

# Differentiable programming accross the PDE/ML divide

---

David A. Ham<sup>1</sup>, Nacime Bouziani<sup>1</sup>, India Marsden<sup>2</sup>, and Reuben Nixon-Hill<sup>1</sup>

November 2023

<sup>1</sup>Department of Mathematics, Imperial College London <sup>2</sup>Mathematical Institute, University of Oxford



- solving nonlinear PDEs
- computing sensitivities
- data assimilation
- design optimisation
- training neural nets
- ...



Not a new idea ...

*LM Beda, LN Korolev, NV Sukkikh, and TS Frolova.* Programs for automatic differentiation for the machine besm.(in russian) technical report, institute for precise mechanics and computation techniques, 1959

by 1981 there are textbooks:

*Louis B Rall. Automatic differentiation: Techniques and applications.*

**Springer, 1981**

So automatic/algorithmic differentiation is nearly as old as computing.



*"Constructing neural networks using pure and higher-order differentiable functions and training them using reverse-mode automatic differentiation is unsurprisingly called Differentiable Programming."*

Erik Meijer. Behind every great deep learning framework is an even greater programming languages concept (keynote).

**In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/FSE 2018*, page 1, New York, NY, USA, 2018. Association for Computing Machinery**



**Abstract** Define symbolic representations for numerical objects and algorithms.

**Compose** Form larger algorithms by plugging together smaller ones.



## Claims:

1. There is a useful extension of this concept of differentiable programming to encompass simulation.
2. FEniCS and Firedrake + pyadjoint are examples<sup>1</sup>.
3. Using this insight, we can naturally extend packages such as this to interact better with external (non-PDE) processes and data.

---

<sup>1</sup>Wilkinson Prize 2015. 2011 prize was to Waechter and Laird for IPopt.

## So you want to solve a PDE using finite elements



1. Write down a residual, boundary/initial conditions, forcings, parametrisations.
2. Choose suitable finite element paces and quadrature rules.
3. Choose a suitable (non)-linear solver and preconditioning strategy.
4. Derive and implement the loops over elements, facets, basis functions, and quadrature points.
5. Implement parallel communication.
6. Implement and compose solvers and preconditioners.
7. Now do it all again for the adjoint.
8. ...

## So you want to solve a PDE using finite elements



1. Write down a residual, boundary/initial conditions, forcings, parametrisations.
2. Choose suitable finite element paces and quadrature rules.
3. Choose a suitable (non)-linear solver and preconditioning strategy.
4. ~~Derive and implement the loops over elements, facets, basis functions, and quadrature points.~~
5. ~~Implement parallel communication.~~
6. ~~Implement and compose solvers and preconditioners.~~
7. ~~Now do it all again for the adjoint.~~
8. ...

About 20 years ago FEniCS worked out how to do this (Kirby & Logg 2006).

10 years later Firedrake joined the party (Rathgeber et al. 2016).





Burgers Equation:

$$\frac{\partial u}{\partial t} + (u \cdot \nabla)u - \nu \nabla^2 u = 0 \quad (1)$$

$$(n \cdot \nabla)u = 0 \text{ on } \Gamma \quad (2)$$

in weak form: find  $u \in V$  such that

$$\int_{\Omega} \frac{\partial u}{\partial t} \cdot v + ((u \cdot \nabla)u) \cdot v + \nu \nabla u \cdot \nabla v \, dx = 0 \quad \forall v \in V_0. \quad (3)$$

For simplicity, use backward Euler in time. At each timestep find  $u^{n+1} \in V_0$  such that:

$$\int_{\Omega} \frac{u^{n+1} - u^n}{dt} \cdot v + ((u^{n+1} \cdot \nabla)u^{n+1}) \cdot v + \nu \nabla u^{n+1} \cdot \nabla v \, dx = 0 \quad \forall v \in V_0. \quad (4)$$



```
1  from firedrake import *
2  n = 30
3  mesh = UnitSquareMesh(n, n)
4  V = VectorFunctionSpace(mesh, "CG", 2)
5  u_ = Function(V, name="Velocity")
6  u = Function(V, name="VelocityNext")
7  v = TestFunction(V)
8  x = SpatialCoordinate(mesh)
9  ic = project(as_vector([sin(pi*x[0]), 0]), V)
10 u_.assign(ic)
11 u.assign(ic)
12 nu = 0.0001
13 timestep = 1.0/n
14 F = (inner((u - u_)/timestep, v) + inner(dot(u,nabla_grad(u)), v) + nu*inner(grad(u), grad(v)))*dx
15 t = 0.0
16 end = 0.5
17 while (t <= end):
18     solve(F == 0, u) # <= all the magic happens here.
19     u_.assign(u)
20     t += timestep
```

