



Approaching fast, differentiable, calibrated chemistry models for methane and tropospheric oxidation

Paul Griffiths





GEOPHYSICAL RESEARCH LETTERS, VOL. 21, NO. 9, PAGES 801-804, MAY 1, 1994

Lifetimes and eigenstates in atmospheric chemistry

Michael J. Prather

Earth System Science, University of California at Irvine

Abstract. The time scales and mode of the atmosphere's response to chemical perturbations are defined by the eigenvalues and eigenvectors of the system. The eigenstates of a simplified one-box CH_4 -CO-OH system are analyzed. The longest time constant (smallest eigenvalue) always exceeds the lifetime defined by the steady-state loss frequency for CH_4 , the longest lived gas. Thus, the extent of a CH_4 perturbation—the methane response time—is always longer than predicted by the steady-state lifetime and is independent of size of the perturbation in the linear limit. This lengthening of the atmospheric recovery time can be diagnosed by how close we are to a chemically unstable troposphere, i.e., how much OH production exceeds that minimum needed to oxidize just the global emissions of CH_4 , CO, and other hydrocarbons and species.

[Isaksen and Hov, 1987] in which OH feedback affects the lifetime of CH_4 . Sensitivity analysis here shows that one key factor explains the extended "lifetime" for CH_4 perturbations: the excess source of tropospheric OH relative to the source of CH_4 and CO that must be oxidized by OH. This factor provides a simple diagnostic that may be used to understand the more complex 3-D tropospheric chemical models for which the calculation of such eigenstates would be exceedingly difficult.

Chemical System

Atmospheric chemistry is complex and four dimensional, varying with time over latitude, longitude, and altitude. The true chemical eigenstates of the atmosphere depend not only on local chemical rates but also on the rates of transport that




Figure 12. Box model studies of interactions between (left) CH_4 , (middle) CO, and (right) OH. Brown/black lines: Scenarios 1 and 4 are plotted from Figure 2. Green line: Scenario 5: Increased OH source and CH_4 emissions increased to 1,170 Tg. Blue line: Scenario 6 based on RCP8.5.

Methane Emissions in a Chemistry-Climate Model: Feedbacks and Climate Response


I. Heimann¹, P. T. Griffiths^{1,2}, N. J. Warwick^{1,2}, N. L. Abraham^{1,2}, A. T. Archibald^{1,2}, and J. A. Pyle^{1,2}

¹Chemistry Department, Cambridge University, Cambridge, UK, ²National Centre for Atmospheric Science, Cambridge University, Cambridge, UK

Peatlands are important carbon reservoirs





Cors Caron bog



Page and Baird, 2016, Annu. Rev. Environ. Resour.

- **600 gigatonnes of carbon** (600 Gt C)
- **One-third** of the global soil carbon storage
- **Twice** as much carbon as the world's forests
- **~370 million tonnes CO_2 /yr** sequestered (≈ 150 million cars) (*IUCN, 2021*)
- **~200 million tonnes CH_4 /yr** emitted per year, largest natural CH_4 source (*Kirschke et al., 2013, Nat. Geosci.*)

Peatland evolution is ecologically and climatically important



First contact with civilisation in Bristol – the Organic Geochemists

⑩ Two constraints, two unknowns:

⑩ $[CH_4]$ fixes the total source against the sink



⑩ $\delta^{13}C$ fixes the flux-weighted source signature

- “ $\delta^{13}C$ provides relatively weak constraints on individual CH_4 sources because of overlapping source $\delta^{13}C$ signatures and temporal variations in $\delta^{13}C$ of the wetland source, as well as uncertainties about CH_4 emissions from biomass burning.”
— Petrenko et al., Nature (2017)

⑩ The ice-core record is not a clean proxy for wetlands:


- “Changes in the ice-core methane record are likely the result of changes in both the source and the sink. Changes in the sink become especially important when changes in methane mixing ratio and/or climate are large.” — Harder, Shindell, Schmidt & Brook, Global Biogeochem. Cycles (2007)

⑩ **Their GCM:** CH_4 lifetime 11.7 yr (Holocene) → 16.5 yr (LGM)
— a ~30% weaker sink in the cold climate, from climate alone.



