

FEniCS Course

Lecture 13: Introduction to dolfin-adjoint

Contributors

Simon Funke

Patrick Farrell

Marie E. Rognes

Computing sensitivities

So far we focused on solving PDEs.

But often we are also interested the sensitivity with respect to certain parameters, for example

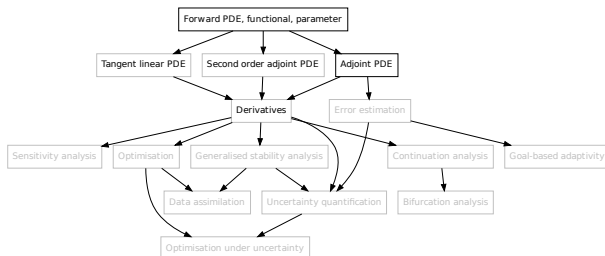
- initial conditions,
- forcing terms,
- unknown coefficients.

Computing sensitivities

So far we focused on solving PDEs.

But often we are also interested the sensitivity with respect to certain parameters, for example

- initial conditions,
- forcing terms,
- unknown coefficients.

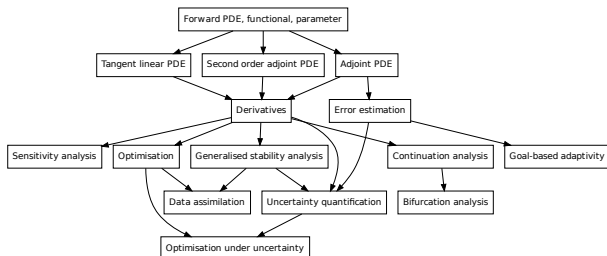


Computing sensitivities

So far we focused on solving PDEs.

But often we are also interested the sensitivity with respect to certain parameters, for example

- initial conditions,
- forcing terms,
- unknown coefficients.



Example

Consider the Poisson's equation

$$\begin{aligned} -\nu \Delta u &= m && \text{in } \Omega, \\ u &= 0 && \text{on } \partial\Omega, \end{aligned}$$

together with the *objective functional*

$$J(u) = \frac{1}{2} \int_{\Omega} \|u - u_d\|^2 dx,$$

where u_d is a known function.

Goal

Compute the sensitivity of J with respect to the *parameter* m : dJ/dm .

Comput. deriv. (i) General formulation

Given

- Parameter m ,
- PDE $F(u, m) = 0$ with solution u .
- Objective functional $J(u, m) \rightarrow \mathbb{R}$,

Goal

Compute dJ/dm .

Reduced functional

Consider u as an implicit function of m by solving the PDE.
With that we define the *reduced functional* R :

$$R(m) = J(u(m), m)$$

Comput. deriv. (i) General formulation

Given

- Parameter m ,
- PDE $F(u, m) = 0$ with solution u .
- Objective functional $J(u, m) \rightarrow \mathbb{R}$,

Goal

Compute dJ/dm .

Reduced functional

Consider u as an implicit function of m by solving the PDE.
With that we define the *reduced functional* R :

$$R(m) = J(u(m), m)$$

Comput. deriv. (i) General formulation

Given

- Parameter m ,
- PDE $F(u, m) = 0$ with solution u .
- Objective functional $J(u, m) \rightarrow \mathbb{R}$,

Goal

Compute dJ/dm .

Reduced functional

Consider u as an implicit function of m by solving the PDE.
With that we define the *reduced functional* R :

$$R(m) = J(u(m), m)$$

Comput. deriv. (ii) Reduced functional

Reduced functional:

$$R(m) \equiv J(u(m), m).$$

Taking the derivative of with respect to m yields:

$$\frac{dR}{dm} = \frac{dJ}{dm} = \frac{\partial J}{\partial u} \frac{du}{dm} + \frac{\partial J}{\partial m}.$$

Computing $\frac{\partial J}{\partial u}$ and $\frac{\partial J}{\partial m}$ is straight-forward, but how handle $\frac{du}{dm}$?

