Appendix S1. Sensitivity to methodological assumptions

Our work focuses on three-dimensional species-species spatial associations. We estimate the location of individual crowns and determine who shades who in the Luquillo Forest Dynamic Plot in Puerto Rico. Crowns are modeled as spheres. The radius of each crown is estimated from the tree DBH through allometries. The location of the crown depends on this allometry for the radius and another allometry for the total tree height (height of the center of the crown = total height of the tree - crown radius). The {x,y} location of the crowns are assumed to be the {x,y} locations of the mapped bases of the trees. *I.e.*, we assume trunks to be perfectly straight. When two trees are close to each other and their crowns overlap, the taller individual adjacent to a shorter individual is credited with the all-or-none overtopping overlap, while the shorter individual is relegated to having no overtopping overlap. In other words, crowns are assumed to be flat horizontal disks when deciding who shades who.

All these assumptions can have consequences for our estimates of aggregated directional overlap between any two species. This document compares the results derived from our assumptions with those derived by alternative assumptions, possibly more realistic, and checks the robustness of our results to error in the allometries that underlie the calculations.

All the alternative assumptions were studied with all the trees in the 16-ha Luquillo Forest Dynamic Plot, at their mapped empirical {x, y} locations.

(1) Noise in heights and crown radiuses

We used local allometries developed for Luquillo. Allometries reflect forest-level averages. We have no reasons to suspect systematic bias. However, we know there is variability between individuals. This can have consequences in any analysis that relies on allometries.

