

1        **SUPPLEMENTARY MATERIALS: THE F-1 ALGORITHM FOR**  
2        **EFFICIENT COMPUTATION OF THE HESSIAN MATRIX OF AN**  
3        **OBJECTIVE FUNCTION DEFINED IMPLICITLY BY THE**  
4        **SOLUTION OF A STEADY-STATE PROBLEM\***

5                    BENOÎT PASQUIER<sup>†</sup> AND FRANÇOIS PRIMEAU<sup>†</sup>

6        **SM1. Global marine phosphorus-cycling model.** Here, we describe the dis-  
7        cretization of the phosphorus-cycling model that we use to demonstrate the perfor-  
8        mance of the F-1 algorithm. Specifically, we detail how the physical and biogeochem-  
9        ical mechanisms are converted into the discretized generic equation (1.1) by explicitly  
10        defining the state function  $\mathbf{F}(\mathbf{x}, \mathbf{p})$ . In practice, we use the **AIBECS** package [SM20]  
11        to download the ocean circulation and generate the state function,  $\mathbf{F}$ , the mismatch  
12        function  $f$ , and their derivatives.

13        **SM1.1. The ocean circulation.** We use the control (CTL) circulation output  
14        from the Ocean Circulation Inverse Model (OCIM1) of DeVries [SM3]. The OCIM1  
15        optimizes a steady-state circulation by assimilating oceanographic tracer data (see  
16        [SM5]), which include potential temperature, salinity, radiocarbon, and CFC-11. The  
17        resulting circulation thus effectively represents an estimate of the mean state of the  
18        global ocean circulation and is formulated as a sparse matrix, denoted here by  $\mathbf{T}_{\text{DIP}}$ ,  
19        to facilitate rapid simulation of biogeochemical tracers. The circulation is embedded  
20        in an Arakawa B-grid with a  $2^\circ \times 2^\circ$  resolution and 24 depth levels with thicknesses  
21        increasing from  $\sim 25$  m at the surface to  $\sim 500$  m for the deepest layer. (This amounts  
22        to 200 160 ocean grid cells.) In practice, we use the **AIBECS** package [SM20] to  
23        download the OCIM matrix.

24        With the state of the system represented by the column vector  $\mathbf{x} = \begin{bmatrix} \mathbf{x}_{\text{DIP}} \\ \mathbf{x}_{\text{POP}} \end{bmatrix}$ ,  
25        the OCIM1 transport matrix is applied to the DIP vector only,  $\mathbf{x}_{\text{DIP}}$ . That is, the  
26        matrix–vector product,  $\mathbf{T}_{\text{DIP}} \mathbf{x}_{\text{DIP}}$ , yields the flux divergence of DIP due to the ocean  
27        circulation. In other words,  $\mathbf{T}_{\text{DIP}} \mathbf{x}_{\text{DIP}}$  is the discrete equivalent of  $\nabla_{\mathbf{r}} \cdot [\mathbf{u} - \mathbf{K}\nabla] \mathbf{x}_{\text{DIP}}$   
28        in (3.1).

29        **SM1.2. The sinking particles.** We now describe how  $\mathbf{T}_{\text{POP}}$ , a sparse matrix  
30        representing the discrete equivalent of  $\nabla \cdot \mathbf{w}$  in (3.1), is created. At the bottom of  
31        each model box, the sinking velocity is hence defined by

32        (SM1.1)                     $\mathbf{w} \equiv w' \mathbf{z}_{\text{bot}} + w_0,$

33        where  $\mathbf{z}_{\text{bot}}$  is a vector of the depths of the bottom of each grid box and  $\mathbf{w}$  is the  
34        vector of the magnitudes of the downward particulate velocity. Note that the  $\mathbf{w}$   
35        symbol defined in (SM1.1) is different from the  $\mathbf{w}$  symbol used for (3.1), which was  
36        the 3D velocity vector. The flux-divergence operator,  $\mathbf{T}_{\text{POP}}$ , is defined such that  
37         $\mathbf{T}_{\text{POP}} \mathbf{x}_{\text{POP}}$  is the divergence of the sinking POP flux. The flux of sinking POP is  
38        approximated by  $\mathbf{w} \mathbf{x}_{\text{POP}}$  at the bottom of each model box.

\*Submitted to the the SIAM Journal of Scientific Computing for review on 2019-07-09.

**Funding:** This work was funded by the US Department of Energy grant DE-SC0016539 and the National Science Foundation grant 1658380.

<sup>†</sup>Department of Earth System Science, University of California, Irvine, CA, United States (pasquieb@uci.edu, fprimeau@uci.edu).

39 The flux divergence is then approximated by the difference between the flux at  
 40 the top and at the bottom of each box. Specifically, the sinking transport operator  
 41  $\mathbf{T}_{\text{POP}}$  is defined by

$$42 \quad (\text{SM1.2}) \quad \mathbf{T}_{\text{POP}} \equiv \Delta \mathbf{Z}^{-1} (\mathbf{U}_{\mathbf{x}_{\text{POP}}} - \mathbf{I}_{\mathbf{x}_{\text{POP}}}) \mathbf{W},$$

43 where  $\Delta \mathbf{Z}$  is a diagonal matrix with diagonal  $\Delta \mathbf{z}$ , which is the vector of the depth  
 44 thickness of each box,  $\mathbf{I}_{\mathbf{x}_{\text{POP}}}$  is the identity matrix, and  $\mathbf{W}$  is the diagonal matrix with  
 45 diagonal  $\mathbf{w}$ . The matrix  $\mathbf{U}_{\mathbf{x}_{\text{POP}}}$  consists of 1's on a superdiagonal similar and is thus  
 46 upper triangular. It is defined similarly to a shift matrix, such that it shifts indices to  
 47 the boxes above, so that  $(\mathbf{U}_{\mathbf{x}_{\text{POP}}} - \mathbf{I}_{\mathbf{x}_{\text{POP}}}) \mathbf{W} \mathbf{x}_{\text{POP}}$  effectively yields the vector of the  
 48 differences between the flux of POP at the top and at the bottom of each grid box.  
 49 Thus,  $\mathbf{T}_{\text{POP}}$ , is a function of the parameters  $w'$  and  $w_0$  only.

