

1 **eddy4R 0.2.0: A DevOps model for community-extensible**
2 **processing and analysis of eddy-covariance data based on**
3 **R, Git, Docker and HDF5**

4
5 **Stefan Metzger^{1,2}, David Durden¹, Cove Sturtevant¹, Hongyan Luo¹, Natchaya**
6 **Pingitha-Durden¹, Torsten Sachs³, Andrei Serafimovich³, Jörg Hartmann⁴,**
7 **Jiahong Li⁵, Ke Xu², Ankur R. Desai²**

8 ¹National Ecological Observatory Network, Battelle, 1685 38th Street, Boulder, CO 80301,
9 USA

10 ²University of Wisconsin-Madison, Dept. of Atmospheric and Oceanic Sciences, 1225 West
11 Dayton Street, Madison, WI 53706, USA

12 ³GFZ German Research Centre for Geosciences, Telegrafenberg, 14473 Potsdam, Germany

13 ⁴Alfred Wegener Institute – Helmholtz Centre for Polar and Marine Research, Am
14 Handelshafen 12, 27570 Bremerhaven, Germany

15 ⁵LI-COR Biosciences, 4647 Superior Street, Lincoln, NE 68504, USA

16 Correspondence to: Stefan Metzger (smetzger@battelleecology.org)

17
18 **Keywords:** computing, container, continuous development, continuous integration, devOps,
19 eddy4R, eddy-covariance, EddyPro, EdiRe, EddyUH, image, NEON, REddyProc,
20 reproducibility, science code, TK3

21

22 **Abstract**

23 Large differences in instrumentation, site setup, data format, and operating system stymie the
24 adoption of a universal computational environment for processing and analyzing eddy-
25 covariance (EC) data. This results in limited software applicability and extensibility in addition
26 to often substantial inconsistencies in flux estimates. Addressing these concerns, this paper
27 presents the systematic development of portable, reproducible, and extensible EC software
28 achieved by adopting a Development and Systems Operation (DevOps) approach. This
29 software development model is used for the creation of the eddy4R family of EC code packages
30 in the open-source R Language for Statistical Computing. These packages are community-
31 developed, iterated via the Git distributed version control system, and wrapped into a portable
32 and reproducible Docker filesystem that is independent of the underlying host operating
33 system. The HDF5 hierarchical data format then provides a streamlined mechanism for
34 highly compressed and fully self-documented data ingest and output.

35 The usefulness of the DevOps approach was evaluated for three test applications. First, the
36 resultant EC processing software was used to analyze standard flux tower data from the first
37 EC instruments installed at a National Ecological Observatory (NEON) field site. Second,
38 through an aircraft test application we demonstrate the modular extensibility of eddy4R to
39 analyze EC data from other platforms. Third, an intercomparison with commercial-grade
40 software showed excellent agreement ($R^2=1.0$ for CO_2 flux). In conjunction with this study,
41 a Docker image containing the first two eddy4R packages and an executable example
42 workflow, as well as first NEON EC data products are released publicly. We conclude by
43 describing the work remaining to arrive at the automated generation of science-grade EC fluxes,
44 and benefits to the science community at large.

45 This software development model is applicable beyond EC, and more generally builds the
46 capacity to deploy complex algorithms developed by scientists in an efficient and scalable
47 manner. In addition, modularity permits meeting project milestones while retaining
48 extensibility with time.

49

50 **1 Introduction**

51 Answering grand challenges in earth system science and ecology requires combining
52 information from hierarchies of environmental observations (tower, aircraft, satellite; Raupach
53 et al., 2005; Running et al., 1999; Turner et al., 2004). Eddy-covariance (EC) measurements
54 serve as crucial observations in this hierarchy to study landscape-scale surface-atmosphere
55 exchange processes that both inform and anchor earth system models. Networks of EC towers
56 such as FLUXNET (Baldocchi et al., 2001), AmeriFlux (Law, 2007), ICOS (Sulkava et al.,
57 2011), and others are vital for providing the necessary distributed observations covering the
58 climate space, with the longest running towers now reaching two decades of observations.

59 A current challenge for EC tower networks in informing regional and continental scale
60 processes is instrument and computational compatibility. The computations involved in EC
61 processing are complex and developmentally dynamic, making code portability, extensibility,
62 and documentation paramount. Much progress has been made in developing community
63 standards for processing algorithms and workflows (Aubinet et al., 2012; Papale et al., 2006).
64 Many authors have included code in publication, or have developed sharable tools (e.g.
65 EddyPro and TK3 by Fratini and Mauder (2014), EddyUH by Mammarella et al. (2016), EdiRe
66 by Clement et al. (2009), despite the significant and often unfunded effort required to
67 adequately document and generalize code. Still, large differences in instrumentation, site setup,
68 data format, and operating systems stymie the adoption of a universal EC processing
69 environment: one that is portable, reproducible, and extensible to allow tailored workflows that
70 incorporate additional data streams, to automate and scale processing across large compute
71 facilities, or to inject additional algorithms that address specific needs or synergistic research
72 questions. In 50% of published scientific code, one cannot even replicate the necessary software
73 dependencies (Collberg et al., 2014), and even widely used and well-documented EC
74 processing software packages have shown substantial inconsistencies in flux estimates (e.g.
75 Fratini and Mauder, 2014). A universal EC processing environment that enables these
76 capabilities would better allow research groups to tailor existing software to their needs (and
77 contribute new algorithms) instead of re-creating code or kludging together multiple software
78 outputs to realize an algorithmic chain for their data analytics.

79 The U.S.-based National Ecological Observatory Network (NEON), once fully operational, will
80 represent the largest single-provider EC tower network globally, with a standardized
81 measurement suite designed explicitly for cross-site comparability and analysis of continental-
82 scale ecological change (Schimel et al., 2007). This capability is accompanied by a strong need
83 for a flexible and scalable processing framework that can incorporate specific data streams, take
84 advantage of close alignment of hardware and software for problem tracking and resolution,
85 provide traceability and reproducibility of outputs, and seamlessly integrate distributed and
86 dynamic community-developed code (written by multiple people in multiple places) within
87 existing cyberinfrastructure (CI). In sum, NEON needs what the EC community is currently
88 lacking.

89 The question we ask in this paper is: How do we collaboratively create portable, reproducible,
90 open-source, scalable, and extensible software that improves reliability and comparability of

91 EC data products? Here, we describe and demonstrate a developmental model that enables these
92 capabilities by embracing a Development and Systems Operation (DevOps) approach. DevOps
93 is a philosophy arising from the software development community that emphasizes
94 collaboration among developers and operators to continuously iterate the development, building,
95 testing, packaging, and release of software (Erich et al., 2014; Loukides, 2012). Tools are
96 adopted that control and automate these processes, allowing distributed development and rapid
97 iteration. Applied to the scientific community, developers are the multitude of scientists
98 creating and improving the scientific algorithms that form the developmentally dynamic
99 community standard. Operators are those deploying the algorithms to process and analyze data,
100 and can be the same or different people as those creating the algorithms. A key aspect of
101 DevOps is the recipe- or script-based generation and packaging of computation environments
102 rather than abstracted documentation, which improves accessibility, extensibility, and
103 reproducibility of scientific software (Boettiger, 2015; Clark et al., 2014). The recipe automates
104 the loading of the software including all dependencies so that the most significant hurdle of
105 reproducing the computational environment is overcome. At the same time, the recipe serves
106 as explicit documentation, and can be easily extended (added to or changed), shared, and
107 versioned. The entire computational environment including any necessary data are packaged
108 into Docker images that work identically across different computers and operating systems, can
109 be deployed at scale, and archived for ultimate reproducibility.


110 In the following we present this framework and demonstrate its success in producing EC data
111 products via a family of modular, open-source R packages wrapped in Docker images. We
112 emphasize that this paper is not a presentation of EC processing software (although this is the
113 ultimate application). Rather, it is a presentation of the development model that facilitates
114 portability, reproducibility, and extensibility of EC processing software. In the following,
115 Sect. 2 describes the DevOps framework, and Sect. 3 provides three core tests of the
116 applicability of this framework: 1) processing tower-based flux data, including NEON's first
117 set of EC data, 2) processing and footprint modeling of aircraft-based flux data, and 3) a
118 software cross-validation. Section 4 summarizes the work remaining to operationally produce
119 EC fluxes from 47 NEON sites, and provides an outlook on future capabilities and science
120 community benefits. Code and data availability information is provided in Sect. 5.

121 **2 The development and operations (DevOps) model**

122 DevOps promotes collaboration and tight integration between software development, testing,
123 and operational deployment by following a core workflow (e.g., Wurster et al., 2015): Plan,
124 Create, Verify, Package, Release, Configure, and Monitor. The text below describes these
125 stages and **Error! Reference source not found.** shows the general sequence and overlap of
126 these stages between software developers (Dev) and operators (Ops).

127 **Plan** involves focusing and prioritizing new software features or capabilities based on their
128 enhancement of value. **Create** is the activity of designing and writing the code that delivers a
129 new feature. **Verify** tests the new software feature against established standards for accuracy
130 and performance (e.g. does it unexpectedly alter the output of pre-existing features? Does it
131 produce the expected result?). **Package** involves the compilation of the code once it is ready

132 for deployment, including all data and software dependencies, and gathers necessary approvals.
133 The **Release** stage deploys the software into production. **Configure** involves supplying and
134 configuring the computational infrastructure required to operate the code at scale, including
135 storage, database operations, and networking. Finally, **Monitor** observes and tracks the use,
136 performance, and end-user impact of the release.
137

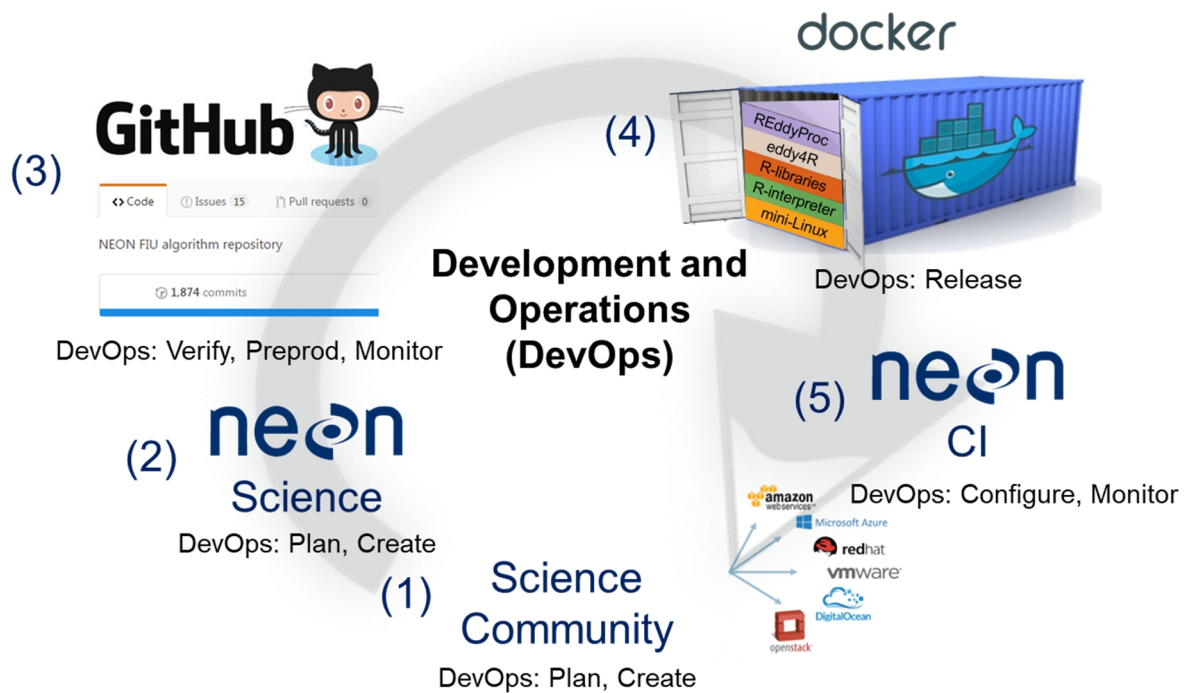


138
139 Figure 1. Stages of the general DevOps workflow (source: Kharnagy via Wikimedia
140 Commons [CC BY-SA 4.0]).
141

142 Variants of this workflow exist (e.g., Chen, 2015), but the general components and sequence
143 are retained. In addition, there is no single set of tools accompanying the DevOps approach.
144 Rather, many tools exist that facilitate the execution of one or more of these workflow steps,
145 often through automation.