A differentiable 1-box model, in Julia – thanks ICCS training

- 1-box atmosphere, 18 kyr, 8 ODE state variables ($^{12}\text{CH}_4$ and $^{13}\text{CH}_4$ for fen / bog / remainder, plus CO and OH)
- Solved with a stiff ODE integrator (Rosenbrock23 / Rodas5P)
- **The point of Julia:** the whole solve is differentiable.
- ForwardDiff.jl carries dual numbers through the solver, interpolators and loss
- One pass yields the loss and its gradient w.r.t. all 362 parameters
- L-BFGS then steps all parameters at once, using curvature from the history
- **Cost:** $\sim 5\text{--}15\times$ a forward solve, replacing thousands of finite-difference runs
- $\mathcal{L} = w_c \cdot \text{MSE}[\text{CH}_4] + w_i \cdot \text{MSE}[\delta^{13}\text{C}] + w_s \cdot \Sigma(\Delta\text{SOH}, k)^2$




181 OH nodes + 181 remainder-signature nodes, interpolated through the ODE

A differentiable 1-box model, in Julia

- 1-box atmosphere, 18 kyr, 8 ODE state variables ($^{12}\text{CH}_4$ and $^{13}\text{CH}_4$ for fen / bog / remainder, plus CO and OH)
- Solved with a stiff ODE integrator (Rosenbrock23 / Rodas5P)
- **The point of Julia:** the whole solve is differentiable.
- ForwardDiff.jl carries dual numbers through the solver, interpolators and loss
- One pass yields the loss and its gradient w.r.t. all 362 parameters
- L-BFGS then steps all parameters at once, using curvature from the history
- **Cost:** $\sim 5\text{--}15\times$ a forward solve, replacing thousands of finite-difference runs
- $\mathcal{L} = w_c \cdot \text{MSE}[\text{CH}_4] + w_i \cdot \text{MSE}[\delta^{13}\text{C}] + w_s \cdot \Sigma(\Delta\text{SOH}, k)^2$

Differentiable model: gradient of a 362-parameter loss for free



*ForwardDiff carries dual numbers through the solver: loss + full gradient in one pass
(finite differences would need ~ 724 model runs per step)*

Setting up the differentiable box model in Julia

Model

- $\text{OH} + \text{CH}_4$ $k_{13} = k_{12}/1.005$
- $\text{CO} + \text{OH}$. CO source 537 Tg.
- OH source term

Data


- Fen/bog fluxes as linear interpolants;
- ice-core $[\text{CH}_4]$, $\delta^{13}\text{C}$ cleaned
- All mapped onto a 50-yr grid, ages as $(18 - \text{ka}) \times 1000$ yr so time runs forward from the LGM.

3. Control nodes.

- $S_{\text{OH}}(t)$ and $\delta^{13}\text{C}_{\text{rem}}(t)$: 181 nodes each, linearly interpolated.
- **Remainder flux** is prescribed (ramp + YD dip)


4. Parameters & bounds.

- Box bounds $S_{\text{OH}} \in [0.5, 2.5]$, $\delta^{13}\text{C}_{\text{rem}} \in [-85, -10]\text{‰}$.




Using the differentiable box model in Julia

- **First: integrate the box with a given parameter set.**
 - solve the 8-species stiff ODE (Rosenbrock23) forward 18ky
- **Then: optimise those parameters.**
 - $\text{MSE}[\text{CH}_4] + w \cdot \text{MSE}[\delta^{13}\text{C}] + \text{'smoothness'}$ and its gradient, update the parameters
- Re-integrate — until the loss stabilises
- **Iterates across: all 362 parameters**
 - 181 S_{OH} + 181 $\delta^{13}\text{C}_{\text{rem}}$, stepped together every iteration by L-BFGS (box-constrained)
- **100–200 iterations**
 - Then final reintegration with the optimum parameters




Fit the transients of both methane and isotopes with a constant remainder




First attempt — fit to the end points of the isotope record

- First approach.
- Simulate ice core methane mole fraction.
- Three components: fen and bog and an aggregate non-boreal-peat source ('remainder')
- Vary OH source term over the transient
- Additional constraint of $\delta^{13}\text{CH}_4$ at 18kya and PI/PD
- A single fixed remainder signature ($\approx -45\text{‰}$) cannot reproduce the $\delta^{13}\text{C}$ transient.