UFL and the FEniCS language were created by the FEniCS project. See Logg et al. 2012  
*Automated Solution of Differential Equations by the Finite Element Method*



$$\int_{\Omega} \frac{u^{n+1} - u^n}{dt} \cdot v + ((u^{n+1} \cdot \nabla)u^{n+1}) \cdot v + \nu \nabla u^{n+1} \cdot \nabla v \, dx$$

```
(inner((u - u_)/timestep, v) + inner(dot(u,nabla_grad(u)), v) \
+ nu*inner(grad(u), grad(v)))*dx
```



We solve PDEs with Newton-like methods:

$$u_{\text{next}} = u_{\text{cur}} - \left( \frac{\partial F(u_{\text{cur}})}{\partial u} \right)^{-1} F(u_{\text{cur}})$$

So our solver is the composition of a Newton-like algorithm with functions that assemble the residual  $F$  and the Jacobian  $\partial F/\partial u$ .



If  $V$  is a real Hilbert space with inner product  $\langle \cdot \cdot \rangle_V$  then  $V^*$  is the space of bounded linear functionals  $V \rightarrow \mathbb{R}$ .

The form, given  $u \in V$ :

$$\int_{\Omega} u \cdot v dx \quad \forall v \in V \quad (5)$$

is a function  $V \rightarrow V^*$ . This is exactly the form of the residual in a steady PDE.



This is exactly Meijer's conception of differentiable programming.  $F$  is a differentiable operator and the Newton solver is a higher order function.

We can get technical with their signatures:

$$F : V \rightarrow V^* \tag{6}$$


$$\text{Newton} : (V \rightarrow V^*) : V \rightarrow V \tag{7}$$



We need to differentiate our residual,  $F$  with respect to  $u$ . How does a computer do that? Take the nonlinear term from Burgers' equation as an example. You write:


```
inner(dot(u,nabla_grad(u)), v)
```

But the computer sees:









$$\frac{\partial(u \cdot \nabla u) \cdot v}{\partial u} \cdot \tilde{u} = ?$$










# A little symbolic differentiation



# A little symbolic differentiation



# A little symbolic differentiation



$$\frac{\partial(u \cdot \nabla u) \cdot v}{\partial u} \cdot \tilde{u} = (\tilde{u} \cdot \nabla u + u \cdot \nabla \tilde{u}) \cdot v$$



Turns out scientists and engineers usually solve inverse problems:

- Sensitivity analysis,
- Parameter estimation,
- Design optimisation,
- Data assimilation.

Common to all of these is a requirement to differentiate the model.

This work rests on Pyadjoint by Mitusch, Funke and Dokken, and Dolfin-adjoint by Farrell, Funke, Ham and Rognes.



At each timestep find  $u^{n+1} \in V_0$  such that:

$$\int_{\Omega} \frac{u^{n+1} - u^n}{dt} \cdot v + ((u^{n+1} \cdot \nabla)u^{n+1}) \cdot v + \nu \nabla u^{n+1} \cdot \nabla v \, dx = 0 \quad \forall v \in V_0. \quad (8)$$

Let's write a simple functional:

$$J(u) = \int_{\Omega} u_{t=\text{end}}^2 + u_{t=0}^2 \, dx$$

and assume that we want to differentiate this with respect to the initial condition  $u_{t=0}$ .  
What would that do to our code?



**Symbolic differentiation** Derivative of all outputs with respect to all inputs at an arbitrary state.

**Forward mode/tangent linear model** Derivative of all outputs with respect to one input at a single given state.


**Reverse mode/adjoint/backpropagation** Derivative of a single output with respect to all inputs at a single given state.




```
1 from firedrake import *
2 from firedrake.adjoint import *
3 continue_annotation()
4 n = 30
5 mesh = UnitSquareMesh(n, n)
6 V = VectorFunctionSpace(mesh, "CG", 2)
7 u_ = Function(V, name="Velocity")
8 u = Function(V, name="VelocityNext")
9 v = TestFunction(V)
10 x = SpatialCoordinate(mesh)
11 ic = project(as_vector([sin(pi*x[0]), 0]), V)
12 u_.assign(ic)
13 u.assign(ic)
14 nu = 0.0001
15 timestep = 1.0/n
16 F = (inner((u - u_)/timestep, v) + inner(dot(u,nabla_grad(u)), v) + nu*inner(grad(u), grad(v)))*dx
17 t = 0.0
18 end = 1.0
19 while (t <= end):
20     solve(F == 0, u)
21     u_.assign(u)
22     t += timestep
23 J = assemble(u*u*dx + ic*ic*dx)
    compute_gradient(J, Control(ic))
```