50 **SM1.3. The state function.** The rate of change of the system is defined by

$$51 \quad (\text{SM1.3}) \quad \mathbf{F}(\mathbf{x}, \mathbf{p}) \equiv \begin{bmatrix} -\mathbf{T}_{\text{DIP}} \mathbf{x}_{\text{DIP}} - \mathbf{U} + \mathbf{R} + \mathbf{C} \\ -\mathbf{T}_{\text{POP}} \mathbf{x}_{\text{POP}} + \mathbf{U} - \mathbf{R} \end{bmatrix},$$

52 where the column vectors  $\mathbf{U}$  and  $\mathbf{R}$  are the uptake and remineralization rates, respec-  
 53 tively, converting DIP to POP and back ( $U$  and  $R$  in (3.1)).

54 The discrete equivalent of the uptake, remineralization, and geological restoring  
 55 rates, as defined in (3.2), can be defined compactly by

$$56 \quad (\text{SM1.4}) \quad \begin{cases} \mathbf{U} \equiv \frac{\mathbf{x}_{\text{DIP}}^+}{\tau} \frac{\mathbf{x}_{\text{DIP}}^+}{\mathbf{x}_{\text{DIP}}^+ + k} (\mathbf{z} \leq z_0) \\ \mathbf{R} \equiv \kappa \mathbf{x}_{\text{POP}} \\ \mathbf{C} \equiv \frac{\langle x^{\text{geo}} \rangle - \mathbf{x}_{\text{DIP}}}{\tau_{\text{geo}}} \end{cases},$$

57 where  $\mathbf{x}_{\text{DIP}}^+ = \mathbf{x}_{\text{DIP}}$  if  $\mathbf{x}_{\text{DIP}} \geq 0$  and  $\mathbf{x}_{\text{DIP}}^+ = 0$  otherwise, ensuring that uptake only  
 58 occurs for positive concentrations. (This is to avoid problems due to the hyperbolic  
 59 term in the uptake when numerical noise generates negative concentrations.) Simi-  
 60 larly, the  $(\mathbf{z} \leq z_0)$  term ensures that uptake only occurs in the euphotic zone. ( $\mathbf{z}$   
 61 is the vector of box depths, oriented positively downwards, and  $z_0 \equiv 75$  m is the  
 62 approximate depth of the bottom of the euphotic layer in the OCIM1 grid.)

63 Note that we used two shortcuts to simplify notation in (SM1.4) that we use  
 64 throughout this section: (i) Boolean expressions are assumed to convert automati-  
 65 cally to the real-valued 0 or 1, and (ii) binary vector operations are assumed to be  
 66 elementwise. Thus, for example, in (SM1.4), each entry of  $\mathbf{x}_{\text{DIP}}^+$  is divided by the  
 67 corresponding entry in  $(\mathbf{x}_{\text{DIP}}^+ + k)$ , and  $(\mathbf{z} \geq z_0)$  is a vector of 0's and 1's with its  
 68 entries multiplied by the corresponding entries of  $\mathbf{x}_{\text{DIP}}^+ / (\mathbf{x}_{\text{DIP}}^+ + k)$ .

69 **SM2. Derivatives with respect to the state.** The scope of the F-1 algorithm  
 70 may appear restricted by the fact that it requires the user to supply the Jacobian  
 71 function,  $\nabla_{\mathbf{x}} \mathbf{F}$ . However, the Jacobian function,  $\nabla_{\mathbf{x}} \mathbf{F}$ , which is required for Newton-  
 72 like inner solvers, may be easy to derive analytically. Additionally, because many  
 73 discretizations for PDEs use local low-order finite difference schemes that produce  
 74 sparse Jacobian matrices,  $\nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x}, \mathbf{p})$  is likely sparse, so that it can be computed  
 75 numerically in only a few dual passes using a graph-coloring algorithm (see, e.g.,  
 76 [SM10]). Furthermore, it is likely that the Jacobian of the nonlinear part of  $\mathbf{F}$  is

77 even sparser, e.g., diagonal or made of blocks that are diagonal, allowing for it to be  
 78 computed in even less dual passes.

79 For instance, in the context of the global marine phosphorus cycling model (sec-  
 80 tion SM1), the Jacobian of the state function  $\mathbf{F}$  defined by (SM1.3) can be evaluated  
 81 in just two dual evaluations. Consider the partition of  $\mathbf{F}(\mathbf{x}, \mathbf{p})$  into the sum of a linear  
 82 part (with respect to  $\mathbf{x}$ ), denoted by  $-\mathbf{T}(\mathbf{p})\mathbf{x}$ , where

$$83 \quad (\text{SM2.1}) \quad \mathbf{T}(\mathbf{p}) \equiv \begin{bmatrix} \mathbf{T}_{\text{DIP}} & 0 \\ 0 & \mathbf{T}_{\text{POP}} \end{bmatrix}$$

84 is a large block-diagonal sparse matrix, and a local part,  $\mathbf{G}(\mathbf{x}, \mathbf{p})$ , which consists of  
 85 the uptake,  $\mathbf{U}$ , remineralization,  $\mathbf{R}$ , and geological restoring,  $\mathbf{C}$  (see section SM1).  
 86 The Jacobian of the local part,  $\nabla_{\mathbf{x}}\mathbf{G}(\mathbf{x}, \mathbf{p})$ , is then made of blocks that are diagonal,  
 87 i.e.,

$$88 \quad (\text{SM2.2}) \quad \nabla_{\mathbf{x}}\mathbf{G}(\mathbf{x}, \mathbf{p}) \equiv \begin{bmatrix} \nabla_{\text{DIP}}\mathbf{C} - \nabla_{\text{DIP}}\mathbf{U} & \nabla_{\text{POP}}\mathbf{R} \\ \nabla_{\text{DIP}}\mathbf{U} & -\nabla_{\text{POP}}\mathbf{R} \end{bmatrix}$$

89 where  $\nabla_{\text{DIP}}\mathbf{C}$ ,  $\nabla_{\text{DIP}}\mathbf{U}$ , and  $\nabla_{\text{POP}}\mathbf{R}$  are the derivatives of  $\mathbf{C}$ ,  $\mathbf{U}$ , and  $\mathbf{R}$  with respect  
 90 to  $\mathbf{x}_{\text{DIP}}$  or  $\mathbf{x}_{\text{POP}}$ , as indicated, and are diagonal matrices because  $\mathbf{C}$ ,  $\mathbf{U}$ , and  $\mathbf{R}$   
 91 are local rates. Because these are diagonals,  $\nabla_{\mathbf{x}}\mathbf{G}(\mathbf{x}, \mathbf{p})$  can be computed in two  
 92 dual evaluations, one where  $\varepsilon$  is added to each entry of  $\mathbf{x}_{\text{DIP}}$ , and one where  $\varepsilon$  is  
 93 added to each entry of  $\mathbf{x}_{\text{POP}}$ . The full Jacobian,  $\nabla_{\mathbf{x}}\mathbf{F}(\mathbf{x}, \mathbf{p})$  is then simply given by  
 94  $-\mathbf{T}(\mathbf{p}) + \nabla_{\mathbf{x}}\mathbf{G}(\mathbf{x}, \mathbf{p})$ . (This is for example implemented under the hood in the AIBECS  
 95 package to generate  $\nabla_{\mathbf{x}}\mathbf{F}$  automatically from  $\mathbf{T}_{\text{DIP}}$ ,  $\mathbf{T}_{\text{POP}}$ , and the local sources and  
 96 sinks [SM20].)

