϕ being a valid point not leading to a crash of the forward simulator for the given LV geometry $\mathcal{H}$.

Re 3: partial error surrogates

Why? The objective functions given as a sum of error terms:

$$f(\phi, \mathcal{H}) = \sum_{i=1}^{K} f^{(i)}(\phi, \mathcal{H}).$$

Standard approach: approximate $f(\mathbf{x})$ using a **single surrogate**.

Potential improvement: approximate the partial errors $f^{(i)}$ using K **surrogates**.

Adjusted EI: based on the conditional posterior mean and variance of the full target $f(\phi, \mathcal{H})$ – given as sums of partial means and variances.

Applications

Set-up

Two studies:

1. Klotz-curve study: with 4 healthy volunteers (HVs)
2. PCA study: one HV + LV geometry reconstructed with different no. of PCA components

Comparison:

- Compare BO with full target and partial surrogates with the state-of-the-art algorithm of Gao et al. (2015, 2018).
- Evaluation based on:
  - speed of convergence (no. of simulator invocations),
  - the final value of the objective function.

Results: convergence of the objective function

Study 1: four HVs

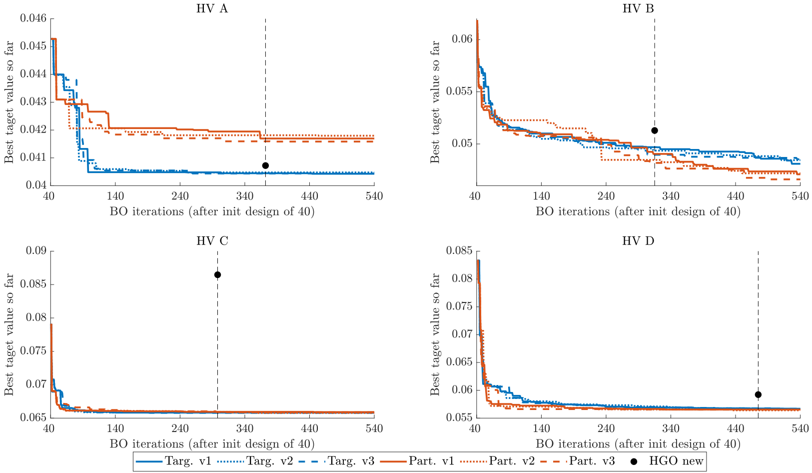

Results: convergence of the objective function

Study 2: one subject + LV geometry reconstructed with different no. of PCA components.

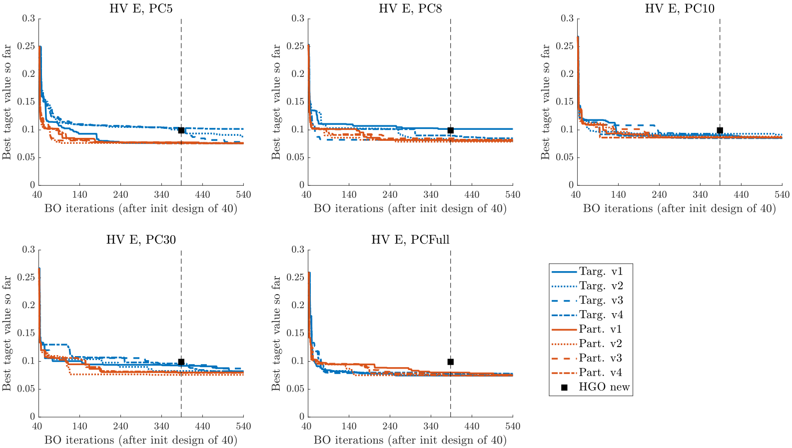

Discussion

Conclusions

- An accurate and efficient Bayesian optimisation–based framework for parameter inference in a cardiac mechanic model of the LV.
- BO converges to lower values of the objective function and requires less invocations of the associated forward simulator than the state-of-the-art multi-step algorithm of Gao et al. (2015, 2018).
- Partial error surrogates: a new approach to minimising a target function given as a sum of error terms.

Discussion

- Better specifications for AF than EI?
    - Information-based policies, e.g. Entropy Search.
    - Portfolios of AFs.

- BO likely to still be too time-consuming to provide a viable tool for the clinical practice
  (optimisation independently for each subject).

  Multi-task BO (Swersky et al., 2013) could address this issue: leveraging prior knowledge from optimisations for previous subjects.

References I

Gao, H, W. G. Li, L. Cai, C. Berry, and X. Y. Luo (2015), "Parameter estimation in a Holzapfel–Ogden law for healthy myocardium." *Journal of Engineering Mathematics*, 95, 231–248.

Gao, H., K. Mangion, C. Berry, and X. Y. Luo (2018), "Mathematical modelling acute myocardial infarction using in vivo magnetic resonance imaging." *Proceeding of virtual physiological human conference, Zaragoza Spain.*

Gelbart, M. A., J. Snoek, and R. P. Adams (2014), "Bayesian optimization with unknown constraints." In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, 250–259.

Holzapfel, G. A. and R. W. Ogden (2009), "Constitutive Modelling of Passive Myocardium: a Structurally Based Framework for Material Characterization." *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 367, 3445–3475.

Jones, D. R., M. Schonlau, and W. J. Welch (1998), "Efficient global optimization of expensive black-box functions." *Journal of Global optimization*, 13, 455–492.

Klotz, S., I. Hay, M. L. Dickstein, G. H. Yi, J. Wang, M. S. Maurer, D. A. Kass, and D. Burkhoff (2006), "Single-beat estimation of end-diastolic pressure-volume relationship: a novel method with potential for noninvasive application." *American Journal of Physiology-Heart and Circulatory Physiology*, 291.

Rasmussen, C. E. and C. K. Williams (2006), *Gaussian Processes for Machine Learning*, volume 1. MIT press Cambridge.

Shahriari, B., K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas (2016a), "Taking the Human Out of the Loop: A Review of Bayesian Optimization." *Proceedings of the IEEE*, 104.

References II

Shahriari, B., K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas (2016b), "Taking the human out of the loop: a review of Bayesian optimization." *Proceedings IEEE*, 104, 148–175.

Snoek, J. R. (2013), *Bayesian otimization and semiparametric models with applications to assistive technology*. Ph.D. thesis, University of Toronto, Toronto, Canada.

Swersky, Kevin, Jasper Snoek, and Ryan P Adams (2013), "Multi-task Bayesian optimization." In *Advances in Neural Information Processing Systems*, 2004–2012.

Wang, H. M., H. Gao, X. Y. Luo, C. Berry, B. E. Griffith, R. W. Ogden, and T. J. Wang (2013), "Structure Based Finite Strain Modelling of the Human Left Ventricle in Diastole." *International Journal for Numerical Methods in Biomedical Engineering*, 29, 83–103.