146 NEON’s DevOps framework consists of a periodic sequence (Figure 2) that incorporates these
147 workflow steps. For this purpose we define NEON Science as personnel working directly on
148 the NEON project, and the Science Community, regardless of whether they also work on the
149 NEON project, as anyone producing or using data, algorithms, or research products related to
150 the NEON data themes (Atmosphere; Biogeochemistry; Ecohydrology; Land Cover and
151 Processes; Organisms, Populations, and Communities): The science community contributes
152 algorithms and best practices (1). Implicitly or explicitly, this embodies the DevOps: Plan stage
153 – the algorithms most valued by the community are being incorporated. Together with NEON
154 Science (2), these algorithms are coded in the open-source R computational environment
155 (DevOps: Create stage). DevOps: Verify (testing) and Package (packaging) are performed as
156 the code is compiled into eddy4R packages via the GitHub distributed version control system
157 (3). NEON Science releases an eddy4R version from GitHub, which automatically builds an
158 eddy4R-Docker image on DockerHub as specified in a “Dockerfile” (4; DevOps: Release stage).
159 The eddy4R-Docker image is immediately available for deployment by NEON CI (5; DevOps:
160 Configure & Monitor stages), the Science Community (1) and NEON Science (2) alike. Here
161 the DevOps: Configure (computational resource allocation) & Monitor stages occur.

162 Monitoring of end-user experience is also performed in GitHub (3) via issue-tracking. This
 163 DevOps cycle can be repeated for continuous development and integration of requests and
 164 future methodological improvements by the scientific community, resulting in the next release.
 165 Two principal types of releases are provided: stable versions are tagged with “0.2.0”, “0.2.1”
 166 etc., and the most recent development built is tagged with “latest”. Thus, the DevOps model
 167 serves as the framework within which the scientific community can efficiently and robustly
 168 collaborate to produce, manage, and iterate software. Through choosing appropriate tools to
 169 implement the DevOps workflow steps, the reproducibility, scalability and extensibility needs
 170 of software development communities (including EC) can be met.
 171



172
 173 Figure 2. NEON-specific DevOps workflow. DevOps workflow steps are called out in
 174 parentheses. Please see text in Sect. 2 for detailed explanation.

175
 176 In the following we describe the key components and tools of this NEON-specific DevOps
 177 model, namely the eddy4R family of code packages (Sect. 2.1), Git-based distributed code
 178 development (Sect. 2.2), packaging of the computational environment in Docker images
 179 (Sect. 2.3), hierarchical data formats (Sect. 2.4), integration with NEON’s CI (Sect. 2.5), and
 180 installation and deployment (Sect. 2.6).

181 **2.1 The eddy4R family of R-packages (DevOps: Plan & Create)**

182 eddy4R is a family of open-source packages for EC raw data processing, analyses and modeling
 183 in the R Language for Statistical Computing (R Core Team, 2016). Forming the DevOps: Plan
 184 & Create stages, it is being developed by NEON scientists with wide input from the

185 micrometeorological community (e.g., De Roo et al., 2014; Kohnert et al., 2015; Lee et al.,
186 2015; Metzger et al., 2012; Metzger et al., 2013; Metzger et al., 2016; Sachs et al., 2014; Salmon
187 et al., 2015; Serafimovich et al., 2013; Starkenburg et al., 2016; Vaughan et al., 2015; Xu et al.,
188 2017). eddy4R currently consists of four packages eddy4R.base, eddy4R.qaqc, eddy4R.turb,
189 and eddy4R.erf. Of these, eddy4R.base and eddy4R.qaqc are published here in conjunction with
190 NEON's release of EC Level 1 data products ([https://w3id.org/smetzger/Metzger-et-
191 al_2017_eddy4R-Docker/portal/0.2.0](https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/portal/0.2.0)): descriptive statistics of calibrated instrument output. In
192 addition, previews of eddy4R.turb and eddy4R.erf are provided, which will be published along
193 NEON's upcoming release of EC Level 4 data products (derived quality-controlled fluxes and
194 related variables). Development of two additional R-packages eddy4R.stor and eddy4R.ucrt has
195 started, which provide functionalities for storage flux computation and uncertainty
196 quantification, respectively. These packages are not covered here, and will be published once
197 available.

198 Each eddy4R package consists of a hierarchical set of reusable definition functions, wrapper
199 functions and workflows. Following best practices, eddy4R is written in controlled and strictly
200 hierarchical terminology consisting of base names and modifiers, which ensures modular
201 extensibility over time. Interactive documentation is provided through the use of Roxygen tags
202 (<http://roxygen.org/>) during development, and follows the Comprehensive R Archive Network
203 (CRAN; <https://cran.r-project.org/>) guidelines for package dissemination. In addition,
204 expanded documentation is available in the form of Algorithm Theoretical Basis Documents
205 from the NEON data portal ([https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-
206 Docker/portal/0.2.0](https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/portal/0.2.0)).

207 EC data processing consists of employing a sequence of model algorithms. These often
208 originate from scientific sub-fields with corresponding publications, and eddy4R provides an
209 integrative, yet modular and extensible framework for their concerted application and continued
210 development: eddy4R.base provides natural constants and basic functions for usability,
211 regularization, transformation, lag-correction, aggregation and unit conversion ensuring
212 consistency of internal units at any point in the workflow. Next, eddy4R.qaqc provides the
213 general quality assurance and quality control (QA/QC) tests of Taylor and Loescher (2013),
214 along the Smith et al. (2014) model for tracking quality information in large datasets, and
215 functions for de-spiking (Brock, 1986; Fratini and Mauder, 2014; Mauder et al., 2013; Mauder
216 and Foken, 2015; Metzger et al., 2012; Vickers and Mahrt, 1997). eddy4R.turb provides
217 standard, Reynolds-decomposed turbulent flux calculation (Foken, 2017), accompanied by
218 models for planar fit transformation (Wilczak et al., 2001) and spectral correction (Nordbo and
219 Katul, 2012). Additional functionalities include Fourier transform, the determination of
220 detection limit (Billesbach, 2011), integral length scales and statistical sampling errors
221 (Lenschow et al., 1994), and flux-specific QA/QC models (Foken and Wichura, 1996; Vickers
222 and Mahrt, 1997). Also, basic scaling variables, atmospheric stability and roughness length
223 (Stull, 1988), as well as the flux footprint (Kljun et al., 2015; Kormann and Meixner, 2001;
224 Metzger et al., 2012) can be determined. Lastly, eddy4R.erf provides time-frequency de-
225 composed flux processing and data mining functionalities to determine an environmental

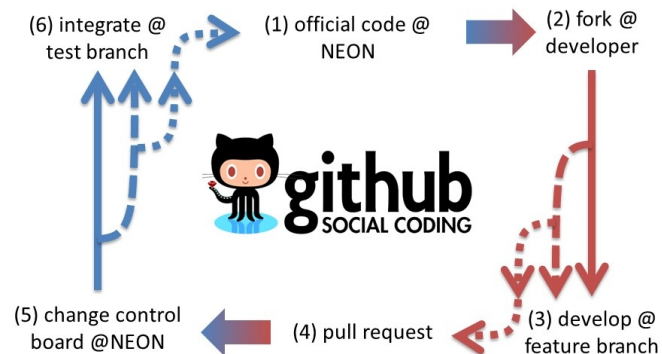
226 response function model and project the flux fields underlying the EC observations (Metzger et
227 al., 2013; Xu et al., 2017).

228 eddy4R can be used with a fully adaptive single-pass workflow (Sect. 3.1), which makes it
229 computationally efficient compared to the multiple passes required by other flux processing
230 schemes. In addition, eddy4R is fully parallelized and memory efficient leveraging R's snowfall
231 parallelization (<https://cran.r-project.org/package=snowfall>) and ff file-backed object
232 (<https://cran.r-project.org/package=ff>) facilities, respectively. This makes eddy4R seamlessly
233 scalable from local laptop development to deployment across massively parallel computing
234 facilities. Lastly, its unique modularity permits straightforward adjustments (extensibility) and
235 versioning as science and/or hardware progresses.

236 2.2 Git distributed version control (DevOps: Verify & Package)

237 The eddy4R source code resides on a version-controlled Git repository on the hosting service
238 GitHub (<https://github.com/>). In general, a developer community uses a version control system
239 to manage and track different states of their works over time. GitHub provides distributed
240 version control and has become widely used by scientific research groups because it is free,
241 open-source, and provides several features that make it useful for managing artifacts of
242 scientific research (Ram, 2013).

243



244

245 Figure 3. NEON's Git workflow. Please see text in Sect. 2.2 for detailed explanation.

246

247 Git allows multiple users and developers to simultaneously access and collaborate on a remote
248 repository by means of independent 'forks' or replicas of the entire repository (Paarsch and
249 Golyaev, 2016). Figure 3 shows NEON's Git workflow: At any given time (1) the official,
250 stable eddy4R source code resides on NEON's GitHub repository. A user can install the eddy4R
251 packages directly from there, and (2) a developer can 'fork' or copy the repository and create
252 'branches' for modification. After (3) 'committing' or creating a new feature, the developer (4)
253 can propose the feature for inclusion in the official eddy4R source code by issuing a pull request
254 to (5) NEON's change control board. After (6) thorough review and all prior test cases
255 reproducing benchmark results (DevOps: Verify stage), the feature can be 'merged' or
256 integrated into the next release of (1) the official, stable eddy4R source code (DevOps: Package

257 stage). This cycle can be repeated to accommodate requests and future developments, resulting
258 in subsequent releases. Including a test case for new code is strongly encouraged to ensure
259 sustainability over time, but is not mandatory. The developers can periodically update their
260 'forks' from the remote repository, ensuring that they always work on basis of the most recent
261 eddy4R source code.

262 The ultimate advantages of Git are provenance, reproducibility and extensibility: Every copy
263 of the code repository includes the complete history of all changes and authorship that can be
264 viewed and searched by anyone (Ram, 2013). This allows developers to build from any stage
265 of the versioned project and makes it easy to collaborate as an integrated scientific community.
266 We note that the DevOps workflow is robust to the business viability of the particular tools
267 used for implementation. Git is simply one instance of a version control system, which could
268 be replaced with another similar tool should Git fail at some point in the future.

269 **2.3 Docker image build and deployment (DevOps: Release)**

270 Facilitating the DevOps: Release stage, Docker images ([https://www.docker.com/what-](https://www.docker.com/what-docker)
271 [docker](https://www.docker.com/what-docker)) wrap a piece of software in a complete filesystem that contains only the minimal
272 context an application needs to run: code, runtime, system libraries and tools. This guarantees
273 that it always performs the same, regardless of the compute environment it is deployed in (i.e.
274 ultimate reproducibility). Compared to the similar but more cumbersome virtual machine
275 approach, a Docker image is an order or magnitude smaller (eddy4R-Docker: 2 GB without
276 example data). Also, by running as native processes it bypasses the virtual machine overhead.
277 Docker is used by many organizations (e.g., National Center for Atmospheric Research,
278 National Snow and Ice Data Center, NSF Agave API), and widely supported across large-scale
279 cloud compute environments (e.g., Amazon EC2 Container Service, Google Container
280 Engine, NSF Xsede). It is particularly well suited to NEON's DevOps strategy: combining
281 development, operation and quality assurance to enable creating, testing, deploying and
282 updating scientific software rapidly and reliably (Figure 2).

283 Docker can build images automatically by reading the instructions from a Dockerfile. A
284 Dockerfile is a text document that contains all the instructions a user would call on the command
285 line to assemble an image. Using e.g. a cloud hosting platform like DockerHub
286 (<https://hub.docker.com/>), the image build, versioning and distribution can be automated. This
287 is realized through executing the series of command-line instructions defined in the Dockerfile
288 whenever a new eddy4R source code version is available on GitHub. A key feature of eddy4R-
289 Docker is that it builds upon "Rocker" pre-built Docker images, maintained by the rOpenSci
290 group (<https://ropensci.org/>). This ensures access to stable, up-to-date base images containing
291 R and a variety of packages commonly used. The eddy4R-Docker image (0.2.0) released in this
292 study was built based on the rocker/ropensci/latest image containing R (3.4.0;
293 <https://hub.docker.com/r/rocker/ropensci/builds/>). As specified in the eddy4R Dockerfile, our
294 R packages eddy4R.base (0.2.0) and eddy4R.qaqc (0.2.0) and their dependencies were
295 automatically built on top of this base image. To complete the eddy4R-Docker processing,
296 analysis and modeling environment, the NEON data portal API Client nneo (0.1.0) as well as

297 the REdDyProc (1.0.0) high-level utilities for aggregated EC data were also included. In
298 addition, the user can install any desired R packages to customize the environment.

