Manuscript hess-2017-147 entitled “Derived Optimal Linear Combination Evapotranspiration (DOLCE): a global gridded synthesis ET estimate”

We would like to thank the reviewers for their constructive comments on our manuscript. This document outlines our point-by-point responses to the reviewers’ comments and the improvements made to the manuscript.

Response to Reviewer #1

General comments

1. **There is obviously a question regarding the representativeness of the FLUXNET stations to their regions (grid box average or dominant surface conditions), which the paper addresses fairly well. However, one specific issue jumps out at me: what about rare representatives, particularly in the tropics? FLUXNET is woefully thin on stations at low latitudes. The out-of-sample tests are like an OSSE, but there is little sampling in the tropics to play with here. What was the specific effect of denying low-latitude stations? Can we use this study to make case to prioritize FLUXNET expansion into the tropics, S/SW Asia and the Southern Hemisphere? I would like to see this issue discussed more, where the global maps are presented and also in the Discussion. Motivating targeted FLUXNET network expansion would be a good "broader impact" of this study.**

This is a great point. We have extended the discussion to try to address this:

There are relatively few towers located in the Southern Hemisphere and the tropics (14 out of 159 sites) and none located in the dry climates over South West Asia and North Africa. The weighting was therefore mostly driven by the ability of products to match sites located in the temperate and cold zones of the Northern Hemisphere, so that performance in climate zones with low FLUXNET site density was under-represented when deriving DOLCE. This might raise questions about the performance of DOLCE in the tropics and the Southern Hemisphere. To evaluate DOLCE in these areas we calculated the four site metrics separately for two groups of sites, 1) those located in the Northern Hemisphere excluding Tropics and 2) sites located in the Tropics and/or Southern Hemisphere. We excluded the two sites ID-Pag and AU-Fog in this exercise since both are wetland sites, and so would complicate a determination of whether these two groups had notable behavioural differences.

If systematic behavioural differences did exist between these two groups, we would expect relatively poorer performance of DOLCE at group2 sites compared to group1 sites. The results, shown in Fig. 10, appear inconclusive. DOLCE performed marginally worse at group2 sites overall, however with the limited number of sites in group2, the validation of the performance of DOLCE in the tropics and Southern Hemisphere remains somewhat uncertain. The uneven distribution of eddy-covariance sites between the Northern and Southern Hemisphere and across the different climates might also explain why much of the largest seasonal differences DOLCE-MPI and DOLCE-LANDFLUX-EVAL shown in Fig. 6 and Fig. 7 reside in the low latitudes (tropics) and the Southern Hemisphere and the persistent differences between DOLCE and LANDFLUX-EVAL in the tropics throughout the year. The expansion of the FLUXNET network into these areas that are lacking observations is clearly something that would improve DOLCE and LSM evaluation more broadly.

2. **What about the uncertainty/error in the FLUXNET observations? They certainly contain random errors and systematic biases. For a measured quantity that has a red-noise spectrum (following a Markov process) the random error can be estimated from the behavior of lagged**

autocorrelations (cf. http://dx.doi.org/10.1175/JHM-D-15-0196.1). Random error will systematically degrade correlation and affect other statistics (cf. http://dx.doi.org/10.1175/JHM-D-15-0063.1). Systematic error is more difficult to identify, but the energy balance corrections in FLUXNET2015 give some clue (see specific comments below). All these mean that using FLUXNET for cal/val is itself flawed and imperfect. On the other hand, model ET is inherently precise (no random error) and can be used to estimate/differentiate this type of error from others by noting its statistical differences from the instrument records. This issue should be acknowledged and discussed, including how the assumptions in DOLCE regarding errors (that they are uncorrelated) affect results. In other words, more discussion about uncertainties.

Yes. We have extended the discussion to address this point, at least to some degree:

Many studies have analysed the systematic and random errors of latent heat flux in FLUXNET measurements (Dirmeyer et al., 2016; Göckede et al., 2008; Richardson et al., 2006). These studies have detected errors of magnitudes that cannot be neglected. A recent study (Cheng et al., 2017) showed that the computed eddy-covariance fluxes have errors in the applied turbulence theory that lead to the underestimation of fluxes, and that this is likely to be one of the causes of the lack of surface energy closure. In this study, we 1) used the flag assigned to the observed flux, to filter out the low quality data and 2) used energy-balance-corrected FLUXNET data which has higher per-site mean values than the raw data at most of the sites (85% of them). We expect that filtering together with the use of corrected data will reduce the magnitude of the uncertainty in the observational data used here and compensate to a certain extent for the underestimation due to the systematic errors. However, we have not formally explored a range of approaches to addressing this. The possibility of systematic biases in FLUXNET data remains, and this could clearly lead to systematic biases in DOLCE.

We have also assumed that error across sites is uncorrelated, which, given the distribution of sites, is unlikely to be true, meaning that the effective number of sites is probably somewhat smaller than those shown in Figure 2. Given this dependence is likely to vary depending on a

range of time varying factors, we have left the job of attempting to disentangle this issue for future work.

3. **In the Supplement: There are captions for Fig S1 and Fig S2, but the PDF does not contain any figures!! Please recheck the rendering.**

   We thank the reviewer for spotting this, we now added the missing Figures in the Supplement.

Specific comments:

4. **P3 L6: Change "85 FLUXNET tower data" to "85 FLUXNET towers"**

   We've corrected this in the manuscript.

5. **P3 L8: "ground-truthed" is not an appropriate verbification for this context. Change '...gridded ET data sets are "ground-truthed" using flux tower data from FLUXNET...' to '...tower data from FLUXNET provide ground truth for gridded ET data sets...'**

   Thanks for picking this up, we've made the change.

6. **P5 L16: Change "where" to "were"**

   Thanks for picking this up, we've corrected this in the manuscript.

7. **P7 L22: RSD cannot convey whether the variance is too large or too small. So for the changes shown in Fig 3b - we cannot tell whether the improvement is due to an increase or decrease in standard deviation. Furthermore, what about locations where mean ET (denominator) is near zero - does that explain some of the very large values and changes?**

   The reviewer is right; as a result of this comment and a comment by Carlos Jimenez, we've replaced the relative standard deviation metric (RSD) with a modified relative standard deviation MRSD defined as $\frac{\sigma_{\text{dataset or observation}}}{max(mean(observation),\ q)}$. This addresses the two issues and removes the potential

for improvement in the RSD metric simply because the mean has improved. We have also added a fourth panel to this Figure showing improvement in the mean, for reference. We added the text below to explain the new metric:

We use a modified relative standard deviation metric MRSD that measures the variability of latent heat flux relative to the mean of the flux measured at each site. This ensures that a comparison between MRSD for a product and observations can tell us whether a product's variability is too large or too small (unlike relative standard deviation). The term 'q' is a threshold representing the 2nd percentile of the distribution of observed mean flux (i.e. temporal mean ET) across all sites (about 13 W/m2), which guarantees that MRSD calculated across many sites is not dominated by sites where the mean flux (denominator in MRSD Equation above) approaches zero. We looked at the bias in MRSD for each product considered- i.e. $|MRSD_{dataset} - MRSD_{observation}|$, and showed the performance improvement of the weighted mean.

8. **Fig 3: So the spread is across 5000, and each of those is an average of 25% out-of sample stations, right? It is not 5000x172x0.25 points. Please make clear.**

   The reviewer is right. We now clarified this in the figure caption:

9. **Also, Fig 4 suggests individual o-o-s stations frequently fare worse. Transferability of calibrations appears to be kind of weak, which is not a surprise. Calibration transferability is a very difficult enterprise. This should also be acknowledged, either here or in the Discussion section (wishing PILPS-San Pedro had been completed as it would have really shone a light on this problem).**

   Yes, this is indeed worth mentioning. We have modified the results section to read:

   In both cases it is important to note that many individual sites agree poorly with the weighted product compared to some other products. The distinction between the results shown in Fig. 3 versus Fig. 4 serve to highlight that DOLCE, and indeed any other large scale gridded ET product,

is not suitable for estimation of an individual site's fluxes, even if prediction over many sites shows notable improvement.

10. **P9 and Fig 4: "widespread out-of-sample improvement that this approach offers over existing gridded ET products" seems like a bit of an overstatement - there are frequently locations that fare worse, and sometimes much worse. In Fig 4, there is typically a tremendous range, and usually the central two quartiles encompass the zero line. For RMSE it appears that _20-49% of the time the estimates are worse than other products and methods, 30-55% for RSD, and 20-55% for COR. While there is definitely a net (and welcome) improvement in almost all cases, often the DOLCE estimate is worse. This should be acknowledged.**

    This is indeed a fair criticism. We have modified this to read

    "*On the basis of the aggregate out-of-sample improvement that this approach offers over existing gridded ET products*".

11. **P9 L26: "Standard Deviation (SD) difference" - this is clearly not the same as RSD defined earlier - please define, which minus which?**

    The reviewer is right, this is not the same as RSD. We've clarified this in the manuscript:

    Standard Deviation (SD) difference (i.e. $\sigma_{DOLCE} - \sigma_{observation}$)

12. **Table 2 has little accompanying discussion, and what is there is very shallow adding little to comprehension. Please discuss more or remove the Table if it does not warrant discussion.**

    We thank the reviewer for his suggestion, we removed this table from the manuscript.

13. **Fig 6: Here it could be said you use DOLCE to estimate MPI, and most places have a positive bias. But the energy-balance-corrected fluxnet data, which close the surface energy balance by construct conserving Bowen ratio, consistently increases ET compared to the raw measurements (at 107 of 122 FLUXNET2015 Tier-1 stations during JJA by my quick**

calculation). There is recent independent indication that tower sites bias low because of errors in the turbulence theory applied to estimate fluxes (see: http://dx.doi.org/10.1002/2017GL073499). Could FLUXNET instrument error (and the simplicity of the energy closure correction) contribute to systematic biases in DOLCE? Please discuss here or in the Discussion section.

Yes, of course it could. We have added to the discussion to make this clearer:

Many studies have analysed the systematic and random errors of latent heat flux in FLUXNET measurements (Dirmeyer et al., 2016; Göckede et al., 2008; Richardson et al., 2006). These studies have detected errors of magnitudes that cannot be neglected. A recent study (Cheng et al., 2017) showed that the computed eddy-covariance fluxes have errors in the applied turbulence theory that lead to the underestimation of fluxes, and that this is likely to be one of the causes of the lack of surface energy closure. In this study, we 1) used the flag assigned to the observed flux, to filter out the low quality data and 2) used energy-balance-corrected FLUXNET data which has higher per-site mean values than the raw data at most of the sites (85% of them). We expect that filtering together with the use of corrected data will reduce the magnitude of the uncertainty in the observational data used here and compensate to a certain extent for the underestimation due to the systematic errors. However, we have not formally explored a range of approaches to addressing this. The possibility of systematic biases in FLUXNET data remains, and this could clearly lead to systematic biases in DOLCE.

We have also assumed that error across sites is uncorrelated, which, given the distribution of sites, is unlikely to be true, meaning that the effective number of sites is probably somewhat smaller than those shown in Fig. 2. Given this dependence is likely to vary depending on a range of time varying factors, we have left the job of attempting to disentangle this issue for future work.

14. Re Fig 7: Much of the largest differences are at low latitudes where there is little FLUXNET data for calibration. Please discuss, as I see this as a major issue (more with FLUXNET station distribution than your methods, but the problem of representative ET estimates in the tropics is an ongoing concern).

We have addressed this issue in our response to the first point raised above.

15. **P10 L21: "ET doesn't exhibit any seasonal change over Greenland and the deserts in North Africa..." - not expected in the absolute, because mean values are tiny. It would be more informative to also show relative (percentage) changes, which are more relevant for local water balances.**

A good point. We have now added two extra plot in Fig.8 that show the seasonal variability of 1) ET estimates and 2) uncertainty estimates, we changed the plot titles, made the caption clearer:

and commented on the plots in the text:

The spatial distribution of DOLCE mean ET and its seasonal variability (standard deviation) over the austral Summer (Dec–Feb) and Winter (Jun–Aug) from 2000 to 2009 is shown in Fig. 8 (a) and (b) respectively. The seasonal variability of ET is larger in the warm season but is always small over Antarctica, Greenland and the deserts in North Africa (Sahara), the middle east (Arabian Peninsula desert) and Asia (i.e. Gobi, Takla Makan and Thar). The average uncertainty shown in if Fig. 8 (c) is bigger in the warm season, this is in agreement with the relatively large size of the flux in the warm season, and its seasonal variability shown in Fig. 8 (d) is also in agreement with the seasonal variability of the flux.

16. **Discussion §: Please also speculate whether some spatial variability in weighting could improve estimates further, even if only in 2 or 3 categories of weights.**

We agree this is worth exploring. That's why we performed clustered weighting where each cluster had a different set of weights. We accept this might not have been clear enough in the manuscript, we clarified this in the discussion:

In this study, we sought a single weight for each product to apply globally. But we have a reason to believe that different products are likely to perform better in different environments, so that different weights in different climatic circumstances might well improve the result of weighting

overall. A similar suggestion was made in the studies of (Ershadi et al., 2014) and (Michel et al., 2016) who highlighted the need to develop a composite model, where individual models are assigned weights based on their performance across particular biome types and climate zones. We therefore tried to cluster flux tower sites into groups (such as vegetation type) so that each group maintains enough members to allow the in- and out-of-sample testing approach used above. We tried clustering by vegetation type, climate zone and aridity index, and implemented the same one site out-of-sample testing approach as above, but this time, in each cluster different sets of weights will be assigned to the weighting products.

17. **Fig 9: Only stations in the tropics are 2 EBF stations; the savannah stations are in southern Africa. These are not o-o-s results, right? I would have expected these types to stand out more, but perhaps the samples are biased to the extratropical stations in the categories. Thoughts?**

We have added this point to the discussion:

The EBF box and whisker plot in Fig. 9 (d) shows the correlation of DOLCE at eight EBF sites, out of which two sites are located in the tropics. The lowest correlation seen in this biome type is at the tropical sites ID-Pag (0.26) and BR-Sa3 (0.62). This suggests that DOLCE tends to represent ET at the extratropical sites better than the tropics, and this is not surprising since most of the sites that were used to calibrate DOLCE were extratropical sites.

18. **Also Fig 9: Crops are tricky. The category is a catchall that is unsatisfactory because there is such variability in phenology, seasonality, stomatal resistances, etc. Are any of the CRO stations rice (which acts very different because of seasonal flooding). Not surprising many of the biggest errors are there.**

Good point. In the site description provided by FLUXNET the crop type is not always clear. There is only one rice site that we are sure about, but we excluded this site from the weighting because it was found in the list of irrigated sites. We excluded all irrigated sites. We've added the text below to explain the reasons:

In (6), we expect that some of the weighting models will largely underestimate the flux at irrigated sites, a result of a missing irrigation module in their scheme (Miralles, 2011; Jung, 2011). Because of this, the error bias of these models at the irrigated sites will modify the mean error bias (i.e. mean bias across all the sites) significantly, which will affect the weighting in favour of the products that can represent better irrigation. We excluded these sites as we do not want the products to be weighted for their inclusion/non-inclusion of physical processes.

We expanded our analysis as suggested by reviewer #2 point 10 to test the performance of DOLCE in three irrigated sites, we've added the text below to show this:

We tested the performance of DOLCE at three irrigated sites that were excluded from the weighting, for reasons explained earlier by computing the four statistics. A description of these sites and the results are shown in Table 2. The results show that the performance of DOLCE is reasonable at US-Ne1 and US-Ne2 and low at US-Twt. These results are discussed further below. DOLCE has also shown a weak performance at US-Twt, which is an irrigated rice paddy. This site gets flooded in spring and drains in early fall, then the rice is harvested. Only 9 months were available for this site, which coincide with the flood and drain period between spring and fall. DOLCE could not depict the flooding and draining event, probably because none of the weighting products can represent such phenomena, so it is expected that the effects of seasonal flooding are not represented in DOLCE.

19. **P13 L1: "...for example anthropogenic water management..." - but it was stated earlier that irrigated sites were excluded. Please expand on this comment.**

    We addressed this concern in the previous point.

20. **Conclusion: Please also state plans for updates, future versions, perhaps (hopefully) covering a longer time period!?**

    We thank the reviewer for his suggestion. We have added future improvement of DOLCE in the conclusion.

21. **References: There seem to be a lot of redundancies in author lists for papers with names showing up 2 or 3 times for the same entry. Please check.**

Thanks for picking this up, we have now removed the redundancies and made the appropriate corrections.

Response to Reviewer #2

1. **I have questions about why do you choose to produce DOLCE at 0.5 deg. resolution? Not 0.25 deg or 0.05 deg (the possibilities provided by MOD16, GLEAM ET), even as high as 1km (the possibility will discussed later in my report). I need a strong rebuttal from the authors here.**

   It would be fantastic to produce a meaningful 0.05 degree global ET product, but among all the weighting products, only MOD16 has information at this scale. It could be possible to produce an ET product at 0.25 degree, but both PML and MPI have no information at this scale. One of the benefits of DOLCE is that it combines the strengths of many estimates, one of the limitations is that it requires many estimates in order to provide better information. So, given the very small number of datasets available at high-resolution DOLCE would likely not provide better information.

2. **In addition, before producing DOLCE, are the weighting ETs resampled to 0.5 degree? Or how did you combine ET at different resolution into 0.5 degree DOLCE?**

   We thank the reviewer for their comment, this was indeed not mentioned in the manuscript. Yes, we resampled the 3 GLEAM products to 0.5 degree. We've included this information in the manuscript:

   The 3 GLEAM products were resampled to 0.5° using bilinear interpolation to match the spatial resolution of DOLCE

3. **My understanding is that ET product at either 0.5, or 1 deg., most of the flux site grid is in-homogeneity, thinking about the land surface covers can varies a lot in 50km*50km resolution. The more higher resolution of ET, the more higher chance could the site**

measurement represent information for a gird. Please check the reference: (Anderson et al., 2012).

The reviewer is of course correct, there is a huge mismatch between the fetch of the flux tower (i.e. 1km^2 or less) and the area of the grid cell containing the site, and a grid cell with an area close to the flux site fetch is more likely to be homogeneous than a 50km*50km grid cell. Unfortunately, as we responded to point 1 above, creating DOLCE at a finer resolution was not possible. We also dedicated a considerable portion of the manuscript to exploring the representativeness question (as noted by Reviewer 1), by testing whether the weighted DOLCE product performed well at sites that were not used to derive the weights (the results shown in Figures 3 and 4, and in particular the two columns of plots in Figure 4). We have also added text to make it clear that DOLCE will *not* perform better at every site, but only perform better across a wide range of sites considered collectively:

In both cases it is important to note that many individual sites agree poorly with DOLCE compared to some other products. The distinction between the results shown in Fig. 3 versus Fig. 4 serve to highlight that DOLCE, and indeed any other large scale gridded ET product, is not suitable for estimation of an individual site's fluxes, even if prediction over many sites shows notable improvement.

4. **DOLCE in this work does not make an extension for spatial resolution, or temporal resolution, even time coverage (10 years, 2000-2009). It is not encouraged for publication, which is similar to already published by Mueller et al. 2015. However, the method used to merge ET products is very useful. But the way of using it at 0.5 deg is not an optimal one.**

As discussed in point 1 above, with currently available ET products, it is not possible to produce DOLCE at a higher resolution, however it is always possible to expand the time coverage of DOLCE as its component products evolve, indeed this is the aim for future versions of DOLCE. We added a statement to this effect in the conclusion.

Expanding DOLCE over longer time periods and incorporating more diagnostic ET datasets (such as PT-JPL, CSIRO-global, GLEAM-V3B, SEBS, WECANN and HOLAPS) will be carried out in future versions.

5. **I would say whether the site is reported homogeneous in the evaluations is related to the scale. To asses site homogeneity by ET producer is not fare for the operation of these flux sites. I understand site PIs would seek to ensure the site represent one typical land surface. All the flux tower can be taken as homogeneous sites when the evaluated ET is at 10 meter resolution. Move to 0.5 deg grid, all the flux is located in an in-homogeneous grid. Thus you cannot say 'homogeneous sites', but 'homogeneous site grid'. Or most likely homogeneous due to good matching between 0.5 deg grid and flux site.**

This is a good point that we do need to clarify, here we use HOM and HET to describe whether the land cover attributed to the grid box (based on the IGBP map at 0.5 grid) matches the land cover at the flux tower. We clarified this in the manuscript and replaced HOM site and HET site with "HOM-case" and "HET-case" respectively.

6. **This also intrigue my interests, if you use the method to combine and calibrate existing ET to be a ET at 1km or higher resolution, which makes all available flux sites homogenous at this scale and can be used to calibrate weighting ET, more land covers can also be used not only DOLCE Tier 1, 2 and 3, say as IGBP 18 land covers. Why don't you use IGBP land covers to replace tier 1, 2 and 3 to derive weight for the 18 land covers? More classification or tiers can also help you produce DOLCE at higher resolution. You may say weighting ET cannot provide information at lower resolution.**

We thank the reviewer for their suggestion. The difference between Tier1, Tier2 and Tier3 is the size of the ensemble of ET products: Tier1 is derived from 6 ET products, Tier2 from 5 and Tier3 from 2 products. We produced Tier2 and Tier3 which have less products than Tier1 to overcome the limitation of the spatial coverage in the excluded products and ensure a global coverage of DOLCE (i.e. composite of Tier1, Tier2 and Tier3), as explained in Section 2.1.

The idea of producing a weight for each land cover is a great idea, but as explained in the manuscript we didn't have flux tower covering all biome types, and for some we had very few sites, not enough to calibrate the weights without over-fitting becoming an issue. We did try to group similar biome types so that each group maintained enough members to allow the in- and out-of-sample testing approach – this was outlined in the discussion with results presented in supplementary material. The results showed that clustering by biome type doesn't improve the weighting (figure S2).

7. **But the weighting ET can be calibrated with flux site at higher spatial resolution. Each weighting ET can be resampled and calibrated to e.g. 1km resolution. Then you can further combine them into a DOLCE high resolution ET.**

   This point is clearly very similar to the first point raised above. Calibrating low resolution products with site observations to produce high resolution product might seem to be a great idea, however the vast majority of the products weighted here do not have information at fine scales, so while there would be data at high resolution, there would be no more information than the low-resolution product. This might lead a user of the product to believe it provided more information than it actually does.

8. **This work will make the DOLCE more useful. The current 0.5 deg. ET does not differentiate from other fused/merged ET at coarse resolution, e.g. Landflux.**

   We thank the reviewer for their suggestion. The out-of-sample tests show that DOLCE @0.5degree) is performing better than Landflux (@1degree) in all the diagnostics we xamined.

9. **In addition, I also think about if your collection of flux tower data is not enough, which may lead to an biased weight for DOLCE calculation. One reviewer also pointed out the limitation of flux tower for the tropical region. Especially, the flux tower play an important role in your method.**

   We have addressed this point in detail in our response to point 1 raised by Reviewer 1.