97 A similar potential restriction to the scope of the F-1 algorithm is that it requires  
 98 the user to supply  $\nabla_{\mathbf{x}}f$ . However, an analytical formula for  $\nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{p})$  is straightfor-  
 99 wardly available if  $f(\mathbf{x}, \mathbf{p})$  is expressed as the sum of the quadratics, as is the case for  
 100 the objective function of our phosphorus-cycling model. That is, the state and  
 101 parameter parts of  $f(\mathbf{x}, \mathbf{p})$  are usually separable, so that  $f(\mathbf{x}, \mathbf{p}) = f(\mathbf{x}) + f(\mathbf{p})$ , and  
 102 the state part,  $f(\mathbf{x})$ , usually takes the generic form  $f(\mathbf{x}) = \frac{1}{2}\delta\mathbf{x}^{\top}\mathbf{\Omega}\delta\mathbf{x}$ , where the  
 103 precision matrix  $\mathbf{\Omega}$  is diagonal, so that

$$104 \quad (\text{SM2.3}) \quad \nabla_{\mathbf{x}}f(\mathbf{x}, \mathbf{p}) = \delta\mathbf{x}^{\top}\mathbf{\Omega}$$

105 can be readily used as the analytical formula supplied to the F-1 algorithm. (Here  
 106 we have abused notation to distinguish the state and parameter parts of  $f(\mathbf{x}, \mathbf{p})$   
 107 depending on its arguments, similar to the multiple-dispatch paradigm of the Julia  
 108 language.) Hence, we argue that in many scientific applications, the F-1 algorithm  
 109 can be used as a fully automatic differentiation tool with little additional effort, e.g.,  
 110 compared to the HYPER or FD2 algorithms.

111 **SM3. Comment on the computation time partition.** We note that in Fig-  
 112 ure 4, the time spent by the inner solver is attributed to time spent on the gradient.  
 113 This is because the Newton-Trust-Region algorithm we use for the optimization com-  
 114 puts the gradient first, then the Hessian, and finally the objective, in that specific  
 115 order. Thus, the inner solver is always invoked by the gradient first, making  $\mathbf{s}(\mathbf{p})$  (and  
 116 the factors of  $\mathbf{A}$  in the case of the F-1 algorithm) available for the subsequent Hessian  
 117 and objective computations. The objective function is computed last because it is not  
 118 needed in Newton’s method for minimization and is merely evaluated for recording  
 119 the progress of the optimization. Hence, the time attributed to the objective func-  
 120 tion,  $\hat{f}(\mathbf{p})$ , is merely the time of a single evaluation of  $f(\mathbf{s}, \mathbf{p})$ , which is negligibly

121 small relative to the other computations (the corresponding bars are not visible in  
122 [Figure 4](#)).

123 **SM4. Extending real numbers for differentiation.**

124 **SM4.1. Complex numbers.** Since its discovery [\[SM12\]](#), and despite its ele-  
125 gance and efficiency, it took nearly 30 years before the complex-step algorithm for  
126 differentiation was described in the literature, by Squire and Trapp [\[SM24\]](#). More  
127 recently, Cleve Moler wrote a simple [blog post](#) [\[SM16\]](#) describing the technique on the  
128 MathWorks website in 2013, and Martins et al. [\[SM13\]](#) wrote an article dedicated to  
129 it. Consider a simple generic real-valued function  $g$ , which depends on a real-valued  
130 scalar variable. The Taylor expansion of  $g$  at  $x$  in the  $ih$  direction, where  $i$  is the  
131 imaginary unit and  $h$  is a small real-valued scalar, is given by

$$132 \text{ (SM4.1)} \quad g(x + ih) = g(x) + ih\nabla g(x) + \underbrace{\frac{(ih)^2}{2}\nabla^2 g(x) + \dots}_{\rightarrow 0 \text{ if } h \rightarrow 0}$$

133 In practice, if  $h$  is small enough, floating-point arithmetic would ensure that only the  
134 first two terms remain (specifically, if  $g(x)$  is indistinguishable from  $g(x) - h^2\nabla^2 g(x)/2$   
135 for the finite-precision machine, assuming the higher-order terms are even smaller).

136 This way, one can approximate the first derivatives to numerical precision, by  
137 taking the imaginary part of  $g(x + ih)$ , i.e.,

$$138 \text{ (SM4.2)} \quad \nabla g(x) \approx \frac{\Im[g(x + ih)]}{h},$$

139 where  $\Im(x)$  is the imaginary part of  $x$ . In particular, because there is no difference  
140 involved and because the relative sizes of the real and imaginary parts do not interfere  
141 within the complex-number type, there is no need to worry about truncation error,  
142 so that one can take an arbitrarily small  $h$  (e.g.,  $h = 10^{-100}$ ).

143 The complex-step algorithm is straightforward to implement in many scientific  
144 computing languages, including Julia, MATLAB, Octave, Scilab, Python, IDL, and  
145 Fortran, which all have a built-in complex-number type. And we note that Martins  
146 et al. [\[SM13\]](#) suggested that complex-step differentiation is equivalent to algorithmic  
147 differentiation (also known as automatic differentiation).

148 **SM4.2. Dual Numbers.** Dual numbers were first introduced by Clifford [\[SM2\]](#),  
149 with the “dual” term being coined by Study [\[SM25\]](#). A dual unit number, denoted  
150  $\varepsilon$ , is introduced, and defined by the simple rule that  $\varepsilon^2 = 0$  (instead of  $i^2 = -1$  for  
151 complex numbers) and is used here as an efficient numerical tool to compute first  
152 derivatives (see, e.g., [\[SM17\]](#)).