We have not measured the height of all the trees, so we do not know how much noise there is around the central estimate implicit in the allometry, for a tree of a given diameter. For that reason, we simulated the impact of 9 different levels of noise, from $\pm$ 10% to $\pm$ 90% of the "allometric height". In particular, for each tree, and for a given level of noise l, we took a random number x between $-l$ to $+l$, and we added x to the tree height. Then, we re-calculated the location of the centers of the crowns and all pairwise directional overlaps between species. Fig. S1 compares these new results with the overlaps obtained directly from the allometries with no extra noise, as it was done in our manuscript. We proceeded in the same way to study the impact of adding noise to the crown radiuses estimated by the allometries (Fig. S2). Adding random noise has almost no impact on the results, when these are aggregated at the level of pairs of species. We conclude that our results are fairly robust to error in the allometries, particularly to the crown radius allometry.
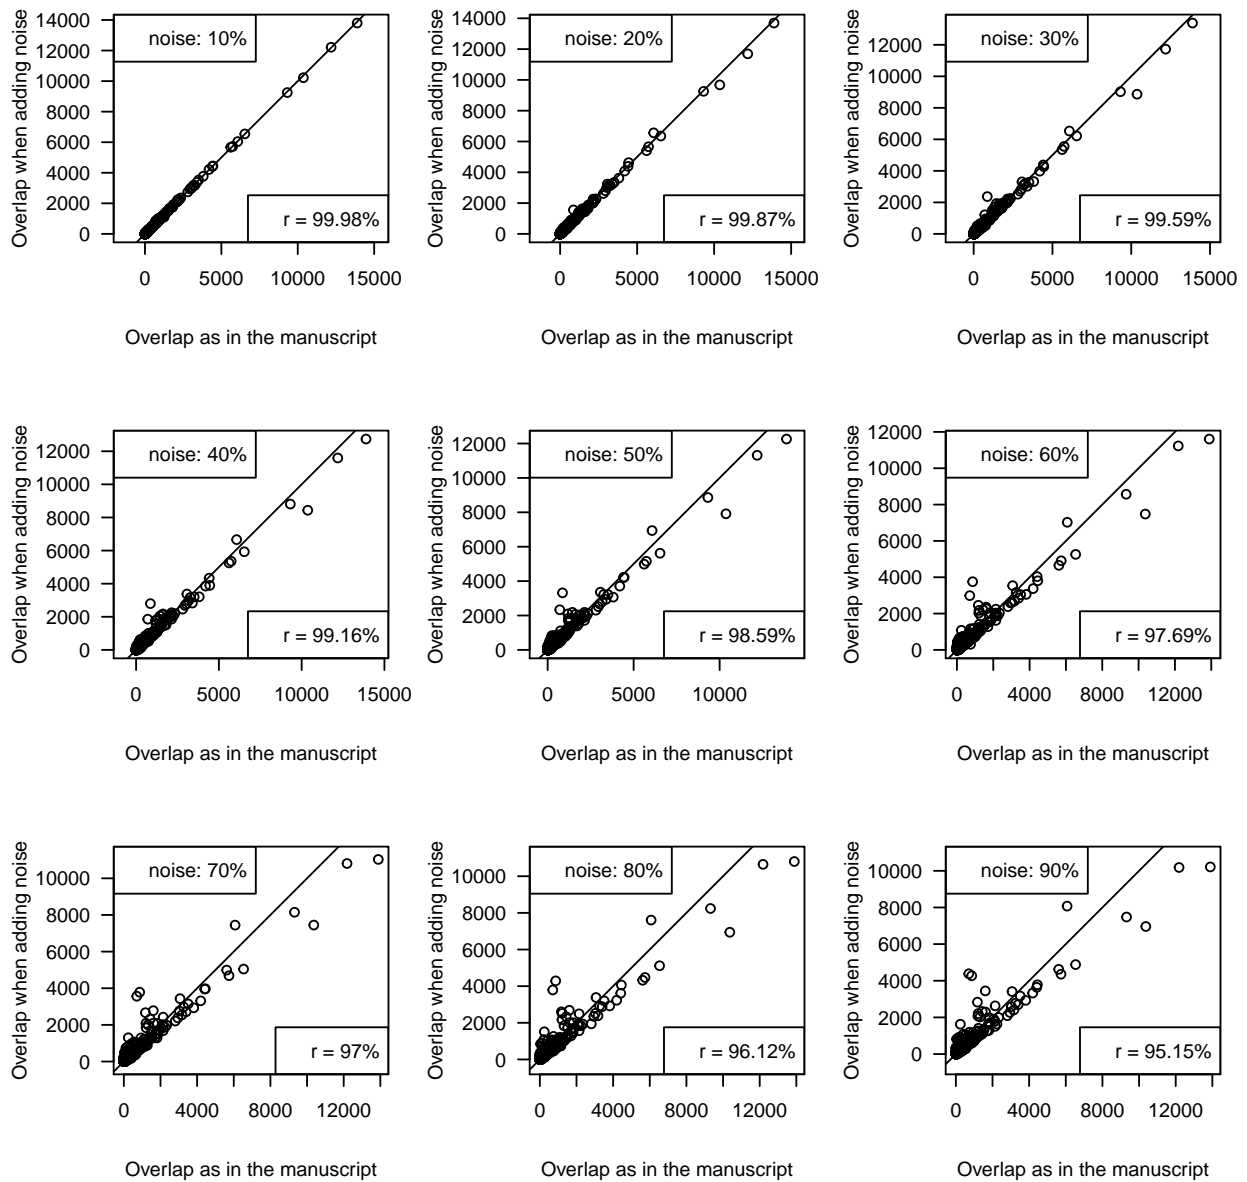
Figure S1. Congruence between results assuming no noise in individual tree heights (exactly what the allometry says) and some level of noise (10%, 20%, ..., 90%) added to the allometric estimation. Each point represents one pair of species. The 1:1 diagonal is shown as a reference of "perfect congruence". The bottom right corner of each panel contains the Pearson's r correlation for the results obtained using both approaches.
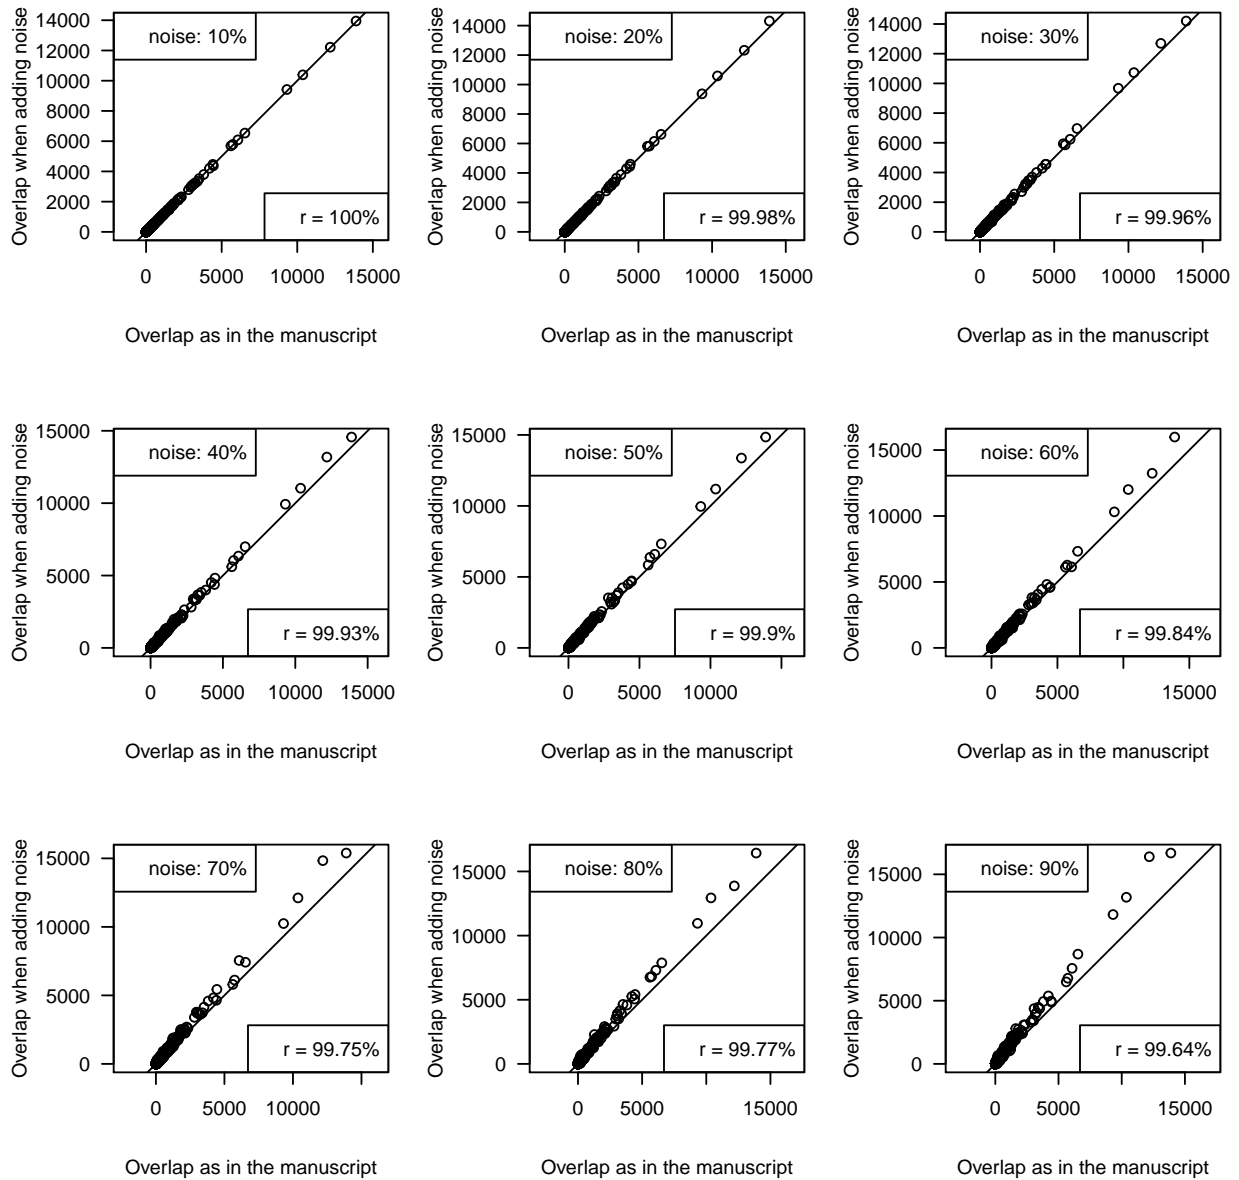
Figure S2. Congruence between results assuming no noise in individual crown radiuses (exactly what the allometry says) and some level of noise (10%, 20%, ..., 90%) added to the allometric estimation. Each point represents one pair of species. The 1:1 diagonal is shown as a reference of "perfect congruence". The bottom right corner of each panel contains the Pearson's r correlation for the results obtained using both approaches.