10. You mentioned that 'irrigated sites' was excluded. What's the purpose of doing this?

We added the text below in the manuscript to clarify this point:

We expect that some of the weighting models will largely underestimate the flux at irrigated sites, a result of a missing irrigation module in their scheme (Jung et al., 2011; Miralles et al., 2011). Because of this, the error bias of these models at the irrigated sites will modify the mean error bias (i.e. mean bias across all the sites) significantly, which will affect the weighting in favour of the products that can represent better irrigation. We excluded these sites as we do not want the products to be weighted for their inclusion/non-inclusion of physical processes.

11. is it believed that all the weighting ET product cannot estimate ET for irrigated crops? If yes, this means DOLCE can also has a big errors for irrigated regions or human influenced regions. Then this need pointed out or at lease add discussion about shortage of this dataset. If no, I would be interested to look at the performance of weighting ET products and DOLCE at irrigated flux sites, since irrigated crops may also influence global water balance. We cannot blind to this issues when producing a global ET. Can't we? I agree if you remove irrigated sites or HT sites, this will makes your ET or paper looks better, but in reality it also expose the shortage of the method and products.

Testing how DOLCE performs at irrigated sites is a good idea. We expanded our analysis to evaluate how DOLCE performs at the three excluded irrigation sites. We displayed the results in Table 2 and added the text below in the Result section:

We tested the performance of DOLCE at three irrigated sites that were excluded from the weighting for reasons explained earlier (section 2.2), by computing the four statistics. A description of these sites and the results are shown in Table 2. The results show that the performance of DOLCE is reasonable at US-Ne1 and US-Ne2 and low at US-Twt. These results are discussed further below.

We also discussed the results by including this text in the Discussion section:

DOLCE has also shown a weak performance at US-Twt, which is an irrigated rice paddy, this site gets flooded in spring and drains in early fall, then the rice is harvested. Only 9 months were available for this site, which coincide with the flood and drain period between spring and fall. DOLCE couldn't depict the flooding and draining event, probably because none of the weighting products can represent such phenomena, so it is expected that representing seasonal flooding is a shortage in DOLCE.

More specific comments list below.

12. **There are a lot of errors in the Table 1. This has been pointed out by Carlos Jimenez.**

We thank the reviewer for spotting this, we have corrected these.

13. **In addition, I also found some ET not included. This is also a 0.05 deg. monthly ET. You may find here: http://en.tpedatabase.cn/portal/MetaDataInfo.jsp?MetaDataId=249454**

A great suggestion. Yes, SEBS is a remote sensing product and looks a good addition to the weighting ensemble. In the current version of DOLCE, we considered global ET products that cover the whole period 2000-2009, whereas SEBS misses a few months in 2000. In the future versions of DOLCE, this will not be an issue and we will consider adding SEBS to the weighting ensemble.

14. **Table3. Weighting ET products have negative or positive mean bias, if you give 0.5 weight to MPIBGC (3.837 mean bias) and GLEAM-v2B (-3.571) respectively, then DOLCE tier 1 will have a mean bias of 0.266 W/mˆ2. I also calculate mean bias with the weight and mean bias in table 3 for tier 1, then I get a mean bias of 0.9021 (0.041*3.756+0.495*3.837+(-0.026*6.180+....=0.9021), why do you think the weight provided in table 3 is better than my suggested weight of the two ETs?**

The bias represents the mean error of each product with respect to the observation. In the weighting technique, we bias correct the product first, then apply the weighting. This was perhaps not clear enough in the manuscript, so we've changed the table caption to read:

Table 3: (1) Bias of weighting products and (2) weights assigned to the bias corrected products in the case of each of the three DOLCE tiers, and the number of flux tower sites used to feed the weighting.

So, to answer this question (a) we are not optimizing for mean bias, and (b) we are testing out of sample to make sure the weighting approach works well.

15. Page 11 line28-30, please see my comments above

Thanks, this point has been discussed above.

16. Page 5, at what time resolution did you do the energy balance correction in the two equations or methods? Monthly or half-hourly? I would not expect 'no qualitative differ-ences' between bowen and residual term correction at half-hourly flux data. Secondly, most of the flux sites have no ground heat flux but soil heat flux at 5cm.

We added the text below in the manuscript to clarify this point:

We used daily averages of latent heat flux represented by "LE_CORR" in FN dataset. In LT dataset, we employed the components of energy imbalance and their associated flags (in brackets), represented by G_f (G_fqcOK), for soil heat flux, H_f (HFqcOK) for sensible heat flux, Rn_f (Rn_fqcOK) for surface net radiation and LE_f (LE_fqcOK) for latent heat flux.

17. What's your consideration of G used in the correction.

We applied quality control and filtering for G as highlighted in steps (2) and (3), section 2.2. We added the text below to clarify this point:

Applying a correction technique for energy imbalance at LaThuile sites required applying (2) and (3) for the other components of energy imbalance (i.e. R_n, G and H), which means that the sites that had to undergo a correction for the energy imbalance, should have monthly estimates for all the fluxes of the energy budgets, where each monthly value has been calculated from at least 15 daily mean flux values. *Because of this constraint, many sites were disregarded from the analysis.*

18. Fig.4 LFD or LDF*,LFA or LFA*?

The reviewer is right, we now changed LDF label to LFD and we modified the caption to clarify why we added * to the four products. The modified caption of Fig.4 reads:

*Figure 4: In (a), (b), (c) and (d), as for Fig. 3 but showing the one site out-of-sample test. Box and whisker plots are generated through selecting one site to be out sample and are repeated for all 138 sites. Products marked with * have limited spatiotemporal availability relative to the diagnostic ensemble, and testing against the LFA, LFD, CS and PT products was limited to 110, 108, 108 and 72 sites respectively. In (e), (f), (g) and (h), the one site out-of-sample test is trained by HOM-case sites data only..*

19. Figure 7, there is no values for the Sahel desert. This is due to non-values from DOLCE or LandFlux? Then I found DOLCE has ET estimate for Sahel desert in Fig. 8. Please explain it.

We created a mask from the spatial intersection of DOLCE, MPI and Landflux so that when we perform the comparison of DOLCE with MPI (fig. 6) and LandFlux (fig. 7) we look at the same areas. This eliminated the Sahel desert from the comparison, as it is not covered by MPI.

20. Fig. 8. How did you say reliability of uncertainty is low? DOLCE has uncertainty with monthly temporal resolution, am I right? I have difficulty understanding 'seasonal variability of global mean ET and its associated uncertainty'. It's better say 'spatial

distribution of a) global ET and (b) its associated uncertainty (standard deviation)in Winter and Summer,'.

We've added two extra plots that show the seasonal variability of 1) ET estimates and 2) uncertainty estimates, we changed the plot titles and rewrote the caption to read:

Figure 8: Seasonal (a) global mean ET and (b) its variability (standard deviation), (c) time average of uncertainty (the standard deviation uncertainty $\sqrt{\sigma_e^2}$ *shown in Equation 11) (d) standard deviation of uncertainty over time (e) reliability, defined as high (*$\frac{Uncertainty\ SD}{mean\ ET} \leq 1$ *in blue), medium(*$|mean\ ET| \leq 5$, $Uncertainty\ SD < 10$ *and* $\frac{Uncertainty\ SD}{mean\ ET} \geq 1$ *in green) and low (in red). DJF is shown in the left column and JJA in the right column.*

21. **Line 10. 'Together with the reasonable density .. are reasonably well constrained.' need rephrase.**

We have rewritten this to read:

These datasets together with in-situ surface observations have provided constraint on the reanalysis products that provide the basis of global gridded LSM forcing products.

22. **Page 1, Line 16, 'point-based estimates of flux towers provide information at the grid scale of these products.', are you sure point flux tower can provide information at 50*50 km pixel? Please check my comments above.**

As noted above in our response to point 3, we also think this is an interesting question that we have gone to some lengths to address – please see our response above for more detail. The fact that overall, the weighting is succeeding out-of-sample shows that the point-based measurements used to weight DOLCE do indeed provide useful information at the 0.5 degree grid cell scale. So yes, within the caveats noted in the revised manuscript, we are reasonably certain that this is the case.

23. Page 1, Line 21, These ET products differ in their data requirements, the approaches used to derive them and their estimates (Wang and Dickinson, 2012). This is well known. No need citation here.

We acknowledge that many in the community do know this, but are reasonably certain that others do not, and so have left the reference as is – there is no cost to this.

24. Line 24, we provide an even stronger vindication of this relationship. Which part show this? Please give some explanation.

In section 2.3 we showed that combining products will always derive a better performing product at the in-sample sites. We now refer the reader to Section 2.3.

25. Please use either 'Time-space step' or 'Space-time step';

We thank the reviewer for spotting this. We now made the necessary changes and chose to use "Time-space".

26. Acknowledgements Where did you use ERA-inerim by saying 'The ERA-Interim reanalysis data are provided by ECMWF and processed by LSCE'.

The reviewer is right, we haven't used ERA-Interim reanalysis data, but part of the term and conditions of the use of FLUXNET data is to use the exact statement included in http://fluxnet.fluxdata.org/data/data-policy/. We assume ERA-Interim is included in the statement because it was used for gap filling.

27. References: Annan, ...n/a-n/ Fisher, JB has been listed two times. Miralles, D.G, also listed two times. Please check the standard format for references used on HESS website

Thanks for picking this up, we have removed the duplicates and made the appropriate corrections.

28. Table s1, please also add RMSE, correlation and mean bias values for each site. This information is also important for DOLCE dataset users.

We obviously have all these values and they are all included in the plots (Figure 5 and Figure 9). The table is already dense but we are very happy to include them if the Editor thinks that fitting these into the table will not make the table unpublishable.

5

Response to Short Comments

Product Selection

While acknowledging the difficulties of finding the rights products to derive a synthesis product, I am a bit surprised about some of the choices.

As the authors state, product selection should follow the criteria of product diversity, so ideally single algorithms with different strengths and weaknesses should be combined together. In that sense, I would have not considered an already synthesis product such as LandFlux-Eval as a possible candidate for the merge.

In my opinion, a more valid alternative regarding LandFlux efforts could have been the three single products publicly available based on different ET algorithms (https://hydrology.kaust.edu.sa/Pages/GEWEX_Landflux.aspx). Also, I do not see much interest in combining obsolete versions of products with the current, and presumably better, product, as it has been done for GLEAM. Just GLEAM V3A (and perhaps GLEAM V3B) seems to me a better option. An interesting product is MPI. This is a global extrapolation of the tower fluxes, the same tower fluxes that are used to decide on the merging weights, and quite different to the other products, which we could consider more "physically" based and less "calibrated" with the tower data. It is not surprise then that MPI is by a large margin the more weighted product. I do not deny that it can be a valid product for the merge, although much less independent than the other products with respect to the tower data. In that sense, it could have been very informative also to see how the merging works, and how the weights are distributed, when that product is left out

Yes, this is indeed a reasonable question. One key distinction of the weighting approach here is that it accounts not only for the performance differences between products, but also the error covariance between them (as noted in Section 2.3). So, if a product were added that were a near copy of another, it would not degrade the performance of the weighting at all. While variants of

GLEAM are indeed likely to be similar, small structural differences might mean that there is in fact an advantage to using a nominally inferior version in addition to the latest version. This might explain why GLEAM-V2A was assigned a negative coefficient when the three GLEAM products participated in the weighting (tier1). By assigning a negative weight, the weighting was removing redundancies or data that was not adding any information to the weighting. It was not possible to include GLEAM v3B in the current version, due to the limitation in the covered time period which doesn't include 2000-2003.

Testing how the weighting would perform *without* MPI is an interesting idea. We therefore did exactly that - removed MPI from the weighting as a separate experiment and included this in the manuscript. Performance, perhaps surprisingly, was very similar:

it is not surprising that MPIBGC was the most weighted product since it is highly calibrated with flux tower data. In a further analysis, we left MPIBGC out and we performed the out of sample tests using the five remaining products. We wanted here to test how the weighting will perform without MPIBGC. The plots in Fig. 11 and 12 show the results of 25 % out-of-sample test and one site out-of-sample test respectively. Overall, without including MPIBGC, the weighting offers a smaller performance improvement than that offered when MPIBGC is a member of weighting ensemble (Fig. 3 and Fig. 4 (a-d)). The distribution of the weights when MPIBGC is absent from the weighting is 0.3 for both PML and GLEAM_v3A, 0.2 for GLEAM_v2B, 0.13 for MOD16 and 0.07 for GLEAM_v2A.

The new plots Fig.11 and 12 are shown below

Spatial and temporal resolution

The resolutions given in the Table 1 seem wrong for a number of products (e.g. MODIS original resolution is 1 km, Zhang 2010 is 0.05 deg, MPI and Zhang 2015 0.5 deg). Also, the

periods of available data should also be revised for GLEAM (I think GLEAM V2A covered 1980-2011 and V3A 1980-2014).

Thanks for spotting this, we have corrected the errors in Table 1

As far as I can see, the only datasets limiting the study period to 2000 is MOD16 and GLEAM V2B. Perhaps the products going into merging could have been separated in two groups, a bit similar to what has been done regarding geographical coverage: from a much earlier year than 2000 will all but MODIS and GLEAM V2B, and from 2000 including MODIS and GLEAM V2B. That would have resulted in a much longer DOLCE dataset, presumably based on a larger collection of tower data and more ET products. A shorter time period than monthly will result is a more useful product. Daily will a be a better objective for future developments, although it will require a different selection of products, possibly more based on the "physically" based diagnostic ET products, where daily is a common time scale. I would suspect a more complex merging exercise, given the larger amount of time variability that needs to be captured by the merge product ET and FluxNet datasets

Great comments and suggestions. In this work, limiting DOLCE to be derived from diagnostic (as opposed to overtly model-based) products only made the choices of temporal and special resolution limited. As suggested, extending the period of DOLCE by using different products for different time slices of DOLCE is a great idea, and will likely be an objective for further development in future versions of DOLCE. Applying the same merging technique on daily diagnostic ET products requires more complex analysis at both the site scale (observation) and grid scale (gridded products), although this is worth investigating for future versions. Thanks!

Tier 3 (i.e, Greenland and Antarctica) is just a very close weighting of an obsolete version of GLEAM and a newer version. It seems a bit awkward to distribute that as part of a synthesis product. It may have been better to just remove those regions from the synthesis product given that nearly no one dares to estimate ET over there

(understandably).

Yes, agreed. We however wanted to produce DOLCE with global coverage (requiring tier1, tier2 and tier3), and did make an effort to show the level of reliability of each tier in Fig. 8. We published tier1, tier2 and tier3 in separate files in order to give the users of DOLCE the flexibility of selecting the most appropriate product.

Some of the ET datasets considered are based on algorithms that estimate separately interception, evaporation from the canopy, and evaporation from the soil (e.g, MOD16, PT-JPL, and GLEAM). Under the assumption that routine EC observations perform very poorly for rainy conditions, in principle interception is not captured by the tower observations. In some recent ET evaluations of these products the tower data has been filtered to remove rainy periods and the interception component has not been evaluated. A discussion about this could have been interesting, given that, as far as I can see, the tower data is not filtered for precipitation conditions, and the merged product is a total ET product.

Yes, we haven't filtered the sites for rainy periods, especially since we're working with monthly data. Something to explore in the future, especially if we investigate using daily data.

Regarding the energy closure issue, the text may give the impression that fluxes correction is always possible, but a large number of stations do not measure Rn and/or G. Given that only corrected fluxes are used in the study, I imagine that a number of stations have to be discarded as they corrected fluxes were not available (FLUXNET-2015) or could not be estimated (LaThuille-2007). This may be worth mentioning

We only corrected LaThuile sites for energy imbalance at sites that had measurements of all the component fluxes. We applied quality control and filtering for G, H and Rn as highlighted in steps

(2) and (3), section 2.2. This is perhaps not clear enough in the manuscript, we therefore added the text below to clarify this point:

Applying a correction technique for energy imbalance at LaThuile sites required applying (2) and (3) for the other components of energy imbalance (i.e. R_n, G and H), which means that the sites that had to undergo a correction for the energy imbalance, should have monthly estimates for all the fluxes of the energy budgets, where each monthly value has been calculated from at least 15 daily mean flux values. Because of this constraint, many sites were disregarded from the analysis.

Merging technique

If I understand the method correctly, the weights are global (i.e. one value for all pixels), time-invariant (i.e., an annual value), and the bias-correction is what Bishop, 2013 calls a "global bias" correction (i.e., is a single value per product using all towers in the in-sample training dataset). If this is true, if Product A was performing better than product B over some biomes and/or at some periods, the method cannot be used to weight more or less the products to reflect that difference in performance. If I am correct, I wonder if there is a way to modify the weighting to take into account those differences. We typically see that ET products perform differently at different biomes and/or seasons, so it may be advantageous to capture this in the weighting. Bishop, 2013 was quite illustrative about this. Given that for that temperature example the "truth" was quasi-global (i.e., not over a very few pixels like for the ET), just a "per-cell bias" correction, even without weighting, outperformed the weighted product with a prior "global bias" correction. Of course, the "per-cell bias" correction and weights cannot be be applied here, given the limited geographical coverage of the flux "truth", so the problem is more complex.

I guess your first take at this is your sites clustering based on vegetation type, presented in some detail in the Supplement. You conclude that it did not improve overall the DOLCE performance based on a global analysis, but I would be interested to see the results (in the Supplement I can download the figures summarizing the analysis are missing). Perhaps

other schemes that better cluster the flux behavior are worth investigating for future versions of DOLCE

Yes, we calculated a single bias term, and assigned one weight to each ET product to be applied globally. As you also note, we did apply clustered weighting, where we derived cluster dependent sets of bias terms and weights for weighting the ET products. We tried clustering by 1) vegetation type, and also 2) climate zone and 3) aridity index (2 and 3 not shown in this study), and we implemented the same one site out-of-sample test, but this time separately in each cluster. We calculated different sets of weights for the ET products based on their performance differences in the different clusters, but none of the clustering succeeded in deriving a better weighted product overall. We added the text below to clarify the clustered weighting.

We show the results of one site out-of-sample test where different weights were assigned to products at each biome type (in supplementary material). We apologize for not putting the figure in the supplementary material, this has now been added.

In this study, we sought a single weight for each product to apply globally. But we have reason to believe that different products are likely to perform better in different environments, so that different weights in different climatic circumstances might well improve the result of weighting overall. A similar suggestion was made in the studies of Ershadi et al. (2014) and Michel et al. (2016) who highlighted the need to develop a composite model where individual models are assigned weights based on their performance across particular biome types and climate zones. We therefore tried to cluster flux tower sites into groups (such as vegetation type) so that each group maintains enough members to allow the in- and out-of-sample testing approach used above. We tried clustering by vegetation type, climate zone and aridity index, and implemented the same one site out-of-sample testing approach as above with different sets of weights in each cluster.

Results

It is stated that the MSE plot in Fig. 3a shows the MSE of the weighted product being better than the ensemble mean, but I do not see it. The central line of the first whisker box in Fig

3a is at the zero line. Perhaps I am missing something regarding how to read these plots. As a side note, it may be good to say what the end of the whiskers represents. In most occasions it is used to represent the max and min of the data, but it is not always the case.

We have modified both the plots and their description in the manuscript. As noted in our response to Reviewer 1, we have modified the RSD metric and added the performance of DOLCE on mean values as well in a fourth panel. We have also clarified the description of the box plots as follows:

We display the results of performance improvement datasets calculated in (a-c) above in 12 box and whisker plots. In each boxplot, the lower and upper hinges represent the first (Q_1) and third (Q_3) quartiles respectively of the performance improvement datasets, and the line located inside the boxplot represents the median value. The extreme of the lower whisker represents the minimum of 1) max(dataset) and 2) (Q_3 + IQR), while the extreme of the upper whisker is the maximum of 1) min(dataset) and 2) (Q_3+ IQR)), where IQR is the interquartile range of the performance improvement dataset. If the median performance improvement is positive, this indicates that the weighting offers an improvement in more than half of the data presented by the boxplot.

Note that (a-c) are explained in section 2.4

In the same Fig. 3a, the whisker boxes for the individual products make me think again about the MPI product. Based on the "heavy" calibration of MPI with the same tower data used to derive the weights, I would speculate that if the MPI product was removed from the merging, the percentage improvements of this new weighted product (i.e., without MPI) over the individual products will be much smaller. This may give a different perspective of the exercise regarding the skill of the tower-based merging to combine the more "tower-independent" physically based products.

Please see our response to a similar point above earlier in this document (in Product selection).

If we just concentrate on the improvements of the weighted product with respect to

the equally weighted product (i.e., first whisker box in Fig 3 a-b-c), the gain in performance of the weighted product seems small. Again, if I read these plots correctly, the gain for MSE and COR is minimal, only the RSD shows some improvement. But given the definition of the RSD metric, I wonder is this is mostly associated to the bias correction. If I understand this correctly, after the bias correction mean-dataset and mean-observation will be equal, so RSD is abs(sigma-dataset –sigma-observation). In other words, I am wondering about a comparison of the equally weighted product, but with a bias-correction first, and the weighted product. I am assuming here that the equally weighted mean did not involve a prior bias correction, as nothing was stated in the paper, but I may be wrong.

Yes, the plots in Fig 3 a,b and c show that overall, the improvement of the weighted product is noticeable with respect to the reference products and marginal with respect to the equally weighted mean. As a result of this comment and those by Reviewer 1, we've replaced the relative standard deviation metric (RSD) with a modified relative standard deviation MRSD defined as $\frac{\sigma_{\text{dataset or observation}}}{max(mean(observation),\ q)}$. This removes the potential for improvement in the RSD metric simply because the mean has improved. We have also added a fourth panel to this Figure showing improvement in the mean, for reference. We added the text below to explain the new metric:

We use a modified relative standard deviation metric MRSD that measures the variability of latent heat flux relative to the mean of the flux measured at each site. This ensures that a comparison between MRSD for a product and observations can tell us whether a product's variability is too large or too small (unlike relative standard deviation). The term 'q' is a threshold representing the 2nd percentile of the distribution of observed mean flux (i.e. temporal mean ET) across all sites (≈ 13 W/m2), which guarantees that MRSD calculated across many sites is not dominated by sites where the mean flux (denominator in MRSD Equation above) approaches zero. We looked at the bias in MRSD for each product considered- i.e. |MRSDdataset - MRSDobservation|, and showed the performance improvement of the weighted mean.

It is also true that the equally weighted mean used here doesn't involve any bias correction. Of course, some of the improvement offered by the weighting is owing to the bias correction and some comes from the weighting. We have now separated these effects in Figure S3 in the supplementary material, and referred to it in the results section:

Part of the success of the weighting approach relative to the multi-product mean is due to the bias correction applied before the weighting. Figure S3 in supplementary material separates the effect of each step.

In the HOM and HET comparison, I see very small improvements in MSE, larger for RSD, and not much for COR (the median of the whisker box is for MSE and COR is at the zero percentage line). And I wonder if the separation into HOM and HET sites may have also implied a separation in land covers, so the improvements we see are more related to the weighted product working better for some specific biomes. One may think that land covers such as forested areas are more likely to be represented in the HOM class, compared with e.g. croplands. I wonder if this has been checked, i.e., that the biome representation in HOM and HET classes does not change too much.