153 Using dual numbers to evaluate a derivative is cleaner and simpler than the  
154 complex-step differentiation of [\(SM4.2\)](#), as

$$155 \text{ (SM4.3)} \quad \nabla g(x) = \mathfrak{D}[g(x + \varepsilon)],$$

156 where  $\mathfrak{D}(x)$  is the dual part (the  $\varepsilon$  part) of  $x$ . The advantage of using dual numbers for  
157 our approach lies in the fact that Taylor expansions of the first order in the non-real  
158 direction are exact, i.e.,

$$159 \text{ (SM4.4)} \quad g(x + \varepsilon) = g(x) + \varepsilon\nabla g(x) + \underbrace{\frac{\varepsilon^2}{2}\nabla^2 g(x) + \dots}_{=0}$$

160 is a strict equality. (A property that is not true for Taylor expansions in the complex  
 161 plane.) In other words, dual numbers provides a robust mathematical tool to represent  
 162 infinitesimal quantities with  $\varepsilon$ .

163 Another advantage compared to the complex step algorithm is that there is no  
 164 need to introduce a small step size,  $h$ , although we note that the behavior of dual  
 165 numbers can be reproduced using complex numbers by choosing  $h$  sufficiently small  
 166 for the complex Taylor series to automatically truncate terms of order greater than  
 167 1. However, we note that the dual-step algorithm offers a cleaner, more modern  
 168 numerical differentiation tool that can be applied to a broader set of functions (e.g.,  
 169 functions that internally invoke complex numbers).

170 We note the dual numbers have a formal algebraic interpretation, whereby they  
 171 extend the real numbers to an algebra over the reals of dimension 2, where  $\varepsilon$  is the  
 172 image of the polynomial  $X \mapsto X$  in the quotient of the univariate polynomial ring with  
 173 real-valued coefficients,  $\mathbb{R}[X]$ , over the ideal generated by the  $X \mapsto X^2$  polynomial. In  
 174 other words, the dual numbers can be identified with  $\mathbb{R}[X]/(X^2)$  just like the complex  
 175 numbers can be identified with  $\mathbb{R}[X]/(X^2 + 1)$ .

176 **SM4.3. Hyper Dual Numbers.** Hyper dual numbers are an extension of dual  
 177 numbers used to evaluate second derivatives (e.g., [SM8, SM6, SM7, SM17]). The  
 178 nonreal components are defined by  $\varepsilon_1$  and  $\varepsilon_2$  such that  $\varepsilon_1^2 = \varepsilon_2^2 = 0$  and  $\varepsilon_1\varepsilon_2 \neq 0$ .  
 179 Consider  $g(x, y)$ , a real-valued function of two variables. Its Taylor expansion in the  
 180  $(\varepsilon_1, \varepsilon_2)$  direction is given by

181 (SM4.5)  $g(x + \varepsilon_1, y + \varepsilon_2) = g(x, y) + \varepsilon_1 \nabla_x g(x, y) + \varepsilon_2 \nabla_y g(x, y) + \varepsilon_1 \varepsilon_2 \nabla_{xy} g(x, y),$

182 and its Taylor expansion in the  $(\varepsilon_1 \varepsilon_2, 0)$  direction is given by

183 (SM4.6)  $g(x + \varepsilon_1, y + \varepsilon_2) = g(x, y) + \varepsilon_1 \nabla_x g(x, y) + \varepsilon_2 \nabla_x g(x, y) + \varepsilon_1 \varepsilon_2 \nabla_{xx} g(x, y),$

184 Hence, the terms of the Hessian of  $g$  can be evaluated via

185 (SM4.7)  $\nabla^2 g(x, y) = \mathfrak{H} \left( \begin{bmatrix} g(x + \varepsilon_1 + \varepsilon_2, y) & g(x + \varepsilon_1, y + \varepsilon_2) \\ g(x + \varepsilon_2, y + \varepsilon_1) & g(x, y + \varepsilon_1 + \varepsilon_2) \end{bmatrix} \right),$

186 where  $\mathfrak{H}(x)$  is the hyperdual part of  $x$  (the  $\varepsilon_1 \varepsilon_2$  coefficient). Note that hyperdual  
 187 numbers can also be used to evaluate the terms of the gradient just like dual numbers,  
 188 e.g., via

189 (SM4.8)  $\nabla g(x, y) = \mathfrak{H}_1 \left( \begin{bmatrix} g(x + \varepsilon_1, y) \\ g(x, y + \varepsilon_1) \end{bmatrix}^\top \right),$

190 where  $\mathfrak{H}_1(x)$  is the "first hyperdual" part of  $x$  (the  $\varepsilon_1$  part). Generalizing these  
 191 formulas to the case of functions of more than two variables is straightforward.

192 We note that the hyperdual numbers also have an algebraic interpretation. They  
 193 form an algebra over the reals of dimension 4, as the quotient of the multivariate  
 194 polynomial ring with real-valued coefficients,  $\mathbb{R}[X, Y]$ , over the ideal generated by  
 195 the  $(X, Y) \mapsto X^2$  and  $(X, Y) \mapsto Y^2$  polynomials. In that quotient,  $\varepsilon_1$  and  $\varepsilon_2$  are  
 196 the images of the polynomials  $(X, Y) \mapsto X$  and  $(X, Y) \mapsto Y$ . In other words, the  
 197 hyperdual numbers can be identified with  $\mathbb{R}[X, Y]/(X^2, Y^2)$ .

198 **SM5. Solving dual and hyperdual linear systems.** As noted in [section 2](#)  
 199 and [section 3](#), the CSD, DUAL, and HYPER algorithms must invoke the inner solver

200 with complex-valued, dual-valued, and hyperdual-valued state and parameters. How-  
 201 ever, assuming the inner solver invokes Newton-type steps, it must solve complex,  
 202 dual, or hyperdual linear systems, respectively. This Appendix describes an efficient  
 203 strategy for solving dual-valued and hyperdual-valued linear systems, and illustrates  
 204 its execution in Julia using packages that were developed specifically for this study:  
 205 `DualMatrixTools` [SM18] and `HyperDualMatrixTools` [SM19].