“Second” guess — optimise simultaneously to XCH4 and delta-13CH4

- A single fixed remainder signature ($\approx -45\text{‰}$) cannot reproduce the $\delta^{13}\text{C}$ transient.
 - Remainder flux initialised at 100 Tg/yr and optimised towards present day
 - During fitting the remainder increases from 100 $\rightarrow$ 200 Tg/yr across the deglaciation (14 $\rightarrow$ 9 ka). The remainder signature relaxes to a plausible $\sim -46\text{‰}$.
- “Major increases in the rates of initiation of circum-Arctic and Alaskan peatlands ... lagged increases in atmospheric methane ... other climatically sensitive sources or sinks must have been poised to rapidly switch on — but which ones?” — Behl, PNAS (2011)*
- Fits both records well, and the change in remainder signature to is now a few per mil
 - Something strange happening during the Younger Dryas?




Setting up the problem

- Let the remainder ramp 100 → 200 Tg/yr across the deglaciation (14→9 ka). The remainder signature relaxes to a plausible ~46%.

“Major increases in the rates of initiation of circum-Arctic and Alaskan peatlands ... lagged increases in atmospheric methane ... other climatically sensitive sources or sinks must have been poised to rapidly switch on — but which ones?” — Behl, PNAS (2011)

- The record demands a sharp CH₄ drop through the YD. We impose a remainder dip.


“Geological methane emissions were no higher than 15.4 Tg yr⁻¹ ... the rapid ~50% increase at the Younger Dryas–Preboreal event was driven by contemporaneous methane from sources such as wetlands.” — Petrenko et al., Nature (2017)



Adding in the Younger Dryas







- The record demands a sharp CH_4 drop through the YD. We impose a remainder dip back to 18kya.

“Geological methane emissions were no higher than 15.4 Tg yr^{-1} ... the rapid $\sim 50\%$ increase at the Younger Dryas–Preboreal event was driven by contemporaneous methane from sources such as wetlands.” — Petrenko et al., Nature (2017)








Varying the peatland isotope signatures

- Are the prescribed fen/bog signatures (−62.6 / −74.0‰) special? Let them float within their measured compilation ranges.
- Result:** on the ramp baseline, fitted fen/bog land at −62.9 / −75.0‰ — essentially the compilation centroids — and the fit is unchanged.








And what about the chlorine sink?

- **Equifinality:** heavy remainder + OH-only from light remainder + OH + Cl?
- Add tropospheric Cl. $\alpha_{\text{Cl}} \approx 1.066$ ($\epsilon_{\text{Cl}} \approx 66\text{‰}$) vs $\alpha_{\text{OH}} \approx 1.005$ ($\epsilon_{\text{OH}} \approx 5.4\text{‰}$). Cl fractionates $\sim 12\times$ stronger than OH.
- $\tau_{\text{Cl}} \approx 100$ yr **$\sim 8\%$ of total CH_4 loss** so a small sink with large isotopic leverage
- Fit essentially unchanged: $\text{RMS}(\text{CH}_4)$ 4.0 $\rightarrow$ 4.3 ppb, $\text{RMS}(\delta^{13}\text{C})$ 0.042 $\rightarrow$ 0.061 ‰ .
- Recovered OH history barely moves: glacial $\rightarrow$ Holocene ratio $\times 2.4 \rightarrow \times 2.6$, same shape, same $\sim -45\%$ deglacial trough.
- The **only** thing that responds is the remainder *delta13C-CH4 signature*: $\sim 7\text{‰}$ lighter, $-45\text{‰} \rightarrow -52\text{‰}$, moving toward physically sensible values?