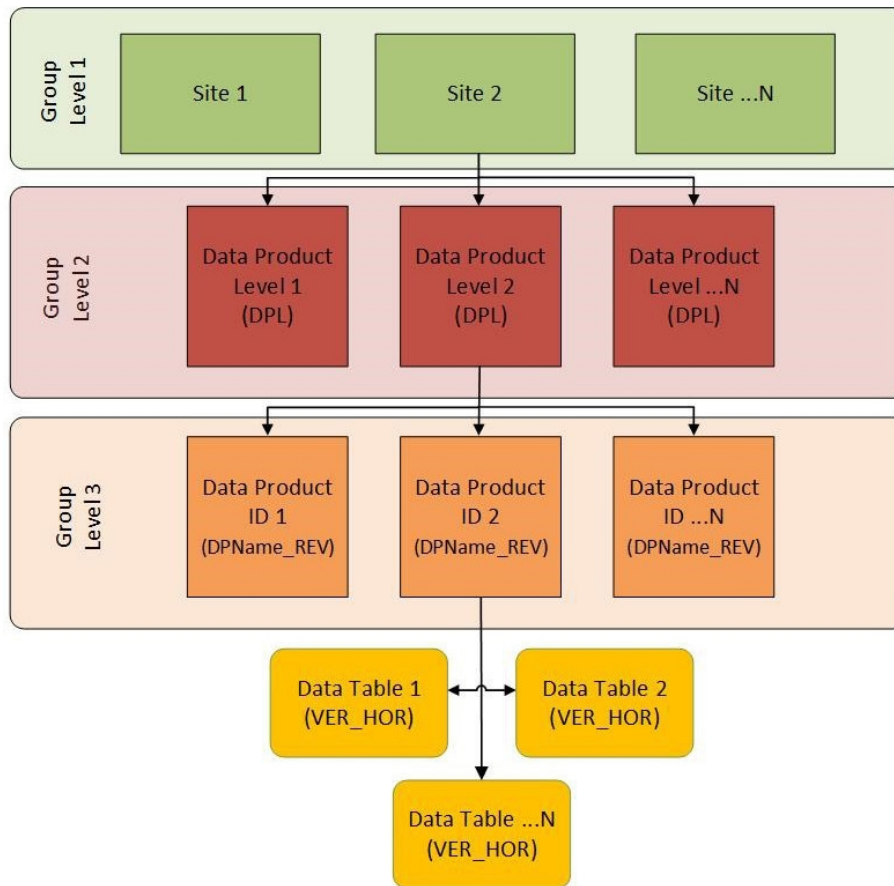
299 Docker’s benefits to scientific software development are described in detail in Boettiger (2015).
300 For NEON’s purposes, several Docker properties are particularly important:

- 301 • **Portability:** Docker images are portable and independent of the underlying operating
302 system. This enables scientists to develop code on local computers or virtual machines
303 without worrying about the deployment architecture.
- 304 • **Reproducibility:** The DevOps principles are ingrained into the Docker build process,
305 thus ensuring a fully traceable and documented Docker image.
- 306 • **Streamlined interface between Science and CI:** Defined inputs, outputs and
307 instructions provide an ideal framework to isolate and package algorithmic services for
308 operational deployment.
- 309 • **Continuous development and integration:** Docker provides a modular and extensible
310 framework, permitting NEON’s data processing to remain up-to-date with the latest
311 algorithmic developments. As shown by the nneo and REdDyProc examples, it enables
312 directly leveraging community-developed code. In this way eddy4R-Docker is
313 functionally extensible, while making it easy for the community to incorporate NEON-
314 developed code into their own data processing.

315 **2.4 Hierarchical Data Format version 5 (DevOps: Configure)**

316 The capability to process large data sets is reliant upon efficient input and output of data, data
317 compressibility to reduce compute resource loads, and the ability to easily package and access
318 metadata. The Hierarchical Data Format (HDF5) is a file format that can meet these needs, and
319 is a key tool aiding the DevOps: Configure (computational resource allocation) stage. A NEON
320 standard HDF5 file structure and metadata attributes allow users to explore larger data sets in
321 an intuitive “directory-like” structure that is based upon the NEON data product naming
322 convention (see Figure 4). Group level 1 separates data by site and site level metadata are
323 attributed at that level. Group level 2 separates data by data product level (DPL) and DPL
324 metadata are attributed at that level, where DPLs correspond to the amount of processing
325 performed. DPL1 are calibrated descriptive statistics, DPL2 are temporally interpolated, DPL3
326 are spatially interpolated, and DPL4 are further-derived quantities. Group level 3 are the
327 individual data products, for instance CO₂ concentration. Lastly, replicates in the horizontal and
328 vertical are separated as individual data tables.

329 This provides a streamlined data-delivery mechanism for the eddy4R-Docker processing
330 framework. For the tower datasets analyzed in this study, including sonic anemometer, infrared
331 gas analyzer and mass flow controller data, file sizes ranged from 1 GB for the uncompressed
332 data in comma-delimited ASCII files to 0.1 – 0.2 GB in HDF5 format, depending on the amount
333 of missing data. The HDF5 files can be written in a simple format where data are stored as
334 single 1-dimensional arrays to maximize compression and efficiency, or the data can be stored
335 as compound datatables that allow multiple datatypes to be written together in columnar format
336 for ease of navigation when data size is not an issue.



337

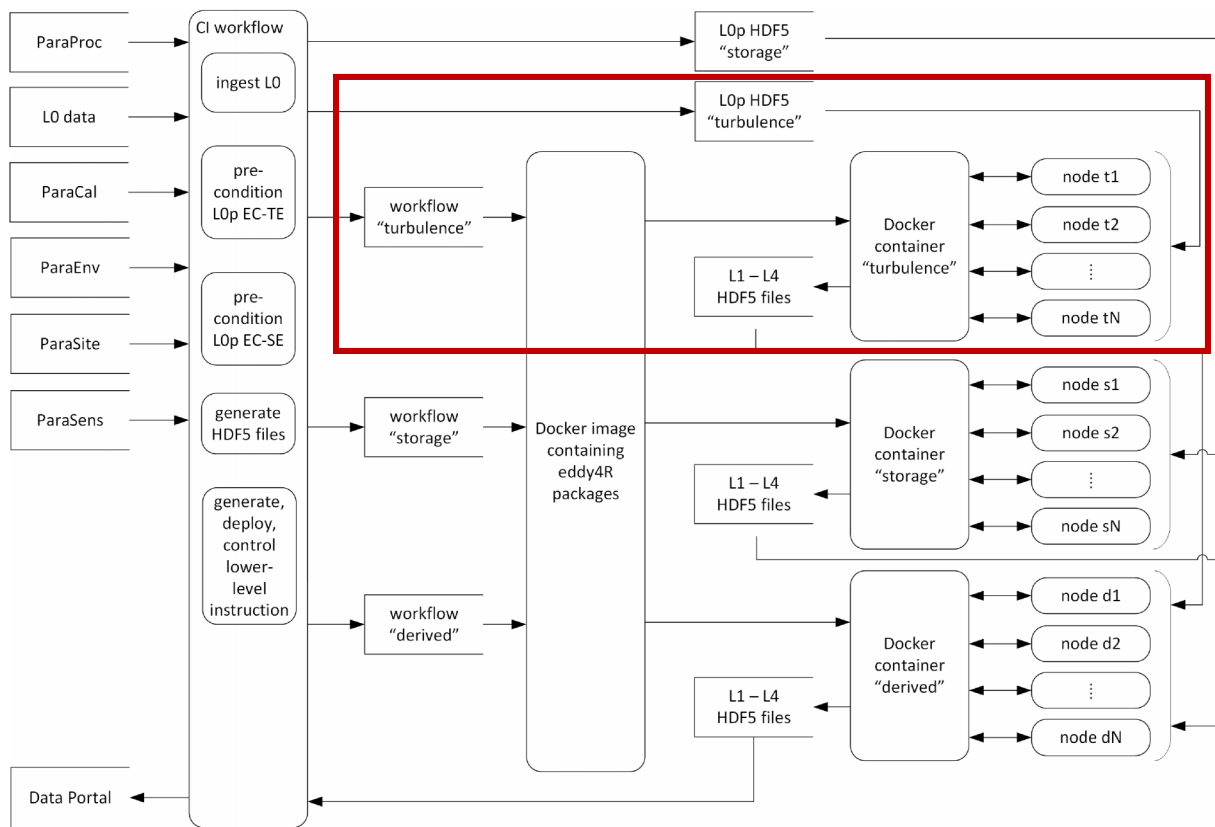
338 Figure 4 The NEON HDF5 file structure based on the NEON data product naming convention.

339

340 Another important function of the HDF5 file format is the ability to attach metadata as attributes,
 341 further promoting reproducibility. The data in this study has the units and variable names as
 342 metadata attached to the data tables in the HDF5 file. Additional metadata are attributed to
 343 various hierarchical groups throughout the file, including environmental parameters, sensor
 344 metadata, and processing parameters. As a result, HDF5 and similar self-documenting
 345 hierarchical data formats are gaining traction in a community that has traditionally relied on
 346 ASCII text column or comma-delimited files, especially as tools for viewing, manipulating, and
 347 extracting data from HDF5 become more commonplace. The utility of HDF5 file format is
 348 demonstrated in the executable example workflow that accompanies this manuscript (see
 349 Sects. 2.6, 5).

350 **2.5 Modular compatibility with existing compute infrastructure (DevOps:** 351 **Configure & Monitor)**

352 To perform a defined series of processing steps, a Docker image is called with a workflow file,
 353 resulting in a running instance called Docker container (Figure 5). Through this mechanism, an
 354 arbitrary number of Docker containers can be run simultaneously performing identical or
 355 different services depending on the workflow file. This provides an ideal framework for scaled
 356 deployment using e.g. high-throughput compute architectures, cloud-based services etc.



357

358 Figure 5. NEON’s eddy4R-Docker EC processing framework. The red box visualizes the scope
 359 of the present study, and individual components are described in the text.

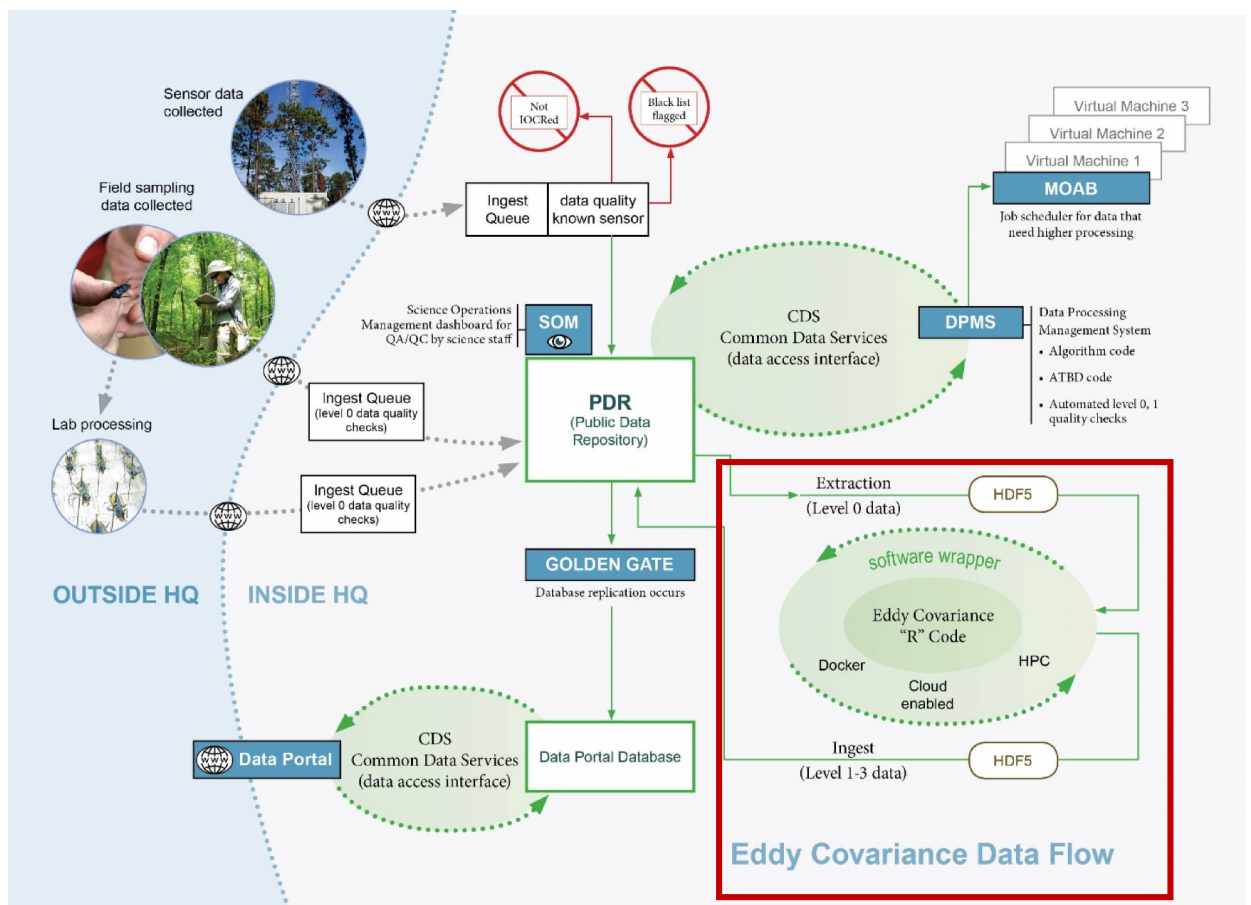
360

361 Embodying the DevOps: Configure stage, NEON’s eddy4R-Docker EC processing framework
 362 begins with ingesting information from various data sources on a site-by-site basis (Figure 5
 363 top left). This includes EC raw data (Level 0, or L0 data) alongside contextual information on
 364 measurement site (ParaSite), environment (ParaEnv), sensor (ParaSens), calibration (ParaCal),
 365 as well as processing parameters (ParaProc). Next, the raw data is preconditioned and all
 366 information is hierarchically combined into a compact and easily transferable HDF5 file
 367 (Figure 5 panel “CI workflow”). Each file contains the calibrated raw data (L0 prime, or L0p)
 368 and metadata for one site and one day, either for EC turbulent exchange or storage exchange.
 369 In this manuscript, we focus on demonstrating the turbulence data process and analysis in the
 370 red box of Figure 5. Together with the “turbulence” workflow file the HDF5 L0p data file is
 371 passed to the eddy4R-Docker image, where a running Docker container is spawned that scales
 372 the computation over a specified number of compute nodes (Figure 5 top right). The resulting
 373 higher-level data products (Level 1 – Level 4, or L1-L4) are collected from the compute nodes
 374 and, together with all contextual information, are combined into a daily L1-L4 HDF5 data file
 375 that is served on the data portal (Figure 5 bottom left). In addition to the daily output files,
 376 monthly concatenated files are also available for download from the NEON data portal
 377 (https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/portal/0.2.0). This sequence
 378 is performed analogously for different combinations of workflows and data, and it is possible
 379 for the workflow instruction sets to interact with each other. For example, the “turbulence” and

380 “storage” containers are processing in parallel, and starting the “derived” container once all
 381 intermediary results are available (Figure 5 bottom right). It should be noted that the
 382 “turbulence”, “storage” and “combined” Docker containers (Figure 5 right) are all spawned
 383 from the same eddy4R-Docker image (Figure 5 center): each container includes the same
 384 underlying functionality (eddy4R packages), but serves a different purpose by being fed the
 385 “turbulence”, “storage” or “combined” workflow files.