Yes, the weighting calibrated by HOM sites offers only a marginal improvement over the weighing using all the sites (HOM and HET). We looked at the separation of Biomes in HOM and HET cases and we found a clear separation of land covers and their distribution across the HOM and HET case:

1) cropland constitutes more than 20% of the HOM-case sites and 7.6% of the HET-case sites,

2) about 23% of HOM-case sites are forests (EBF, DBF and MF), while more than half of the HET-case sites are forests, the big number of forest sites in the HET case is because most of the MF sites are located on grid boxes identified as EBF and DBF.

3) WET and WSA sites are found only in the HET-case

We added table S2 below in the supplementary material and we modified the text.

These results nevertheless lead us to expect that if we construct DOLCE by incorporating HOM-case sites only, we might get a better product, but the small number of sites satisfying this property, the fact that the separation of sites into HOM-case and HET-case can lead to a separation of land covers, and the difficulty in defining a meaningful definition for expected flux homogeneity are limiting factors.

Regarding the boxplots of Figure 5, it is true that the end of the whiskers are larger for the HET sites, but the RMSE and correlation median and percentiles look slightly better for the HET class. I think this is also worth discussing, as it can potentially indicate again that the differences in performance between HOM and HET classes are small, with just a few HET sites having bad statistics. Given that this is based on one site out-of-sample, I wonder if the bad performance at some individual sites may be nothing to do with homogeneity, but with the fact that the site is one of a kind, so any weights derived from a sample without that site are not informative. Given the location of some the sites, it would not be a surprise.

An interesting point. It could certainly be the case that the reason for some sites in the HET group showing poor performance of DOLCE might simply be about measurement quality or site uniqueness. The box plots in Figure 5 show the performance of DOLCE in the HOM and HET cases separately. Yes, while the whiskers are larger for the HET sites but we cannot infer from these plots that DOLCE performs better in any of the two groups of sites. We modified the text in the Results and Discussion to make this point clear.

There is some expectation that the weighting will show better performance if it is trained with HOM-case sites only, although HOM-case sites consist of about one-thirds of the total number of sites used in this study.

And

Determining whether DOLCE performs better at HOM-case sites or HET-case sites is inconclusive. Even that the worst performance of DOLCE was achieved in HET-case sites, the boxplots in Fig. 5 (a), (d) show that the value of the median, lower and upper quartiles are better in HET-case for two metrics (i.e. RMSE and COR). While we expect that calibrating the weighting with HOM-case could lead to a better product, we don't expect to see DOLCE overall performing better in any of the two groups.

We might expect that (3) would decrease the overall performance of the HET-case sites (see Fig.3). On the other hand, it is very likely that(1) and (2) increase the performance of HET-case sites and eventually compensate the decrease of performance due the WET sites.

These results nevertheless lead us to expect that if we construct DOLCE by incorporating HOM-case sites only, we might get a better product, but the small number of sites satisfying this property, the fact that the separation of sites into HOM-case and HET-case can lead to a separation of land covers, and the difficulty in defining a meaningful definition for expected flux homogeneity are limiting factors. Determining whether DOLCE performs better at HOM-case sites or HET-case sites is inconclusive. Even that the worst performance of DOLCE was achieved in HET-case sites, the boxplots in Fig. 5 (a), (d) show the value of the median, lower and upper quartiles are better in HET-case for two metrics (i.e. RMSE and COR). While we expect that calibrating the weighting with HOM-case could lead to a better product, we don't expect to see DOLCE overall performing better in any of the two groups.

MPIBGC is the larger contributor to DOLCE, with a weight close to ~0.5. Perhaps it is not surprising than the differences of MPI with DOLCE sown in Figure 6 are smaller than with LandFlux-EVAL-Diag in Figure 7, which is not part of DOLCE.

I acknowledge that it is quite difficult to come out with the right classes to try to illustrate the reliability of the weighted product. But when I look at figure 8-c what I broadly see is that arid places, Greenland and Antarctica have low reliability, snowed places medium, and

the rest high. Perhaps just the uncertainty over the mean ET would have been more informative.

Great suggestion. We've expanded the plot and showed the seasonal variability of 1) ET estimates and 2) uncertainty estimates, and we've changed the plot titles and rewritten the caption to read:

Figure 8: Seasonal (a) global mean ET and (b) its variability (standard deviation), (c) time average of uncertainty (the standard deviation uncertainty $\sqrt{\boldsymbol{\sigma_e^2}}$ *shown in Equation 11) (d) standard deviation of uncertainty over time (e) reliability, defined as high (*$\frac{\boldsymbol{Uncertainty\ SD}}{\boldsymbol{mean\ ET}} \leq \mathbf{1}$ *in blue), medium(*$|\boldsymbol{mean\ ET}| \leq \mathbf{5}, \boldsymbol{Uncertainty\ SD} < \mathbf{10}$ *and* $\frac{\boldsymbol{Uncertainty\ SD}}{\boldsymbol{mean\ ET}} \geq \mathbf{1}$ *in green) and low (in red). DJF is shown in the left column and JJA in the right column.*

Discussion

The discussion about a possible selection of HOM sites to construct DOLCE is quite appropriate. Being very strict, and considering that the typical tower fetch is of the order of hundreds of meters, while the grid cell has an area of ~2500 km2, possibly none of the sites will truly qualify as HOM. But if we want to keep using tower fluxes, we need to live with this. Perhaps a simple measure to reduce the effect of this tower-fetch and cell-scale mismatch is to try to work at finer resolutions in the future. GLEAM, MOD16 and Zhang, 2010 are already at resolutions 0.25 deg. PT-JPL will be soon available at 5 km. Perhaps a couple of products will never be available at 0.25 deg (e.g., MPI), but they could be downscaled to 0.25 deg and merged with the other products (e.g., by a nearest-neighbor interpolation if we do not want to add any extra information). The lack of success of any clustering of the sites based on vegetation, climate, etc, possibly is more indicative of the limitations of the tower flux, rather than limitation of the conceptual idea. Even if not terribly successful, a look at how the relative weights of the different products change for different clusters can be informative regarding the products performance for different conditions.

Yes, both are good suggestions for future work. As noted in our response to Reviewer 2, when globally gridded ET products allow us to derive DOLCE at higher resolutions, there are obviously many benefits. More tower data will also obviously help, especially with this second point – understanding which gridded products are better suited to different conditions.

As I mentioned before, using MPI as part of the weighted product is perfectly valid. But I think the main feature defining MPI for this study is not the fact that it is a statistical product, but more that it is a global extrapolation of the same tower fluxes used to derive the weights. As clearly stated in the paper, the tower fluxes have limitations, which will have an impact on the weighted product. Presumably, the MPI product will suffer from the same limitations than the tower fluxes, but these limitations will not become apparent when looking at differences with the tower fluxes, so it will not be reflected in the weights. An example could be the ET estimation for those biomes and moments when interception can be a relevant component of the fluxes. If the tower fluxes are not properly capturing this component, the physical methods that try to do so (e.g., GLEAM, PT-JPL, MOD16) may be penalized in the weight derivation, compared with MPI, but their fluxes may be more correct. Of course, the problem is how to merge or validate the ET products away from the tower flux data. Given that DOLCE is a global product with a relatively long time series, at least the annually integrated ET could be compared with basin-integrated differences of precipitation and runoff, which may shed more light into the merits of the individual products, the equally weighted mean product, and DOLCE.

Yes. We hope that our addition of DOLCE without the MPI product to the manuscript (as described above), and the fact that it still performs well, goes some way to alleviating this concern. At the annual scale, we could indeed neglect the change in water storage, and compare evapotranspiration with the difference (Precipitation - Runoff). This difference could serve as a test dataset for DOLCE, the equally weighted mean and the reference ensemble, away from flux data, and is independent from both MPI and DOLCE. However, this validation exercise requires observational datasets for precipitation and runoff. While there is indeed a big network for

observed precipitation, gridded products can vary by surprising amounts, and currently there are no global time varying observational estimates of runoff that we are aware of. Year-to-year storage changes are also not necessarily negligible. Nevertheless, validating DOLCE as a component of the water balance is part of our next piece of research… so we clearly agree this is an aim worth pursuing.

Derived Optimal Linear Combination Evapotranspiration (DOLCE): a global gridded synthesis ET estimate

Sanaa Hobeichi[1,2], Gabriel Abramowitz[1,2], Jason Evans[1] and Anna Ukkola[2]

[1] Climate Change Research Centre, University of New South Wales, Sydney, NSW 2052, Australia
[2]ARC Centre of Excellence for Climate System Science, University of New South Wales, Sydney, NSW 2052, Australia

Correspondence to: Sanaa Hobeichi (s.hobeichi@student.unsw.edu.au)

Abstract. Accurate global gridded estimates of evapotranspiration (ET) are key to understanding water and energy budgets, as well as being required for model evaluation. Several gridded ET products have already been developed which differ in their data requirements, the approaches used to derive them and their estimates, yet it is not clear which provides the most reliable estimates. This paper presents a new global ET dataset and associated uncertainty with monthly temporal resolution for 2000–2009. Six existing gridded ET products are combined using a weighting approach trained by observational datasets from 159 FLUXNET sites. The weighting method is based on a technique that provides an analytically optimal linear combination of ET products compared to site data, and accounts for both the performance differences and error covariance between the participating ET products. We examine the performance of the weighting approach in several in-sample and out-of-sample tests that confirm that point-based estimates of flux towers provide information at the grid scale of these products. We also provide evidence that the weighted product performs better than its six constituent ET product members in four common metrics. Uncertainty in the ET estimate is derived by rescaling the spread of participating ET products so that their spread reflects the ability of the weighted mean estimate to match flux tower data. While issues in observational data and any common biases in participating ET datasets are limitations to the success of this approach, future datasets can easily be incorporated and enhance the derived product.

1 Introduction

Improving the accuracy and understanding of uncertainties in the spatial and temporal variations of Evapotranspiration (ET) globally is key to a number of endeavours in climate, hydrological and ecological research. Its estimation is critical to water resource, heatwave and ecosystem stress prediction and it provides constraint on the energy, water and carbon cycles. For these reasons, it is also useful for evaluating the performance of land surface models (LSMs). To identify LSM performance issues and diagnose their probable causes, the observed values of LSM inputs and outputs need to be known with sufficient accuracy and precision. At the site scale, FLUXNET (Baldocchi et al., 2001; Baldocchi, 2008), a global group of tower sites that measures the exchange of energy, water and carbon between the land surface and the atmosphere, provides direct observations of ET and most of the drivers and outputs of LSMs with soundly quantifiable uncertainty and sufficient accuracy to make diagnostic evaluation possible. This fact has led to a range of model evaluation, comparison and benchmarking studies using

FLUXNET data (e.g. Abramowitz, 2005, 2012; Best et al., 2015; Chen et al., 1997; Stöckli et al., 2008; Wang et al., 2007). Haughton et al. (2016) clearly showed that the flux tower data quality is good enough to provide diagnostic constraint on LSMs at the site scale. However, the point scale is not the spatial scale at which these models are typically used. In either climate projections or weather forecast applications, gridded model estimates are required at the regional or global scale, with grid cell surface areas of order 25 km^2–40 000 km^2. The relevance of model diagnostic information at the site scale to broader scale gridded simulations remains unclear.

The advancement of remote sensing technology and satellite image processing techniques has dramatically improved the estimation of many components of LSM forcing (e.g. surface air temperature, precipitation, radiation components and vegetation properties) (Loew et al., 2016). These datasets t~~T~~ogether with ~~the reasonable density of~~ in-situ surface observations ~~, this has meant that~~have provided constraint on the reanalysis products ~~– usually the basis of gridded land surface forcing data sets at regional and global scales – are reasonably well constrained~~that provide the basis of global gridded LSM forcing products. On the other hand, very few LSM outputs (such as evapotranspiration or sensible heat flux) can be directly observed by remote sensing, and the only currently available way to derive them is by the use of modelling schemes or empirical formulations that use satellite based datasets (Jung et al., 2011; Fisher et al., 2008; Miralles et al., 2011~~b~~; Martens et al., 2016; Mu et al., 2007; 2011; Su, 2002; Zhang et al., 2016; 2010) . This approach leads to large and usually unquantifiable uncertainties. As a result, gridded LSM evaluation products show considerable differences (Ershadi et al., 2014; Jiménez et al., 2011; McCabe et al., 2016 ; Michel et al., 2016; Vinukollu et al., 2011). If an accurate description of these uncertainties was available, then it might be possible to evaluate gridded LSM simulations in certain circumstances, but until now very few studies quantify LSM uncertainties at regional or global scales (Badgley et al., 2015; Loew et al., 2016; Zhang et al., 2016).

Several gridded ET products have already been developed (see Table 1). These ET products differ in their data requirements, the approaches used to derive them and their estimates (Wang and Dickinson, 2012). So far, it is not at all clear which product provides the most reliable estimates. Recently, inter-comparison studies have aimed to evaluate and compare the available gridded ET products. For instance, Vinukollu et al. (2011), ran simulations of the Surface Energy Balance System SEBS, Penman–Monteith scheme by Mu (PM-Mu) and modified Priestley–Taylor (PT–JPL) forced by a compiled gridded dataset. The monthly model ET estimates were then compared with ET observations from 12 eddy-covariance towers. The root mean square deviation (RMSD) and biases for all the models fell within a small range with the highest correlation with tower data for PT–JPL, and the lowest for SEBS. These results somewhat disagree with the work of Ershadi et al. (2014) who inspected the performance of SEBS, the single-source Penman–Monteith, advection-aridity (AA) complementary method (Brutsaert and Stricker, 1979) and PT–JPL using a high quality forcing dataset from 20 FLUXNET stations and found that the models can be ranked from the best to the worst model as: PT–JPL followed closely by SEBS then PM and finally AA. In the same study, a more detailed analysis revealed that no single model was consistently the best across all landscapes. More recently, in the Energy and Water cycle Experiment (GEWEX) LandFlux project, McCabe et al. (2016) ran simulations of the algorithms in SEBS, PT–JPL, PM-Mu and Global Land Evaporation Amsterdam Model (GLEAM) from common global-scale gridded

forcing data, as well as site-based forcing data, and assessed their response relative to ET measurements from forty-five globally distributed FLUXNET towers. The results indicated that PT–JPL achieved the highest statistical performance, followed closely by GLEAM, whereas PM-MU and SEBS tended to under- and over- estimate fluxes respectively. Furthermore, in the Water Cycle Multi-mission Observation Strategy for Evapotranspiration (WACMOS-ET) project-part 1, Michel et al. (2016) carried out validation experiments for PT–JPL, PM-MU, SEBS and GLEAM against in situ measurements from 85 FLUXNET towers ~~data~~ and found that there was no single best performing model.

In each of the evaluation studies described above, tower data from FLUXNET provide ground truth for gridded ET datasets~~gridded ET data sets are "ground-truthed" using flux tower data from FLUXNET~~ by comparing grid cell values to those measured at the site scale. Most gridded ET products have a 0.5-degree resolution, so that each grid cell can represent an area of around 2500 km^2. The fetch of flux tower measurements varies depending on terrain, vegetation and weather, but is typically under 1 km^2 (Burba and Anderson, 2010) . None of these studies directly address this obvious scale mismatch, and the degree to which surface heterogeneity might nullify any information that flux towers provide about fluxes at these larger scales. Indeed, all the evaluation studies that compared gridded estimates against flux tower observations highlighted this point as a limitation in the evaluation approach.

Zhang et al. (2016) pointed out that merging an ensemble of gridded ET products using a sophisticated data fusion method is likely to generate a better ET product with reduced uncertainty. Ershadi et al. (2014) and McCabe et al. (2016) noted in their studies that the multi-product mean produces improved estimates relative to individual ET products. Similarly, the analysis result of the LandFlux-EVAL project (Mueller et al., 2013) suggested that deriving ET values using the mean of multiple datasets outperforms the ET values from individual datasets. These results do suggest that flux towers likely provide some information at the gridded scale, since there are solid theoretical reasons to expect the mean to outperform individual estimates (Annan and Hargreaves, 2010; Bishop and Abramowitz, 2013), although this fact was not noted in these studies. That is, we would expect that the mean of a range of relatively independent approaches – whose errors are somewhat uncorrelated – would provide a better estimate to observations if those observations were of the same system being simulated. In this study, we provide an even stronger vindication of this relationship in section 2.3 below.

This paper introduces a new method for deriving globally gridded ET as well as its spatio-temporal uncertainty by combining existing gridded ET estimates. The method is based on the ensemble weighting and rescaling technique suggested by Bishop and Abramowitz (2013). The technique provides an analytically optimal linear combination of ensemble members that minimizes mean square error when compared to an observational dataset, and so accounts for both the performance differences and error covariance between the participating ET products (for example, caused by the fact that different ET datasets might share a forcing product). In this way, at least in-sample, the optimal linear combination is insensitive to the addition of redundant information. We use a broad collection of FLUXNET sites as the observational constraint.

We examine the performance of the weighting approach in several in-sample and out-of-sample tests that confirm flux towers do indeed provide information at the grid scale of these products. Having confirmed this, we note that this "Derived Optimal

Linear Combination Evapotranspiration" (DOLCE) is more observationally constrained than the individual participating gridded ET estimates and therefore provides a valuable addition to currently available gridded ET estimates.

Sections 2.1 and 2.2 present the gridded ET datasets and the flux tower ET used to derive the weighted ET product. Sect. 2.3 describes the weighting approach used to calculate the linear combination ET and its uncertainty. In Sect. 2.4 we explain our methods for testing the weighting approach and the representativeness of the point scale to the grid scale. We then present and discuss our results in Sect. 3 and 4 before concluding.

2 Methods, data and experimental setup

To build the DOLCE product, six global gridded ET products, all classified as diagnostic datasets based on Mueller et al. (2011), are weighted based on their ability to match site-level data from 159 globally distributed flux tower sites. DOLCE is derived at 0.5° spatial resolution and monthly temporal resolution for 2000–2009.

2.1 ET datasets

This study employs monthly values from ten existing gridded ET datasets (Table 1). Only six of these datasets – referred to in this study as the Diagnostic Ensemble – are used for building DOLCE, while all of the ten are used to evaluate its performance (we'll refer to these as the Reference Ensemble). The reasons for restricting the Diagnostic Ensemble are (a) to maximize the time period covered by DOLCE (see Table 1), (b) maximize the number of flux tower sites that can inform the weights (noting that datasets have different spatial coverage), and (c) to avoid LSM-based estimates in the final DOLCE product, so that its validity for LSM evaluation is clearer. For these reasons, CSIRO-global, LandFlux-EVAL-ALL, LandFlux-EVAL-DIAG and PT–JPL were removed from the Diagnostic Ensemble.

The weighting approach becomes increasingly effective as more products are included in the weighted mean, but since different products have different spatial coverage, there is a need to strike a balance between maximizing the number of products included in DOLCE and maximizing spatial coverage of DOLCE. To resolve this issue, three different subsets of the Diagnostic Ensemble were chosen to derive DOLCE "tiers". The three DOLCE tiers differ in their spatial coverage of ET and uncertainty, the number of ET products included, and the number of flux towers employed to compute them. DOLCE tier 1 was derived by employing all the products in the Diagnostic Ensemble and 138 sites, and has a spatial coverage equal to the intersection of the global land covered by these products, as shown in Red in Fig. 1. In DOLCE tier 2, GLEAM v2B was excluded and the remaining products were used to compute ET and uncertainty, increasing the spatial coverage and allowing 151 observational flux tower sites. The newly added regions in DOLCE tier 2 are shown in orange in Fig. 1. Finally, ET and uncertainties in the regions shown in yellow were computed using only two ET products, GLEAM V2A and GLEAM V3A, to create DOLCE tier 3, using 159 flux sites. With this approach, we manage to create a global product without sacrificing data quality in regions where more information is available.

2.2 Flux tower data

The flux tower data used in this work is a composite of daily values from the FLUXNET2015-Tier 1 (FN) and LaThuile2007 (LT) Free Fair Use databases ~~(~~(Baldocchi, 2008; Baldocchi et al., 2001; Papale et al., 2012; ~~(~~ORNL DAAC, 2015). We applied quality control and filtering to the site data as follows:

1. Omit the observations that don't fall within the temporal coverage of DOLCE.
2. Omit the daily ET observations if less than 50% of half-hourly ET was observed on that day (as opposed to gap-filled).
3. Omit the monthly ET aggregates that have been calculated from less than 15 daily mean ET values.
4. Correct the remaining site-months of ET observations from the LT dataset for energy balance non-closure on a site-by-site basis (energy closure corrected FN daily data was used)
5. Keep only the sites that are located within the spatial coverage of products in any of the three tiers of the Diagnostic Ensemble (Fig. 2).
6. Exclude irrigated sites or those that have only one monthly record.

Applying (1), (2) and (3) reduced the number of tower sites that were available at the beginning of the analysis from 246 to 172 sites, and produced a total of 7891 site-months (out of a possible 18468). Applying (5) and (6) reduced the number of sites against which the weighting is tested to 138 sites in DOLCE tier 1, increasing to 151 sites in DOLCE tier 2 and 159 in DOLCE tier 3. In (4), two different correction techniques w~~h~~ere applied for energy balance non-closure at ~~LaThuile~~ LT sites. Both involve ensuring that $R_n - G = H + LE$ where R_n is the net radiation, G, H and LE are the ground heat flux, sensible heat flux and latent heat flux respectively, at either the monthly timescale or over the length of the entire site record. Applying a correction technique for energy imbalance at L~~aThuile~~T sites required applying (2) and (3) for the other components of energy imbalance (i.e. R_n, G and H), which means that the sites that had to undergo a correction for the energy imbalance, should have monthly estimates for all the fluxes of the energy budgets, where each monthly value has been calculated from at least 15 daily mean flux values. Because of this constraint, many sites were disregarded from the analysis. One of the two correction method for energy imbalance at ~~LaThuile~~ LT sites distributes the residual errors in heat fluxes according to the Bowen ratio (BR) (Bowen, 1926; Sumner and Jacobs, 2005; Twine et al., 2000) such that:

$$BR = \frac{H_{uncorrected}}{LE_{uncorrected}} \quad (1)$$

is applied using

$$F_{corrected} = \frac{R_n - G}{H_{uncorrected} + LE_{uncorrected}} \times F_{uncorrected} \quad (2)$$

where F is either H or LE, so that BR is unchanged between corrected and uncorrected fluxes. At first, the Bowen Ratio technique was applied on monthly basis, but this occasionally seemed to give erratic results in the corrected flux values. However, when a single correction per site was applied across all the monthly records, the results were more consistent.