Comput. deriv. (ii) Reduced functional

Reduced functional:

$$R(m) \equiv J(u(m), m).$$

Taking the derivative of with respect to m yields:

$$\frac{dR}{dm} = \frac{dJ}{dm} = \frac{\partial J}{\partial u} \frac{du}{dm} + \frac{\partial J}{\partial m}.$$

Computing $\frac{\partial J}{\partial u}$ and $\frac{\partial J}{\partial m}$ is straight-forward, but how handle $\frac{du}{dm}$?

Comput. deriv. (ii) Reduced functional

Reduced functional:

$$R(m) \equiv J(u(m), m).$$

Taking the derivative of with respect to m yields:

$$\frac{dR}{dm} = \frac{dJ}{dm} = \frac{\partial J}{\partial u} \frac{du}{dm} + \frac{\partial J}{\partial m}.$$

Computing $\frac{\partial J}{\partial u}$ and $\frac{\partial J}{\partial m}$ is straight-forward, but how handle $\frac{du}{dm}$?

Comput. deriv. (iii) Computing $\frac{du}{dm}$

Taking the derivative of $F(u, m) = 0$ with respect to m yields:

$$\frac{dF}{dm} = \frac{\partial F}{\partial u} \frac{du}{dm} + \frac{\partial F}{\partial m} = 0$$

Hence:

$$\frac{du}{dm} = - \left(\frac{\partial F}{\partial u} \right)^{-1} \frac{\partial F}{\partial m}$$

Final formula for functional derivative

$$\frac{dJ}{dm} = - \overbrace{\frac{\partial J}{\partial u} \left(\frac{\partial F}{\partial u} \right)^{-1} \frac{\partial F}{\partial m}}^{\text{adjoint PDE}} + \frac{\partial J}{\partial m},$$

tangent linear PDE

Comput. deriv. (iii) Computing $\frac{du}{dm}$

Taking the derivative of $F(u, m) = 0$ with respect to m yields:

$$\frac{dF}{dm} = \frac{\partial F}{\partial u} \frac{du}{dm} + \frac{\partial F}{\partial m} = 0$$

Hence:

$$\frac{du}{dm} = - \left(\frac{\partial F}{\partial u} \right)^{-1} \frac{\partial F}{\partial m}$$

Final formula for functional derivative

$$\frac{dJ}{dm} = - \overbrace{\frac{\partial J}{\partial u} \left(\frac{\partial F}{\partial u} \right)^{-1} \frac{\partial F}{\partial m}}^{\text{adjoint PDE}} + \frac{\partial J}{\partial m},$$

tangent linear PDE

Comput. deriv. (iii) Computing $\frac{du}{dm}$

Taking the derivative of $F(u, m) = 0$ with respect to m yields:

$$\frac{dF}{dm} = \frac{\partial F}{\partial u} \frac{du}{dm} + \frac{\partial F}{\partial m} = 0$$

Hence:

$$\frac{du}{dm} = - \left(\frac{\partial F}{\partial u} \right)^{-1} \frac{\partial F}{\partial m}$$

Final formula for functional derivative

$$\frac{dJ}{dm} = - \overbrace{\frac{\partial J}{\partial u} \left(\frac{\partial F}{\partial u} \right)^{-1} \frac{\partial F}{\partial m}}^{\text{adjoint PDE}} + \frac{\partial J}{\partial m},$$

tangent linear PDE

Dimensions of a finite dimensional example

$$\frac{dJ}{dm} = \underbrace{\left[-\frac{\partial J}{\partial u} \right]}_{\text{discretised adjoint PDE}} \times \underbrace{\left[\left(\frac{\partial F}{\partial u} \right)^{-1} \right]}_{\text{discretised tangent linear PDE}} \times \left[\frac{\partial F}{\partial m} \right] + \left[\frac{\partial J}{\partial m} \right]$$

The tangent linear solution is a matrix of dimension $|u| \times |m|$ and requires the solution of m linear systems. The adjoint solution is a vector of dimension $|u|$ and requires the solution of one linear systems.

Adjoint approach

- 1 Solve the adjoint equation for λ

$$\frac{\partial F^*}{\partial u} \lambda = -\frac{\partial J^*}{\partial u}.$$

- 2 Compute

$$\frac{dJ}{dm} = \lambda^* \frac{\partial F}{\partial m} + \frac{\partial J}{\partial m}.$$

The computational expensive part is (1). It requires solving the (linear) adjoint PDE, and its cost is independent of the choice of parameter m .