386 This eddy4R-Docker EC processing framework modularly integrates into pre-existing data
 387 processing pipelines, such as NEON’s CI (Figure 6): in NEON’s pre-existing framework the
 388 CI group encoded simple algorithms (e.g. temporal means) in Java, based on algorithm
 389 documentation provided by Science staff. The key difference of the eddy4R-Docker EC
 390 processing framework is that instead of algorithm documentation, NEON Science staff now
 391 provide documented algorithms that perform a complex series of processing steps, which can
 392 be directly deployed by CI. Not only does this adoption of the NEON-DevOps workflow
 393 (Figure 2) streamline end-to-end operational implementation and efficiency, it empowers the
 394 Science community at large by putting the key to the scientific algorithms into the hand of
 395 scientists.

396



397

398 Figure 6. NEON’s CI for streaming data processing. The red box visualizes the eddy4R-Docker
 399 EC processing framework within the overall CI.

400

401 To address the DevOps: Monitor stage, the computational resource load and performance
402 statistics of operating eddy4R-Docker can easily be monitored with standard profiling
403 procedures within NEON’s CI or other compute infrastructures. eddy4R-Docker further utilizes
404 the R logging package (0.7-103) to provide hierarchical logging, multiple handlers and
405 formattable log records. Finally, end-user experience is monitored via the Issues feature in
406 GitHub, where users can report code bugs, deployment problems, etc.

407 **2.6 Installation and operation**

408 One source of resistance to reproducible research is the initial burden of learning a new
409 workflow. The eddy4R-Docker image aims to reduce the initial setup effort and learning
410 requirements. This is achieved by providing a computational environment that contains all the
411 necessary software dependencies, the Rstudio graphical development environment
412 (<https://www.rstudio.com/>), and a code base consisting of example workflows and easily
413 accessible functions. Combined with a simple and thoroughly documented installation
414 procedure it provides a similar feel to working locally.

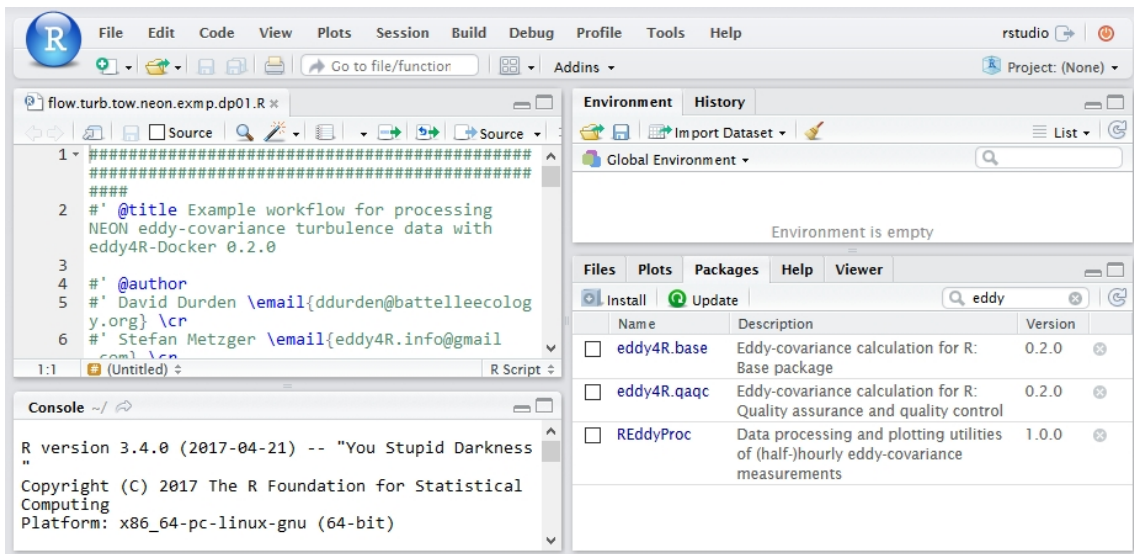
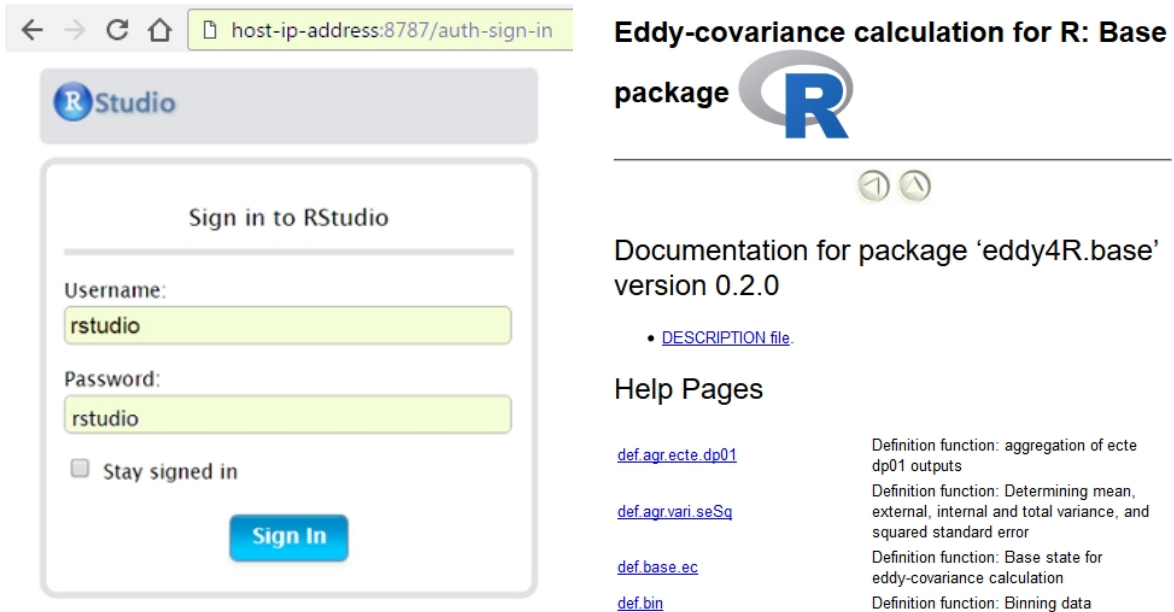
415 To work with the eddy4R-Docker image, one first needs to sign up at DockerHub
416 (<https://hub.docker.com/>) and install the Docker host software following the Docker installation
417 instructions (https://docs.docker.com/engine/getstarted/step_one/). Next, the download of the
418 eddy4R-Docker image and subsequent creation of a container can be performed by two simple
419 commands in an open shell (Linux/Mac) or the Docker Quickstart Terminal (Windows):

```
420     docker login  
421     docker run -d -p 8787:8787 stefanmet/eddy4r:0.2.0
```

422 The first command will prompt for the user’s DockerHub ID and password. The second
423 command will download the latest eddy4R-Docker image and start a Docker container that
424 utilizes port 8787 for establishing a graphical interface via web-browser. The release version of
425 the Docker image can be specified, or alternatively the specifier `latest` provides the most up-
426 to-date development image. In addition, it is possible to download and run a specific digest
427 using the `docker run stefanmet/eddy4r@sha256:` command. If data is not directed from/to
428 cloud hosting, a physical file system location on the host computer (`local/dir`) can be
429 mounted to a file system location inside the Docker container (`docker/dir`). This is achieved
430 with the `docker run` option `-v local/dir:docker/dir`.

431 The interactive Rstudio Server session running inside the Docker container can then be accessed
432 via a web browser at `http://host-ip-address:8787`, using the IP address of the Docker host
433 machine. The IP address of the Docker host can be determined by typing `localhost` in a shell
434 session (Linux/Mac) or by typing `docker-machine ip default` in `cmd.exe` (Windows).
435 Lastly, in the web browser the user can log into the RStudio session with username and
436 password `rstudio` (see Figure 7 top left panel). Figure 7 also shows the Rstudio integrated
437 development environment and interactive help for the `eddy4R.base` package in the bottom and
438 top left panels, respectively. Additional information about the use of Rstudio and eddy4R
439 packages in Docker containers can be found on the `rocker-org/rocker` website

440 (<https://github.com/rocker-org/rocker/wiki/Using-the-RStudio-image>) and the eddy4R Wiki
 441 pages.
 442



443 Figure 7 Docker-based Rstudio server session login via web-browser. Top left panel: Sign-in
 444 screen with highlighted areas showing information to input by the user. Top right panel:
 445 Interactive help for the eddy4R.base package. Bottom panel: Integrated development
 446 environment with workflow template, R console, Git staging area and eddy4R packages.
 447

448 To demonstrate the ease of “Docker-assisted” data analysis and provide a template for potential
 449 eddy4R-Docker users, an executable example workflow and data are included in the eddy4R-
 450 Docker image. Once the eddy4R container is started, the example workflow, input data (NEON
 451 dp0p HDF5 file) and output data (NEON dp01 HDF5 file) are available from the Docker-

452 internal directory `/home/eddy/`. The example workflow is located at
453 `/home/eddy/flowExmp/flow.turb.tow.neon.exmp.dp01.R`, and provides a selection of the
454 processing steps that yield the EC dp01 data on the NEON data portal
455 (https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/portal/0.2.0). The example
456 workflow is fully documented to guide readers through the various processing steps, and
457 employs key functionalities of the `eddy4R.base` and `eddy4R.qaqc` packages. These include data
458 and metadata import from the input HDF5 file, data assignment to file-backed objects,
459 processing of 1 minute and 30 minute data statistics and data quality, and writing the output
460 HDF5 file. In addition, outputs from the quality flag and quality metric model are visualized.

461 As described above, the eddy4R Docker image can be used for code development (DevOps:
462 Create stage) through accessing a running eddy4R Docker container via a web browser.
463 Alternatively, the eddy4R Docker image can be used from the command line to perform scaled
464 batch processing (DevOps: Configure & Monitor stages). Deployment from the command line
465 consists of passing the R workflow file to the Docker image. This is achieved by using the
466 `docker run` command with the additional argument `Rscript docker/dir/filename.R`, with
467 `filename.R` being the desired workflow. Thus, the eddy4R Docker image can be used to
468 simultaneously deploy multiple Docker containers to process data for multiple days or sites to
469 the capacity of the computational platform.

470 **3 Test applications**

471 In the following we present three test applications of eddy4R-Docker to evaluate whether the
472 NEON DevOps model can indeed produce collaborative, portable, reproducible, and extensible
473 EC software. Code development, packaging, release, and operation followed the NEON
474 DevOps model presented in this paper. Code modules have been contributed by order 10
475 individuals, distributed across multiple institutions and utilizing various computer systems.
476 Nevertheless, each contributor achieved identical results per validation scripts during the
477 DevOps: Verify & Package stage (Sect. 2.2), emphasizing the achieved portability and
478 reproducibility. The majority of the calculations presented here were performed on 12 Intel
479 Xeon X5550 2.67GHz CPUs, 32 GB memory with 10 Mbit interconnects and 10 Mbit access
480 to 8 TB storage on an Oracle Zettabyte File System. The software specifications were CentOS 7
481 (3.10.0-327.el7.x86_64) with docker-engine (1.11.0). In Sect. 3.1, results of processing 12 days
482 of EC data from a fixed tower at a NEON field site are shown. Next, in Sect. 0, we present the
483 processing of EC fluxes from a 1-hour recording of a moving platform: airborne observations
484 in a convectively mixed boundary layer. Lastly, a validation via software intercomparison is
485 provided in Sect. **Error! Reference source not found.**

486 **3.1 Tower eddy-covariance measurements**

487 Here, we use tower EC measurements to test a typical implementation of the eddy4R processing
488 framework. The Smithsonian Environmental Research Center (SERC) in Edgewater, MD, USA
489 is located on the Rhode and West Rivers, and hosts the NEON SERC tower (38°53'24.29" N,
490 76°33'36.04" W; 30 m a.s.l.). The ecosystem at SERC is a closed-canopy hardwood deciduous

491 forest dominated by tulip popular, oak and ash, with a mean canopy height of approximately
492 38 m (Figure 8). EC turbulent flux sensors are mounted at the tower top at 62 m above ground
493 or 24 m above the forest canopy.