Cl leaves the fit and OH alone; only the remainder level shifts

Beyond two end-points



An isotope mass-balance gap

source	Tg/yr	$\delta^{13}\text{C}$
boreal peat (fen+bog)	80	-62.6 / -74
tropical wetland	100	-59
termites / animals	35	-60 to -63
biomass burning	18	-23
geologic	12	-40
ocean+lakes	10	-52

- Plugging these in gives CH_4 5‰ too light with OH sink alone;
- Flux-weighted source -58‰ which with OH-only fractionation, give atmosphere -53‰.
- Observed PI atmosphere is -47‰ — a 6‰ gap.

sink	ϵ_{app}	predicted atmosphere	vs observed -47.7‰
OH only	5.4‰	-52.7‰	misses by 5‰
OH + CI ($\tau=100$ yr)	10.1‰	-48.0‰	-0.3‰ ✓

- a more strongly fractionating sink like CI does better. A range of CI lifetimes from CI 60–200 yr and budget $\pm 2\%$ fit the budget and the observations.
- The recovered remainder signature gets heavier by $\sim +2\%$ from glacial to Holocene. A doubling of biomass burning in the PI budget increases methane emissions from 18 to 36 Tg/yr and would shift the source +2.3‰ — matching the observed shift. [various approaches, e.g. *Fischer et al., Nature (2008)*]

Differentiable model complements the Monte Carlo approach (I hope)

Differentiable box modelling makes this (362-parameter) joint inversion problem tractable.

The non-boreal peatland emissions flux must **ramp** across the deglaciation (isotopes)


Given the right baseline the **peatland data fit** to observations - reassuring.

Fitted OH is **~30–40% weaker at the deglaciation** - encouraging similarity

Biomass burning ~doubling covers the pyrogenic shift - feasible?

Absolute source-vs-sink partitioning is **equifinal**. A fixed bottom-up budget is 5‰ too light with OH alone; **CI reconciles budget to isotopes**

	Monte Carlo ensemble	Differentiable / optimisable
what it does	samples the prior space	Can change one thing, calculate the response
If degeneracy?	maps it (broad posterior)	localises each sensitivity
cost	thousands of forward runs	~one gradient per step



Also known as

“Lessons learned from ATom missions?”

“What is the minimal viable set of variables for atmospheric chemistry?”


“Chemistry is expensive. Can we replace it with ML?”

“Chemistry is boring. Can we replace it with a shell script?”


Getting OH without solving the chemistry: four recent answers

- Pimlott et al. 2022: steady state driven by satellite sources and sinks. Their simplified form is one source over three sinks, with no HO₂ term. Against ATom, $r = 0.78$ and magnitude within about 30 per cent
- Baublitz et al. 2023: a reduced-form proxy for OH variability in the remote marine troposphere, built on ATom, with HO₂ supplied as a measured input. Within-bin $r^2 = 0.90$
- Anderson et al. 2022: gradient-boosted regression trees trained on a full-chemistry chemistry-climate model, designed to be integrated back into that model so the OH feedbacks stay nonlinear
- Anderson et al. 2026: primary OH production, the O(1D) plus water term, constrained from space between 60S and 60N using machine learning with satellite and meteorological data

- A new Earth system model, built to learn from data
- Schneider et al. 2017: models that learn from observations and from targeted high-resolution simulations
- Julia throughout, differentiable, GPU native
- Resolve what you can afford, calibrate the rest
- Calibrate the parameters, emulate the expensive forward map, sample the posterior
- Currently physics only. Or it was
- First step: CO transported and oxidised, the loss driven by a 3-D OH lookup table regridded onto the cubed sphere



- A new Earth system model, built to learn from data
- Schneider et al. 2017: models that learn from observations and from targeted high-resolution simulations
- Julia throughout, differentiable, GPU native
- Resolve what you can afford, calibrate the rest
- Calibrate, emulate, sample
- Calibrate the parameters, emulate the expensive forward map, sample the posterior
- Currently physics only. Or it was
- Second step: SO₂ and H₂SO₄ tracers with CEDS emissions, with OH as the rate-limiting oxidant driven by a 3-D OH lookup table regridded onto the cubed sphere