(2) Noise and bias in {x, y} coordinates

In our analyses, we assume that crowns are located directly above the {x, y} coordinates of the (empirical, mapped) bases of the trunks. However, this is not perfectly realistic, as many trees show some degree of leaning. Here, we explore two ways of leaning: "dumb" leaning, and "smart" leaning. Dumb leaning is just leaning in random directions, *i.e.* noise added to the {x, y} coordinates. Smart leaning, on the other hand, moves crowns towards gaps or available vertical light. Both cases were simulated independently and with different algorithms. However, both were based on the same approximated distribution of degrees of leaning

in tropical forests (Figure 1h in Zuleta et al. *in press*). This distribution reflects degrees of leaning of >30,000 trees of all sizes in several tropical forest sites. It groups trees into 9 levels of leaning: 70% of trees lean 0°-10°, 13% of trees lean 10°-20°, 4% of trees lean 20°-30°, 3% of tees lean 30°-40°; the rest of leaning levels (40°-50°, 50°-60°, 60°-70°, 70°-80° and 80°-90°) contain 2% of the trees each.

In both simulations (dumb leaning and smart leaning), we assigned a specific degree of leaning to each tree by (1) assigning each tree to one level of leaning according to the proportion of trees within each class; and (2) picking one random number within the range of the leaning class.

For the simulation of dumb leaning, we translated the random leaning into distances of horizontal displacement. Then we displaced the crowns horizontally in a random direction. Then we re-calculated all overlaps between individual trees and aggregated the results for each pair of species.