494 An enclosed infrared gas analyzer (IRGA, LI-COR Biosciences, Lincoln, NE, USA, model: LI-
495 7200, firmware version 7.3.1.) was used to measure the turbulent fluctuations of H₂O and CO₂.
496 A mass flow controller (Alicat Scientific, Burlington, VT, USA, model: MCRW-20 SLPM-DS-
497 NEON) was used to maintain a constant flow rate of 12 SLPM through the IRGA cell. A sonic
498 anemometer (Campbell Scientific, Logan, UT, USA, model: CSAT3, firmware version 3) was
499 used to measure the 3-dimensional turbulent wind components. Data from the IRGA and the
500 sonic anemometer was synchronized using triggering and network timing protocol, and
501 collected simultaneously at 20 Hz sampling rate.

502



503 Figure 8. Left panel: Ecosystem at the NEON SERC tower (credit: Stephen Voss Photography;
504 <http://www.stephenvoss.com/blog/neon-tower-smithsonian>). Right panels: EC instrumentation
505 on top of the NEON SERC tower. Right top panel: Campbell Scientific CSAT-3 three-
506 dimensional sonic anemometer (front) and LI-COR Biosciences LI-7200 infrared gas analyser
507 (back) on the retracted tower-top boom. Right bottom panel: Same instrumentation but with the
508 tower-top boom extended at 230° from true north.

509 Here, data from April 22 to May 3, 2016 were used. The mean temperature during this time
510 period was 15°C, with a maximum temperature of 29°C and a minimum of 8°C. A total of
511 15 mm of precipitation was observed at nearby Annapolis Naval Academy.

512 **3.1.1 Algorithm settings and profiling**

513 The eddy4R workflow file was configured to ingest on the order of 50 data streams at 20 Hz,
514 including 3-D wind components, sonic temperature, and H₂O and CO₂ concentrations. The data
515 were processed to half-hourly L1 data products and turbulent fluxes. The L1 data products are
516 essentially state variables (wind, temperature, concentrations) with basic statistical products
517 derived, i.e. mean, minimum, maximum, standard error of the mean and variance. The
518 algorithmic processing for the L4 flux calculations requires additional scientific and procedural
519 complexity to test assumptions of the EC theory. The resultant fluxes represent half-hourly
520 vertical turbulent exchanges between the earth's surface and the atmosphere corresponding to
521 these state variables.

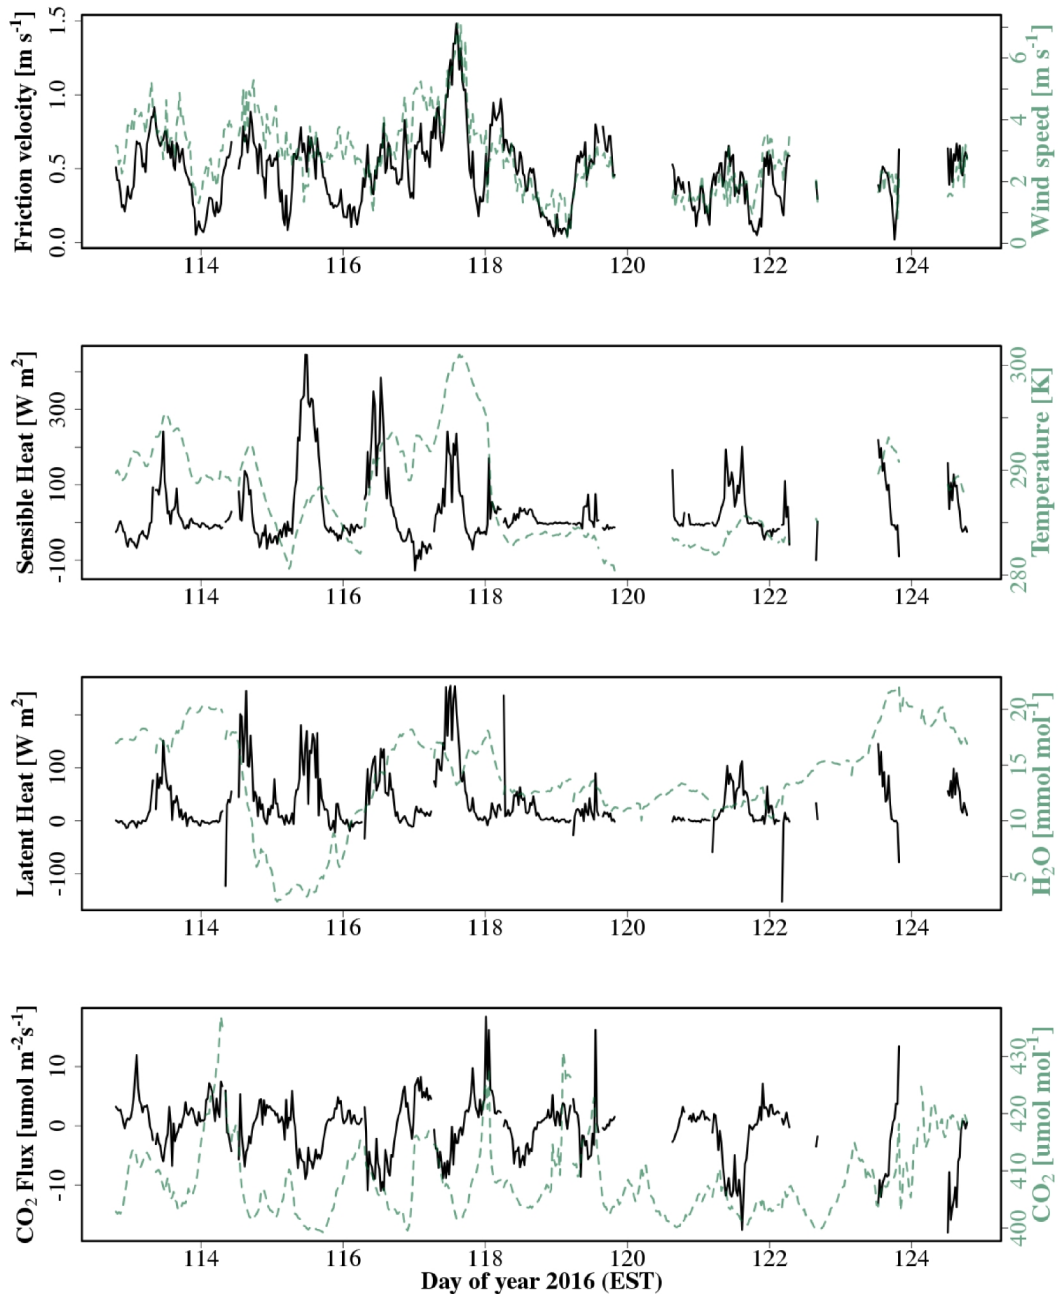
522 For the datasets analyzed in this study, the L0p input file sizes ranged from 0.1 – 0.2 GB in
523 HDF5 format depending on the amount of missing data, with metadata attached as attributes.
524 We used the simple data format for our HDF5 files, as opposed to compound data type, this
525 resulted in reduced read in time from 60 seconds to 3 seconds for 20 Hz IRGA data. Elementary
526 testing indicates that in this framework 6 CPU-minutes were required to process 1 day of 20 Hz
527 L0 data, and 1.2 CPU-minutes per 1 day of L0p data (100,000,000 observations). This
528 difference arises mainly from application of plausibility tests per Taylor and Loescher (2013)
529 in the transition from L0 to L0p. No reduction in efficiency was observed between direct
530 software deployment and its Docker implementation. Once flux QA/QC and uncertainty budget
531 is implemented, the computational expense will likely increase by a factor of two to three. This
532 suggests that eddy4R performs comparably to other flux processors. Memory usage is kept
533 below 2 GB through the use of fast access file-backed objects, enabling more sophisticated
534 scientific analyses through access to multiple days of data without overloading random access
535 memory (RAM) resources. Additionally, the snowfall R package allows for logical
536 parallelization frameworks to be implemented in the processing framework, even at low-level
537 analysis steps.

538 **3.1.2 Results and discussion**

539 The time series ranging from April 22 to May 3, 2016 was processed to deliver both state (L1)
540 and flux (L4) quantities; however, the initial eddy4R package release will only contain
541 functions necessary to report state variables or L1 data products in the NEON data product
542 description. During the processing of the proof-of-concept results averaging periods with >10%
543 missing data (incl. bad sensor diagnostic flags) were removed, and dedicated flux QA/QC and
544 uncertainty quantification were disabled.

545 Figure 9 shows the resultant time series of shear stress (friction velocity), sensible heat, latent
546 heat and CO₂ flux. The derived values fall into typical ranges for mid-latitude hardwood forests
547 in spring. As expected, fluxes follow the general trends in the scalar quantities. Good data

548 coverage can be seen for the LI-7200 measurements even during the rainy period at the end of
549 the analysis. A footprint analysis revealed that 90% of the flux measurement signals were
550 sourced within 800 m from the tower, and 80% were within 500 m from the tower at our site.
551 Data coverage was reduced after day of year (DOY) 120 due to inclement weather conditions.
552



553
554 Figure 9. Time-series of turbulent fluxes derived from EC measurements atop the NEON SERC
555 tower. Top to bottom: Vertical turbulent exchange of shear (friction velocity) and wind speed,
556 sensible heat and temperature, latent heat and H₂O dry mole fraction and CO₂ flux and CO₂ dry
557 mole fraction.

558 The spiky results preceding and following periods with >10% invalid data highlight the need
559 for enabling the full flux QA/QC and uncertainty budget to subset science-grade fluxes. This
560 implementation of eddy4R in a Docker image, as it will interact with NEON CI, clearly
561 demonstrates the applicability of the DevOps model for generating EC L1-L4 data products.

562 **3.2 Aircraft eddy-covariance measurements**

563 Here, we use aircraft EC measurements to test more advanced scientific capabilities of the
564 eddy4R processing framework. Airborne turbulent flux observations were performed along
565 more than 3100 km of low level (i.e. 50 m above ground level) flights across the North Slope
566 of Alaska, USA in July 2012, using the research aircraft Polar 5 (Tetzlaff et al., 2015). The
567 example data used in this manuscript were recorded during a SSW-NNE flight line near the
568 village of Atqasuk, Alaska, above tundra dominated by sedges and emerging herbaceous
569 wetland vegetation. Large, often oriented lakes and the meandering Meade River characterize
570 the surrounding landscape.

571 The aircraft was equipped with a 3 m nose boom holding a 5-hole probe for wind measurements,
572 an open wire Pt100 in an unheated Rosemount housing for air temperature measurements, and
573 an HMT-330 (Vaisala, Helsinki, Finland) in a Rosemount housing for relative humidity.
574 Sample air was drawn from an inlet above the cabin at about 9.7 l s^{-1} , analysed in an RMT-200
575 cavity-ringdown trace gas sensor (Los Gatos Research Inc., Mountain View, California, USA)
576 and recorded at 20 Hz. Aircraft position, velocity and attitude was provided by several Global
577 Positioning Systems (NovAtel Inc., Calgary, Alberta, USA) and an Inertial Navigation System
578 (Laseref V, Honeywell International Inc., Morristown, New Jersey, USA). Height above ground
579 was determined by a radar altimeter (KRA 405B/ Honeywell International Inc., Morristown,
580 New Jersey, USA) and a laser altimeter (LD90/ RIEGL Laser Measurements Systems GmbH,
581 Horn, Austria). The input data used in this study included the pre-derived 3-D wind vector from
582 5-hole probe and aircraft position, velocity and attitude. After spike removal the sampling
583 frequency of the original data was reduced from 100 Hz to 20 Hz resolution using block
584 averaging. These steps were performed prior to import into eddy4R processing, but could
585 equally well be performed therein.