206 At first, it would appear that in order to solve the linear systems involving a dual-  
 207 valued matrix of the type  $\mathbf{M} = \mathbf{A} + \varepsilon\mathbf{B}$ , the factorization of the dual-valued matrix  
 208  $\mathbf{M}$  is necessary. However, doing so is not straightforward and is likely not efficient:  
 209 (i) there is no guarantee that the underlying package that performs factorizations  
 210 can handle dual-valued sparse matrices (e.g., UMFPACK can only handle real and  
 211 complex numbers), and (ii) performing dual-valued factorizations is computationally  
 212 more expensive than performing real-valued factorizations. We note that this is also  
 213 the case for solving complex-valued linear systems, and that a similar technique —  
 214 avoiding complex-valued factorizations — can be used for solving complex-valued  
 215 linear systems, as long as the imaginary part is small (as is the case for the complex-  
 216 step differentiation method).

217 Instead, a better strategy is to only factorize  $\mathbf{A}$ , the real part of  $\mathbf{M}$ . Because we  
 218 are using dual numbers, the Taylor expansion of the inverse of  $\mathbf{A} + \varepsilon\mathbf{B}$  is given by

$$219 \quad (\text{SM5.1}) \quad (\mathbf{A} + \varepsilon\mathbf{B})^{-1} = (\mathbf{I}_x - \varepsilon\mathbf{A}^{-1}\mathbf{B})\mathbf{A}^{-1},$$

220 where  $\mathbf{I}_x$  is the identity matrix (with its size,  $n \times n$ , indicated by the  $x$  superscript).  
 221 Equation (SM5.1) allows to solve dual-valued linear systems using only the factors of  
 222 the real part,  $\mathbf{A}$ , via two real-valued forward and back substitutions.

223 For the DUAL algorithm to work in our implementation, i.e., for the DUAL  
 224 algorithm to be able to invoke the inner solver with dual-valued state and parameters,  
 225 we developed `DualMatrixTools` [SM18], to perform the shortcut provided by (SM5.1)  
 226 under the hood. That is, it allows to factorize the real part only of  $\mathbf{M}$  (using the  
 227 `factorize` function) and to solve the corresponding dual-valued linear system (using  
 228 the backslash function, `\`) via (SM5.1) in a single line of code for each of those two  
 229 operations. We note that this package can be used in other applications and, again,  
 230 emphasize its ease of use.

231 In the case of the HYPER algorithms, we provide a hyperdual equivalent of  
 232 (SM5.1), which shows that only the inverse of  $\mathbf{A}$  is required for computing the inverse  
 233 of a hyperdual-valued matrix  $\mathbf{M} = \mathbf{A} + \varepsilon_1\mathbf{B} + \varepsilon_2\mathbf{C} + \varepsilon_1\varepsilon_2\mathbf{D}$ . In fact, the inverse of  
 234  $\mathbf{M}$  is given by

$$235 \quad (\text{SM5.2}) \quad \mathbf{M}^{-1} = [\mathbf{I}_x - \varepsilon_1\mathbf{A}^{-1}\mathbf{B} - \varepsilon_2\mathbf{A}^{-1}\mathbf{C} - \varepsilon_1\varepsilon_2\mathbf{A}^{-1}(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B} - \mathbf{B}\mathbf{A}^{-1}\mathbf{C})]\mathbf{A}^{-1}.$$

236 Just like for dual-valued linear systems, we developed a Julia package, `Hyper-`  
 237 `DualMatrixTools` [SM19], to solve hyperdual-valued systems under the hood. In a  
 238 single line of code, one can factorize the real part of  $\mathbf{M}$  using `factorize` or solve a  
 239 hyperdual-valued linear system using backslash (`\`). This package has been optimized  
 240 to compute hyperdual-valued forward and back substitutions in just four real-valued  
 241 forward and back substitutions, as can be derived from (SM5.2).

242 **SM6. The pitfalls of black-box approaches.** In this appendix, we describe  
 243 what happens when invoking the solver with complex, dual or hyperdual parameters  
 244 as a black box. Here we show that not only is this naive approach slower, it directly  
 245 introduces numerical errors. That is, we describe what happens when using the

246 DUAL, HYPER, and CSD algorithms. We emphasize that when using an iterative  
 247 algorithm that was written for real-valued numbers, it is imperative to set additional  
 248 tolerances on the non-real part(s) of the iterates. Setting these tolerances, which  
 249 control when the loop terminates, can be complicated. But before diving in the  
 250 details, we first recall Newton’s method in the real case.

251 Starting from an initial estimate  $\mathbf{x}_0$ , Newton’s method is based on the recursion  
 252 relation

$$253 \quad (\text{SM6.1}) \quad \mathbf{x}_{l+1} = \mathbf{x}_l - \nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x}_l, \mathbf{p})^{-1} \mathbf{F}(\mathbf{x}_l, \mathbf{p}),$$

254 where  $\nabla_{\mathbf{x}} \mathbf{F}(\mathbf{x}_l, \mathbf{p})$  is the Jacobian evaluated at the  $l$ th iterate of the state,  $\mathbf{x}_l$ . If all  
 255 goes well, the iterates converge to the steady state, i.e.,  $\lim_{l \rightarrow \infty} \mathbf{x}_l = \mathbf{s}(\mathbf{p})$ .

256 In practice, when invoked for real-valued parameters, the solver we use applies  
 257 Shamanskii’s method, which improves on Newton’s method by not updating the Jaco-  
 258 bian at each iterate, increasing computational performance. Additionally, the solver  
 259 we use also performs an Armijo line search if the Shamanskii step overshoots the  
 260 solution (the solver we use was adapted from the `nsold` solver from Kelley [SM11]).  
 261 That is, the Armijo line search looks for the minimum of the norm of  $\mathbf{F}(\mathbf{x}, \mathbf{p})$  along  
 262 the Shamanskii-step direction via a quadratic approximation. However, for the sake  
 263 of clarity, here we only detail the classical Newton’s method. (The solver we actually  
 264 use is available in the `AIBECS` package [SM20].)

265 We now explore the convergence of the Newton’s method when the parameters are  
 266 dual-valued (i.e., what happens when using the DUAL algorithm). We thus replace  
 267  $\mathbf{p}$  with  $\mathbf{p} + \varepsilon \mathbf{e}_j$ . Although we start from the real-valued  $\mathbf{x}_0 = \mathbf{s}(\mathbf{p})$ , the iterates  $\mathbf{x}_l$   
 268 include a potentially non-trivial dual part for  $l > 0$ , since the dual part of  $\mathbf{p} + \varepsilon \mathbf{e}_j$  will  
 269 spread to  $\mathbf{x}_l$  at each Newton iteration. We thus partition  $\mathbf{x}_l$  into its real and dual  
 270 parts, such that

$$271 \quad (\text{SM6.2}) \quad \mathbf{x}_l = \mathbf{a}_l + \varepsilon \mathbf{b}_l,$$

272 where we omit the  $\mathbf{p} + \varepsilon \mathbf{e}_j$  argument for brevity.