For the simulation of smart leaning, we had to apply a more complicated algorithm. It was not feasible to optimize the direction of leaning/displacement for >26,000 trees simultaneously, so we used an iterative algorithm. First, we ranked all trees from the tallest to the shortest. Second, we re-assigned all the leanings in that same order, so tall trees are straight and short trees lean more. This gives more freedom to shorter trees to move around. Then, starting with the tallest tree, we adjusted the direction of leaning/displacement to maximize the exposure to vertical light. We iterated through all the trees and adjusted the direction of leaning one at a time, evaluating exhaustively all possible directions of leaning (360°). In this process, the adjustment of the target tree takes into account all the (already adjusted) positions of all the crowns taller than the target tree; the crowns of shorter individuals cannot interfere with its exposure to vertical light. This iterative algorithm guarantees that the total accumulated exposure to vertical light is maximized in the forest, subject to the constraints (empirical {x,y} locations in Luquillo and a realistic distribution of trunk leaning in tropical forests). As in the other case, after displacing the crowns, we re-calculated all overlaps between individual trees and then aggregated the results for each pair of species.
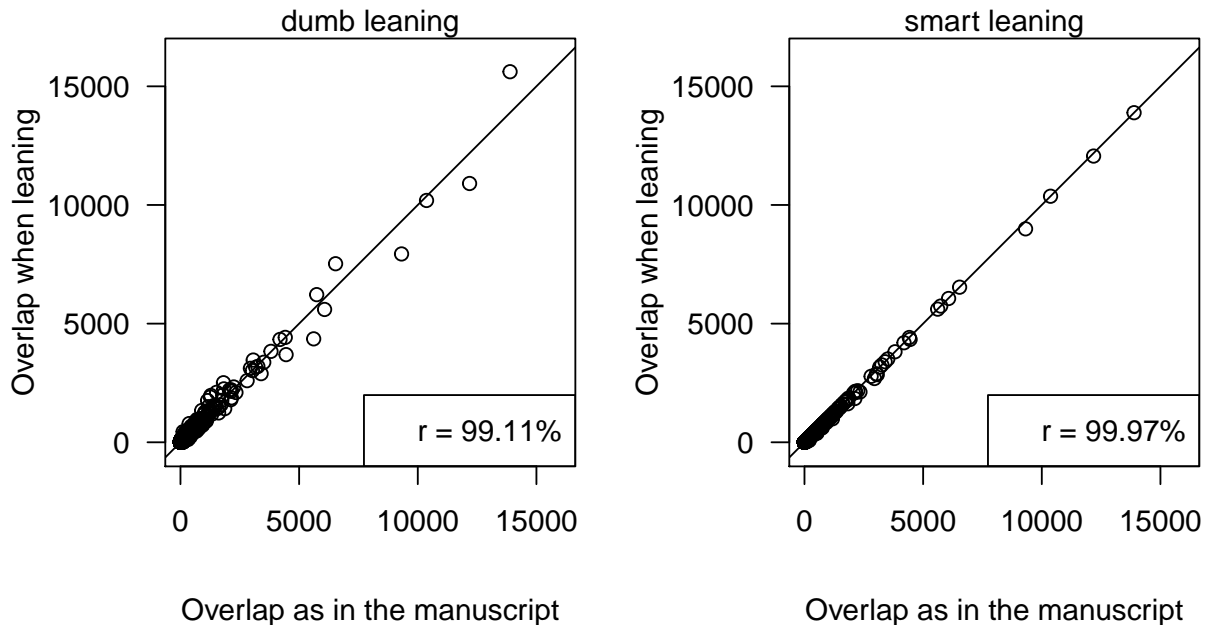
Figure S3. Congruence between results assuming crowns located exactly above the base of the tree (horizontal axis) and more realistic leaning. The left panel includes results for leaning in random directions ("dumb leaning"). The right panel includes results for leaning in optimized directions that maximize the total exposure to vertical light ("smart leaning"). Each point represents one pair of species. The 1:1 diagonal is shown as a reference of "perfect congruence". The bottom right corner of each panel contains the Pearson's *r* correlation for the results obtained using both approaches.

4

Moving the crowns in dumb or smart directions barely changes the aggregated results (Fig. S3). Most overlaps between pairs of species are very similar to those obtained under the assumption of crowns directly above the mapped {x,y} coordinates. Interestingly, the "smart leaning" algorithm has almost no impact whatsoever, suggesting that mapped {x,y} crown locations already represent highly optimized locations, in terms of exposure to vertical light.

Note that simulating realistic levels of leaning just means displacing a bit the {x,y} coordinates of the center of the crowns. This process is the same that one would use to simulate asymmetrical crown growth (dumb or smart) and moderate error in mapping the bases of the trees. Therefore, our results also apply to this other sources of error. Overall, we can conclude that our results are robust to error in {x, y} locations.

