Fig. S1. Empirical distributions of the component metrics A) *S*; B) *C.I. range*; and C) R^2_{BG} for the analysis of resting metabolic rate in Great Cormorants (raw data thinned using `thinData(CormorantData, by=3)`, and a call to `rankLocReg` using `alpha=0.1`). Note the extremely long tail in the distribution of R^2_{BG} (panel C), and the large discrepancy between the $L_\%$ benchmark and the other metrics.
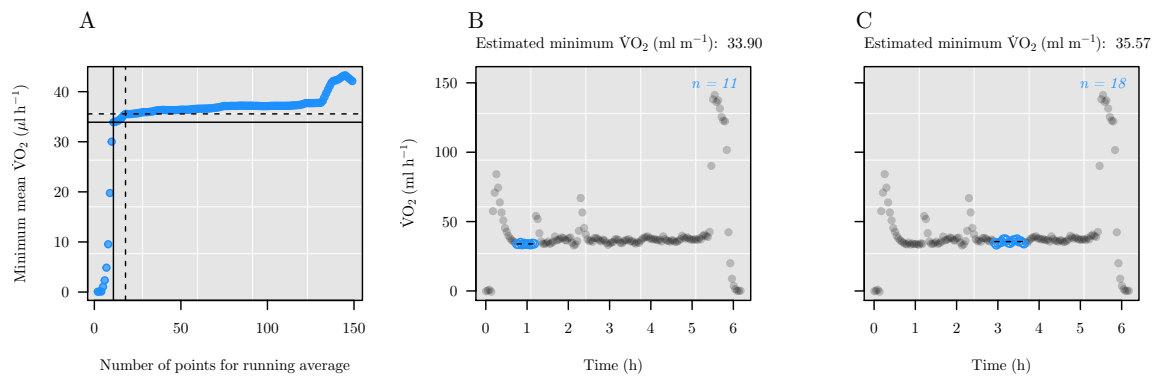
Fig. S2. A graphical summary of conventional methods of estimating resting metabolic rate from $\dot{V}O_2$ time series data (after Withers, 2001). Panel A shows the minimum estimated $\dot{V}O_2$ as a function of the number of adjacent observations used to calculate the running average. After a rapid increase in the minimum $\dot{V}O_2$ with increasing observations up to $n = 11$ (solid black cross), there is a brief region with an intermediate slope, before the plot plateaus at $n = 18$ (dotted cross). Using these methods, a researcher could potentially justify using the running average associated with either of these two points as an estimate of resting metabolic rate. Panels A and B show the full time series, with the subsets associated with $n = 11$ and $n = 18$ respectively highlighted in blue. Note that both subsets are much smaller than the number of observations included in the $L_\%$ rank 1 local regression identified by `rankLocReg` (Fig. 3).
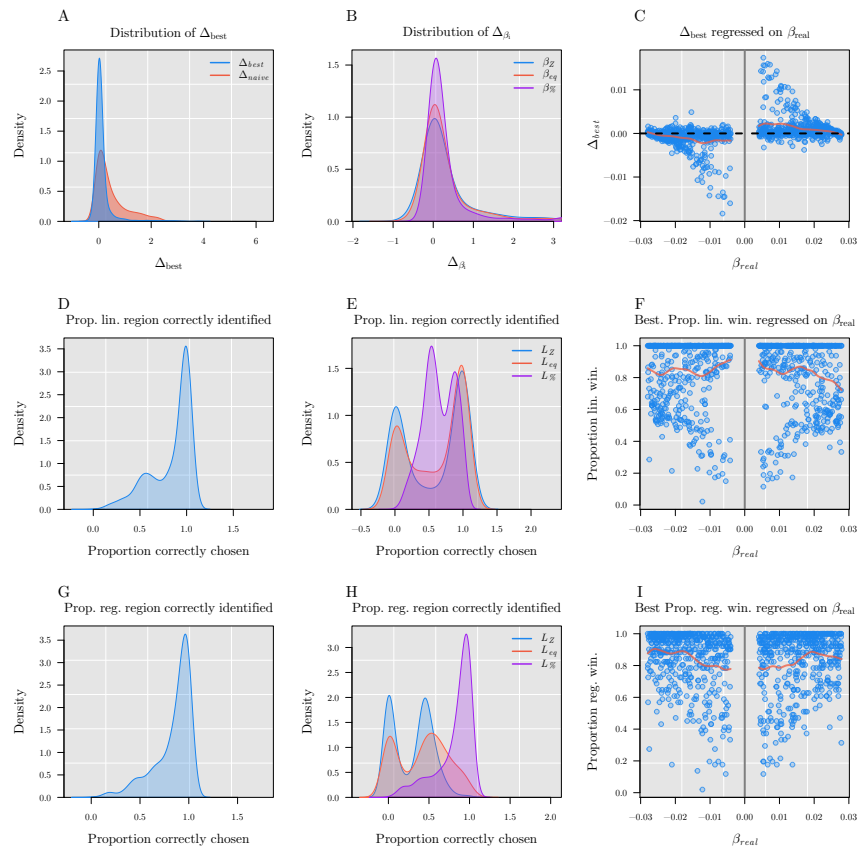
Fig. S3. Graphical summary of the performance of `LoLinR` methods using certain types of simulated data. Panels A-C show results for the performance metric Δ; panels D-F show results for the proportion of the linear subset of the data that are correctly included in the local regressions identified by `rankLocReg`; and panels G-I show results for the proportion of the local regression identified by `rankLocReg` that correctly includes the linear subset of the simulated data. The first column of panels (A,D,G) show the distributions for each performance metric of the 'best' local regression identified by `rankLocReg` (the local regression with the smallest Δ). Note that the x-axes are scaled in units of β_{real}; thus, in panel A, a Δ value of 2 indicates that Δ_{best} was twice as large as the real slope of the simulated data (β_{real}. Panel A shows a comparison between the performance of `rankLocReg` and a naive regression by comparing the distributions of Δ_{best} and Δ_{naive}. As expected, the slopes identified by naive regression are systematically biased by the non-linearity present in the simulated data, while `rankLocReg` is better able to identify slopes that are closer to the β_{real}. The second column of panels (B,E,H) show the distributions of each performance metric for each of the three L metric methods provided in `rankLocReg`. The third column of panels (C,F,I) show scatter plots of each performance metrics regressed on the actual slope of the simulated data (β_{real}), with smoothing splines overlaid to help visualize any trends. This is to visualize that, as expected, `rankLocReg` performs better as the slope of the linear region is further from 0. Note that in panel B long right tails have been truncated to better visualize the peaks of the distributions

Click here to Download Script S1

Click here to Download Script S2