The chemical mechanism is something we might not afford

- MOZART_T1 is 161 species, UKCA StratTrop 85 species and 252 reactions
- Each call is a stiff ODE solve, and it would be one per column or more per chemistry step
- A closure should evaluate in microseconds from simple inputs
- **Schneider 2017:** closure parameters “necessarily arise in any such coarse graining of nonlinear governing equations” – e.g. Turbulence and clouds are carried by EDMF, an eddy-diffusivity mass-flux scheme with calibrated parameters
 - They “can then be learned from a broad range of observations and high-resolution simulations”
 - Learned inside the known governing equations, so symmetries and conservation laws are respected
 - Structured parameterisations “likely have greater out-of-sample predictive power than unstructured parameterization schemes, such as neural networks”

“Out-of-sample predictive power will be crucial if high-resolution simulations performed in selected locations and under selected conditions are to provide information globally and in changed climates”

Why a closure is the right form for OH

- OH lives under a second, so it is always in quasi-steady state with its sources and sinks at every model step.
- Various SSA approaches in the literature. An algebraic balance: primary production, recycling from e.g. HO₂, termination
- So, tracer transport not critical as fast chemistry, almost diagnostic.
- Closure scheme introduces α_J, R_{VOC} , termination parameter (Γ)
- Consistent with current ClimaAtmos chemical tracers

$$[OH] = \frac{Production}{Loss(1 - f)}$$


$$f = \frac{recycling}{recycling + termination}$$

$$Loss = k_{CO}[CO] + k_{CH_4}[CH_4] + R_{VOC}$$

$$Prod = \alpha_J 2J_{O_1D}[O_3][H_2O]$$

First closure, data-driven - calibrated on ATom data



- 4 parameters, photostationary HOx steady state, MAP against the observed OH
- Posteriors 28 to 44 times tighter than prior on the photolysis slopes, so the observations constrain the parameters that carry the climate response
- Differentiable, microseconds to evaluate, and calibratable rather than tuned



Four-parameter closure against ATom observed OH

Three candidate targets for training


- ATom observations. Real data, at least, but confined to the clean daylight remote troposphere - promising, but limited in coverage: never continental, never the polluted surface, never dark
- Start somewhere with a direct mechanism. Derive new training states *from somewhere other than observations*
- Is it accurate?
- Everybody loves box models, and they could provide the right training data.
- The UKCA StratTrop (ST) mechanism (similar to other CTMs) performs adequately against the ATom data.
- Gives atmospheric chemistry mechanism data that covers the whole parameter space, but at the cost of carrying the mechanism's own biases.



Fitting to a Gaussian Process emulator – develop UKCA-GP

- The SVGP emulates StratTrop KPP data and the closure is fitted against the SVGP.
- Ten state variables in, $\log_{10}(\text{OH})$ (and $\log_{10}(\text{HO}_2)$ if required) returned. Fitted to the 26,241 KPP-solved states, bias corrected to observations
- SVGP rather than a full GP inversion of a 26,241 square matrix
- Variational in as much as the inducing points and the hyperparameters are optimised together
- Held out against the mechanism: R^2 0.947 for OH, 0.989 for HO₂
- Use the SVGP to develop the closure – captures all atmospheric states, and is accurate w.r.t. A_{Tom} where appropriate.

Sparse: M inducing points carry the posterior, $M \ll N$
Cost falls from N^3 to NM^2 , and the band widens where the points thin out




Sparse GP: M inducing points summarise N data ($M \ll N$)

The mechanism carries a bias against the observations

- To calibrate/remove bias inside the mechanism we would be adjusting rate constants and branching ratios, hundreds of them, none separately identifiable from an OH measurement.

Leave that for DeTOX!