586 **3.2.1 Algorithm settings and profiling**

587 Here, aircraft-measured vertical wind speed and CH_4 dry mole fraction were analysed to
588 determine CH_4 emissions by means of a time-frequency-resolved version of the EC method
589 (Metzger et al., 2013). For this purpose a combination of settings were chosen in the eddy4R
590 workflow file that differ from Sect. 3.1: Initially the small (<1 MB) EC raw data file consisting
591 of 17 variables and 12,800 data points (or 42 km flight data) was read in ASCII Gzip format –
592 standard R capabilities for data ingest can be used to read data in various formats, frequencies
593 and units. Aircraft-measured vertical wind speed and CH_4 dry mole fraction were then
594 correlated using a Wavelet transform (Metzger et al., 2013). This process includes ranging and
595 de-spiking of unphysical raw data values (Mauder et al., 2013; Metzger et al., 2012), fast dry
596 mole fraction derivation (e.g., Burba et al., 2012) and spectroscopic correction (Tuzson et al.,
597 2010) of CH_4 trace gas observations, and high-frequency spectral correction (Ammann et al.,

598 2006) by means of applying a sigmoidal transfer function (Eugster and Senn, 1995) directly in
599 Wavelet space. This permits estimating turbulent fluxes with improved spatial discretization
600 and determining ~100 biophysically relevant surface properties in the flux footprint. The
601 analysis took 56 minutes with 8-fold parallelization and consumed <3 GB RAM thanks to the
602 use of fast access file-backed objects.

603 3.2.2 Results and discussion

604 The resulting Wavelet cross-scalogram (Figure 10) is integrated in frequency over transport
605 scales up to 20 km, and along the flight path over a 1000 m moving window with 100 m step
606 size, similar to the resolution of the land surface data. The result is an in-situ observed space-
607 series of the CH₄ surface-atmosphere exchange at 100 m spatial resolution. Analogously,
608 turbulence statistics characterizing shear stress and buoyancy are determined for characterizing
609 the atmospheric transport between the emitting land surface and the aircraft position.
610



611
612 Figure 10. Wavelet cross-scalogram of the CH₄ flux equivalent to a time (x-axis) frequency (y-
613 axis) resolved version of EC. For each combination of aircraft position and eddy size, blue and
614 red areas indicate transport toward and away from the surface, respectively.

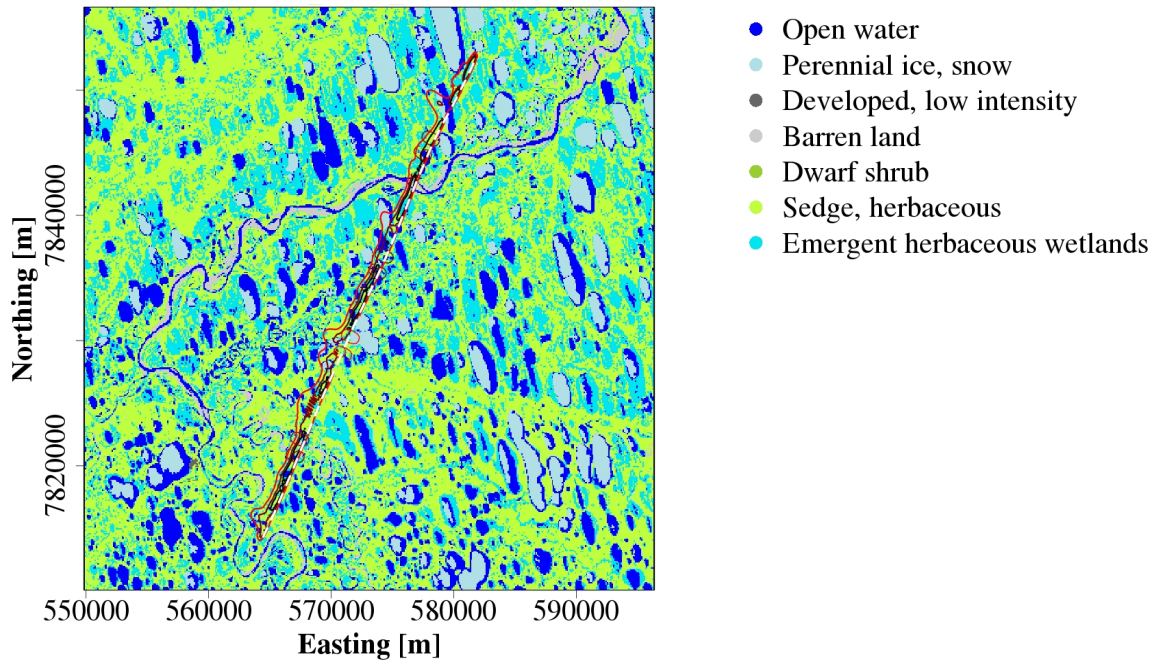
615
616 Corresponding systematic and random statistical errors are calculated following Lenschow and
617 Stankov (1986) and Lenschow et al. (1994), and the flux detection limit is calculated after
618 Billesbach (2011).

619 The relationship between the aircraft-observed CH₄ surface-atmosphere exchange and land
620 surface properties is established through an atmospheric transport operator, the so-called flux
621 footprint function (e.g., Schmid, 1994). Here we use a computationally efficient one-
622 dimensional parameterization of a Lagrangian particle model for the along-wind footprint
623 extent (Kljun et al., 2002; Kljun et al., 2004), combined with an analytical approach to
624 determine cross-wind surface contributions to each 100 m aircraft measurement, depending on
625 aircraft position (Figure 11; Metzger et al., 2012).

626 For each 100 m observation of the CH₄ surface-atmosphere exchange an individual footprint
627 weight matrix derived from the footprint parameterization is convolved with the land surface


628 drivers. The results are space-series of land surface contributions accompanying the CH₄
 629 measured surface-atmosphere exchange (Figure 12).

630



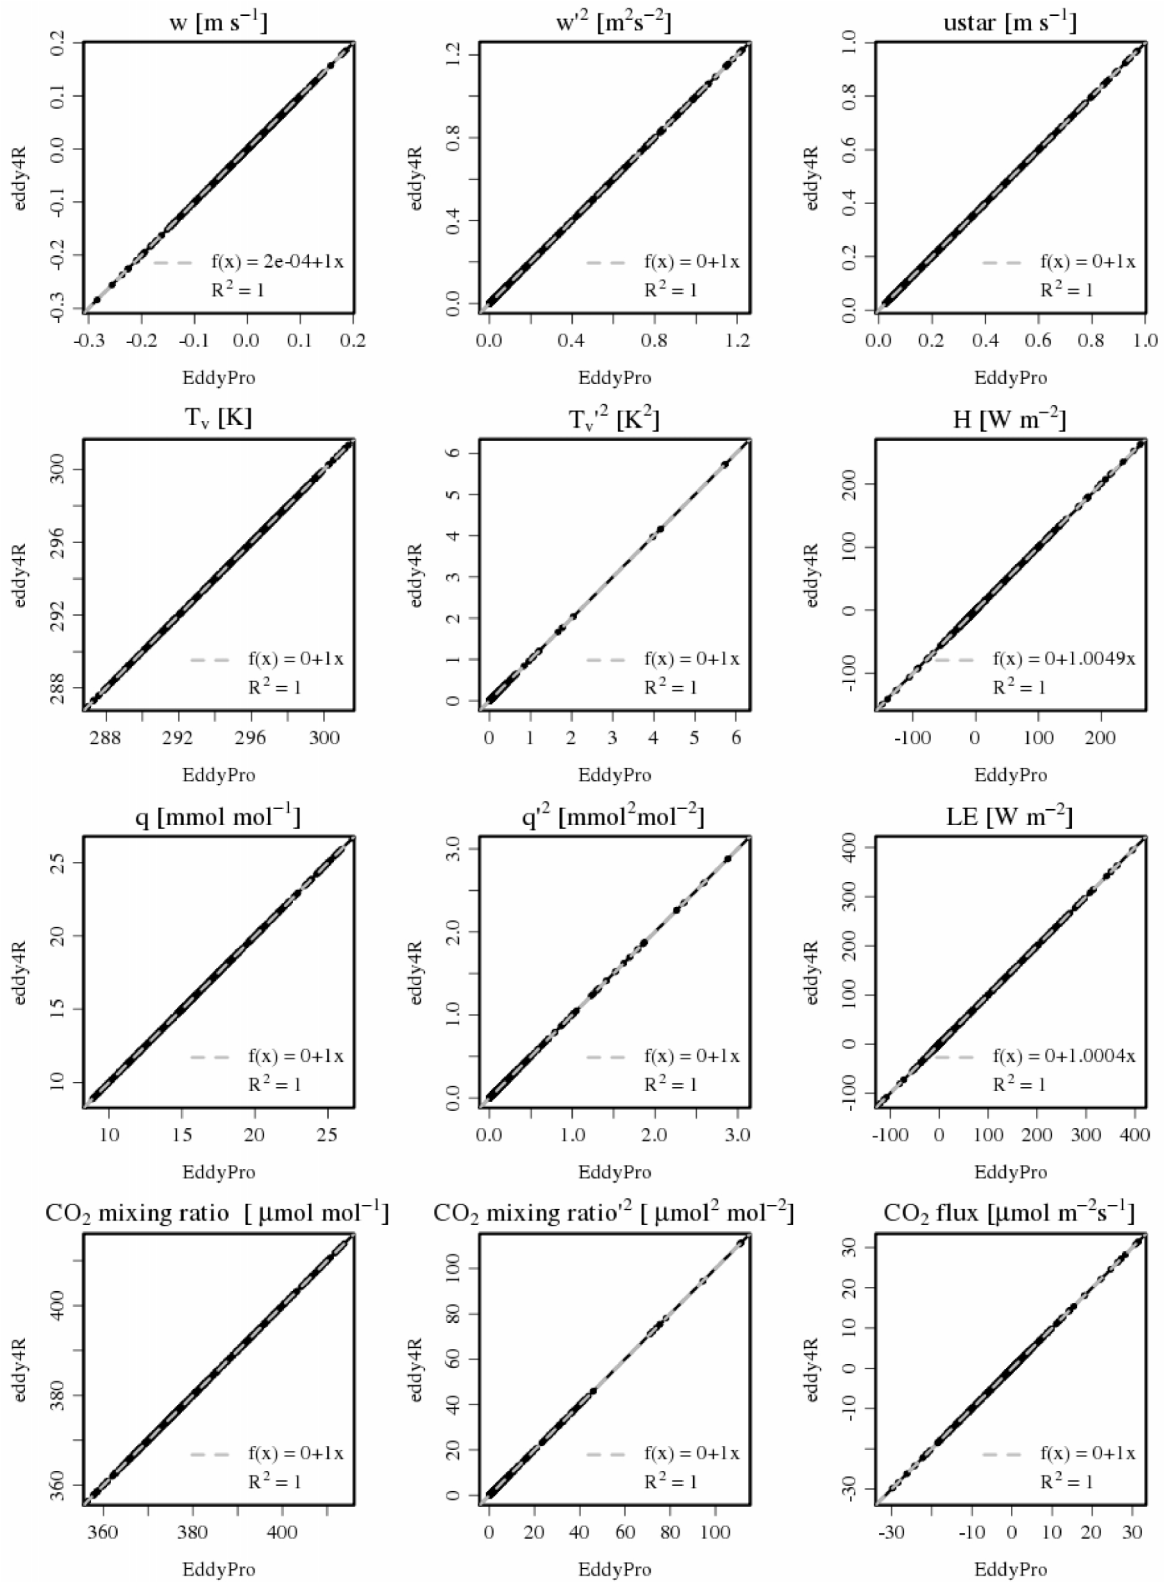
631 Figure 11. The composite flux footprint along the flight line (30 %, 60 %, 90% contour lines)
 632 superimposed over the National Land Cover Database. The white dashed line represents the
 633 aircraft flight track.

634



635

636 Figure 12. Space-series of 399 CH₄ concentration (purple line) and flux (black line)
 637 observations each 100 m, averaged over 1000 m windows. The random sampling errors are
 638 indicated by the shaded areas enveloping each line, and the flux detection limit is shown as
 639 salmon envelope around the abscissa. Circles indicate the dominating land cover in the footprint
 640 of each observation (Figure 11) with full circles corresponding to ‘pure’ fluxes (>80% surface
 641 contribution).



642

643 Figure 13. Scatterplot of means (vertical wind speed, w ; sonic temperature, T_v ; H₂O dry mole
 644 fraction (q); and CO₂ dry mole fraction), variances and fluxes (friction velocity, u_{star} ; sensible
 645 heat flux, H ; latent heat, LE ; CO₂ flux). Data are generated from 2011 July to Aug WLEF data
 646 in EddyPro and eddy4R. Each point represents a one-hour averaging period. Black lines are 1:1
 647 lines, and dashed lines are robust regressions (Salibian-Barrera and Yohai, 2006).

648 The successful application of eddy4R-Docker to both, basic tower and advanced aircraft EC
649 data analyses, highlights how the DevOps model promotes modular extensibility.

650 **3.3 Validation and verification**

651 eddy4R includes a verification script which automatically processes subsets of the tower and
652 aircraft data introduced in Sect. 3.1 and Sect. 3.2, and verifies the results against a reference,
653 e.g. generated with a different software.