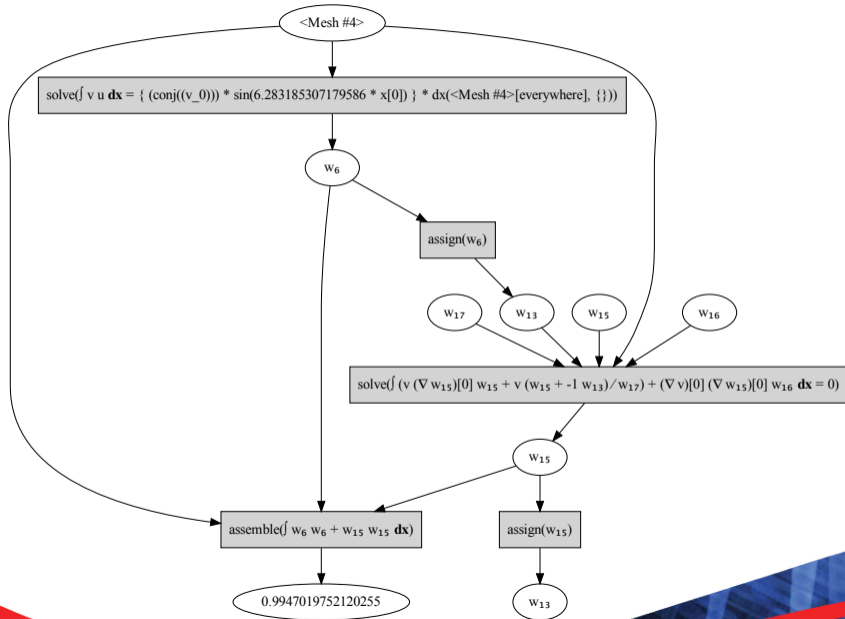




$$\text{solve}(\int[\text{rest of domain}] (v (\text{grad } w_{17})[0] w_{17} + v (w_{17} + -1 w_{15})/w_{19}) + (\text{grad } v)[0] (\text{grad } w_{17})[0] w_{18} \, dx = 0)$$



$$\text{solve}(\int[\text{rest of domain}] (v (\text{grad } w_{17})[0] w_{17} + v (w_{17} + -1 w_{15})/w_{19}) + (\text{grad } v)[0] (\text{grad } w_{17})[0] w_{18} \, dx = 0)$$





For each operation:

$$y = f(x) \tag{9}$$

We compute the adjoint (transpose derivative):

$$x' = \left( \frac{\partial f}{\partial x}(x) \right)^* y' \tag{10}$$



Solve:

$$F(u; v) = 0 \quad (11)$$

the implicit function theorem gives us:

$$u' = \left( \frac{\partial F}{\partial u}(u; \tilde{u}, v) \right)^{-*} \lambda \quad (12)$$

$$= \left( \frac{\partial F}{\partial u}(u; v, \tilde{u}) \right)^{-1} \lambda \quad (13)$$



Back to the function signatures, expanded for unsteady:

$$F : \underbrace{V}_{u_{\text{old}}} \times \underbrace{V}_{u_{\text{new}}} \rightarrow V^* \quad (14)$$

$$\text{Newton} : (V \times V \rightarrow V^*) : \underbrace{V}_{u_{\text{old}}} \rightarrow \underbrace{V}_{u_{\text{new}}} \quad (15)$$

$$\text{Newton}^* : (V \rightarrow V^*) : \underbrace{V}_{u_{\text{old}}} \times \underbrace{V}_{u_{\text{new}}} \times \underbrace{V^*}_{u'_{\text{new}}} \rightarrow \underbrace{V^*}_{u'_{\text{old}}} \quad (16)$$



Suppose what I need is find  $u \in V$  such that:

$$F(u, N(u); v) = 0 \quad \forall v \in V^* \quad (17)$$

Where  $F$  is a PDE residual but  $N$  is not.  $N$  could be a parametrisation whose value is give by e.g.:

1. Solving an ODE at each point.
2. Solving an algebraic equation at each point.
3. Evaluating a neural net.



$$\min_{u \in V} \|u - u_{\text{obs}}\| + N(u) \quad (18)$$

Subject to:

$$F(u; v) = 0 \quad \forall v \in V \quad (19)$$

Where  $N$  is a regularisation term e.g. found by evaluating a neural net.