The second approach calculates ET as the Residual term in the Energy Balance Equation $LE_{corrected} = R_n - H - G$. The two correction methods were tested separately but no qualitative differences were noticed. All results below use the BR technique. In (6), we expect that some of the weighting models will largely underestimate the flux at irrigated sites, a result of a missing irrigation module in their scheme (Jung et al., 2011; Miralles et al., 2011). Because of this, the error bias of these models at the irrigated sites will modify the mean error bias (i.e. mean bias across all the sites) significantly, which will affect the weighting in favour of the products that can represent better irrigation. We excluded these sites as we do not want the products to be weighted for their inclusion/non-inclusion of physical processes.

We used daily averages of latent heat flux represented by "LE_CORR" in FN dataset. In LT dataset, we employed the components of energy imbalance and their associated flags (in brackets), represented by G_f (G_fqcOK), for soil heat flux, H_f (HFqcOK) for sensible heat flux, Rn_f (Rn_fqcOK) for surface net radiation and LE_f (LE_fqcOK) for latent heat flux.

2.3 Weighting approach

The weighting approach presented here was suggested by (Bishop and Abramowitz, 2013). Given an ensemble of ET estimates and a corresponding observational dataset across time and space, the weighting builds a linear combination of the ensemble members that minimizes the mean square difference (MSD) with respect to the observational dataset such that: If x_k^j is the j^{th} time-space step of the k^{th} bias-corrected ET product (i.e. after subtracting the mean error from the product), and a linear combination of the ET estimates is expressed as

$$\mu_e^j = w^T x^j = \sum_{k=1}^{K} w_k x_k^j \quad , \qquad (3)$$

then the weights w^T provide an analytical solution to the minimization of

$$\sum_{j=1}^{J} (\mu_e^j - y^j)^2 \qquad (4)$$

subject to the constraint that $\sum_{k=1}^{K} w_k = 1$, where $j \in [1, J]$ are the ~~space-and-time~~ time-space steps, $k \in [1, K]$ represent the ET products, y^j is the j^{th} observed ~~space-time~~time-space step.

The solution is expressed as

$$w = \frac{A^{-1}1}{1^T A^{-1} 1}, \qquad (5)$$

where $1^T = \overbrace{[1,1,\ldots,1]}^{k \text{ elements}}$ and A is the k×k error covariance matrix of the gridded products. Further details are in (Abramowitz et al., 2015; Bishop and Abramowitz, 2013; Zeller and Hehn, 1996).

As noted above, this weighting approach has two key advantages. 1) It provides an optimal solution to minimizing mean square error differences between the weighted ET estimates and the observational tower data in-sample. 2) It accounts for the error covariance between the participating datasets (e.g. caused by the fact that single datasets may share ET schemes or forcing), that is, they may not provide independent estimates. The analytical solution guarantees that the weighted mean will perform as well or better than the unweighted mean or any individual estimate included in the linear combination, at least on the data

used to train the weights (that is, in-sample). Moreover, the addition of a new gridded ET product to the ensemble will not degrade the performance of the weighted product, even if it consists of a duplicate product or a poor performing product. In order to confirm this in-sample weighting improvement, we performed an in-sample test presented in the supplementary material (Fig. S1). However, there is no guarantee that the weighted product will necessarily perform well out-of-sample, at a collection of sites not used for training. We explore this more below.

We use the ensemble dependence transformation process presented in (Bishop and Abramowitz, (2013) to calculate the spatiotemporal uncertainty of DOLCE. This involves first quantifying the discrepancy between our weighted ET estimate, μ_e, and the flux tower data, expressed as an error variance, s_e^2, over time and space. We then transform our Diagnostic Ensemble so that its variance about μ_e at a given time-space step, σ_e^{2j}, averaged over all time and space steps where we have flux tower data, is equal to s_e^2. This process ensures that the spread of the transformed Diagnostic Ensemble provides a better uncertainty estimate than simply using the spread of all products in the original Diagnostic Ensemble, since the spread of the transformed ensemble now accurately reflects uncertainty in those grid cells where flux tower data is available.

We calculate the discrepancy between observations and our weighted mean as

$$s_e^2 = \frac{\sum_{j=1}^{J}(\mu_e^j - y^j)^2}{J-1} \quad (62)$$

Next, we wish to ensure that

$$\frac{1}{J}\sum_{j=1}^{J}\sigma_e^{2j} = s_e^2 \quad (73)$$

but the variance of our existing Diagnostic Ensemble will not, in general, satisfy this equation. To transform it so that it does, we first modify the coefficients from Eq. (31), so that they're guaranteed to all be positive:

$$\widetilde{w} = \frac{w^T + (\alpha-1)\frac{1^T}{K}}{\alpha} \quad (84)$$

where $\alpha = 1 - K\min(w_k)$ and $\min(w_k)$ is the smallest negative weight (and α is set 1 if all w_k are non-negative). We then transform the ensemble using

$$\tilde{x}_k^j = \mu_e^j + \beta(\bar{x}^j + \alpha(x_k^j - \bar{x}^j) - \mu_e^j) \quad (59)$$

where

$$\beta = \sqrt{\frac{s_e^2}{\frac{1}{J}\sum_{j=1}^{J}\sum_{k=1}^{K}\widetilde{w_k}(\bar{x}^j + \alpha(x_k^j - \bar{x}^j) - \mu_e^j)^2}} \quad (610)$$

If we then define the weighted variance estimate

$$\sigma_e^{2j} = \sum_{k=1}^{K}\widetilde{w}_k(\tilde{x}_k^j - \mu_e^j)^2 \quad (711)$$

we ensure that both Eq. (6~~3~~) above holds, and also that $\mu_e^j = \sum_{k=1}^{K} w_k x_k^j = \sum_{k=1}^{K} \tilde{w}_k \tilde{x}_k^j$ (see Bishop and Abramowitz, 2013, for proofs). We then use $\sqrt{\sigma_e^{2j}}$ as the spatially and temporally varying estimate of uncertainty standard deviation, which we will refer to below as u~~U~~ncertainty ~~SD~~.

2.4 Experimental setup

We employed ~~three~~ four statistical metrics that reflect how well a gridded ET product represents the quality controlled flux tower observations: Mean Square Error MSE=mean(dataset – observation)^2; Mean Bias=mean| dataset – observation |; ~~Relative Standard Deviation RSD = $\left|\frac{\sigma_{dataset}}{mean_{dataset}} - \frac{\sigma_{observation}}{mean_{observation}}\right|$ and~~ Correlation COR=corr(observation, dataset) and Modified Relative Standard Deviation MRSD = $\frac{\sigma_{\text{dataset or observation}}}{max(mean(observation),\ q)}$.

We use a modified relative standard deviation metric MRSD that measures the variability of latent heat flux relative to the mean of the flux measured at each site. This ensures that a comparison between MRSD for a product and observations can tell us whether a product's variability is too large or too small (unlike relative standard deviation). The term '~~,~~q'~~.~~ is a threshold representing the 2nd percentile of the distribution of observed mean flux (i.e. temporal mean ET) across all sites (≈~~about~~ 13 W/m^2), which guarantees that MRSD calculated across many sites is not dominated by sites where the mean flux (denominator in MRSD Equation above) approaches zero. We looked at the bias in MRSD for each product considered- i.e. $|MRSD_{dataset} - MRSD_{observation}|$, and showed the performance improvement of the weighted mean.

In every test, we split the available sites into in-sample and out-of-sample sites (that is, calibration and validation sites). First, we applied the weighting process at the in-sample sites to weight the individual datasets of the Diagnostic Ensemble and derive a weighted product. Then, we calculated the four metrics above using the out-of-sample sites and displayed results by showing:

a. The percentage performance improvement of the derived weighted product compared to the equally weighted mean (~~WmeanD~~mean) of the Diagnostic Ensemble;

b. The percentage performance improvement of the derived weighted product compared to each individual product in the Reference Ensemble, yielding 10 different values of performance improvement.

c. The aggregate (Ragg) of the values of performance improvement calculated across all 10 products in the Reference Ensemble.

We display the results of performance improvement datasets calculated in (a-c) above in 12 box and whisker plots. In each boxplot, the lower and upper hinges represent the first (Q_1) and third (Q_3) quartiles respectively of the performance improvement datasets and the line located inside the boxplot represents the median value. The extreme of the lower whisker represents the minimum of 1) max(dataset) and 2) (Q_3 + IQR), while the extreme of the upper whisker is the maximum of 1) min(dataset) and 2) (Q_3 + IQR)), where IQR is the interquartile range of the performance improvement dataset. If the median performance improvement is positive, this indicates that the weighting offers an improvement in more than half of the data presented in the boxplot.

We first divide sites between the in-sample and out-of-sample groups by randomly selecting 25 % of the sites as out-of-sample. The remaining sites form the in-sample training set, and are used to calculate a scalar bias correction term and a weight for each participating gridded ET product. These bias correction terms and the weights are then applied to the products at the out-of-sample sites. The test was repeated 5000 times with different random selection of sites being out of sample.

Next, we repeat the process with just one site in the out-of-sample testing group. The bias corrections and weights are therefore derived on all sites except one, and tested on the single out-of-sample site. The same test was repeated for all the participating sites.

3. Results

The results for the 25 % out-of-sample test are displayed in the box and whisker plots presented in Fig. 3 (a), (b), (c) and (d). The MSE plot in Fig. 3 (a) indicates that across most of the random selections of the 25 % out-of-sample sites (57 % of them random selections), the MSE of the weighted product is slightly better (with maximum 12 % improvement maximum) than the equally weighted mean (Wmean Dmean) of the Diagnostic Ensemble. The weighting also improved results for almost all random selections of sites when compared to each individual dataset in the Reference Ensemble. The MRSD plot in (Fig. 3 (b)) shows that the weighting succeeded in improving MRSD relative to Wmean Dmean and PT-JPL by no more than 4 % and 36 % respectively most of the products in the reference ensemble across almost all the majority of the combinations of sites, – but didn't improve MRSD relative to the GLEAM V2B and CSIRO other reference products. This is perhaps unsurprising, given that both the weighted and unweighted mean should have decreased variability when the variations in individual products are not temporally coincident. The COR plot in (Fig. 3 (c)) shows that despite the expected drop in variability in the weighted product, the correlation with site data has in fact improved across most of the different random selections of the 25 % out-of-sample sites, relative to all of the Reference Ensemble datasets. The improvement was minimal relative to Wmean Dmean (maximum 2 %) and MPI (maximum 3 %). Similarly, the Mean Bias plot (Fig. 3 (d)) shows a clear improvement relative to Dmean and the Reference Ensemble datasets across most combinations of sites.

The results of the one site out-of-sample tests are displayed in Fig. 4 (a), (b), (c) and (d). These box-plots indicate that across all the three metrics (MSE, MRSD and Mean Bias) the weighted mean clearly outperformed all the products in the Reference Ensemble. However, when the performance improvement of the weighted product is compared against Dmean the improvement was marginal. However, when the performance improvement of the weighted product is compared against Wmean the results were inconclusive.. In MSE, an improvement was achieved at more than half (53 %) more than half of the sites (Fig. 4 (a)). In MRSD, as many 51 % of the sites showed a small improvement relative to Dmean as those that didn't (Fig. 4 (b)). I Finally in COR, the performance decreased at 57 % of the sites, and a minimal improvement over the Reference Ensemble was shown at more than half of sites (Fig. 4 (c)). Finally, in Mean Bias, 58 % of the sites showed performance

improvement over Dmean (Fig. 4 (d)). Part of the success of the weighting approach relative to the multi-product mean is due to the bias correction applied before the weighting. Figure S3 in supplementary material separates the effect of each step.

The results of the one site out-of-sample test suggest that spatial heterogeneity - the discrepancy between site and grid scale estimates - is significant, as the spread of these box-plots is large. These sites are more likely to be where the fetch of the flux site observations cannot represent the data of the entire grid (estimates). A further analysis (not shown here) has indicated that sites that show poor improvement in Fig. 4 (a), (b), ~~and~~ (c) and (d) have a consistently high bias against all of the gridded products.

To investigate the effect this spatial heterogeneity is having on the out-of-sample tests, we also repeat the one site out-of-sample experiments using a subset of the sites within grid cells that are deemed to be relatively homogeneous (Fig. 4 (e~~d~~), (f~~e~~), (g) ~~-~~and (h~~f~~)). Homogeneity is defined in this case by using only those sites that have the same IGBP vegetation type as the grid cell that contains them. Almost all the gridded ET products use the dominant land cover type for computing ET, yet this is not always the same as the land cover type surrounding the flux site. Since this mismatch occurs in heterogeneous terrains, we will denote the sites that have this property as "HET-case" sites, whereas the sites that show agreement with the underlying grid cell vegetation type ~~we~~ are denoted as "HOM-case" sites. We retrieve the IGBP vegetation cover data of the grid cells from MODIS Land Cover for year 2009 at 0.5° spatial resolution ~~((Channan et al., 2014,~~ http://glcf.umd.edu/data/lc/; Friedl et al., 2010). The vegetation type of individual sites is taken from metadata on the FLUXNET website. We show the distribution of HET-case and HOM-case sites by biome type in Table S2 in the supplements. There is some expectation that the weighting ~~DOLCE~~ will show better performance ~~at~~ if it is trained with HOM-case sites ~~compared to HET sites~~only, although ~~HET~~HOM-case sites consist of about ~~two~~one-thirds of the total number of sites used in this study. To further investigate this idea, we carried out the one site out-of- sample test training and testing on HOM-case sites data only. Figures 4 (e~~d~~), (f~~e~~), (g) and (h~~f~~) show that the HOM-case subset of tower sites does indeed improve the result marginally at least in terms of MSE, COR and ~~COR~~ Mean Bias when compared to the original dataset (HOM & HET -case; Fig. 4 (a), (b), (c) and (d~~e~~)). We did not investigate the 25 % out-of-sample tests with HOM-case sites, since the number of in-sample sites would become too small, and likely lead to overfitting. In both cases it is important to note that many individual sites agree poorly with ~~DOLCE~~the weighted product compared to some other products. The distinction between the results shown in Fig. 3 versus Fig. 4 serve to highlight that DOLCE, and indeed any other large scale gridded ET product, is not suitable for estimation of an individual site's fluxes, even if prediction over many sites shows notable improvement.

On the basis of the ~~widespread~~ aggregate out-of-sample improvement that this approach offers over existing gridded ET products, in terms of MSE, MRSD, COR and ~~COR~~ Mean Bias against site data, we now present details of the DOLCE product, which is trained using all site data and derived from the combination of the six diagnostic products. The 3 GLEAM products were resampled to 0.5° using bilinear interpolation to match the spatial resolution of DOLCE.

We calculated four statistics - Mean Bias, RMSE, Standard Deviation (SD) difference (i.e. $\sigma_{DOLCE} - \sigma_{observation}$) and Correlation - to see how well DOLCE performs at each site. The results are displayed in Fig. 5 (a), (b), (c) and (d). At about ~~50 %~~half of sites, DOLCE had values between –6 and 6 W m^{-2}, –5 and 5 W m^{-2} in the Mean bias and SD difference metrics

respectively, and values greater than 0.93 W m^{-2} and less than 14 W m^{-2} in correlation and RMSE respectively. DOLCE metrics vary in the range [–43.4–30.1 W m^{-2}] for mean bias, [–25.8–23.5 W m^{-2}] for SD difference, [0.09–0.99 W m^{-2}] for correlation and [4.2–60.9 W m^{-2}] for RMSE.

We also calculated the on-site metrics (i.e., RMSE, Mean Bias, SD difference, and Correlation-) for DOLCE separately ~~for~~ at each subset of sites (i.e. HOM-case and HET-case) and we displayed the results in two boxplots (Fig. 5 (left column)). These boxplots show that most of the sites at which DOLCE showed low performance are HET sites, since the end of whiskers is larger for the HET-case sites.

We tested the performance of DOLCE at three irrigated sites that were excluded from the weighting for reasons explained earlier (section 2.2), by computing the four statistics. A description of these sites and the results are shown in Table 2. The results show that the performance of DOLCE is reasonable at US-Ne1 and US-Ne2 and low at US-Twt. These results are discussed further below.

~~We calculated the global mean ET for the three DOLCE tiers and for DOLCE (composite of DOLCE tiers as shown in Fig. 1) from 120 monthly means between 2000 and 2009 (Table 2). The computed values vary significantly across the three tiers. The differences in the global monthly means across DOLCE and the three tiers result from the different spatial coverage of the three tiers.~~

We now look at the differences between DOLCE and two widely used ET products. Figures 6 and 7 show~~s~~ the seasonal mean difference ET between MPIBGC and DOLCE and between LandFlux-EVAL-DIAG and DOLCE ~~(Fig. 7)~~respectively. The differences were computed from seasonal means using the same spatial mask over the period 2000–2009 for MPIBGC and 2000–2005 for LandFlux-EVAL-DIAG. The seasonal plots in Fig. 6 show that, overall, DOLCE has lower ET than MPIBGC. DOLCE tends to have higher ET values in the Asian and Australian tropics during the austral summer and autumn. DOLCE has higher ET values in Sahel during Sep–Nov and in the high plateau of Madagascar during Dec–May. In the Amazon, DOLCE exhibits higher ET values between June and November. Higher values of DOLCE are seen in Guiana highlands throughout the year. The Brazilian highlands show higher values for MPIBGC from June to November. The Mid latitudes of the Northern hemisphere show higher values for MPIBGC between March and August but no significant differences are seen in this area between September and February.

In Fig. 7, there are large differences in ET showing higher values in LandFlux-EVAL-DIAG ET over the Amazon, the r~~R~~ainforests of Southeast Asia and in the high plateau of Madagascar throughout the year. DOLCE tends to exhibit higher ET values in the Ethiopian highlands and Myanmar between June and November.

The spatial distribution of DOLCE mean ET and its seasonal variability (standard deviation) over the austral Summer (Dec–Feb) and Winter (Jun–Aug) from 2000 to 2009 is shown in Fig. 8 (a) and (b) respectively. The seasonal variability of ET is larger in the warm season but is always small over Antarctica, Greenland and the deserts in North Africa (Sahara), the middle east (Arabian Peninsula desert) and Asia (i.e. Gobi, Takla Makan and Thar). The average uncertainty shown in if Fig. 8 (c) is bigger in the warm season, this is in agreement with the relatively large size of the flux in the warm season, and its seasonal variability shown in Fig. 8 (d) is also in agreement with the seasonal variability of the flux~~The spatial distribution of DOLCE~~

~~mean ET and uncertainty over the austral Summer (Dec–Feb) and Winter (Jun–Aug) from 2000 to 2009 is shown in Fig. 8 (a) and (b) respectively. The magnitude of the seasonal change of ET is minor in the tropical regions over the Congo basin, the South-eastern part of the Brasilian highlands, the Ethiopian Highland and Southeast Asia. ET doesn't exhibit any seasonal change over Greenland and the deserts in North Africa (Sahara), the middle east (Arabian Peninsula desert) and Asia (i.e. Gobi, Takla Makan and Thar)~~. Figure 8 (~~e~~e) is intended to give some indication of the reliability of DOLCE. Regions in blue show grid cells that satisfy $\frac{\text{Uncertainty SD}}{\text{mean ET}} \leq 1$. Those shown in green have |mean ET|<5 and Uncertainty SD < 10 but don't satisfy $\frac{\text{Uncertainty SD}}{\text{mean ET}} \leq 1$, so that the uncertainty estimates are higher than their associated ET, but both the uncertainty and the ET estimates are very small. All the remaining grid cells, perhaps those least reliable, are shown in red. The global maps in Fig. 8 (~~e~~e) shows that most of the non-reliable uncertainty values are located in the land added in tier 3 (Fig. 2 in yellow). This is not surprising, because uncertainty values have been derived using only two gridded ET datasets that are also not observationally constrained due to the lack of the sites representing the terrestrial ecosystems and environmental conditions in these areas. Some regions show unreliable uncertainty estimates during the cold season only, located in the plateau of Tibet, over the Andes, the Australian deserts, the South-African deserts, American Great basin and Rocky Mountains.

4. Discussion

Many studies have analysed the systematic and random errors of latent heat flux in FLUXNET measurements (Dirmeyer et al., 2016; Göckede et al., 2008; Richardson et al., 2006). These studies have detected errors of magnitudes that cannot be neglected. A recent study (Cheng et al., 2017) showed that the computed eddy-covariance fluxes have errors in the applied turbulence theory that lead to the underestimation of fluxes, and that this is likely to be one of the causes of the lack of surface energy closure. In this study, we 1) used the flag assigned to the observed flux, to filter out the low-quality data and 2) used energy-balance-corrected FLUXNET data which has higher per-site mean values than the raw data at most of the sites (85% of them). We expect that filtering together with the use of corrected data will reduce the magnitude of the uncertainty in the observational data used here and compensate to a certain extent for the underestimation due to the systematic errors. However, we have not formally explored a range of approaches to addressing this. The possibility of systematic biases in FLUXNET data remains, and this could clearly lead to systematic biases in DOLCE.

We have also assumed that error across sites is uncorrelated, which, given the distribution of sites, is unlikely to be true, meaning that the effective number of sites is probably somewhat smaller than those shown in Fig. 2. Given this dependence is likely to vary depending on a range of time varying factors, we have left the job of attempting to disentangle this issue for future work.

Nevertheless, t~~T~~he results of the one site out-of-sample and 25 % out-of-sample tests above suggest that the weighting of gridded ET products in this way produces a more capable ET product overall. Critically, the fact that the weighting improves

out of sample performance suggests that while the representativeness of point-scale measurement for the grid scale may not exist at every single site, it does exist across all these sites as a whole.

It might be reasonable to assume that eliminating sites that poorly represent grid cell properties might further improve these results, yet the results of categorizing sites into HOM-case and HET-case subsets suggested that this is not necessarily a simple process. Our characterisation was based on a relatively simple vegetation characterisation, yet there are a range of soil-vegetation interactions that likely affect ET in each location, that are not captured in this classification. Among HET-case sites, for example, the worst performance occurred at AU-Fog. This poor performance is a reflection of the large discrepancy between the land cover seen by the flux tower which is a wetland landscape in a monsoonal climate, and thus characterised by high evaporation rates throughout the year and the Savannah land cover seen by the satellite products where the vegetation dries or dies in the dry season. Also, we note that none of the six diagnostic products was able to reproduce the flux measured at this site. Perhaps unsurprisingly, DOLCE scored the worst values for c~~C~~orrelation, RMSE and Mean Bias in AU-Fog.