654 Here, we demonstrate such approach at the Park Falls, Wisconsin WLEF very tall tower
655 Ameriflux site (US-PFa). The 447 m tall WLEF television tower (45.946°N, 90.272°W) has
656 been instrumented for EC measurements in 1996, and is part of the AmeriFlux network. Flux
657 measurements at 30 m, 122 m and 396 m sample a mixed landscape of forests and wetlands
658 (Desai et al., 2015). The surrounding forest canopy has approximately 70% deciduous and 30%
659 coniferous trees, and a mean canopy height of 20 m. The site has an interior continental climate.
660 Instrumentation at each level consists of fast response wind speed and temperature from a sonic
661 anemometer (Applied Technologies., Inc., Seattle, USA, ATI Type K). 10 Hz dry mole fraction
662 of CO₂ and H₂O at the 122 m level used here were measured by a closed-path infrared gas
663 analyzer (LI-COR, Inc., Lincoln, USA, LI-6262) located on the tower.

664 A data set from July 27 to August 19, 2011 was used in the intercomparison between eddy4R
665 and the reference software EddyPro (LI-COR, Inc., Lincoln, USA, version 6.2.0). EddyPro was
666 released in April 2011 and is widely used in the EC community.

667 **3.3.1 Algorithm settings**

668 Several preprocessing steps were applied, and the resulting data and settings were used in both,
669 eddy4R and EddyPro: (i) The raw data was pre-cleaned in eddy4R using the Brock (1986) de-
670 spiking algorithm with a filter width of 9 data points for all variables. (ii) EddyPro was used to
671 calculate the planar-fit rotation parameters (Wilczak et al., 2001) over the entire dataset (offset
672 = -0.06 ms^{-1} , pitch = -5.27° , roll = -1.81°). (iii) Time lags for dry mole fractions of CO₂ (0.8 s
673 behind vertical wind) and H₂O (0.1 s behind vertical wind) were calculated in eddy4R using
674 maximum correlation (median lag time over entire dataset).

675 Because CO₂ and H₂O fluxes were calculated from dry mole fractions, the Webb et al. (1980)
676 density correction was not necessary and therefore not applied (Burba et al., 2012). Frequency
677 response correction was not considered in this validation and therefore not applied. Means,
678 variances and fluxes were calculated on the basis of one-hour block averages. Based on
679 Schotanus et al. (1983), sensible heat flux was calculated from point-by-point conversion of
680 sonic temperature in eddy4R, and with the half-hourly statistical correction in EddyPro.

681 **3.3.2 Results and discussion**

682 eddy4R and EddyPro produce nearly identical results (Figure 13), and the gain error is within
683 0.04% for most outputs. Sensible heat flux values produced by eddy4R have slightly larger
684 magnitude compared to EddyPro, by 0.49%. This is likely a result of the different methods
685 applied when converting sonic temperature to air temperature. This intercomparison confirms
686 that applying the DevOps model to scientific EC software achieved results comparable to

687 commercial-grade software. A detailed end-to-end intercomparison considering additional
688 processing steps and EC software is planned for a separate manuscript accompanying NEON's
689 release of flux data products.

690 **4 Summary and conclusions**

691 Adopting a DevOps philosophy has facilitated the creation of a universal processing
692 environment for producing NEON's EC data products. Portable, reproducible, and extensible
693 software is reliably and efficiently created by incorporating the DevOps workflow steps of Plan,
694 Create, Verify, Package, Release, Configure, and Monitor into a NEON-specific DevOps model
695 based on the tools R, Git, HDF5, and Docker. Git-distributed version control facilitates
696 simultaneous internal-external collaboration on scientific algorithms, the outcome being a
697 modular family of open-source R packages. The use of Hierarchical Data Format allows for
698 efficient, self-describing data input and output. Docker images package the entire processing
699 environment for robust, scalable, and portable deployment. The capability of this framework
700 was demonstrated with cross-validated tower and aircraft fluxes.

701 The results presented here are from a file-based implementation of the eddy4R Docker
702 workflow, with EC instrument data accessed directly e.g. from the NEON site and manually
703 processed into the HDF5 ingest format (Sect. 2.5). The subsequent focus is the operational
704 implementation of the eddy4R-Docker workflow for reporting means and variances. This
705 includes: (i) Automated ingest of streaming raw data into the NEON database; (ii) Processing
706 of raw data into the standard, defined inputs required by the eddy4R-Docker in HDF5 format,
707 and (iii) Developing the software and hardware infrastructure to pass data and instructions back
708 and forth to the eddy4R-Docker workflow, and control program execution in a distributed
709 computing framework.

710 Remaining scientific algorithms are being integrated into eddy4R-Docker for producing
711 turbulent exchange data products. These algorithms include on-the-fly de-spiking, lag
712 correction, planar-fit and spectral correction, flux QA/QC, and uncertainty budget estimation.
713 Finally, eddy4R-Docker is being expanded to include "storage" and "derived" workflows
714 (Figure 6) for generating reproducible net ecosystem exchange data products in 2018. Lessons
715 learned here will profit the community at large, e.g. through enabling streaming processing
716 directly at an EC site or over cellular modems with the same eddy4R-Docker open-source
717 software as used for sophisticated analyses (Sect. 3.2). Already now, the executable example
718 workflow and data included in eddy4R-Docker image invite the reader to realize their own end-
719 to-end data analysis and apply it to their data (Sects. 2.6, 5).

720 While our sole focus in developing and implementing this model has been to generate EC data
721 products with the unique capabilities and constraints of NEON, it has become clear that the
722 NEON DevOps model enables the implementation of a suite of complex processing algorithms,
723 such as temporal gap filling of sensor time series data or modeling re-aeration rates. There exist
724 many potential synergies between NEON, other tower networks, and the user community for
725 producing high level EC data products. We hope this framework can serve as a model for
726 implementing community-sourced, distributed-development scientific code while combatting

727 the deficiencies of current computational frameworks that limit accessibility, reproducibility,
728 and extensibility.

729 **5 Code and data availability**

730 The source code packages eddy4R.base (0.2.0) and eddy4R.qaqc (0.2.0) used in this study are
731 archived at https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/code/0.2.0, under
732 the GNU Affero General Public License (GNU AGPLv3). Similarly, the corresponding
733 eddy4R-Docker image (0.2.0), including an executable example workflow and data, is available
734 at https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/docker/0.2.0. In addition, a
735 data supplement is provided at [https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-](https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/data/0.2.0)
736 [Docker/data/0.2.0](https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/data/0.2.0), including an extended abstract and all NEON SERC raw data used in this
737 study, accompanied by variable documentation. Lastly, NEON EC data products generated with
738 eddy4R-Docker are available at [https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-](https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/portal/0.2.0)
739 [Docker/portal/0.2.0](https://w3id.org/smetzger/Metzger-et-al_2017_eddy4R-Docker/portal/0.2.0).

740 **Acknowledgements**

741 Many colleagues at Battelle Ecology supported this study. In particular, Santiago Bonarrigo
742 provided pre-parsed high-frequency data from the SERC site, and Andrew Fox (now: National
743 Center for Atmospheric Research), Mike SanClements and David Hulslander commented on
744 an earlier version of the manuscript. Henry Loescher and Leslie Goldman designed Figure 6,
745 and Andrea Thorpe (now: Washington Natural Heritage Program), Thomas Gulbransen and
746 Michael Kuhlman helped shepherding this study and its publication through required
747 administrative procedures. Special thanks goes to Timothy Brown at the National Oceanic and
748 Atmospheric Administration, for numerous discussions and invaluable advice on scientific
749 computing. The National Ecological Observatory Network is a project sponsored by the
750 National Science Foundation and managed under cooperative agreement by Battelle Ecology,
751 Inc. This material is based upon work supported by the National Science Foundation under the
752 grant DBI-0752017. Any opinions, findings, and conclusions or recommendations expressed in
753 this material are those of the author(s) and do not necessarily reflect the views of the National
754 Science Foundation. Ankur Desai acknowledges support from NSF DBI-1457897 and DOE
755 Office of Science Ameriflux Network Management Project core site support to the ChEAS
756 cluster. Torsten Sachs and Andrei Serafimovich are supported by the Helmholtz Association of
757 German Research Centres through a Helmholtz Young Investigators Group grant to Torsten
758 Sachs (grant VH-NG-821).

759 **References**

760 Ammann, C., Brunner, A., Spirig, C., and Neftel, A.: Technical note: Water vapour
761 concentration and flux measurements with PTR-MS, *Atmos. Chem. Phys.*, 6, 4643-4651,
762 doi:10.5194/acp-6-4643-2006, 2006.

763 Aubinet, M., Vesala, T., and Papale, D., (Eds.): *Eddy covariance: A practical guide to*
764 *measurement and data analysis*, Springer, Dordrecht, Heidelberg, London, New York, 438 pp.,
765 2012.

766 Baldocchi, D., Falge, E., Gu, L., Olson, R., Hollinger, D., Running, S., Anthoni, P., Bernhofer,
767 C., Davis, K., Evans, R., Fuentes, J., Goldstein, A., Katul, G., Law, B., Lee, X., Malhi, Y.,
768 Meyers, T., Munger, W., Oechel, W., U, K., Pilegaard, K., Schmid, H., Valentini, R., Verma,
769 S., Vesala, T., Wilson, K., and Wofsy, S.: FLUXNET: A new tool to study the temporal and
770 spatial variability of ecosystem-scale carbon dioxide, water vapor, and energy flux densities,
771 *Bull. Am. Meteorol. Soc.*, 82, 2415-2434, doi:10.1175/1520-
772 0477(2001)082<2415:FANTTS>2.3.CO;2, 2001.

773 Billesbach, D. P.: Estimating uncertainties in individual eddy covariance flux measurements:
774 A comparison of methods and a proposed new method, *Agric. For. Meteorol.*, 151, 394-405,
775 doi:10.1016/j.agrformet.2010.12.001, 2011.

776 Boettiger, C.: An introduction to Docker for reproducible research, with examples from the R
777 environment, *Operating Systems Review*, 49, 71-79, doi:10.1145/2723872.2723882, 2015.

778 Brock, F. V.: A nonlinear filter to remove impulse noise from meteorological data, *J. Atmos.*
779 *Oceanic Technol.*, 3, 51-58, doi:10.1175/1520-0426(1986)003<0051:anftri>2.0.co;2, 1986.

780 Burba, G., Schmidt, A., Scott, R. L., Nakai, T., Kathilankal, J., Fratini, G., Hanson, C., Law,
781 B., McDermitt, D. K., Eckles, R., Furtaw, M., and Velgersdyk, M.: Calculating CO₂ and H₂O
782 eddy covariance fluxes from an enclosed gas analyzer using an instantaneous mixing ratio,
783 *Global Change Biol.*, 18, 385-399, doi:10.1111/j.1365-2486.2011.02536.x, 2012.

784 Chen, L.: Continuous delivery: Huge benefits, but challenges too, *IEEE Softw.*, 32, 50-54,
785 doi:10.1109/ms.2015.27, 2015.

786 Clark, D., Culich, A., Hamlin, B., and Lovett, R.: BCE: Berkeley's common scientific compute
787 environment for research and education, *Proceedings of the 13th Python in Science Conference*
788 *(SCIPY 2014) Austin, USA, 2014.*

789 Clement, R. J., Burba, G. G., Grelle, A., Anderson, D. J., and Moncrieff, J. B.: Improved trace
790 gas flux estimation through IRGA sampling optimization, *Agric. For. Meteorol.*, 149, 623-638,
791 doi:10.1016/j.agrformet.2008.10.008, 2009.

792 Collberg, C., Proebsting, T., Moraila, G., Shankaran, A., Shi, Z., and Warren, A. M.: Measuring
793 reproducibility in computer systems research, *University of Arizona, Department of Computer*
794 *Science, Tucson, USA, 37, 2014.*

795 De Roo, F., Abdul Huq, S. U., Metzger, S., Desai, A. R., Xu, K., and Mauder, M.: On the benefit
796 of driving large-eddy simulation with spatially resolved surface fluxes derived from
797 environmental response functions, *TERENO International Conference, Bonn, Germany, 29*
798 *September - 2 October, 2014.*

799 Desai, A. R., Xu, K., Tian, H., Weishampel, P., Thom, J., Baumann, D., Andrews, A. E., Cook,
800 B. D., King, J. Y., and Kolka, R.: Landscape-level terrestrial methane flux observed from a very
801 tall tower, *Agric. For. Meteorol.*, 201, 61-75, doi:10.1016/j.agrformet.2014.10.017, 2015.

802 Erich, F., Amrit, C., and Daneva, M.: A mapping study on cooperation between information
803 system development and operations, *15th International Conference on Product-Focused*
804 *Software Process Improvement, PROFES 2014, Helsinki, Finland, 2014.*

805 Eugster, W., and Senn, W.: A cospectral correction model for measurement of turbulent NO₂
806 flux, *Boundary Layer Meteorol.*, 74, 321-340, doi:10.1007/bf00712375, 1995.

807 Foken, T., and Wichura, B.: Tools for quality assessment of surface-based flux measurements,
808 *Agric. For. Meteorol.*, 78, 83-105, doi:10.1016/0168-1923(95)02248-1, 1996.

809 Foken, T.: *Micrometeorology*, 2 ed., Springer, Berlin, Heidelberg, 362 pp., 2017.

810 Fratini, G., and Mauder, M.: Towards a consistent eddy-covariance processing: An
811 intercomparison of EddyPro and TK3, *Atmos. Meas. Tech.*, 7, 2273-2281, doi:10.5194/amt-7-
812 2273-2014, 2014.