273 The dual-valued  $\mathbf{F}(\mathbf{a}_l + \varepsilon \mathbf{b}_l, \mathbf{p} + \varepsilon \mathbf{e}_j)$  can be expressed via the Taylor expansion  
 274 of  $\mathbf{F}$  at  $(\mathbf{a}_l, \mathbf{p})$  in the  $(\varepsilon \mathbf{b}_l, \varepsilon \mathbf{e}_j)$  direction, which gives

$$275 \quad (\text{SM6.3}) \quad \mathbf{F}(\mathbf{a}_l + \varepsilon \mathbf{b}_l, \mathbf{p} + \varepsilon \mathbf{e}_j) = \mathbf{F}(\mathbf{a}_l, \mathbf{p}) + \varepsilon \mathbf{A}_l \mathbf{b}_l + \varepsilon \nabla_{\mathbf{p}} \mathbf{F}(\mathbf{a}_l, \mathbf{p}) \mathbf{e}_j,$$

276 where  $\mathbf{A}_l \equiv \nabla_{\mathbf{x}} \mathbf{F}(\mathbf{a}_l, \mathbf{p})$  is a real-valued Jacobian matrix. Similarly, the dual-valued  
 277 Jacobian matrix  $\nabla_{\mathbf{x}} \mathbf{F}(\mathbf{a}_l + \varepsilon \mathbf{b}_l, \mathbf{p} + \varepsilon \mathbf{e}_j)$  can also be expressed by a dual Taylor ex-  
 278 pansion. However, we need not detail all its terms, and instead simply denote its real  
 279 part by  $\mathbf{A}_l$  and its dual part by the unspecified matrix  $\mathbf{B}_l$ , respectively, so that

$$280 \quad (\text{SM6.4}) \quad \nabla_{\mathbf{x}} \mathbf{F}(\mathbf{a}_l + \varepsilon \mathbf{b}_l, \mathbf{p} + \varepsilon \mathbf{e}_j) = \mathbf{A}_l + \varepsilon \mathbf{B}_l.$$

281 Note that  $\mathbf{A}_l = \nabla_{\mathbf{x}} \mathbf{F}(\mathbf{a}_l, \mathbf{p})$  is also the real part of the dual-valued Jacobian.

282 Evaluating the iterations of (SM6.1) with dual-valued parameters, it appears that  
 283 the real part has exactly the same form as the original real-valued iterations, except  
 284 with  $\mathbf{a}_l$  symbols in place of  $\mathbf{x}_l$  symbols. Assuming we start with  $\mathbf{a}_0 = \mathbf{s}(\mathbf{p})$  means  
 285  $\mathbf{a}_l = \mathbf{s}(\mathbf{p})$  throughout. This is expected because the dual-valued call to the solver will  
 286 not improve the real-valued solution.

287 The non-real part however contains additional terms, with

$$288 \quad (\text{SM6.5}) \quad \mathbf{b}_{l+1} = \mathbf{b}_l - \mathbf{A}_l^{-1} \mathbf{A}_l \mathbf{b}_l - \mathbf{A}_l^{-1} \nabla_{\mathbf{p}} \mathbf{F}(\mathbf{a}_l, \mathbf{p}) \mathbf{e}_j \\ + \mathbf{A}_l^{-1} \mathbf{B}_l \mathbf{A}_l^{-1} \mathbf{F}(\mathbf{a}_l, \mathbf{p}),$$

289 where we have used the analytical formula for the inverse of a dual-valued matrix,  
 290 (SM5.1). Note the first and second terms, which should cancel each other. These  
 291 were intentionally left for the reader to appreciate how a black-box approach will ex-  
 292 ecute unnecessary computations, both slowing down the computation and potentially  
 293 introducing additional numerical errors ( $\mathbf{b}_l$  would seldom be strictly exactly equal to  
 294  $\mathbf{A}_l^{-1}\mathbf{A}_l\mathbf{b}_l$  without an infinite precision computer.) Additionally, the last term should  
 295 vanish because  $\mathbf{a}_l = \mathbf{s}(\mathbf{p})$  (by definition of  $\mathbf{s}(\mathbf{p})$ ). However, this term will not vanish  
 296 exactly with finite precision machines, so that the black-box approach will compute  
 297 this term and potentially introduce errors while doing so.

298 In our implementation, the tolerance for the dual part is decided by the ratio of  
 299 the norm of the dual-part of the Newton step to the norm of the dual part. Specifically,  
 300 the solver loop terminates if  $\|\mathbf{b}_l - \mathbf{b}_{l-1}\|/\|\mathbf{b}_{l-1}\|$  is smaller than  $10^{-7}$ . Starting from  
 301 the solution  $\mathbf{s}(\mathbf{p})$ , the dual-type black-box approaches will thus invoke the inner solver  
 302 with one iteration to converge to the dual-valued solution, and another iteration to  
 303 check that the dual-part has converged, thus requiring two iterations overall.

304 We note that in this specific context, one could get away with forgetting to set  
 305 a tolerance for the dual part only because the inner solver effectively converges in a  
 306 single iteration. However, this is not the case for all iterative algorithms taken as a  
 307 black box, as illustrated by the  $\sim 100$ -iterations lag of the non-real part described by  
 308 Martins et al. [SM13]. Therefore we consider that omitting these tolerances is *not*  
 309 a reasonable approach. Conversely, we also note that setting too tight a tolerance  
 310 for the non-real part could potentially result in a large (or even infinite) number  
 311 of iterations for the solver loop to terminate, when in fact it oscillates over a small  
 312 neighborhood of the solution until a small enough dual-valued step is randomly taken.  
 313 We emphasize that all the arguments brought up in this appendix so far also apply to  
 314 the complex-step algorithm. That is, the exact same phenomenon occurs when using  
 315 the CSD algorithm. In comparison, the F-1 algorithm completely alleviates the need  
 316 to set a tolerance for the dual part of the inner solver, since it does not invoke it.