Aside from AU-Fog, three sites showed low correlation (less than 0.5) between DOLCE and observations. These include ID-Pag, ES-LgS and IL-Yat (see Supplementary Table S1). IL-Yat site is an Evergreen Needleleaf Forest site located over a desert (Bare) grid cell and therefore belongs to the HET-~~-~~case which might be the reason for the discrepancy between the satellite derived ET and the observed ET. ID-Pag belongs to the HOM-~~-~~case, but a burning event occurred in a nearby forest between mid-August and late ~~October,~~October 2002, and caused the site to be covered with a dense smoke-haze. Although these three months were excluded from the analysis, it is very likely that the burned region produced extremely different ET compared to the surrounding unburned landscape, making it a HET-case site until the regeneration of the vegetation and might be the reason why the weighting products couldn't reproduce ET in the following couple of months. ES-LgS is an Open Shrubland site that belongs to the HOM-~~-~~case. This site operated from 2007 to 2009 and is currently inactive. The reason of the low correlation of DOLCE with the observed fluxes is not clear, however previous studies (Ershadi et al., 2014; McCabe et al., 2016) stated that the complexity of S~~s~~hrubland landscape is the cause of the poor performance of many models (PM, GLEAM and PT–JPL). In the case of DOLCE, the poor performance was revealed in one metric only (i.e. correlation). These results nevertheless lead us to expect that if we construct DOLCE by incorporating ~~homogeneous~~ HOM-case sites only, we might get a better product, but the small number of sites satisfying this property, the fact that the separation of sites into HOM-case and HET-case can lead to a separation of land covers, and the difficulty in defining a meaningful definition for expected flux homogeneity are limiting factors. Determining whether DOLCE performs better at HOM-case sites or HET-case sites is inconclusive. Even that the worst performance of DOLCE was achieved in HET-case sites, the boxplots in Fig. 5 (a), (d) show that the value of the median, lower and upper quartiles are better in HET-case for two metrics (i.e. RMSE and COR). While we expect that calibrating the weighting with HOM-case could lead to a better product, we don't expect to see DOLCE overall performing better in any of the two groups.

~~DOLCE has also shown a weak performance at US-Twt, which is an irrigated rice paddy. This site gets flooded in spring and drains in early fall, then the rice is harvested. Only nine months were available for this site, which coincide with the flood and drain period between spring and fall. DOLCE could not depict the flooding and draining event, probably because none of the~~

~~weighting products can represent such phenomena, so it is expected that the effects of seasonal flooding are not represented in DOLCE.~~

Figure 9 shows that DOLCE performs differently for different vegetation types, at least as sampled by the flux towers we use here. Previous studies have shown that different ET estimation approaches perform differently over different biomes. For example, Ershadi et al. (2014) found that the PM models outperform PT–JPL across Gg~~g~~rasslands, C~~c~~roplands and D~~d~~eciduous B~~b~~roadleaf F~~f~~orest~~s~~, but PT–JPL has higher performance than PM models over S~~s~~hrublands and E~~e~~vergreen N~~n~~eedleleaf F~~f~~orest biomes. McCabe et al. (2016) observed poor performance of GLEAM, PM models and PT–JPL over S~~s~~hrublands and low performance over the forest biomes and higher performance over short canopies. The number of sites for each vegetation type in our case varies, with four for Open Shrublands (OSH), 39 for E~~e~~vergreen N~~n~~eedleleaf F~~f~~orest (ENF), nine for E~~e~~vergreen B~~b~~roadleaf F~~f~~orest~~s~~ (EBF), 18 for D~~d~~eciduous B~~b~~roadleaf Forest (DBF), 10 for Mixed Forests (MF), four for Woody Savannas (WSA), four for Savannas (SAV), 42 for grassland (GRA), eight for Wetland (WET), one for Cropland/Natural Vegetation Mosaic (VEG) and 19 for Croplands (CRO). Of the 159 sites that are used to weight the tier 3 products (GLEAM v2A and GLEAM v3A), only 142 sites are within the spatial coverage of DOLCE. This decrease is caused by the coarser resolution of DOLCE compared to GLEAM products. The remaining 17 sites were located on grids that were identified as land grid cells in GLEAM's spatial resolution (0.25º) but water body in DOLCE's spatial resolution (0.5º). The EBF box and whisker plot in Fig. 9 (d) shows the correlation of DOLCE at eight EBF sites, out of which two sites are located in the tropics. The lowest correlation seen in this biome type is at the tropical sites ID-Pag (0.26) and BR-Sa3 (0.62). This suggests that DOLCE tends to represent ET at the extratropical sites better than the tropics, and this is not surprising since most of the sites that were used to calibrate DOLCE were extratropical sites.

DOLCE has also shown a weak performance at US-Twt, which is an irrigated rice paddy. This site gets flooded in spring and drains in early fall, then the rice is harvested. Only nine months were available for this site, which coincide with the flood and drain period between spring and fall. DOLCE could not depict the flooding and draining event, probably because none of the weighting products can represent such phenomena, so it is expected that the effects of seasonal flooding are not represented in DOLCE.

In this study, we sought a single weight for each product to apply globally. But we have reason to believe that different products are likely to perform better in different environments, so that different weights in different climatic circumstances might well improve the result of weighting overall. A similar suggestion was made in the studies of Ershadi et al. (2014) and Michel et al. (2016) who highlighted the need to develop a composite model, where individual models are assigned weights based on their performance across particular biome types and climate zones. We therefore tried to cluster flux tower sites into groups (such as vegetation type) so that each group maintains enough members to allow the in- and out-of-sample testing approach used above. We tried clustering by vegetation type, climate zone and aridity index, and implemented the same one site out-of-sample testing approach as above, but this time, in each cluster different sets of weights will be assigned to the weighting products~~We additionally tried to cluster flux tower sites into groups (such as vegetation type) so that each group maintains~~

~~enough members to allow the in- and out-of-sample testing approach used above. We tried clustering by vegetation type, climate zone and aridity index, and implemented the same one site out-of-sample testing approach as above~~. The motivation for this approach was to try to reduce the number of sites that did not show improvement as a result of the weighting, and ideally improve performance of DOLCE overall, yet none of these clustering approaches delivered any improvement, possibly due to the considerably reduced sample size in each cluster (increasing the likelihood of overfitting). It might also be a result of grid cell heterogeneity within the small sample that constitutes each vegetation type. Finally, as noted above, we have no guarantee that vegetation type is necessarily a good proxy for surface flux behaviour. A summary of this investigation and a plot showing the results of the clustered weighting by vegetation type (Fig. S2) are included in the supplementary material.

There are relatively few towers located in the Southern Hemisphere and the tropics (14 out of 159 sites) and none located in the dry climates over South West Asia and North Africa. The weighting was therefore mostly driven by the ability of products to match sites located in the temperate and cold zones of the Northern Hemisphere, so that performance in climate zones with low FLUXNET site density was under-represented when deriving DOLCE. This might raise questions about the performance of DOLCE in the tropics and the Southern Hemisphere. To evaluate DOLCE in these areas we calculated the four site metrics (RMSE, Mean bias, SD difference and Correlation) separately for two groups of sites, 1) those located in the Northern Hemisphere excluding Tropics and 2) sites located in the Tropics and/or Southern Hemisphere. We excluded the two sites ID-Pag and AU-Fog in this exercise since both are wetland sites, and so would complicate a determination of whether these two groups had notable behavioural differences. If systematic behavioural differences did exist between these two groups, we would expect relatively poorer performance of DOLCE at group2 sites compared to group1 sites. The results, shown in Fig. 10, appear inconclusive. DOLCE performed marginally worse at group2 sites overall, however with the limited number of sites in group2, the validation of the performance of DOLCE in the tropics and Southern Hemisphere remains somewhat uncertain. The uneven distribution of eddy-covariance sites between the Northern and Southern Hemisphere and across the different climates might also explain why much of the largest seasonal differences DOLCE-MPI and DOLCE-LANDFLUX-EVAL shown in Fig. 6 and Fig. 7 reside in the low latitudes (tropics) and the Southern Hemisphere and the persistent differences between DOLCE and LANDFLUX-EVAL in the tropics throughout the year. The expansion of the FLUXNET network into these areas that are lacking observations is clearly something that would improve DOLCE and LSM evaluation more broadly.

In this study, DOLCE was derived by combining remote sensing products that use an empirical approach (i.e. MPIBGC) and a more physical approach (i.e. MOD16, GLEAM v2A, GLEAM v2B and GLEAM v3A and PML). Constructing an ET product by combining different approaches can take advantage of their desirable features and reduce their limitations (Zhang et al., 2016). This is indeed reflected by the enhanced performance of DOLCE over all the biome types against the reference products. Table 3 shows that the weights were attributed almost equally to each approach (i.e. ≈ 0.5 for both MPIBGC and the physical ensemble in tier 1 and tier 2), which indicates that the two approaches contributed equally in deriving DOLCE, it is not surprising that MPIBGC was the most weighted product since it is highly calibrated with flux tower data. In a further analysis, we left MPIBGC out and we performed the out of sample tests using the five remaining products. We wanted here to test how

the weighting will perform without MPIBGC. The plots in Fig. 11 and 12 show the results of 25 % out-of-sample test and one site out-of-sample test respectively. Overall, without including MPIBGC, the weighting offers a smaller performance improvement than that offered when MPIBGC is a member of weighting ensemble (Fig. 3 and Fig. 4 (a-d)). The distribution of the weights when MPIBGC is absent from the weighting is 0.3 for both PML and GLEAM v3A, 0.2 for GLEAM v2B, 0.13 for MOD16 and 0.07 for GLEAM v2A.

It is important to mention that the weighting approach offers the possibility of enhancing DOLCE in the future by incorporating any ET dataset that becomes available, (for example, WECANN, (Alemohammad et al., 2016), and HOLAPS v1.0, (Loew et al., 2016).

The limitations of the weighting approach presented here arise from issues in the observational data and limitations in the weighted datasets that are employed. FLUXNET data inevitably has some instrumentation-related problems (Göckede et al., 2008) that affect the data quality of the measured flux in some terrains and under some environmental conditions. Moreover, FLUXNET sites are not globally-representative of all terrestrial ecosystems, and not evenly distributed across the biome types, which might lead to biases when computing the weights of the products in favour of the product that outperforms over the most frequent cover types. Also, in some areas where tower density is relatively high, the information from different towers is not necessarily independent.

Finally, a common imperfection in all the weighted products due to for example anthropogenic water management, will of course lead to the same imperfection in the derived product.

The uncertainty estimate presented here is firmly grounded in the spread of existing gridded ET products, but is better than this spread alone, since this spread has been recalibrated so that the uncertainty of DOLCE where flux tower data exists is precisely the spread of the recalibrated products.

5. Conclusion

In this study, we presented a new global ET product with monthly temporal resolution for 2000–2009 at 0.5° spatial resolution and a calibrated estimate of its uncertainty. The approach used to weight existing gridded ET products accounted for both the performance differences and error covariance between the participating ET products. The DOLCE product performs better than any of its six constituent members overall, as well outperforming other well-known gridded ET products (CSIRO-global, LandFlux-EVAL-ALL, LandFlux-EVAL-DIAG and PT–JPL) across a range of metrics. It was shown that despite the scale mismatch between the flux tower and the grid cell, the ensemble of flux towers as a whole can provide information about the grid cells that contain them. While the representativeness of the point scale for the grid scale is enhanced by only considering sites that lie within homogeneous grid cells we suggest that an optimal definition of homogeneity for flux behaviour be the subject of future investigation. Nevertheless, DOLCE appears to outperform existing gridded ET products overall, and offers the opportunity for improvement as more flux tower data and new gridded ET products become available. Expanding DOLCE

over longer time periods and incorporating more diagnostic ET datasets (such as PT-JPL, CSIRO-global, GLEAM-V3B, SEBS, WECANN and HOLAPS) will be carried out in future versions.

6. Competing interests

The authors declare that they have no conflict of interest.

7. Data availability

DOLCE v1.0 can be downloaded from geonetwork.nci.org.au and its DOI is http://dx.doi.org/10.4225/41/58980b55b0495

8. Acknowledgements

We acknowledge the support of the Australian Research Council Centre of Excellence for Climate System Science (CE110001028). This work used eddy covariance data acquired and shared by the FLUXNET community, including these networks: AmeriFlux, AfriFlux, AsiaFlux, CarboAfrica, CarboEuropeIP, CarboItaly, CarboMont, ChinaFlux, Fluxnet-Canada, GreenGrass, ICOS, KoFlux, LBA, NECC, OzFlux-TERN, TCOS-Siberia, and USCCC. The ERA-Interim reanalysis data are provided by ECMWF and processed by LSCE. The FLUXNET eddy covariance data processing and harmonization was carried out by the European Fluxes Database Cluster, AmeriFlux Management Project, and Fluxdata project of FLUXNET, with the support of CDIAC and ICOS Ecosystem Thematic Center, and the OzFlux, ChinaFlux and AsiaFlux offices. This work also used eddy covariance data acquired by the FLUXNET community and in particular by the following networks: AmeriFlux (U.S. Department of Energy, Biological and Environmental Research, Terrestrial Carbon Program (DE - FG02 - 04ER63917 and DE-FG02 - 04ER63911)), AfriFlux, AsiaFlux, CarboAfrica, CarboEuropeIP, CarboItaly, CarboMont, ChinaFlux, Fluxnet - Canada (supported by CFCAS, NSERC, BIOCAP, Environment Canada, and NRCan), GreenGrass, KoFlux, LBA, NECC, OzFlux, TCOS - Siberia, USCCC. We acknowledge the financial support to the eddy covariance data harmonization provided by CarboEuropeIP, FAO - GTOS - TCO, iLEAPS, Max Planck Institute for Biogeochemistry, National Science Foundation, University of Tuscia, Université Laval and Environment Canada and US Department of Energy and the database development and technical support from Berkeley Water Center, Lawrence Berkeley National Laboratory, Microsoft Research eScience, Oak Ridge National Laboratory, University of California - Berkeley, University of Virginia.

9. References

Abramowitz, G.: Towards a benchmark for land surface models, Geophys. Res. Lett., 32(22), doi:10.1029/2005GL024419, 2005.

Abramowitz, G.: Towards a public, standardized, diagnostic benchmarking system for land surface models, Geosci. Model

Dev., 5(3), 819–827, doi:10.5194/gmd-5-819-2012, 2012.

Abramowitz, G., Bishop, C. H., Abramowitz, G. and Bishop, C. H.: Climate Model Dependence and the Ensemble Dependence Transformation of CMIP Projections, J. Clim., 28(6), 2332–2348, doi:10.1175/JCLI-D-14-00364.1, 2015.

Alemohammad, S. H., Fang, B., Konings, A. G., Green, J. K., Kolassa, J., Prigent, C., Aires, F., Miralles, D. and Gentine, P.: Water, Energy, and Carbon with Artificial Neural Networks (WECANN): A statistically-based estimate of global surface turbulent fluxes using solar-induced fluorescence, Biogeosciences Discuss., 1–36, doi:10.5194/bg-2016-495, 2016.

Annan, J. D. and Hargreaves, J. C.: Reliability of the CMIP3 ensemble, Geophys. Res. Lett., 37(2), doi:10.1029/2009GL041994, 2010.

Badgley, G., Fisher, J. B., Jiménez, C., Tu, K. P. and Vinukollu, R.: On Uncertainty in Global Terrestrial Evapotranspiration Estimates from Choice of Input Forcing Datasets, J. Hydrometeorol., 16(4), 1449–1455, doi:10.1175/JHM-D-14-0040.1, 2015.

Baldocchi, D.: "Breathing" of the terrestrial biosphere: lessons learned from a global network of carbon dioxide flux measurement systems, Aust. J. Bot., 56(1), 1, doi:10.1071/BT07151, 2008.

Baldocchi, D., Falge, E., Gu, L., Olson, R., Hollinger, D., Running, S., Anthoni, P., Bernhofer, C., Davis, K., Evans, R., Fuentes, J., Goldstein, A., Katul, G., Law, B., Lee, X., Malhi, Y., Meyers, T., Munger, W., Oechel, W., Paw, K. T., Pilegaard, K., Schmid, H. P., Valentini, R., Verma, S., Vesala, T., Wilson, K., Wofsy, S., Baldocchi, D., Falge, E., Gu, L., Olson, R., Hollinger, D., Running, S., Anthoni, P., Bernhofer, C., Davis, K., Evans, R., Fuentes, J., Goldstein, A., Katul, G., Law, B., Lee, X., Malhi, Y., Meyers, T., Munger, W., Oechel, W., Paw, K. T., Pilegaard, K., Schmid, H. P., Valentini, R., Verma, S., Vesala, T., Wilson, K. and Wofsy, S.: FLUXNET: A New Tool to Study the Temporal and Spatial Variability of Ecosystem–Scale Carbon Dioxide, Water Vapor, and Energy Flux Densities, Bull. Am. Meteorol. Soc., 82(11), 2415–2434, doi:10.1175/1520-0477(2001)082<2415:FANTTS>2.3.CO;2, 2001.

Belelli Marchesini, L., Papale, D., Reichstein, M., Vuichard, N., Tchebakova, N. and Valentini, R.: Carbon balance assessment of a natural steppe of southern Siberia by multiple constraint approach, Biogeosciences, 4(4), 581–595, doi:10.5194/bg-4-581-2007, 2007.

Best, M. J., Abramowitz, G., Johnson, H. R., Pitman, A. J., Balsamo, G., Boone, A., Cuntz, M., Decharme, B., Dirmeyer, P. A., Dong, J., Ek, M., Guo, Z., Haverd, V., van den Hurk, B. J. J., Nearing, G. S., Pak, B., Peters-Lidard, C., Santanello, J. A., Stevens, L., Vuichard, N., Best, M. J., Abramowitz, G., Johnson, H. R., Pitman, A. J., Balsamo, G., Boone, A., Cuntz, M., Decharme, B., Dirmeyer, P. A., Dong, J., Ek, M., Guo, Z., Haverd, V., Hurk, B. J. J. van den, Nearing, G. S., Pak, B., Peters-Lidard, C., Jr., J. A. S., Stevens, L., Vuichard, N., J Van Den Hurk, B. J., Nearing, G. S., Pak, B., Peters-Lidard, C., Santanello Jr, J. A., Stevens, L. and Vuichard, N.: The Plumbing of Land Surface Models: Benchmarking Model Performance, J. Hydrometeorol., 16(3), 1425–1442, doi:10.1175/JHM-D-14-0158.1, 2015.

Bishop, C. H. and Abramowitz, G.: Climate model dependence and the replicate Earth paradigm, Clim. Dyn., 41(3–4), 885–900, doi:10.1007/s00382-012-1610-y, 2013.

Bowen, I. S.: The Ratio of Heat Losses by Conduction and by Evaporation from any Water Surface, Phys. Rev., 27(6), 779–787, doi:10.1103/PhysRev.27.779, 1926.

Brutsaert, W. and Stricker, H.: An advection-aridity approach to estimate actual regional evapotranspiration, Water Resour. Res., 15(2), 443–450, doi:10.1029/WR015i002p00443, 1979.

Burba, G. and Anderson, D.: A Brief Practical Guide to Eddy Covariance Flux Measurements: Principles and Workflow Examples for Scientific and Industrial Applications, LI-COR Biosciences., 2010.

Chen, T. H., Henderson-Sellers, A., Milly, P. C. D., Pitman, A. J., Beljaars, A. C. M., Polcher, J., Abramopoulos, F., Boone, A., Chang, S., Chen, F., Dai, Y., Desborough, C. E., Dickinson, R. E., Dümenil, L., Ek, M., Garratt, J. R., Gedney, N., Gusev, Y. M., Kim, J., Koster, R., Kowalczyk, E. A., Laval, K., Lean, J., Lettenmaier, D., Liang, X., Mahfouf, J.-F., Mengelkamp, H.-T., Mitchell, K., Nasonova, O. N., Noilhan, J., Robock, A., Rosenzweig, C., Schaake, J., Schlosser, C. A., Schulz, J.-P., Shao, Y., Shmakin, A. B., Verseghy, D. L., Wetzel, P., Wood, E. F., Xue, Y., Yang, Z.-L., Zeng, Q., Chen, T. H., Henderson-Sellers, A., Milly, P. C. D., Pitman, A. J., Beljaars, A. C. M., Polcher, J., Abramopoulos, F., Boone, A., Chang, S., Chen, F., Dai, Y., Desborough, C. E., Dickinson, R. E., Dümenil, L., Ek, M., Garratt, J. R., Gedney, N., Gusev, Y. M., Kim, J., Koster, R., Kowalczyk, E. A., Laval, K., Lean, J., Lettenmaier, D., Liang, X., Mahfouf, J.-F., Mengelkamp, H.-T., Mitchell, K., Nasonova, O. N., Noilhan, J., Robock, A., Rosenzweig, C., Schaake, J., Schlosser, C. A., Schulz, J.-P., Shao, Y., Shmakin, A. B., Verseghy, D. L., Wetzel, P., Wood, E. F., Xue, Y., Yang, Z.-L. and Zeng, Q.: Cabauw Experimental Results from the Project for Intercomparison of Land-Surface Parameterization Schemes, J. Clim., 10(6), 1194–1215, doi:10.1175/1520-0442(1997)010<1194:CERFTP>2.0.CO;2, 1997.

Cheng, Y., Sayde, C., Li, Q., Basara, J., Selker, J., Tanner, E. and Gentine, P.: Failure of Taylor's hypothesis in the atmospheric surface layer and its correction for eddy-covariance measurements, Geophys. Res. Lett., 44(9), 4287–4295, doi:10.1002/2017GL073499, 2017.

Dirmeyer, P. A., Wu, J., Norton, H. E., Dorigo, W. A., Quiring, S. M., Ford, T. W., Santanello, J. A., Bosilovich, M. G., Ek, M. B., Koster, R. D., Balsamo, G. and Lawrence, D. M.: Confronting Weather and Climate Models with Observational Data from Soil Moisture Networks over the United States, J. Hydrometeorol., 17(4), 1049–1067, doi:10.1175/JHM-D-15-0196.1, 2016.

Ershadi, A., McCabe, M. F., Evans, J. P., Chaney, N. W. and Wood, E. F.: Multi-site evaluation of terrestrial evaporation models using FLUXNET data, Agric. For. Meteorol., 187, 46–61, doi:10.1016/j.agrformet.2013.11.008, 2014.

Fisher, J. B., Tu, K. P. and Baldocchi, D. D.: Global estimates of the land-atmosphere water flux based on monthly AVHRR and ISLSCP-II data, validated at 16 FLUXNET sites, Remote Sens. Environ., 112(3), 901–919, doi:10.1016/j.rse.2007.06.025, 2008.

Friedl, M. A., Sulla-Menashe, D., Tan, B., Schneider, A., Ramankutty, N., Sibley, A. and Huang, X.: MODIS Collection 5 global land cover: Algorithm refinements and characterization of new datasets, Remote Sens. Environ., 114, 168–182, doi:10.1016/j.rse.2009.08.016, 2010.

Göckede, M., Foken, T., Aubinet, M., Aurela, M., Banza, J., Bernhofer, C., Bonnefond, J. M., Brunet, Y., Carrara, A., Clement, R., Dellwik, E., Elbers, J., Eugster, W., Fuhrer, J., Granier, A., Grünwald, T., Heinesch, B., Janssens, I. A., Knohl, A., Koeble, R., Laurila, T., Longdoz, B., Manca, G., Marek, M., Markkanen, T., Mateus, J., Matteucci, G., Mauder, M., Migliavacca, M.,

Minerbi, S., Moncrieff, J., Montagnani, L., Moors, E., Ourcival, J.-M., Papale, D., Pereira, J., Pilegaard, K., Pita, G., Rambal, S., Rebmann, C., Rodrigues, A., Rotenberg, E., Sanz, M. J., Sedlak, P., Seufert, G., Siebicke, L., Soussana, J. F., Valentini, R., Vesala, T., Verbeeck, H. and Yakir, D.: Quality control of CarboEurope flux data–Part 1: Coupling footprint analyses with flux data quality assessment to evaluate sites in forest ecosystems, Biogeosciences, 5(2), 433–450, doi:10.5194/bg-5-433-2008, 2008.