(3) Impact of the flat disk assumption

We defined overlap of i over j as the amount of i that can be seen from above within the area of horizontal overlap between the crowns of i and j. The flat disk assumption consist on assigning all the overlap to the tallest of i and j. We used this simplification for computational reasons: the area of non-directional overlap between two circles (= flat projections of two spheres) can be calculated in a vectorized way, and then a simple `ifelse()` statement in R assigns the direction: ifelse(A taller than B?, yes $\rightarrow$ give all to A, no $\rightarrow$ give all to B). A more nuanced model splits the overlap into some portion of i that can be seen from above within that area of overlap and another portion of the area that corresponds to j seen from above within that area of overlap (Fig. S4). This more correct approach is computationally intensive, and it is not an option for our analyses because of the 999 instances of the null model. Anyway, the differences between both approaches are very small when aggregating all the overlaps between individuals into overlaps between species (Fig. S5).
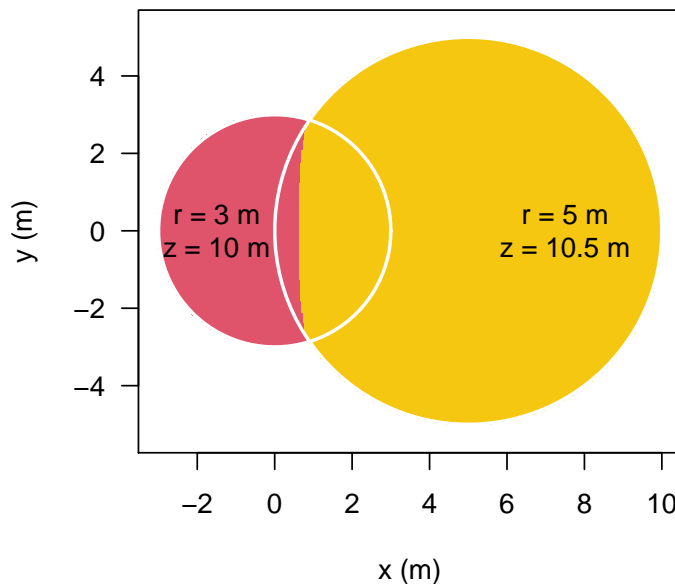

Figure S4. Two colliding spheres seen from above. The flat disk assumption would assign all the area of overlap to the larger sphere, which is not completely correct. The text within the spheres indicate their radius (r) and the height of their centers (z).