UFL External operators:

1. Symbolic behaviour (given by calculus rules).
2. Numerical implementation (not UFL's problem).



UFL is the Unified **Form** Language:

$$V \times W \rightarrow \mathbb{R} \quad (20)$$

but the external operators we're interested in aren't (obviously) forms:

$$V \rightarrow W \quad (21)$$



A really simple idea: we can apply function arguments one at a time:

$$V \times W \rightarrow \mathbb{R} \equiv V \rightarrow (W \rightarrow \mathbb{R}) \quad (22)$$



A really simple idea: we can apply function arguments one at a time:

$$V \times W \rightarrow \mathbb{R} \equiv V \rightarrow (W \rightarrow \mathbb{R}) \quad (22)$$

$$\equiv V \rightarrow W^* \quad (23)$$

So any form is an operator into the dual space of its last argument.

## But we need the other direction



Happily, all the spaces we care about are reflexive:

$$V \Leftrightarrow V^{**} (\equiv V^* \rightarrow \mathbb{R}) \quad (24)$$

by identifying  $v^{**} \in V^{**}$  with  $v \in V$  such that:

$$v^{**}(u^*) = u^*(v) \quad \forall u^* \in V^* \quad (25)$$

## But we need the other direction



Happily, all the spaces we care about are reflexive:

$$V \Leftrightarrow V^{**} (\equiv V^* \rightarrow \mathbb{R}) \quad (24)$$

by identifying  $v^{**} \in V^{**}$  with  $v \in V$  such that:

$$v^{**}(u^*) = u^*(v) \quad \forall u^* \in V^* \quad (25)$$

Hence:

$$V \rightarrow W \equiv V \rightarrow W^{**} \quad (26)$$

$$\equiv V \rightarrow (W^* \rightarrow \mathbb{R}) \quad (27)$$

$$\equiv V \times W^* \rightarrow \mathbb{R} \quad (28)$$



```
1 X = FunctionSpace(...)
2 V = FunctionSpace(...)
3 u = Coefficient(V)
4 m = Coefficient(V)
5 v = TestFunction(V)
6 uhat = TrialFunction(V)
7
8 # N : V x V x V* -> R
9 #     u, m, v* -> N(u, m; v*)
10 N = ExternalOperator(u, m, function_space=X)
11
12 # Define a given form F
13 F = u * N * v * dx
14
15 # Symbolically compute the derivative  $\frac{\partial N(u, m; \hat{u}, v^*)}{\partial u}$ 
16 dNdu = derivative(N, u, uhat)
17
18 # Symbolically compute the derivative  $\frac{dF}{du}$ 
19 dFdu = derivative(F, u, uhat)
```

## External operators obey the right symbolic rules.



UFL expression	ExternalOperator	Derivatives	Argument slots	Output type
$N$	$N(u, m; v^*)$	(0, 0)	$(v^*, )$	Coefficient
$dNdu = \text{derivative}(N, u, \hat{u})$	$\frac{\partial N(u, m; \hat{u}, v^*)}{\partial u}$	(1, 0)	$(v^*, \hat{u})$	Matrix
$dNdm = \text{derivative}(N, m, \hat{m})$	$\frac{\partial N(u, m; \hat{m}, v^*)}{\partial m}$	(0, 1)	$(v^*, \hat{m})$	Matrix
$\text{action}(dNdu, w)$	$\frac{\partial N(u, m; w, v^*)}{\partial u}$	(1, 0)	$(v^*, w)$	Coefficient
$\text{adjoint}(dNdm)$	$\frac{\partial N(u, m; v^*, \hat{m})}{\partial m}$	(0, 1)	$(\hat{m}, v^*)$	Matrix
$\text{action}(\text{adjoint}(dNdm), \tilde{v})$	$\frac{\partial N(u, m; \tilde{v}, \hat{m})}{\partial m}$	(0, 1)	$(\hat{m}, \tilde{v})$	Coefficient