317 We now show that, in our context, the inner solver invoked with hyperdual-  
 318 valued parameters converges in exactly two iterations (i.e., when using the HYPER  
 319 algorithm). We give a brief outline of how the derivation goes, but leave its details out  
 320 for brevity. Again, we consider the hyperdual parameters defined by  $\mathbf{p}_{jk} = \mathbf{p} + \varepsilon_1\mathbf{e}_j +$   
 321  $\varepsilon_2\mathbf{e}_k$ . Starting from the real-valued solution,  $\mathbf{s}(\mathbf{p})$ , just like in the dual case, the real  
 322 part of the hyperdual state remains constant, equal to  $\mathbf{s}(\mathbf{p})$ , throughout the Newton  
 323 iterations. The first iteration permeates hyperdual-values from the parameters into  
 324 the first state iterate,  $\mathbf{x}_1$ . In fact, the  $\varepsilon_1$  and  $\varepsilon_2$  parts of the state converge in that first  
 325 iteration, so that the  $\varepsilon_1$  and  $\varepsilon_2$  parts of  $\mathbf{s}(\mathbf{p}_{jk})$  are the  $\varepsilon_1$  and  $\varepsilon_2$  parts of  $\mathbf{x}_1$ . However,  
 326 a second iteration is required for the  $\varepsilon_1\varepsilon_2$  part to converge, because it requires that  
 327 the previous state (in this case  $\mathbf{x}_1$ ) contains the converged  $\varepsilon_1$  and  $\varepsilon_2$  parts. Thus,  
 328 convergence of all the non-real parts requires exactly two iterations. This is because,  
 329 unlike for the dual case, some of the non-real terms that include non-real parts of the  
 330 state (the terms including  $\mathbf{b}_l$  in (SM6.5)) do not cancel out in the hyperdual case.  
 331 This is also because the  $\varepsilon_1\varepsilon_2$  part of  $\mathbf{x}_{l+1}$  only depends on the real,  $\varepsilon_1$ , and  $\varepsilon_1$  parts of  
 332  $\mathbf{x}_l$ . Just like in the dual case, we emphasize that the inner solver must check that the  
 333 non-real parts have converged, so that the HYPER algorithm requires one additional  
 334 Newton step to terminate, bringing the total number of iterations to three.

335 **Code and data availability.** The state function,  $\mathbf{F}$ , and the mismatch function,  
 336  $f$ , as well as their derivatives, were automatically generated using the AIBECS pack-  
 337 age [SM20] or manually generated in the FastBGCPParameterOptimization GitHub



338 repository. The ocean circulation was taken from the output of the control (CTL)  
339 run of the Ocean Circulation Inverse Model (OCIM1 [SM4]) and is available online  
340 in a MAT file (the MATLAB data format) on the [website](#) of Tim DeVries. The  
341 OCIM1 grid and transport matrix have been converted to Julia to be downloaded  
342 and used by the [AIBECS](#) package [SM20]. The mismatch of the state used the ob-  
343 jectively analyzed mean field of the phosphate concentration from the World Ocean  
344 Atlas 2018 (WOA18 [SM9]). The WOA18 data was downloaded and fitted to the  
345 OCIM grid using the [WorldOceanAtlasTools](#) package [SM22]. The F-1 algorithm is  
346 implemented in the [F1Method](#) package [SM21]. The CSD, FD1, DUAL, HYPER, and  
347 FD2 algorithms, the figures, and the underlying data that are shown in this study  
348 are contained in the [FastBGCPParameterOptimization](#) GitHub repository. Dual- and  
349 hyperdual-valued factorizations and forward and back substitutions of sparse linear  
350 systems were performed using the [DualMatrixTools](#) [SM18] and [HyperDualMatrix-](#)  
351 [Tools](#) [SM19] packages, respectively. The optimizations were performed using the  
352 [Optim](#) package [SM14, SM15]. The concrete types defining the steady-state prob-  
353 lems and the steady-state solutions were built upon the [DiffEqBase](#) package [SM23].  
354 Benchmarks were performed with the [BenchmarkTools](#) and [TimerOutputs](#) packages.  
355 The diagram was created using the [L<sup>A</sup>T<sub>E</sub>X TikZ](#) package [SM26]. Result figures were  
356 created using the [PyPlot](#) package. All the computations were performed using the  
357 [Julia](#) language [SM1].

358

## REFERENCES

- 359 [SM1] J. BEZANSON, A. EDELMAN, S. KARPINSKI, AND V. SHAH, *Julia: A fresh approach to numerical*  
360 *computing*, SIAM Review, 59 (2017), pp. 65–98, <https://doi.org/10.1137/141000671>.
- 361 [SM2] W. CLIFFORD, *Preliminary sketch of bi-quaternions*, Proceedings of London Mathematical  
362 Society, 4 (1873), pp. 361–395.
- 363 [SM3] T. DEVRIES, *The oceanic anthropogenic CO<sub>2</sub> sink: Storage, air-sea fluxes, and transports*  
364 *over the industrial era*, Global Biogeochem. Cycles, 28 (2014), pp. 631–647, [https://doi.](https://doi.org/10.1002/2013GB004739)  
365 [org/10.1002/2013GB004739](https://doi.org/10.1002/2013GB004739).
- 366 [SM4] T. DEVRIES, J.-H. LIANG, AND C. DEUTSCH, *A mechanistic particle flux model applied to*  
367 *the oceanic phosphorus cycle*, Biogeosciences, 11 (2014), pp. 5381–5398, [https://doi.org/](https://doi.org/10.5194/bg-11-5381-2014)  
368 [10.5194/bg-11-5381-2014](https://doi.org/10.5194/bg-11-5381-2014).
- 369 [SM5] T. DEVRIES AND F. PRIMEAU, *Dynamically and observationally constrained estimates of*  
370 *water-mass distributions and ages in the global ocean*, J. Phys. Oceanogr., 41 (2011),  
371 pp. 2381–2401, <https://doi.org/10.1175/JPO-D-10-05011.1>.
- 372 [SM6] J. FIKE AND J. ALONSO, *The development of hyper-dual numbers for exact second-derivative*  
373 *calculations*, in 49th AIAA Aerospace Sciences Meeting including the New Horizons Forum  
374 and Aerospace Exposition, 2011, p. 886.
- 375 [SM7] J. A. FIKE AND J. J. ALONSO, *Automatic differentiation through the use of hyper-dual num-*  
376 *bers for second derivatives*, in Recent Advances in Algorithmic Differentiation, S. Forth,  
377 P. Hovland, E. Phipps, J. Utke, A. Walther, T. J. Barth, M. Griebel, D. E. Keyes,  
378 R. M. Nieminen, D. Roose, and T. Schlick, eds., vol. 87 of Lecture Notes in Computa-  
379 tional Science and Engineering, Springer Berlin Heidelberg, 2012, pp. 163–173, [https:](https://doi.org/10.1007/978-3-642-30023-3_15)  
380 [//doi.org/10.1007/978-3-642-30023-3\\_15](https://doi.org/10.1007/978-3-642-30023-3_15).
- 381 [SM8] J. A. FIKE, S. JONGSMA, J. J. ALONSO, AND E. VAN DER WEIDE, *Optimization with gradient*  
382 *and Hessian information calculated using hyper-dual numbers*, in AIAA paper 2011-3807,  
383 29th AIAA Applied Aerodynamics Conference, 2011.
- 384 [SM9] H. E. GARCIA, K. WEATHERS, C. R. PAVER, I. SMOLYAR, T. P. BOYER, R. A. LOCARNINI,  
385 M. M. ZWENG, A. V. MISHONOV, O. K. BARANOVA, D. SEIDOV, AND J. R. REAGAN,  
386 *World Ocean Atlas 2018*, Volume 4: Dissolved Inorganic Nutrients (phosphate, nitrate  
387 and nitrate+nitrite, silicate) (in preparation). A. Mishonov Technical Ed.
- 388 [SM10] A. H. GEBREMEDHIN, F. MANNE, AND A. POTHEN, *What color is your jacobian? graph*  
389 *coloring for computing derivatives*, SIAM review, 47 (2005), pp. 629–705.
- 390 [SM11] C. T. KELLEY, *Solving Nonlinear Equations with Newton’s Method*, SIAM, 2003, ch. 2.  
391 Finding the Newton Step with Gaussian Elimination, pp. 27–55, <https://doi.org/10.1137/>