5
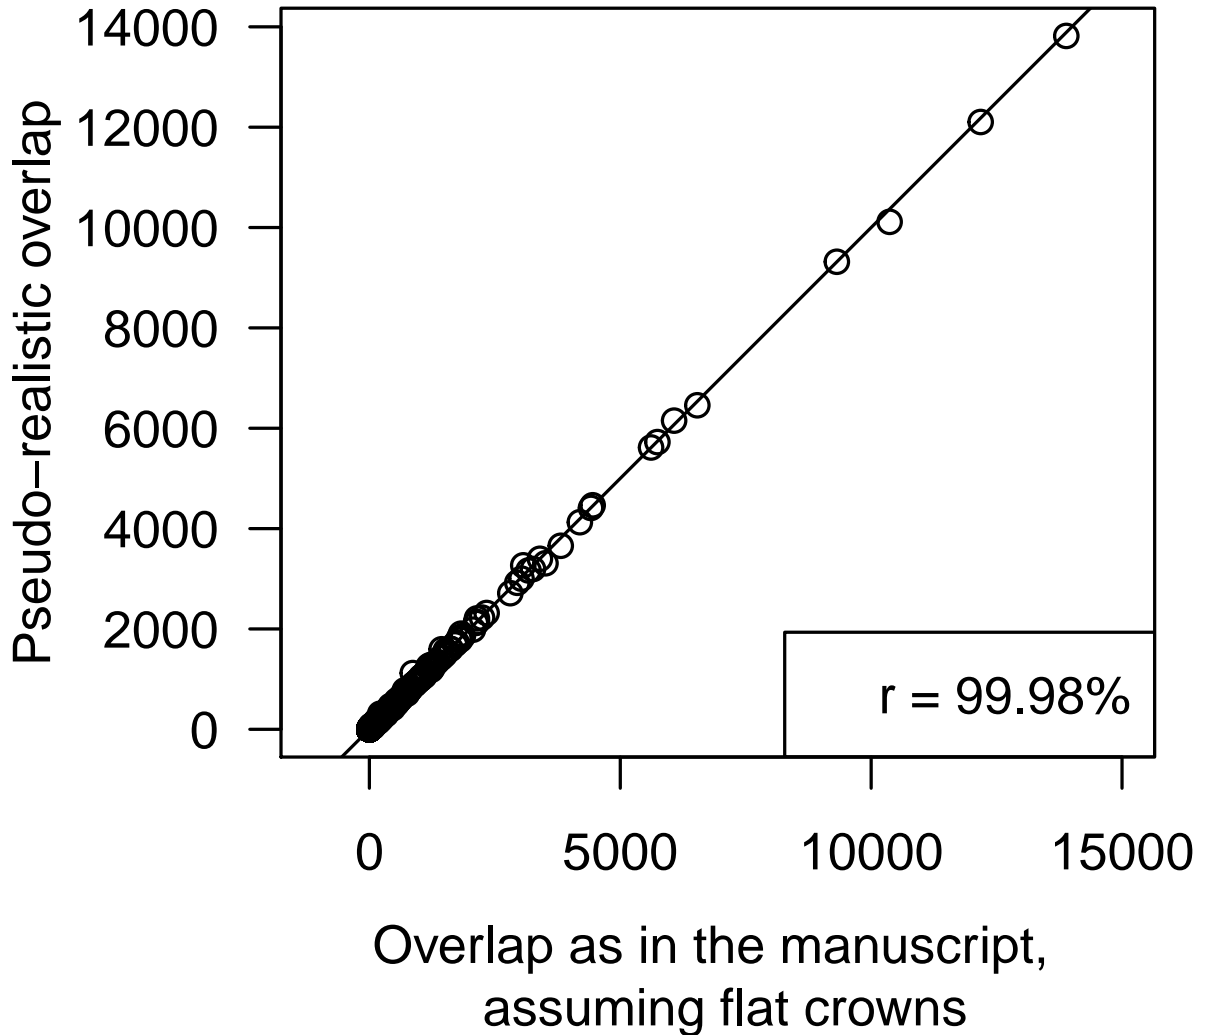
Figure S5. Congruence between results assuming flat crowns and a more realistic model of sphere collision. Each point represents one pair of species. The 1:1 diagonal is shown as a reference of "perfect congruence". The bottom right corner of each panel contains the Pearson's r correlation for the results obtained using both approaches.