Haughton, N., Abramowitz, G., Pitman, A. J., Or, D., Best, M. J., Johnson, H. R., Balsamo, G., Boone, A., Cuntz, M., Decharme, B., Dirmeyer, P. A., Dong, J., Ek, M., Guo, Z., Haverd, V., van den Hurk, B. J. J., Nearing, G. S., Pak, B., Santanello, J. A., Stevens, L. E., Vuichard, N., Haughton, N., Abramowitz, G., Pitman, A. J., Or, D., Best, M. J., Johnson, H. R., Balsamo, G., Boone, A., Cuntz, M., Decharme, B., Dirmeyer, P. A., Dong, J., Ek, M., Guo, Z., Haverd, V., Hurk, B. J. J. van den, Nearing, G. S., Pak, B., Jr., J. A. S., Stevens, L. E. and Vuichard, N.: The Plumbing of Land Surface Models: Is Poor Performance a Result of Methodology or Data Quality?, J. Hydrometeorol., 17(6), 1705–1723, doi:10.1175/JHM-D-15-0171.1, 2016.

Jiménez, C., Prigent, C., Mueller, B., Seneviratne, S. I., McCabe, M. F., Wood, E. F., Rossow, W. B., Balsamo, G., Betts, A. K., Dirmeyer, P. A., Fisher, J. B., Jung, M., Kanamitsu, M., Reichle, R. H., Reichstein, M., Rodell, M., Sheffield, J., Tu, K. and Wang, K.: Global intercomparison of 12 land surface heat flux estimates, J. Geophys. Res. Atmos., doi:10.1029/2010JD014545, 2011.

Jung, M., Reichstein, M., Margolis, H. A., Cescatti, A., Richardson, A. D., Arain, M. A., Arneth, A., Bernhofer, C., Bonal, D., Chen, J., Gianelle, D., Gobron, N., Kiely, G., Kutsch, W., Lasslop, G., Law, B. E., Lindroth, A., Merbold, L., Montagnani, L., Moors, E. J., Papale, D., Sottocornola, M., Vaccari, F. and Williams, C.: Global patterns of land-atmosphere fluxes of carbon dioxide, latent heat, and sensible heat derived from eddy covariance, satellite, and meteorological observations, J. Geophys. Res., 116(G3), G00J07, doi:10.1029/2010JG001566, 2011.

Loew, A., Peng, J. and Borsche, M.: High-resolution land surface fluxes from satellite and reanalysis data (HOLAPS v1.0): evaluation and uncertainty assessment, Geosci. Model Dev, 9, 2499–2532, doi:10.5194/gmd-9-2499-2016, 2016.

Martens, B., Miralles, D., Lievens, H., Van Der Schalie, R., De Jeu, R., Fernández-Prieto, D. and Verhoest, N.: GLEAM v3: updated land evaporation and root-zone soil moisture datasets, Geophys. Res. Abstr. EGU Gen. Assem., 18(23), 2016–4253, 2016.

McCabe, M. F., Ershadi, A., Jimenez, C., Miralles, D. G., Michel, D. and Wood, E. F.: The GEWEX LandFlux project: Evaluation of model evaporation using tower-based and globally gridded forcing data, Geosci. Model Dev., doi:10.5194/gmd-9-283-2016, 2016.

Michel, D., Jiménez, C., Miralles, D. G., Jung, M., Hirschi, M., Ershadi, A., Martens, B., McCabe, M. F., Fisher, J. B., Mu, Q., Seneviratne, S. I., Wood, E. F. and Fernández-Prieto, D.: The WACMOS-ET project Part 1: Tower-scale evaluation of four remote-sensing-based evapotranspiration algorithms, Hydrol. Earth Syst. Sci., 20(2), 803–822, doi:10.5194/hess-20-803-2016, 2016.

Miralles, D. G., Holmes, T. R. H., De Jeu, R. A. M., Gash, J. H., Meesters, A. G. C. A. and Dolman, A. J.: Global land-surface

evaporation estimated from satellite-based observations, Hydrol. Earth Syst. Sci., 15(2), 453–469, doi:10.5194/hess-15-453-2011, 2011.

Monteith J. L.: Evaporation and envrionment, Symp. Soc. Exp. Biol., 19(205–23), 1965.

Mu, Q., Heinsch, F. A., Zhao, M. and Running, S. W.: Development of a global evapotranspiration algorithm based on MODIS and global meteorology data, Remote Sens. Environ., 111(4), 519–536, doi:10.1016/j.rse.2007.04.015, 2007.

Mu, Q., Zhao, M. and Running, S. W.: Improvements to a MODIS global terrestrial evapotranspiration algorithm, Remote Sens. Environ., 115(8), 1781–1800, doi:10.1016/j.rse.2011.02.019, 2011.

Mueller, B., Seneviratne, S. I., Jimenez, C., Corti, T., Hirschi, M., Balsamo, G., Ciais, P., Dirmeyer, P., Fisher, J. B., Guo, Z., Jung, M., Maignan, F., McCabe, M. F., Reichle, R., Reichstein, M., Rodell, M., Sheffield, J., Teuling, A. J., Wang, K., Wood, E. F. and Zhang, Y.: Evaluation of global observations-based evapotranspiration datasets and IPCC AR4 simulations, Geophys. Res. Lett., 38(6), doi:10.1029/2010GL046230, 2011.

Mueller, B., Hirschi, M., Jimenez, C., Ciais, P., Dirmeyer, P. A., Dolman, A. J., Fisher, J. B., Jung, M., Ludwig, F., Maignan, F., Miralles, D. G., McCabe, M. F., Reichstein, M., Sheffield, J., Wang, K., Wood, E. F., Zhang, Y. and Seneviratne, S. I.: Benchmark products for land evapotranspiration: LandFlux-EVAL multi-data set synthesis, Hydrol. Earth Syst. Sci., doi:10.5194/hess-17-3707-2013, 2013.

ORNL DAAC: Home | fluxnetweb.ornl.gov, [online] Available from: https://fluxnet.ornl.gov/ (Accessed 1 December 2016), 2015.

Papale, D., Agarwal, D. A., Baldocchi, D., Cook, R. B., Fisher, J. B. and van Ingen, C.: Database Maintenance, Data Sharing Policy, Collaboration, in Eddy Covariance, pp. 399–424, Springer Netherlands, Dordrecht., 2012.

Richardson, A. D., Hollinger, D. Y., Burba, G. G., Davis, K. J., Flanagan, L. B., Katul, G. G., William Munger, J., Ricciuto, D. M., Stoy, P. C., Suyker, A. E., Verma, S. B. and Wofsy, S. C.: A multi-site analysis of random error in tower-based measurements of carbon and energy fluxes, Agric. For. Meteorol., 136(1–2), 1–18, doi:10.1016/j.agrformet.2006.01.007, 2006.

Stöckli, R., Lawrence, D. M., Niu, G.-Y., Oleson, K. W., Thornton, P. E., Yang, Z.-L., Bonan, G. B., Denning, A. S. and Running, S. W.: Use of FLUXNET in the Community Land Model development, J. Geophys. Res. Biogeosciences, 113(G1), doi:10.1029/2007JG000562, 2008.

Su, Z.: The Surface Energy Balance System (SEBS) for estimation of turbulent heat fluxes, Hydrol. Earth Syst. Sci., 6(1), 85–100, doi:10.5194/hess-6-85-2002, 2002.

Sumner, D. M. and Jacobs, J. M.: Utility of Penman–Monteith, Priestley–Taylor, reference evapotranspiration, and pan evaporation methods to estimate pasture evapotranspiration, J. Hydrol., 308(1–4), 81–104, doi:10.1016/j.jhydrol.2004.10.023, 2005.

Twine, T. E., Kustas, W. P., Norman, J. M., Cook, D. R., Houser, P. R., Meyers, T. P., Prueger, J. H., Starks, P. J. and Wesley, M. L.: Correcting eddy-covariance flux underestimates over a grassland, NASA Publ., 2000.

Vinukollu, R. K., Wood, E. F., Ferguson, C. R. and Fisher, J. B.: Global estimates of evapotranspiration for climate studies

using multi-sensor remote sensing data: Evaluation of three process-based approaches, Remote Sens. Environ., doi:10.1016/j.rse.2010.11.006, 2011.

Wang, K. and Dickinson, R. E.: A review of global terrestrial evapotranspiration: Observation, modeling, climatology, and climatic variability., 2012.

Wang, Y. P., Baldocchi, D., Leuning, R., Falge, E. and Vesala, T.: Estimating parameters in a land-surface model by applying nonlinear inversion to eddy covariance flux measurements from eight FLUXNET sites, Glob. Chang. Biol., 13(3), 652–670, doi:10.1111/j.1365-2486.2006.01225.x, 2007.

Zeller, K. and Hehn, T.: Measurements of upward turbulent ozone fluxes above a subalpine spruce-fir forest, Geophys. Res. Lett., 23(8), 841–844, doi:10.1029/96GL00786, 1996.

Zhang, K., Kimball, J. S., Nemani, R. R. and Running, S. W.: A continuous satellite-derived global record of land surface evapotranspiration from 1983 to 2006, Water Resour. Res., 46(9), doi:10.1029/2009WR008800, 2010a.

Zhang, K., Kimball, J. S. and Running, S. W.: A review of remote sensing based actual evapotranspiration estimation, Wiley Interdiscip. Rev. Water, 3(6), 834–853, doi:10.1002/wat2.1168, 2016.

Zhang, Y., Leuning, R., Hutley, L. B., Beringer, J., McHugh, I. and Walker, J. P.: Using long-term water balances to parameterize surface conductances and calculate evaporation at 0.05° spatial resolution, Water Resour. Res., 46(5), doi:10.1029/2009WR008716, 2010b.

Zhang, Y., Peña-Arancibia, J. L., Mcvicar, T. R., Chiew, F. H. S., Vaze, J., Liu, C., Lu, X., Zheng, H., Wang, Y., Liu, Y. Y., Miralles, D. G. and Pan, M.: Multi-decadal trends in global terrestrial evapotranspiration and its components, Nat. Publ. Gr., 6, 19124, doi:10.1038/srep19124, 2015.

Tables

Table 1: Gridded ET products used in this paper.

ET product and Reference	Abbreviation	Time period & Spatial Resolution	Forcing data source	Calculation Method(s)
CSIRO-global (Zhang et al., 2010a)	CS	1983–2006 0.5° Also available at 8km and 1°	Meteorological observations from flux tower distributed across all global biome types Remote sensing inputs	An extended ET product of CSIRO (Zhang et al., 2010b) that covers a global domain NDVI-based PM model PT equation for open water evaporation
GLEAM-V2A (Miralles et al., 2011b)	G2A	20001980–20092011 0.25°	Remote sensing based observations Gauged based precipitation	PT equation Canopy Interception Model, Soil water module and Stress module
GLEAM-V2B (Miralles et al., 2011b)	G2B	2000–20092011 0.25°	Remote sensing based observations	PT equation Canopy Interception Module, Soil water module and Stress module
GLEAM-V3A (Martens et al., 2016)	G3A	20001980–20122014 0.25°	Satellite based inputs Multi-source precipitation	A revised version of GLEAM V2A in which new satellite-observed geophysical variables have been incorporated and the representation of the surface soil moisture and evaporation has been improved
LandFlux-Eval-Diag	LFD	1989–2005 1°	Simple mean of 5 diagnostic ET datasets	

(Mueller et al., 2011, 2013)				
LandFlux-Eval-All (Mueller et al., 2011, 2013)	LFA		Simple mean of 14 Diagnostic, LSM and Reanalysis datasets.	
MOD16 MODIS global ET products (Mu et al., 2011)	MOD	2000–2014 0.5° also available at 0.0 5°~~1°~~	Global Modeling and Assimilation Office (GMAO) meteorological reanalysis data Remote sensing inputs from MODIS 8-day retrievals	PM formula (Monteith J. L., 1965)
MPIBGC (Jung et al., 2011)	MPI	1982–2011 ~~1~~0.5°	FLUXNET data from 253 sites Remote sensing datasets from (SeaWiFS)	Empirical methods: a Model Tree Ensemble (MTE) Machine learning techniques
PML PM-Leuning model (Zhang et al., 2015)	PML	1981–2012 ~~1~~0.5°	GMAO Reanalysis products	PM Leuning method
PT–JPL (Fisher et al., 2008)	PT	1984–2006 1°	Meteorological reanalysis data from ISLSCP –II Remote sensing based observations from monthly AVHRR data	PT equation

~~**Table 2: Standard Deviation (SD) and mean over space and time of DOLCE and the weighting products, derived from the monthly means over the time period 2000-2010 (a total of 120 months) and within the spatial coverage intersect of DOLCE monthly land cover across all the months. All the values are rounded to the nearest tenth**~~

	Tier 1 spatial coverage		Tier 2 spatial coverage		Tier 3 spatial coverage		Tier 1, Tier 2 and Tier 3 composite	
Product	SD W m^{-2}	Mean W m^{-2}	SD W m^{-2}	Mean W m^{-2}	SD W m^{-2}	Mean W m^{-2}	SD W m^{-2}	Mean W m^{-2}
DOLCE	8.9	41.7	11.6	34.7	5.7	20.2	6.4	21.0
MOD16	8.0	42.9	11.0	38.8				
MPIBGC	10.1	44.9	12.8	39.4				
GLEAM v2A	8.1	43.6	10.6	38.4	6	22.5		
GLEAM v2B	7.6	43						
GLEAM v3A	7.7	33.7	10	29.3	5.5	16.5		
PML	8.8	44.6	10.4	38				

Table 2: Four metrics (RMSE, Mean bias, SD difference and Correlation) of DOLCE at three irrigated sites, and the number of available monthly records for each site.

Site-Code	Longitude	Latitude	Description	RMSE	Mean bias	SD difference	Correlation	Number of months
US-Ne1	-96.4766	41.1651	Rice paddy	16.6	-7.41	-9.25	0.96	103
US-Ne2	-96.4701	41.1649	Mead irrigated continuous maize site	15.8	-5.05	-7.44	0.95	103
US-Twt	-121.6521	38.1055	Mead irrigated maize-soybean rotation site	91.9	-67.39	-55.23	0.49	9

Table 3: Weights and bias terms (1) Bias of weighting products and (2) weights assigned to the weighting bias corrected products in the case of each of the three DOLCE tiers, and the number of flux tower sites used to feed the weighting.

	DOLCE tier 1 138 sites		DOLCE tier 2 151 sites		DOLCE tier 3 159 sites	
Product	weight	~~mean bias~~ Bias W m^{-2}	weight	~~mean~~ Bbias W m^{-2}	weight	~~mean~~ Bbias W m^{-2}
MOD16	0.041	3.756	0.05	4.0		
MPIBGC	0.495	3.837	0.537	3.882		
GLEAM –v2A	–0.026	6.180	0.123	6.098	0.44	5.735
GLEAM-v2B	0.192	–3.571				
GLEAM-v3A	0.171	–5.483	0.151	–5.221	0.56	–6.344
PML	0.127	4.982	0.139	4.668		

Figures

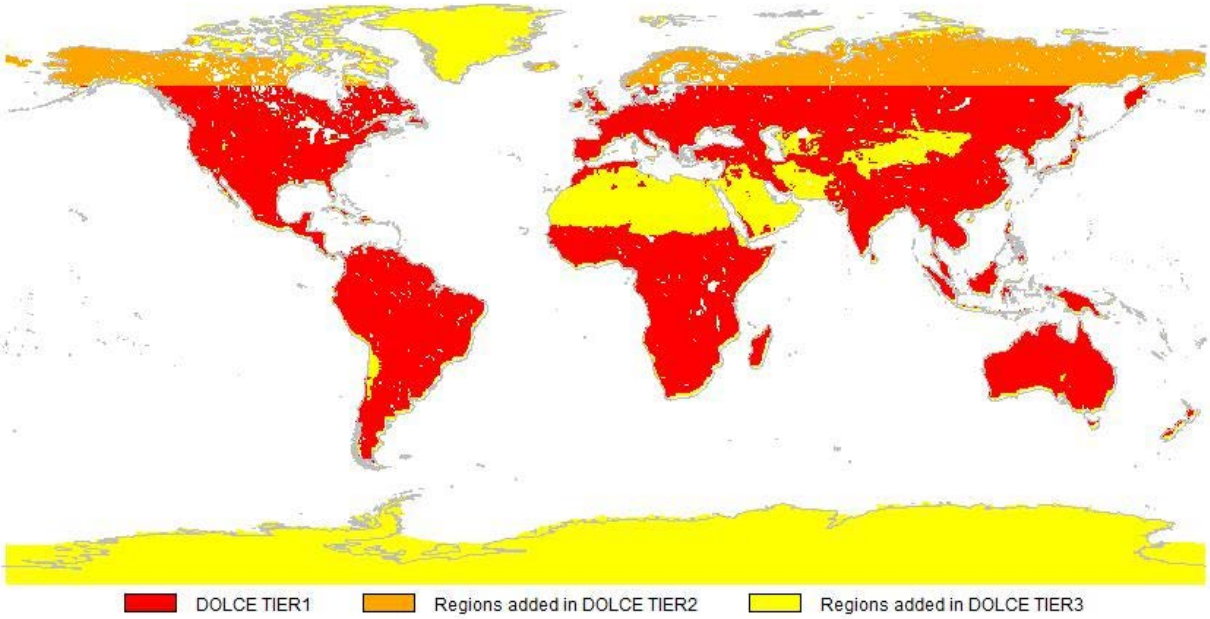

Figure 1: Spatial coverage of DOLCE tier 1 (red), regions added in tier 2 (orange) and tier 3 (yellow). Spatial coverage varies temporally; May 2000 is shown here as a representative example.

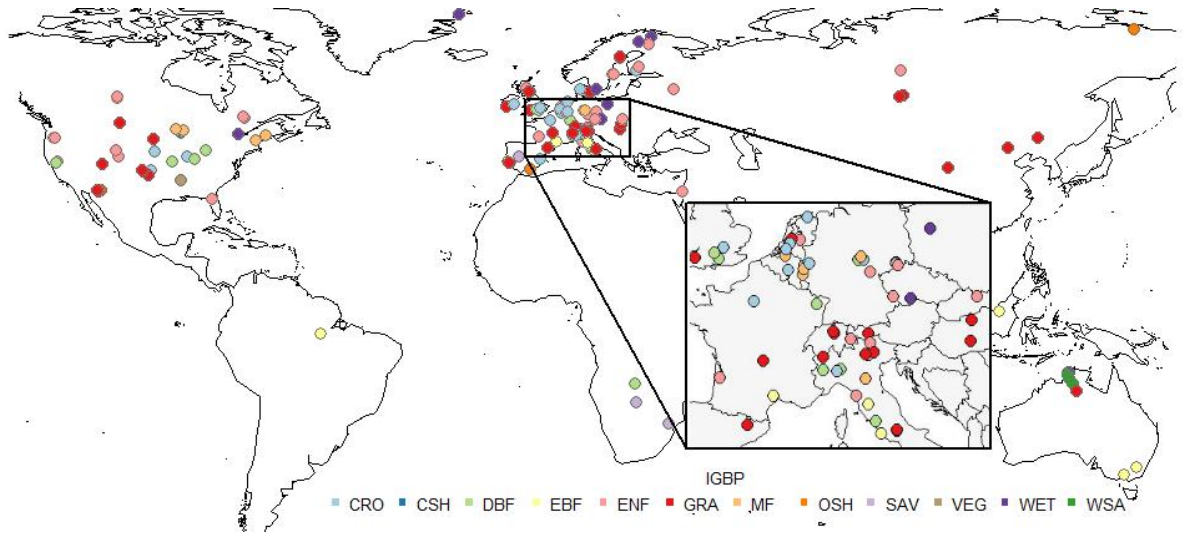

Figure 2: Location of the 159 flux tower sites involved in building DOLCE, classified by vegetation type given by International Geosphere-Biosphere Programme (IGBP): Croplands (CRO), Closed Shrublands (CSH), Deciduous Broadleaf Forest (DBF), Evergreen Broadleaf Forest (EBF), Evergreen Needleleaf Forest (ENF), Grasslands (GRA), Mixed Forests (MF), Open Shrublands (OSH), Savannas (SAV), Cropland/Natural Vegetation mosaic (VEG), Wetlands (WET), and Woody Savannas (WSA).