813 Kljun, N., Rotach, M. W., and Schmid, H. P.: A three-dimensional backward lagrangian
814 footprint model for a wide range of boundary-layer stratifications, *Boundary Layer Meteorol.*,
815 103, 205-226, doi:10.1023/A:1014556300021, 2002.

816 Kljun, N., Calanca, P., Rotach, M. W., and Schmid, H. P.: A simple parameterisation for flux
817 footprint predictions, *Boundary Layer Meteorol.*, 112, 503-523,
818 doi:10.1023/B:BOUN.0000030653.71031.96, 2004.

819 Kljun, N., Calanca, P., Rotach, M. W., and Schmid, H. P.: A simple two-dimensional
820 parameterisation for Flux Footprint Prediction (FFP), *Geosci. Model Dev.*, 8, 3695-3713,
821 doi:10.5194/gmd-8-3695-2015, 2015.

822 Kohnert, K., Serafimovich, A., Metzger, S., Hartman, J., and Sachs, T.: Geogenic sources
823 strongly contribute to the Mackenzie River Delta's methane emissions derived from airborne
824 flux data, 48th AGU annual Fall Meeting, San Francisco, U.S.A., 14 - 18 December, 2015.

825 Kormann, R., and Meixner, F. X.: An analytical footprint model for non-neutral stratification,
826 *Boundary Layer Meteorol.*, 99, 207-224, doi:10.1023/A:1018991015119, 2001.

827 Law, B.: AmeriFlux network aids global synthesis, *Eos, Transactions American Geophysical*
828 *Union*, 88, 286-286, doi:10.1029/2007eo280003, 2007.

829 Lee, J., Vaughan, A., Lewis, A., Shaw, M., Purvis, R., Carlslaw, D., Hewitt, C., Misztal, P.,
830 Metzger, S., Beevers, S., Goldstein, A., Karl, T., and Davison, D.: Spatially resolved emissions
831 of NO_x and VOCs and comparison to inventories, 48th AGU annual Fall Meeting, San Francisco,
832 U.S.A., 14 - 18 December, 2015.

833 Lenschow, D. H., and Stankov, B. B.: Length scales in the convective boundary layer, *Journal*
834 *Of The Atmospheric Sciences*, 43, 1198-1209, doi:10.1175/1520-
835 0469(1986)043<1198:LSITCB>2.0.CO;2, 1986.

836 Lenschow, D. H., Mann, J., and Kristensen, L.: How long is long enough when measuring
837 fluxes and other turbulence statistics?, *J. Atmos. Oceanic Technol.*, 11, 661-673,
838 doi:10.1175/1520-0426(1994)011<0661:HLILEW>2.0.CO;2, 1994.

839 Loukides, M.: *What is DevOps? Infrastructure as Code*, O'Reilly Media, Ebook, Safari Books
840 Online, 15 pp., 2012.

841 Mammarella, I., Peltola, O., Nordbo, A., Järvi, L., and Rannik, Ü.: Quantifying the uncertainty
842 of eddy covariance fluxes due to the use of different software packages and combinations of

843 processing steps in two contrasting ecosystems, *Atmos. Meas. Tech.*, 9, 4915-4933,
844 doi:10.5194/amt-9-4915-2016, 2016.

845 Mauder, M., Cuntz, M., Drüe, C., Graf, A., Rebmann, C., Schmid, H. P., Schmidt, M., and
846 Steinbrecher, R.: A strategy for quality and uncertainty assessment of long-term eddy-
847 covariance measurements, *Agric. For. Meteorol.*, 169, 122-135,
848 doi:10.1016/j.agrformet.2012.09.006, 2013.

849 Mauder, M., and Foken, T.: Eddy-covariance software TK3 [Data set]. Documentation and
850 instruction manual of the eddy-covariance software package TK3 (update). University of
851 Bayreuth, Bayreuth, Germany, doi:10.5281/zenodo.20349, 2015.

852 Metzger, S., Junkermann, W., Mauder, M., Beyrich, F., Butterbach-Bahl, K., Schmid, H. P.,
853 and Foken, T.: Eddy-covariance flux measurements with a weight-shift microlight aircraft,
854 *Atmos. Meas. Tech.*, 5, 1699-1717, doi:10.5194/amt-5-1699-2012, 2012.

855 Metzger, S., Junkermann, W., Mauder, M., Butterbach-Bahl, K., Trancón y Widemann, B.,
856 Neidl, F., Schäfer, K., Wieneke, S., Zheng, X. H., Schmid, H. P., and Foken, T.: Spatially
857 explicit regionalization of airborne flux measurements using environmental response functions,
858 *Biogeosciences*, 10, 2193-2217, doi:10.5194/bg-10-2193-2013, 2013.

859 Metzger, S., Burba, G., Burns, S. P., Blanken, P. D., Li, J., Luo, H., and Zulueta, R. C.:
860 Optimization of an enclosed gas analyzer sampling system for measuring eddy covariance
861 fluxes of H₂O and CO₂, *Atmos. Meas. Tech.*, 9, 1341-1359, doi:10.5194/amt-9-1341-2016,
862 2016.

863 Nordbo, A., and Katul, G.: A wavelet-based correction method for eddy-covariance high-
864 frequency losses in scalar concentration measurements, *Boundary Layer Meteorol.*, 146, 81-
865 102, doi:10.1007/s10546-012-9759-9, 2012.

866 Paarsch, H. J., and Golyaev, K.: A gentle introduction to effective computing in quantitative
867 research: What every research assistant should know, MIT Press, Cambridge, USA, 776 pp.,
868 2016.

869 Papale, D., Reichstein, M., Aubinet, M., Canfora, E., Bernhofer, C., Kutsch, W., Longdoz, B.,
870 Rambal, S., Valentini, R., Vesala, T., and Yakir, D.: Towards a standardized processing of Net
871 Ecosystem Exchange measured with eddy covariance technique: algorithms and uncertainty
872 estimation, *Biogeosciences*, 3, 571-583, doi:10.5194/bg-3-571-2006, 2006.

873 R Core Team: R: A language and environment for statistical computing, R Foundation for
874 Statistical Computing, Vienna, Austria, 2016.

875 Ram, K.: Git can facilitate greater reproducibility and increased transparency in science, *Source
876 Code Biol. Med.*, 8, 7, doi:10.1186/1751-0473-8-7, 2013.

877 Raupach, M. R., Rayner, P. J., Barrett, D. J., DeFries, R. S., Heimann, M., Ojima, D. S., Quegan,
878 S., and Schmullius, C. C.: Model-data synthesis in terrestrial carbon observation: Methods, data
879 requirements and data uncertainty specifications, *Global Change Biol.*, 11, 378-397,
880 doi:10.1111/j.1365-2486.2005.00917.x, 2005.

881 Running, S. W., Baldocchi, D. D., Turner, D. P., Gower, S. T., Bakwin, P. S., and Hibbard, K.
882 A.: A global terrestrial monitoring network integrating tower fluxes, flask sampling, ecosystem
883 modeling and EOS satellite data, *Remote Sens. Environ.*, 70, 108-127, doi:10.1016/S0034-
884 4257(99)00061-9, 1999.

885 Sachs, T., Serafimovich, A., Metzger, S., Kohnert, K., and Hartmann, J.: Low permafrost
886 methane emissions from arctic airborne flux measurements, 47th AGU annual Fall Meeting, San
887 Francisco, U.S.A., 15 - 19 December, 2014.

888 Salibian-Barrera, M., and Yohai, V. J.: A fast algorithm for S-regression estimates, *Journal Of*
889 *Computational And Graphical Statistics*, 15, 414-427, 2006.

890 Salmon, O., Caulton, D., Shepson, P., Brian, S., Metzger, S., and Musinsky, J.: Attributing
891 airborne measurements of forest CO₂ exchange to finer spatial scales, 5th NACP Principal
892 Investigators Meeting, Washington D.C., U.S.A., 26 - 29 January, 2015.

893 Schimel, D., Hargrove, W., Hoffman, F., and MacMahon, J.: NEON: a hierarchically designed
894 national ecological network, *Frontiers in Ecology and the Environment*, 5, 59-59,
895 doi:10.1890/1540-9295(2007)5[59:nahdne]2.0.co;2, 2007.

896 Schmid, H. P.: Source areas for scalars and scalar fluxes, *Boundary Layer Meteorol.*, 67, 293-
897 318, doi:10.1007/bf00713146, 1994.

898 Schotanus, P., Nieuwstadt, F. T. M., and Bruin, H. A. R.: Temperature measurement with a
899 sonic anemometer and its application to heat and moisture fluxes, *Boundary Layer Meteorol.*,
900 26, 81-93, doi:10.1007/BF00164332, 1983.

901 Serafimovich, A., Metzger, S., Kohnert, K., Hartmann, J., and Sachs, T.: The airborne
902 measurements of methane fluxes (AIRMETH) arctic campaign, 46th AGU annual Fall Meeting,
903 San Francisco, U.S.A., 9 - 13 December, 2013.

904 Smith, D. E., Metzger, S., and Taylor, J. R.: A transparent and transferable framework for
905 tracking quality information in large datasets, *PLoS One*, 9, e112249,
906 doi:10.1371/journal.pone.0112249, 2014.

907 Starkenburg, D., Metzger, S., Fochesatto, G. J., Alfieri, J. G., Gens, R., Prakash, A., and
908 Cristóbal, J.: Assessment of de-spiking methods for turbulence data in micrometeorology, *J.*
909 *Atmos. Oceanic Technol.*, doi:10.1175/jtech-d-15-0154.1, 2016.

910 Stull, R. B.: *An Introduction to Boundary Layer Meteorology*, Kluwer Academic Publishers,
911 Dordrecht, The Netherlands, 670 pp., 1988.

912 Sulkava, M., Luysaert, S., Zaehle, S., and Papale, D.: Assessing and improving the
913 representativeness of monitoring networks: The European flux tower network example, *J.*
914 *Geophys. Res.*, 116, G00J04, doi:10.1029/2010jg001562, 2011.

915 Taylor, J. R., and Loescher, H. L.: Automated quality control methods for sensor data: A novel
916 observatory approach, *Biogeosciences*, 10, 4957-4971, doi:10.5194/bg-10-4957-2013, 2013.

917 Tetzlaff, A., Lüpkes, C., and Hartmann, J.: Aircraft-based observations of atmospheric
918 boundary-layer modification over Arctic leads, *Q. J. R. Meteorolog. Soc.*, 141, 2839-2856,
919 doi:10.1002/qj.2568, 2015.

920 Turner, D. P., Ollinger, S. V., and Kimball, J. S.: Integrating remote sensing and ecosystem
921 process models for landscape- to regional-scale analysis of the carbon cycle, *BioScience*, 54,
922 573-584, doi:10.1641/0006-3568(2004)054[0573:irsaep]2.0.co;2, 2004.

923 Tuzson, B., Hiller, R. V., Zeyer, K., Eugster, W., Neftel, A., Ammann, C., and Emmenegger,
924 L.: Field intercomparison of two optical analyzers for CH₄ eddy covariance flux measurements,
925 *Atmos. Meas. Tech.*, 3, 1519-1531, doi:10.5194/amt-3-1519-2010, 2010.

926 Vaughan, A. R., Lee, J., Misztal, P., Metzger, S., Shaw, M. D., Lewis, A. C., Purvis, R., Carslaw,
927 D., Goldstein, A., Hewitt, C. N., Davison, B., Beevers, S. D., and Karl, T.: Spatially resolved
928 flux measurements of NO_x from London suggest significantly higher emissions than predicted
929 by inventories, *Faraday Discuss.*, doi:10.1039/c5fd00170f, 2015.

930 Vickers, D., and Mahrt, L.: Quality control and flux sampling problems for tower and aircraft
931 data, *J. Atmos. Oceanic Technol.*, 14, 512-526, doi:10.1175/1520-
932 0426(1997)014<0512:QCAFSP>2.0.CO;2, 1997.

933 Webb, E. K., Pearman, G. I., and Leuning, R.: Correction of flux measurements for density
934 effects due to heat and water vapour transfer, *Q. J. R. Meteorolog. Soc.*, 106, 85-100,
935 doi:10.1002/qj.49710644707, 1980.

936 Wilczak, J. M., Oncley, S. P., and Stage, S. A.: Sonic anemometer tilt correction algorithms,
937 *Boundary Layer Meteorol.*, 99, 127-150, doi:10.1023/A:1018966204465, 2001.

938 Wurster, L. F., Colville, R. J., and Duggan, J.: Market Trends: DevOps - not a market, but a
939 tool-centric philosophy that supports a continuous delivery value chain, Gartner, Inc.,
940 G00274555, Stamford, U.S.A., 14 pp., 2015.

941 Xu, K., Metzger, S., and Desai, A. R.: Upscaling tower-observed turbulent exchange at fine
942 spatio-temporal resolution using environmental response functions, *Agric. For. Meteorol.*, 232,
943 10-22, doi:10.1016/j.agrformet.2016.07.019, 2017.

944

945