- 392 [1.9780898718898.ch2](https://doi.org/10.1137/0704019).
- 393 [SM12] J. N. LYNES AND C. B. MOLER, *Numerical differentiation of analytic functions*, SIAM  
394 Journal on Numerical Analysis, 4 (1967), pp. 202–210, <https://doi.org/10.1137/0704019>.
- 395 [SM13] J. R. R. A. MARTINS, P. STURDZA, AND J. J. ALONSO, *The complex-step derivative ap-*  
396 *proximation*, ACM Trans. Math. Softw., 29 (2003), pp. 245–262, [https://doi.org/10.1145/](https://doi.org/10.1145/838250.838251)  
397 [838250.838251](https://doi.org/10.1145/838250.838251).
- 398 [SM14] P. K. MOGENSEN AND A. N. RISETH, *Optim: A mathematical optimization package for Julia*,  
399 Journal of Open Source Software, 3 (2018), p. 615, <https://doi.org/10.21105/joss.00615>.
- 400 [SM15] P. K. MOGENSEN, J. M. WHITE, A. N. RISETH, T. HOLY, M. LUBIN, C. STOCKER, A. LEVITT,  
401 C. ORTNER, B. JOHNSON, A. NOACK, Y. YU, K. CARLSSON, D. LIN, T. R. COVERT,  
402 R. ROCK, J. REGIER, B. KUHN, A. WILLIAMS, RYAN, K. SATO, D. SMITH, R. ANANTHARA-  
403 MAN, M. GOMEZ, J. REVELS, I. DUNNING, D. MACMILLEN, C. RACKAUCKAS, B. LEGAT,  
404 AND A. STUKALOV, *JuliaNLSolvers/Optim.jl: Bugfix for Fminbox, better docs, and exper-*  
405 *imental maximize function*, Sept. 2018, <https://doi.org/10.5281/zenodo.1412092>.
- 406 [SM16] C. MOLER, *Cleve's Corner: Cleve Moler on Mathematics and Computing —*  
407 *Complex Step Differentiation*, 2013, [https://blogs.mathworks.com/cleve/2013/10/14/](https://blogs.mathworks.com/cleve/2013/10/14/complex-step-differentiation)  
408 [complex-step-differentiation](https://blogs.mathworks.com/cleve/2013/10/14/complex-step-differentiation) (accessed 2019-04-29).
- 409 [SM17] M. NEUENHOFEN, *Review of theory and implementation of hyper-dual numbers for first and*  
410 *second order automatic differentiation*, CoRR, abs/1801.03614 (2018), [http://arxiv.org/](http://arxiv.org/abs/1801.03614)  
411 [abs/1801.03614](http://arxiv.org/abs/1801.03614), <https://arxiv.org/abs/1801.03614>.
- 412 [SM18] B. PASQUIER, *DualMatrixTools: A Julia package for solving dual-valued linear systems*, Nov.  
413 2018, <https://doi.org/10.5281/zenodo.1493571>.
- 414 [SM19] B. PASQUIER, *HyperDualMatrixTools: A Julia package for solving hyperdual-valued linear*  
415 *systems*, Nov. 2018, <https://doi.org/10.5281/zenodo.1670841>.
- 416 [SM20] B. PASQUIER, *AIBECs.jl: The ideal tool for exploring global marine biogeochemical cycles*,  
417 May 2019, <https://doi.org/10.5281/zenodo.2864051>.
- 418 [SM21] B. PASQUIER, *F1Method.jl: A julia package for computing the gradient and Hessian of an*  
419 *objective function defined implicitly by the solution to a steady-state problem*, May 2019,  
420 <https://doi.org/10.5281/zenodo.2667835>.
- 421 [SM22] B. PASQUIER, *Worldoceanatlastools.jl: Download and process world ocean data in julia*, May  
422 2019, <https://doi.org/10.5281/zenodo.2677666>.
- 423 [SM23] C. RACKAUCKAS AND Q. NIE, *Differentialequations.jl – a performant and feature-rich eco-*  
424 *system for solving differential equations in julia*, Journal of Open Research Software, 5  
425 (2017), <https://doi.org/10.5334/jors.151>.
- 426 [SM24] W. SQUIRE AND G. TRAPP, *Using complex variables to estimate derivatives of real functions*,  
427 SIAM Review, 40 (1998), pp. 110–112, <https://doi.org/10.1137/S003614459631241X>.
- 428 [SM25] E. STUDY, *Geometrie der dynamen: Die zusammensetzung von kräften und verwandte*  
429 *gegenstände der geometrie*, 1903.
- 430 [SM26] T. TANTAU, *The TikZ and PGF Packages*, 2013, <http://sourceforge.net/projects/pgf/>.