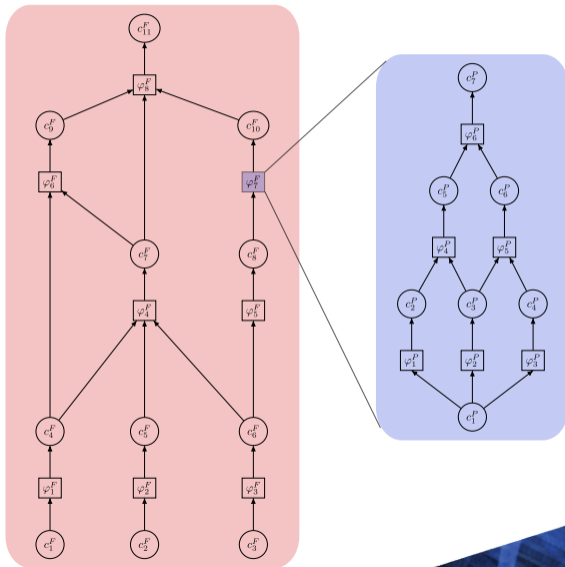


# The user provides the implementation



```
1 class MyExternalOperator(AbstractExternalOperator):
2     def __init__(self, *args, **kwargs):
3         ...
4     @assemble_method((0, 0), (0,))
5     # or @assemble_method(0, (0,))
6     def N(self, *args, **kwargs):
7         """Evaluate my external operator N"""
8         ...
9     @assemble_method((1, 0), (0, 1))
10    def dNdu(self, *args, **kwargs):
11        """Evaluate dNdu"""
12        ...
13    @assemble_method((1, 0), (0, None))
14    def dNdu_action(self, *args, **kwargs):
15        """Evaluate the action of dNdu"""
16        ...
17    @assemble_method((0, 1), (1, 0))
18    def dNdm_adjoint(self, *args, **kwargs):
19        """Evaluate dNdm*"""
20        ...
21    @assemble_method((0, 1), (None, 0))
22    def dNdm_adjoint_action(self, *args, **kwargs):
23        """Evaluate the action of dNdm*"""
24        ...
```

Suppose we want to define a neural net operator.



Suppose we want to define a neural net operator.



PyTorch	Firedrake
$f(x^P; \theta)$	$N(x^F; v^*)$
<code>jacobian(f, x<sup>P</sup>)</code>	$J = \text{derivative}(N, x^F)$ <code>assemble(J)</code>
<code>jvp(f, x<sup>P</sup>, z<sup>P</sup>)</code>	<code>assemble(action(J, z<sup>F</sup>))</code>
<code>vjp(f, x<sup>P</sup>, z<sup>P</sup>)</code>	<code>assemble(action(adjoint(J), z<sup>F</sup>))</code>
<code>hvp(f, x<sup>P</sup>, z<sup>P</sup>)</code>	$H = \text{derivative}(J, x^F)$ <code>assemble(action(H, z<sup>F</sup>))</code>
<code>vhp(f, x<sup>P</sup>, z<sup>P</sup>)</code>	<code>assemble(action(adjoint(H), z<sup>F</sup>))</code>

Turns out PyTorch has all of our operators.



```
1 import torch
2 from torch.autograd.functional import jvp
3 ...
4 class PytorchOperator(MLOperator):
5     ...
6     @assemble_method(0, (0,))
7     def forward(self, *args, **kwargs):
8         V = self.function_space()
9         xF, _ = self.ufl_operands
10        # Convert input to PyTorch
11        xP = self.ml_backend.to_ml_backend(xF)
12        # Forward pass
13        yP = self.model(xP)
14        # Convert output to Firedrake
15        yF = self.ml_backend.from_ml_backend(yP, V)
16        return yF
17    ...
```



$$\min_{c \in P} \frac{1}{2} \|\varphi - \varphi^{obs}\|_V + \alpha \mathcal{R}(c) \quad (29)$$


subject to:

$$F(\varphi, c; v) = 0 \quad \forall v \in V \quad (30)$$



```
1 from firedrake import *
2 from firedrake_adjoint import *
3
4 ...
5 # Get a pre-trained PyTorch model
6 model = ...
7 # Define the external operator from the model
8 pytorch_op = neuralnet(model, function_space=...)
9 N = pytorch_op(vel)
10
11 # Solve the forward problem defined by equation (??)
12 solve(F(c, phi, v) == 0, phi, ...)
13 # Assemble the cost function:
14 J = assemble(0.5*(inner(phi-phi_obs, phi-phi_obs) +
15                 alpha*inner(N, N))*dx)
16
17 # Optimise the problem
18 Jhat = ReducedFunctional(J, Control(c))
19 c_opt = minimize(Jhat, method="L-BFGS-B", tol=1.0e-7,
20                 options={"disp": True, "maxiter" : 20})
```

# The only computational result in this talk



Recovered wave speed  $c$  as a function of position  $(x, z)$ : exact velocity (upper left), without regularisation (upper right), Tikhonov regulariser (lower left), neural network-based regulariser (lower right).

Other composable abstraction layers in and around Firedrake:

**Fireshape** Shape optimisation (Alberto Paganini, Leicester)

**Deflated continuation** Finding multiple solutions to nonlinear PDEs (Patrick Farrell, Oxford)

**Point data operators** Interact with real data and have first class point sources (Reuben Nixon-Hill)





# *Firedrake*



UK Research  
and Innovation



Natural  
Environment  
Research Council



Engineering and  
Physical Sciences  
Research Council