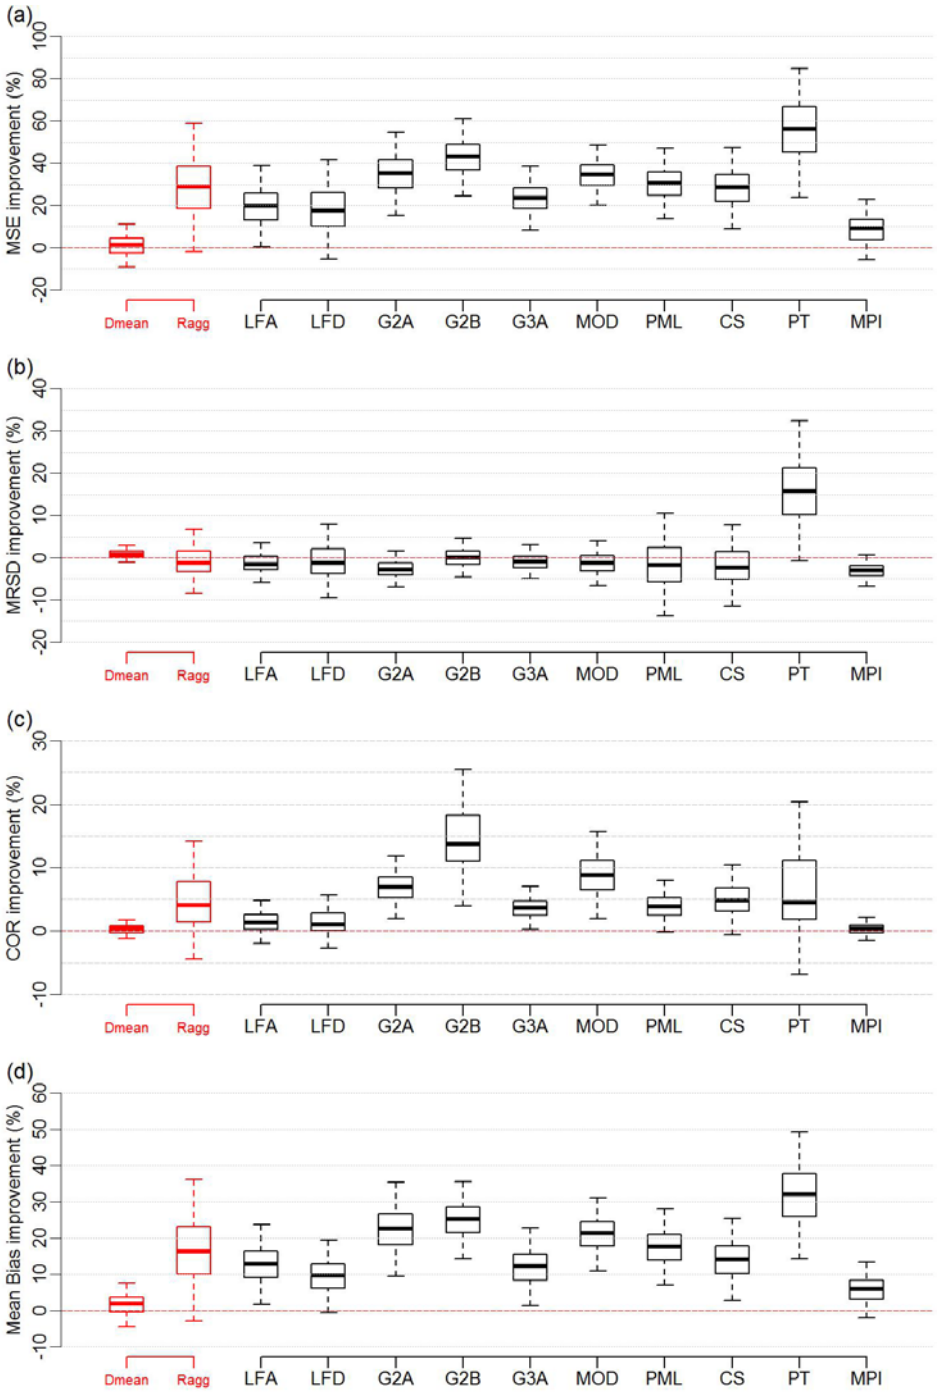

Figure 3: Box and whisker plots displaying the percentage improvement that the weighted product provides in the 25% out-of-sample sites test for ~~three~~ four metrics: MSE (a), MRSD (b), ~~and~~ COR (c) and Mean bias (d), when compared to the equally weighted mean ~~(Wmean)~~ of the Diagnostic Ensemble (Dmean), aggregated Reference Ensemble (Ragg) and each member of the reference

ensemble. Box and whisker plots represents 5000 entries, each entry isare generated through randomly selecting the 25% of sites to be out sample 5000 times.

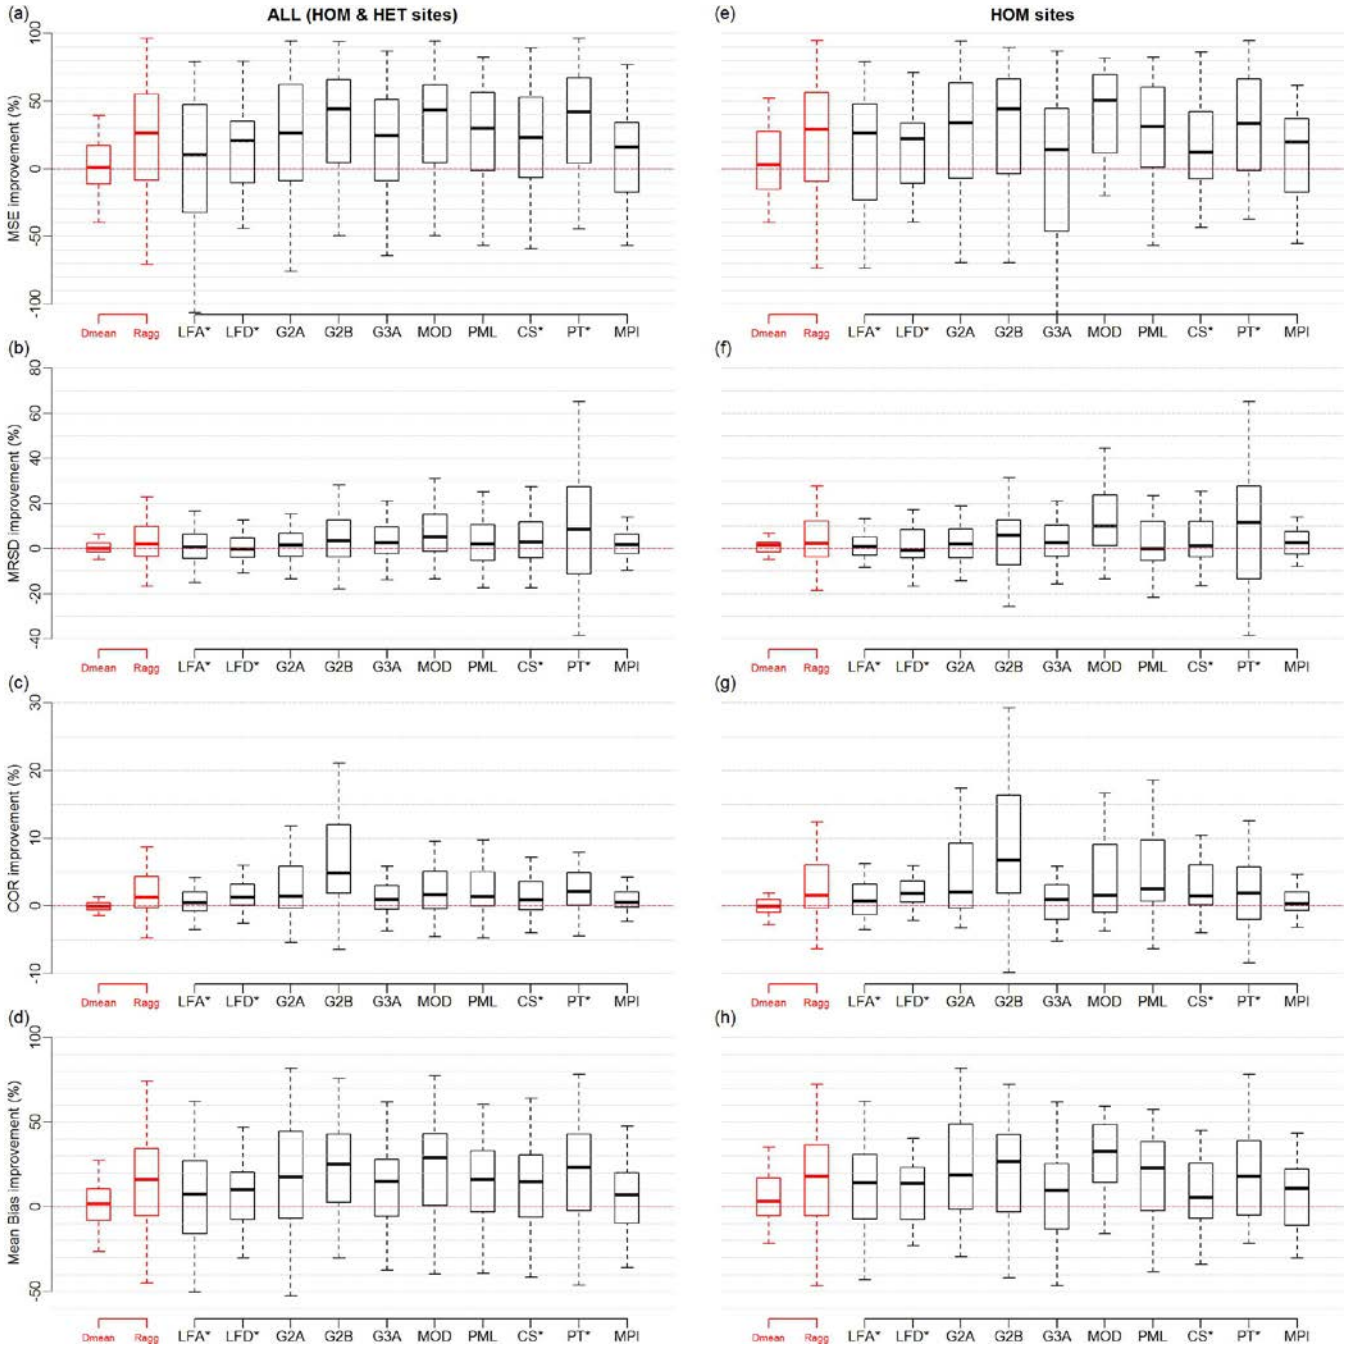

Figure 4: In (a), (b), (c) and (de), as for Fig. 3 but showing the one site out-of-sample tests. Box and whisker plots are generated through selecting one site to be out sample and are repeated for all 138 sites. Products marked with * have limited spatiotemporal availability relative to the diagnostic ensemble, and tTesting against the LFA, LFD, CS and PT products was limited to 110, 108, 108

and 72 sites respectively. ~~due to the limited spatiotemporal availability of these products relative to the Diagnostic Ensemble.~~ In (cd), (fe), (g) and (hf), the one site out-of-sample test is trained by HOM-case sites data only.

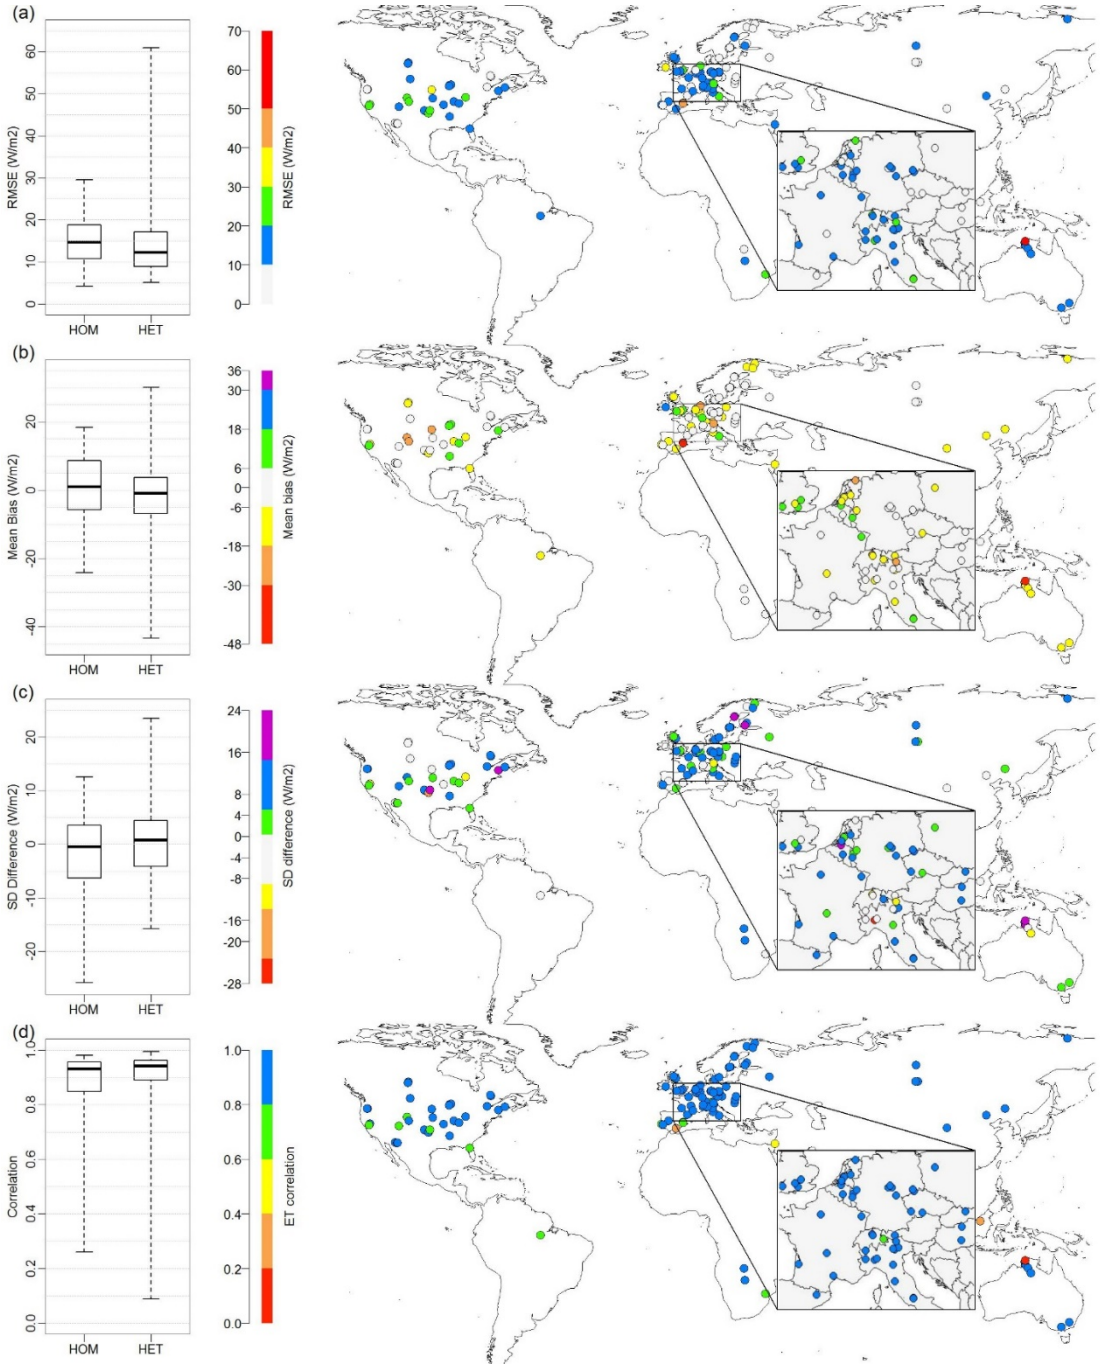

Figure 5: The global maps in the right column show the performance of DOLCE against in-situ measurements at each of the 142 sites by calculating four statistics: (a) RMSE, (b) Mean bias, (c) SD difference and (d) Correlation. The boxplots on the left display the results of these four metrics calculated separately for the sites that are satisfying the HOM-case and the HET-case.

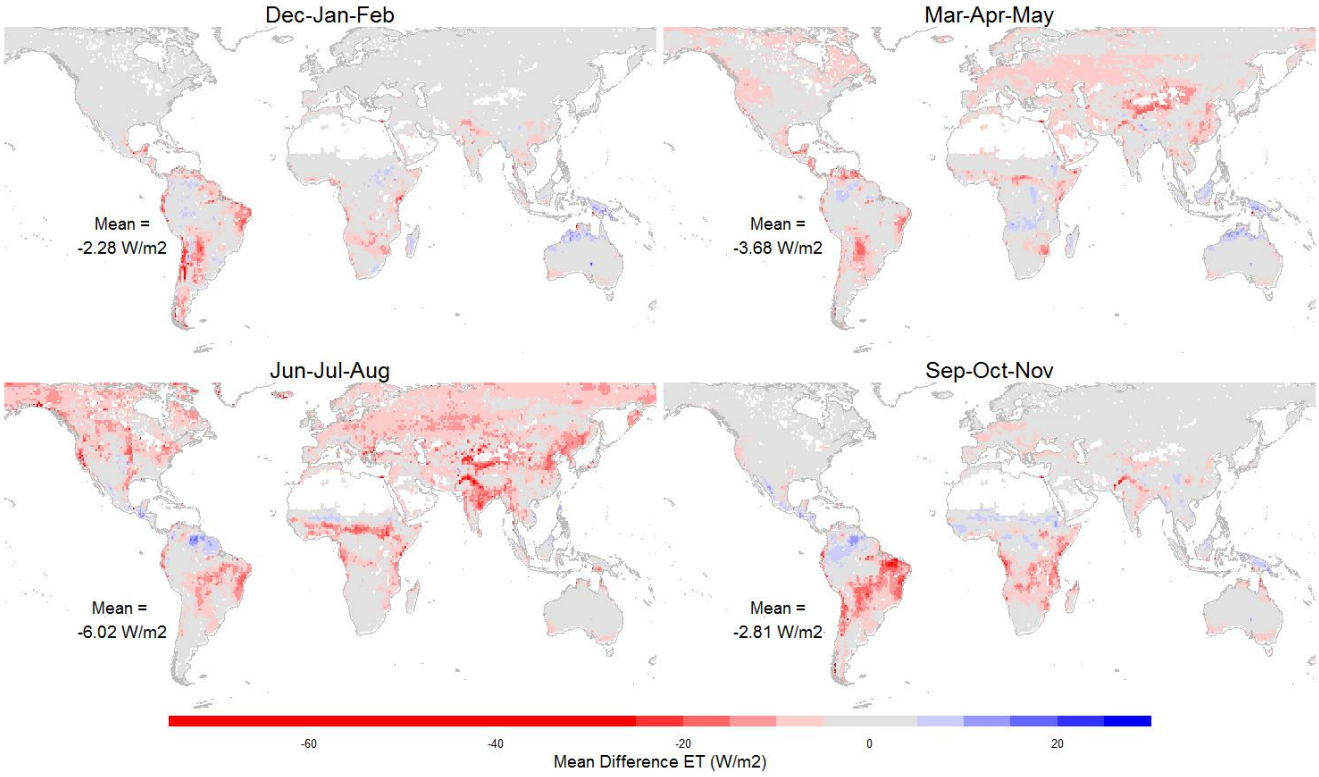

Figure 6: Difference between the seasonal average ET from MPIBGC and DOLCE calculated over the period 2000–2009.

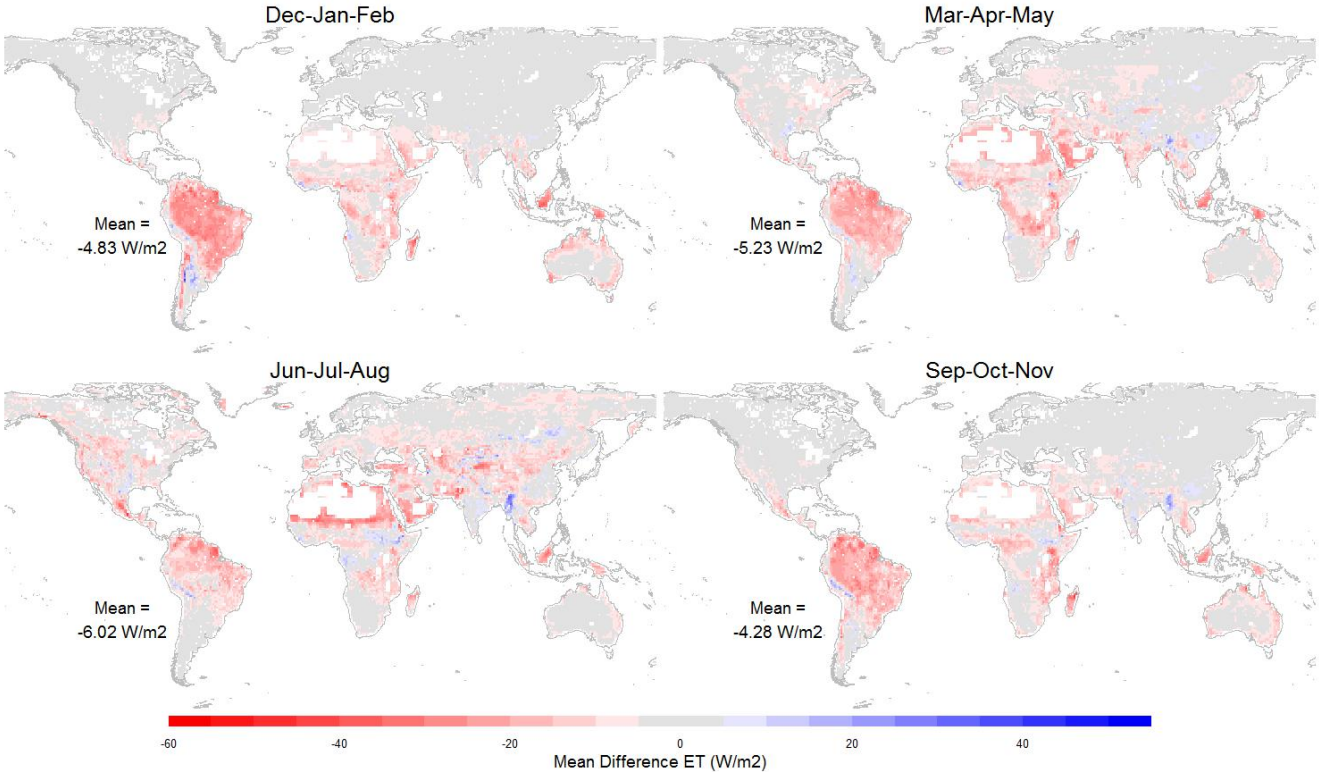

Figure 7: Difference between the seasonal average ET from LandFlux-EVAL-DIAG and DOLCE calculated over the period 2000–2005

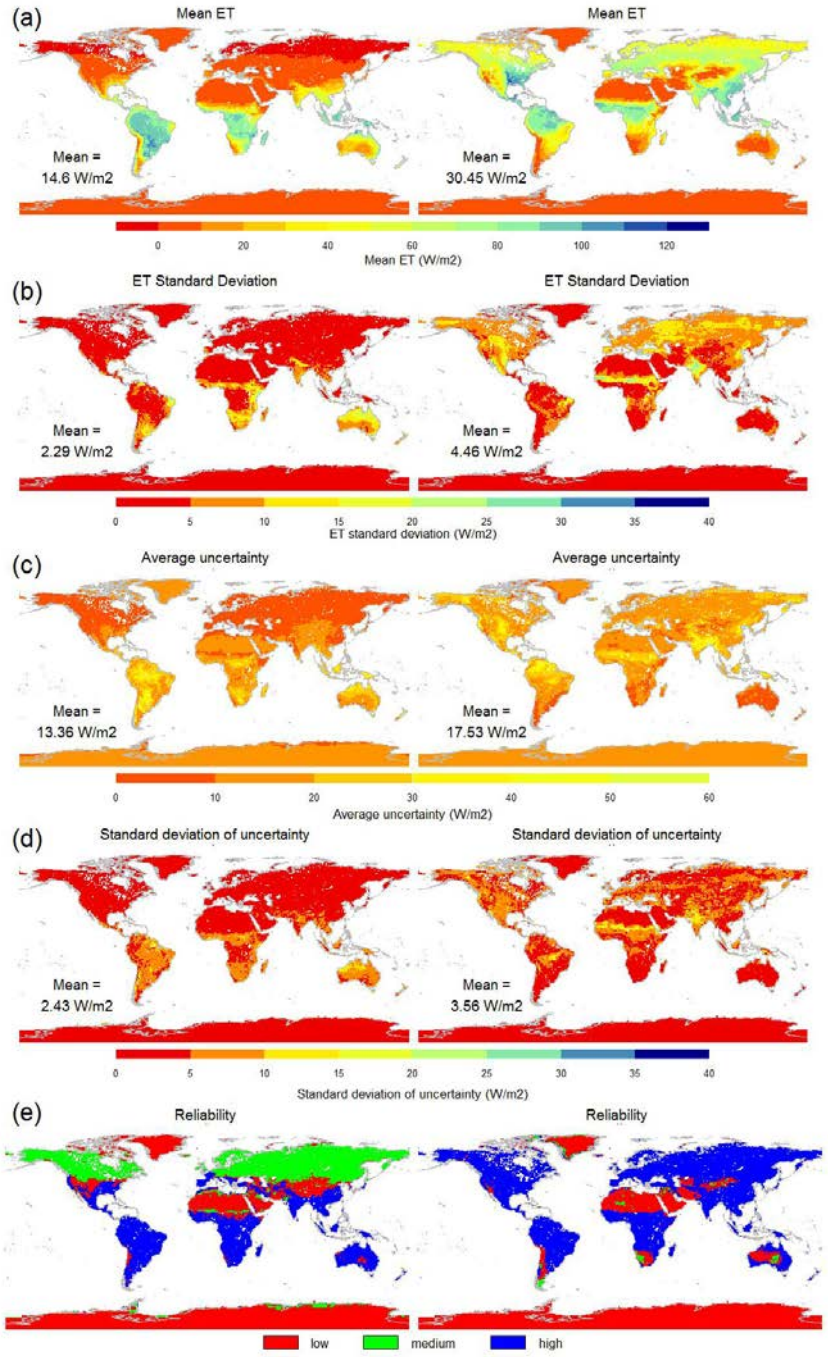

Figure 8: ~~Separate DJF and JJA~~ Seasonal ~~columns showing~~ **(a) global mean ET and (b) its variability (standard deviation), (c) time average of uncertainty (the standard deviation uncertainty $\sqrt{\sigma_e^2}$ shown in Equation 11~~7~~) (d)** ~~uncertainty~~ **standard deviation of uncertainty over time (e) reliability, defined as high ($\frac{Uncertainty\ SD}{mean\ ET} \leq 1$ in blue), medium($|mean\ ET| \leq 5$, $Uncertainty\ SD < 10$ and $\frac{Uncertainty\ SD}{mean\ ET} \geq 1$ in green) and low (in red). DJF is shown in the left column and JJA in the right column.** ~~**Figure 8: (a) Seasonal**~~