What is dolfin-adjoint?


Dolfin-adjoint is an extension of FEniCS for: solving adjoint and tangent linear equations; generalised stability analysis; PDE-constrained optimisation.

Main features

- Automated derivation of first and second order adjoint and tangent linear models.
- Discretely consistent derivatives.
- Parallel support and near theoretically optimal performance.
- Interface to optimisation algorithms for PDE-constrained optimisation.
- Documentation and examples on www.dolfin-adjoint.org.

What has dolfin-adjoint been used for?



Layout optimisation of tidal turbines



- Up to 400 tidal turbines in one farm.
- What are the optimal locations to maximise power production?

What has dolfin-adjoint been used for?

Layout optimisation of tidal turbines



What has dolfin-adjoint been used for?

Layout optimisation of tidal turbines

Python code


```
from dolfin import *
from dolfin_adjoint import *

# FEniCS model
# ...

J = Functional(turbines*inner(u, u)**(3/2)*dx*dt)
m = Control(turbine_positions)
R = ReducedFunctional(J, m)
maximize(R)
```

What has dolfin-adjoint been used for?

Reconstruction of a tsunami wave




Is it possible to reconstruct a tsunami wave from images like this?


¹Image: ASTER/NASA PIA06671

What has dolfin-adjoint been used for?


Reconstruction of a tsunami wave



Correct tsunami wave



Reconstructed tsunami wave



Reconstruction of a tsunami wave

Python code

```
from fenics import *
from dolfin_adjoint import *

# FEniCS model
# ...

J = Functional(observation_error**2*dx*dt)
m = Control(input_wave)
R = ReducedFunctional(J, m)
minimize(R)
```

Other applications

Dolfin-adjoint has been applied to lots of other cases, and works for many PDEs:

Some PDEs we have adjoined

- Burgers
- Navier-Stokes
- Stokes + mantle rheology
- Stokes + ice rheology
- Saint Venant + wetting/drying
- Cahn-Hilliard
- Gray-Scott
- Shallow ice
- Blatter-Pattyn
- Quasi-geostrophic
- Viscoelasticity
- Gross-Pitaevskii
- Yamabe
- Image registration
- Bidomain
- ...

Example

Compute the sensitivity of

$$J(u) = \int_{\Omega} \|u - u_d\|^2 dx$$

with known u_d and the Poisson equation:

$$\begin{aligned} -\nu \Delta u &= m && \text{in } \Omega \\ u &= 0 && \text{on } \partial\Omega. \end{aligned}$$

with respect to m .

Poisson solver in FEniCS

An implementation of the Poisson's equation might look like this:

Python code

```
from fenics import *

# Define mesh and finite element space
mesh = UnitSquareMesh(50, 50)
V = FunctionSpace(mesh, "Lagrange", 1)

# Define basis functions and parameters
u = TrialFunction(V)
v = TestFunction(V)
m = interpolate(Constant(1.0), V)
nu = Constant(1.0)

# Define variational problem
a = nu*inner(grad(u), grad(v))*dx
L = m*v*dx
bc = DirichletBC(V, 0.0, "on_boundary")

# Solve variational problem
u = Function(V)
solve(a == L, u, bc)
plot(u, title="u")
```

Dolfin-adjoint (i): Annotation

The first change necessary to adjoin this code is to import the dolfin-adjoint module *after* importing DOLFIN:

Python code

```
from fenics import *  
from dolfin_adjoint import *
```

With this, dolfin-adjoint will record each step of the model, building an *annotation*. The annotation is used to symbolically manipulate the recorded equations to derive the tangent linear and adjoint models.

In this particular example, the `solve` function method will be recorded.

Dolphin-adjoint (ii): Objective Functional

Next, we implement the objective functional, the square L^2 -norm of $u - u_d$:

$$J(u) = \int_{\Omega} \|u - u_d\|^2 dx$$

or in code

Python code

```
j = inner(u - u_d, u - u_d)*dx
J = Functional(j)
```

Dolfin-adjoint (ii): Control parameter

Next we need to decide which parameter we are interested in. Here, we would like to investigate the sensitivity with respect to the source term m .