We can visualize both approaches to get a more intuitive feeling of their similarities and differences. Fig. S6 contains pseudo-realistic (the best guess) aerial view of Luquillo, compared with the estimated aerial view under the flat disk assumption, and the differences between both approaches. They coincide in 95.72% of the surface. Only when zooming in to specific areas within the plot we can easily observe the differences between both approaches. Overall, we conclude that our results are robust to the flat disk assumption.
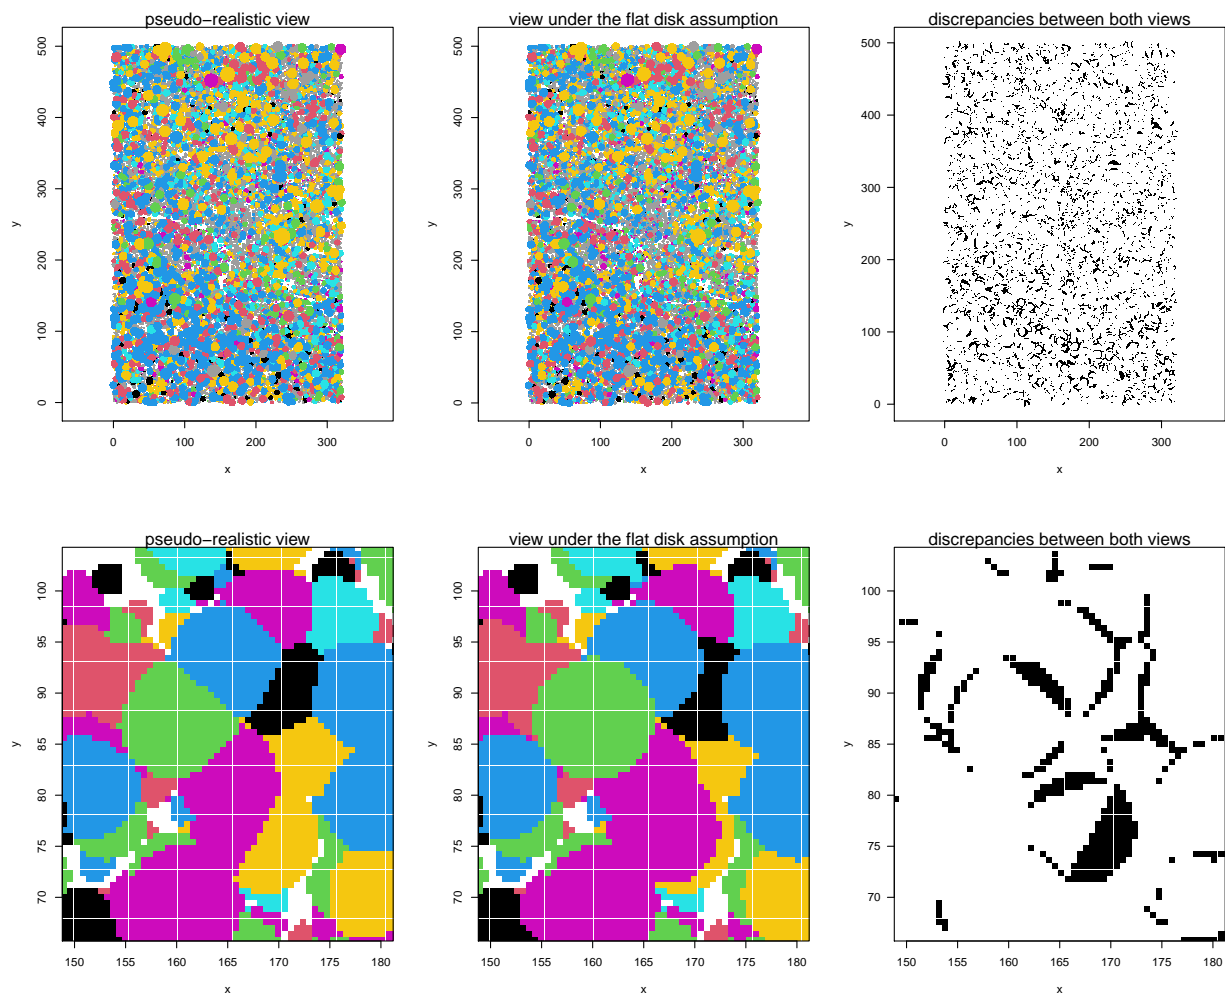
Figure S6. Estimated aerial view of the Luquillo Forest Dynamic Plot in Puerto Rico. The left panels show a pseudo-realistic view that takes into account the three-dimensionality of spheres when they collide. The panels in the center use the flat disk assumption: when two crowns overlap, it assumes that we can see the crown of the tallest tree only. The panels in the right shows the points where both approaches differ in which individual can be seen from above. The panels in the top are views for the 16 hectares, and are colored by species identity. The panels at the bottom take a closer look to few individuals, and are colored by individual identity.

References

Zuleta, D., G. Arellano, H.C. Muller-Landau, S.M. McMahon, S. Aguilar, S. Bunyavejchewin, D. Cardenas, C.H. Chang-Yang, A. Duque, D. Mitre, M. Nasardin, R. Perez, I.F. Sun, T.L. Yao, S.J. Davies. *In press.* Individual tree damage dominates mortality risk factors across six tropical forests. *New Phytologist.*