~~variability of global mean ET and (b) its associated uncertainty (standard deviation) showing their spatial distribution in Winter and Summer, (c) spatial distribution of the reliability of the uncertainty (standard deviation), showing regions where the reliability of uncertainty estimates is high ($\frac{\text{Uncertainty SD}}{\text{mean ET}} \leq 1$ in blue), medium($|\text{mean ET}| \leq 5$, Uncertainty SD < 10 and $\frac{\text{Uncertainty SD}}{\text{mean ET}} \geq 1$ in green) and low (in red).~~

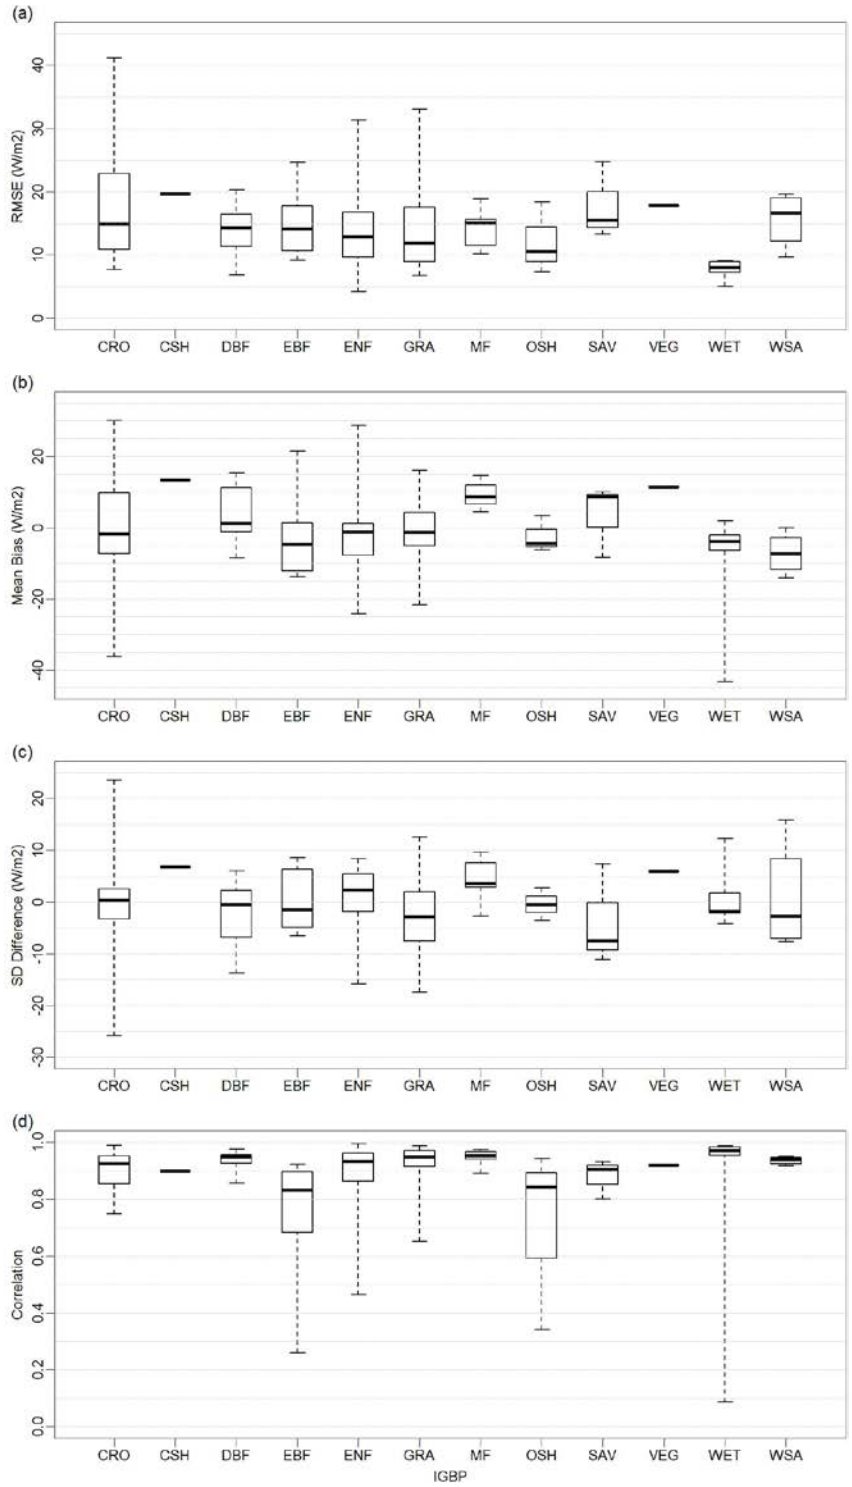

Figure 9: Four statistics, (a) RMSE, (b) Mean bias, (c) SD -difference and (d) Correlation, calculated for DOLCE at 142 flux tower sites and displayed by biome types. See Fig. 2 for biome abbreviations.

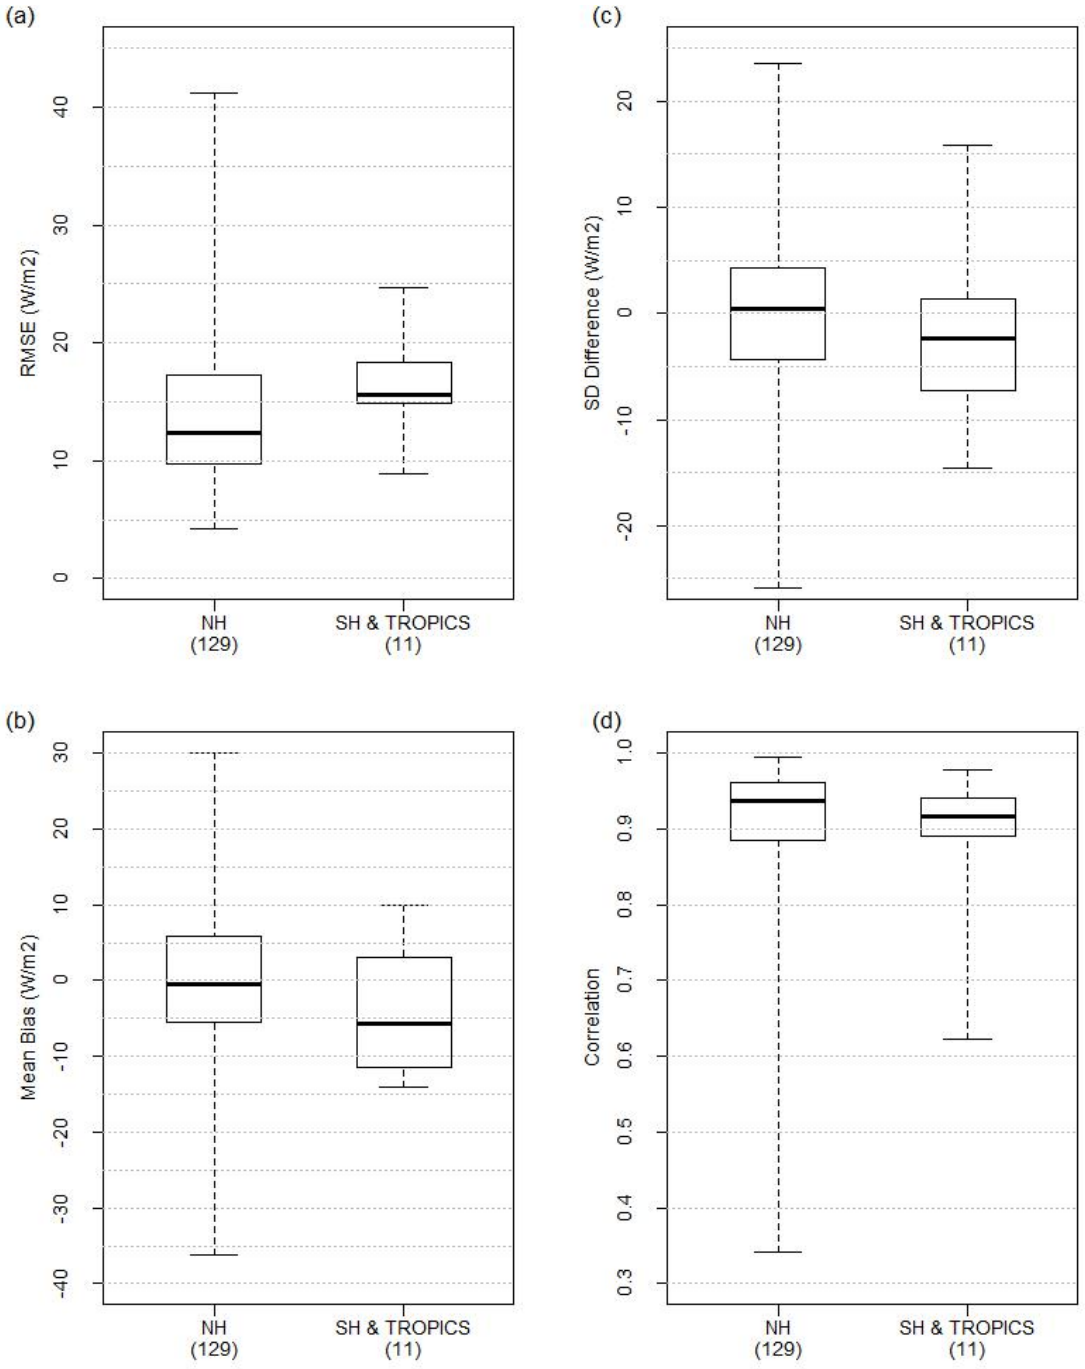

Figure 10: Four statistics, (a) RMSE, (b) Mean bias, (c) ~~SD difference~~SD difference and (d) Correlation, calculated for DOLCE separately at 129 flux towers located at the Northern Hemisphere excluding the tropics (NH) and at ~~the~~ 11 towers in the Southern Hemisphere and the tropics (SH & TROPICS).

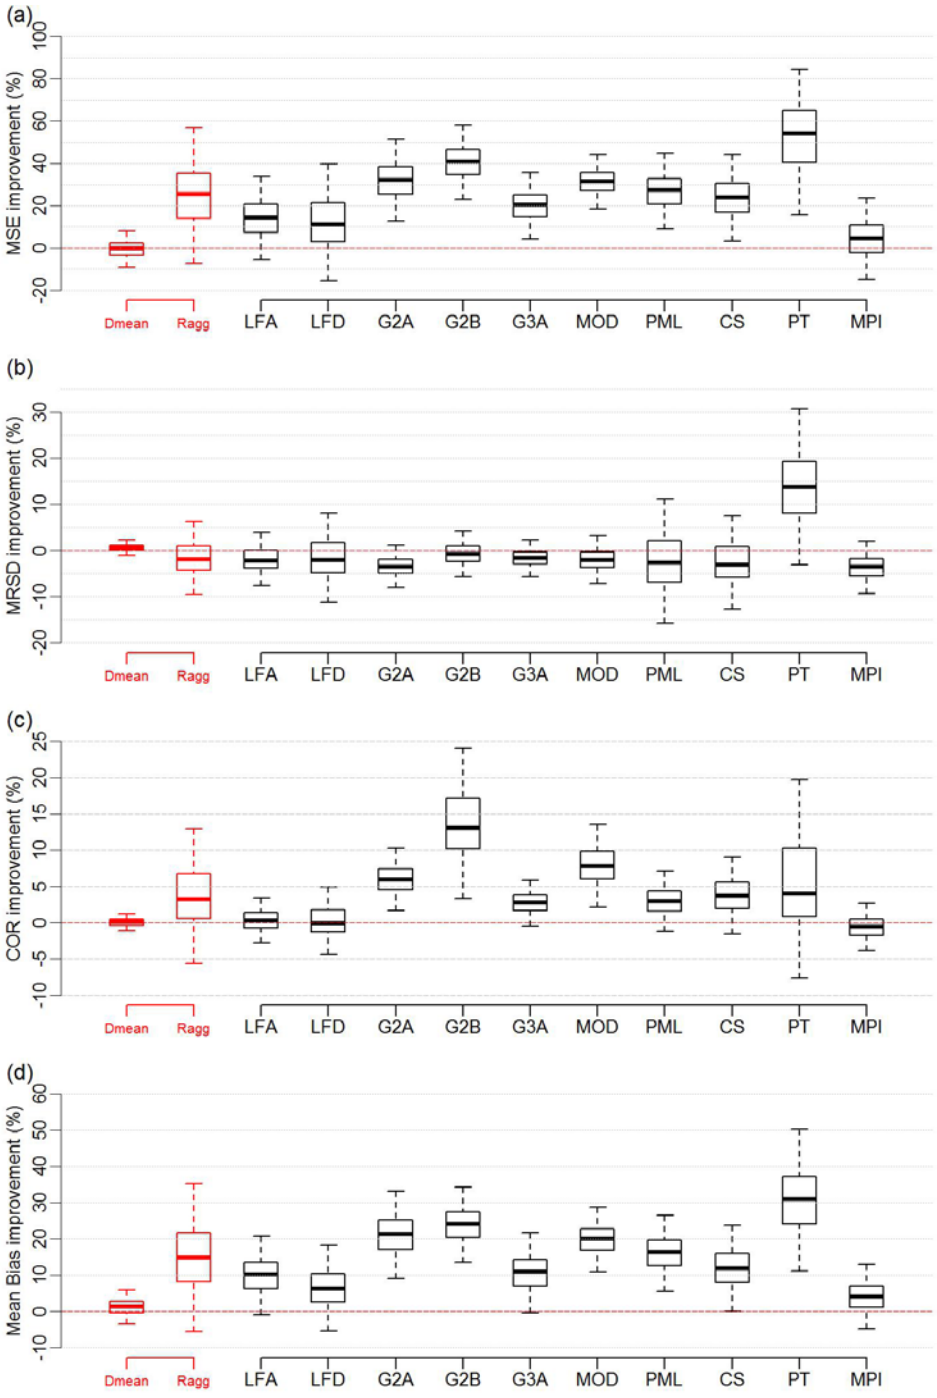

Figure 11: Box and whisker plots displaying the percentage improvement that the weighted product excluding MPIBGC provides in the 25% out-of-sample sites test for four metrics: MSE (a), MRSD (b), COR (c) and Mean Bias (d), when compared to the equally weighted mean (Dmean) of the Diagnostic Ensemble, aggregated Reference Ensemble (Ragg) and each member of the reference

ensemble. Box and whisker plots represents 5000 entries, each entry is generated through randomly selecting 25% of sites to be out sample.

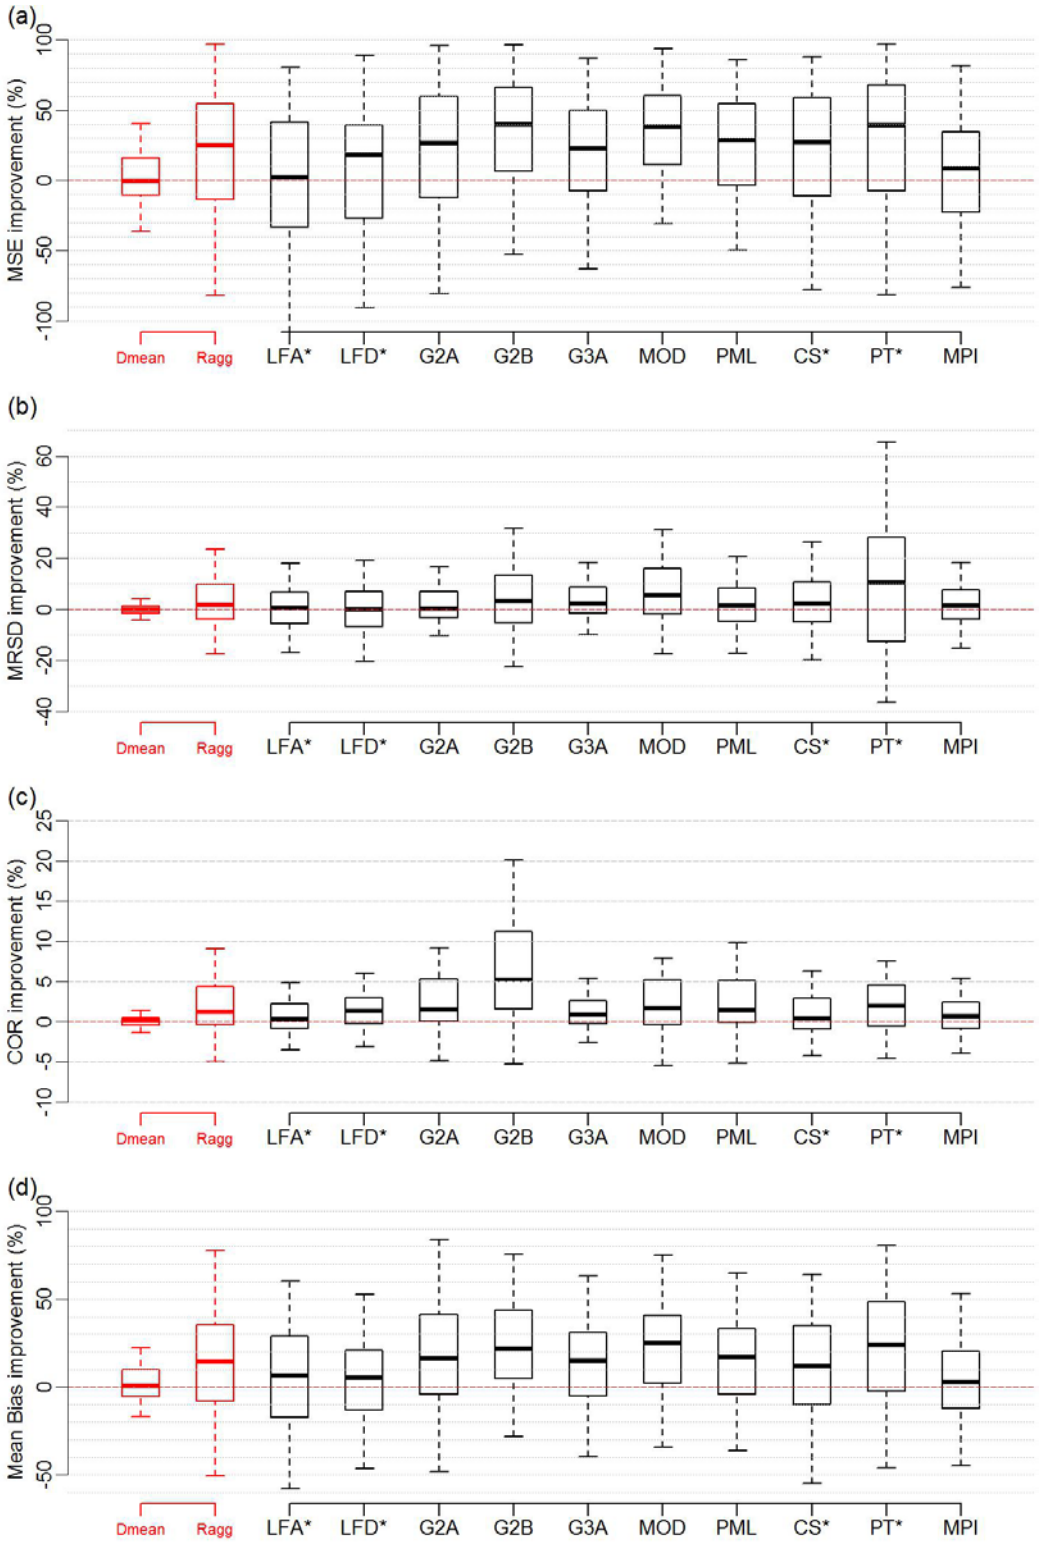
(a)
MSE improvement (%)
100
50
0
-50
-100
Dmean
Ragg
LFA*
LFD*
G2A
G2B
G3A
MOD
PML
CS*
PT*
MPI
(b)
MRSD improvement (%)
60
40
20
0
-20
-40
Dmean
Ragg
LFA*
LFD*
G2A
G2B
G3A
MOD
PML
CS*
PT*
MPI
(c)
COR improvement (%)
25
20
15
10
5
0
-5
-10
Dmean
Ragg
LFA*
LFD*
G2A
G2B
G3A
MOD
PML
CS*
PT*
MPI
(d)
Mean Bias improvement (%)
100
50
0
-50
Dmean
Ragg
LFA*
LFD*
G2A
G2B
G3A
MOD
PML
CS*
PT*
MPI

Figure 12: Box and whisker plots displaying the percentage improvement that the weighted product excluding MPIBGC provides in the one site out-of-sample ~~sites~~ test for four metrics: MSE (a), MRSD (b), COR (c) and Mean Bias (d), when compared to the equally weighted mean (Dmean) of the Diagnostic Ensemble, aggregated Reference Ensemble (Ragg) and each member of the reference ensemble. Products marked with * have limited spatiotemporal availability relative to the diagnostic ensemble, and testing against the LFA, LFD, CS and PT products was limited to 110, 108, 108 and 72 sites respectively.