We inform dolfin-adjoint of this:

Python code

```
m = Control(m)
```

Dolfin-adjoint (iii): Computing gradients

Now, we can compute the gradient with:

Python code

```
dJdm = compute_gradient(J, m, project=True)
```

Dolfin-adjoint derives and solves the adjoint equations for us and returns the gradient.

Note

If you call `compute_gradient` more than once, you need to pass `forget=False` as a parameter. Otherwise you get an error: *Need a value for u_1:0:0:Forward, but don't have one recorded.*

Computational cost

Computing the gradient requires one adjoint solve.

Dolfin-adjoint (iii): Computing gradients

Now, we can compute the gradient with:

Python code

```
dJdm = compute_gradient(J, m, project=True)
```

Dolfin-adjoint derives and solves the adjoint equations for us and returns the gradient.

Note

If you call **compute_gradient** more than once, you need to pass *forget=False* as a parameter. Otherwise you get an error: *Need a value for u_1:0:0:Forward, but don't have one recorded.*

Computational cost

Computing the gradient requires one adjoint solve.

Dolfin-adjoint (iii): Computing gradients

Now, we can compute the gradient with:

Python code

```
dJdm = compute_gradient(J, m, project=True)
```

Dolfin-adjoint derives and solves the adjoint equations for us and returns the gradient.

Note

If you call **compute_gradient** more than once, you need to pass *forget=False* as a parameter. Otherwise you get an error: *Need a value for u_1:0:0:Forward, but don't have one recorded.*

Computational cost

Computing the gradient requires one adjoint solve.

Dolfin-adjoint (iii): Computing Hessians

Dolfin-adjoint can also compute the second derivatives (Hessians):

Python code

```
H = hessian(J, m)
direction = interpolate(Constant(1), V)
plot(H(direction))
```

Computational cost

Computing the directional second derivative requires one tangent linear and two adjoint solves.

Dolfin-adjoint (iii): Computing Hessians

Dolfin-adjoint can also compute the second derivatives (Hessians):

Python code

```
H = hessian(J, m)
direction = interpolate(Constant(1), V)
plot(H(direction))
```

Computational cost

Computing the directional second derivative requires one tangent linear and two adjoint solves.

Verification

How can you check that the gradient is correct?

Taylor expansion of the reduced functional R in a perturbation δm yields:

$$|R(m + \epsilon \delta m) - R(m)| \rightarrow 0 \quad \text{at } \mathcal{O}(\epsilon)$$

but

$$|R(m + \epsilon \delta m) - R(m) - \epsilon \nabla R \cdot \delta m| \rightarrow 0 \quad \text{at } \mathcal{O}(\epsilon^2)$$

Taylor test

Choose $m, \delta m$ and determine the convergence rate by reducing ϵ . If the convergence order with gradient is ≈ 2 , your gradient is probably correct.

The function `taylor_test` implements the Taylor test for you. See `help(taylor_test)`.

Verification

How can you check that the gradient is correct?

Taylor expansion of the reduced functional R in a perturbation δm yields:

$$|R(m + \epsilon \delta m) - R(m)| \rightarrow 0 \quad \text{at } \mathcal{O}(\epsilon)$$

but

$$|R(m + \epsilon \delta m) - R(m) - \epsilon \nabla R \cdot \delta m| \rightarrow 0 \quad \text{at } \mathcal{O}(\epsilon^2)$$

Taylor test

Choose $m, \delta m$ and determine the convergence rate by reducing ϵ . If the convergence order with gradient is ≈ 2 , your gradient is probably correct.

The function `taylor_test` implements the Taylor test for you. See `help(taylor_test)`.

Getting started with Dolfin-adjoint

- 1 Compute the gradient and Hessian of the Poisson example with respect to m .
- 2 Run the Taylor test to check that the gradient is correct.
- 3 Measure the computation time for the forward, gradient and Hessian computation. What do you observe? Hint: Use `help(Timer)`.