- Use BOXMOX+UKCA StratTrop to generate the training data
 - 26,241 states, sampled for coverage of the driver space rather than for its frequency
 - GEOS-Chem supplies the ranges of the drivers, never the states themselves. A model state carries that model's chemistry in its reservoir species
 - Try to enforce couplings based on physically reasonable states: T on P, water on T, the photolysis pair on their ratio



Training data vs ATom


Fitting box model data to a Gaussian Process emulator – UKCA-GP

- The SVGP emulates StratTrop KPP data.
- Ten state variables in, $\log_{10}(\text{OH})$ (and $\log_{10}(\text{HO}_2)$ if required) returned. Fitted to the 26,241 KPP-solved states, bias corrected to observations
- SVGP rather than a full GP inversion of a 26,241 square matrix
- Variational in as much as the inducing points and the hyperparameters are optimised together
- Held out against the mechanism: R^2 0.947 for OH, 0.989 for HO₂
- Can now use the SVGP to develop a closure.



Calibrating a GP based on UKCA data – performance against ATom


- The SVGP emulates StratTrop KPP, and does well against UKCA-StratTrop inside a box model.
- **Performance against ATom data is fine, but some bias across altitudes**
- A second SVGP is then fitted to the residual between the coverage GP and observed OH, evaluated at the ATom states
- By definition, this has zero mean, so recover BOXMOX-trained GP when away from ATom states. Strikes a balance.
- After bias correction, the altitude structure bias of the base flattens from ~ 1.4 times high at the surface and 0.75 aloft, becoming flat.
- The scatter is nearly the same in both, so what the correction removes is the offset and the vertical structure rather than the noise



Calibrating a GP based on UKCA data – performance against ATOm



- The SVGP emulates StratTrop KPP, and does well against UKCA-StratTrop inside a box model.
- Performance against ATOm data is also fine, but some bias across altitudes
- Add second SVGP to fit the residual between the UKCA-ST GP and observed ATOm OH, evaluated at the ATOm states
- By definition, this second GP has zero mean, so recover BOXMOX-trained GP when away from ATOm states. Strikes the right balance?
- **After** bias correction, the altitude structure bias of the base flattens from ~ 1.4 times high at the surface and 0.75 aloft, becoming flat.
- The scatter is nearly the same in both, so what the correction removes is the offset and the vertical structure rather than the noise

Bias-corrected emulator against held-out ATOm OH




Phase 1: a closure in terms of only what ClimaAtmos currently carries

- Which, at the moment, is not very much, so expect this to be a challenge.
- Reads jO(1D), jNO₂, O₃, H₂O, T, P, CO and CH₄
- So, e.g. OH plus NO₂ sink and HO₂ plus NO recycling terms are folded into calibrated effective terms
- **On the scoped remote domain:** median ratio 1.01, arithmetic mean ratio 1.18
- Not too shabby but high biased.
- J-F would tell me (indeed has told me) but, Paaaul, that's a log scale. Fair.
- But competitive on ATom and valid in the dark



Phase 2 closure: train and add NO and HCHO – any better?

- Primary over gross production 0.33, against Lelieveld 0.33, Secondary over primary 2.03, against 1.99




Thank you

A solid red horizontal bar spanning the entire width of the slide at the bottom.



■ SPARE SLIDES

Expanding from three peatlands to global peatland evolution



- 89 peatlands with methane flux measurements (n = 372)
- 82 peatlands with $\delta^{13}\text{CH}_4$ measurements (n = 432)
- 133 peatlands with paleoecology and paleo-pH records, based on plant macrofossils, pollen, spores, and biomarker-inferred pH (n = 133)



Holocene global peatland evolution: timing and patterns



Peatland succession pattern from 39 sites in eastern North America

Yu et al., 2013, *Glob. Biogeochem. Cycles*;
Loisel et al., 2017, *Earth-Sci. Rev.*

Modern peatland CH₄ flux and $\delta^{13}\text{C}$ dataset



- Modern peatland methane flux and $\delta^{13}\text{C}$ values are significantly positively correlated with pH
 - Modern pattern consistent with core biomarkers:
 - pH $\uparrow$ - methanogenesis $\uparrow$ - CH₄ flux $\uparrow$
 - pH $\uparrow$ - methanotrophy $\uparrow$ - methanotrophic lipid $\delta^{13}\text{C}$ $\downarrow$ - residual CH₄ $\delta^{13}\text{C}$ $\uparrow$
- Key point: Spatial gradient $\leftrightarrow$ temporal evolution