World Journal of Educational Studies
Print ISSN: 2959-9989
Online ISSN: 2959-9997
DOI: https://doi.org/10.61784/wjes3219

CASE-DRIVEN REFORM OF APPLIED STATISTICS TEACHING THROUGH QUANTILE REGRESSION OF FOREST TREE GROWTH

HaiYan Gu*, QiZhi Zhang, WenYan Wang, MeiLing Xu
Northeast Forestry University, Harbin 150040, Heilongjiang, China.
**Corresponding Author: HaiYan Gu*

Abstract: In graduate applied-statistics courses, knowing how to run a regression does not necessarily prepare students to choose a model for a forestry problem or explain what its estimates mean. We describe a case-based teaching plan built around the height–diameter relationship in a mixed larch–birch stand. The source project reports measurements for 677 trees in four plots, including diameter at breast height, total height, live-crown height, and crown dimensions. The proposed course asks students to examine data quality, fit ordinary least squares (OLS) and quantile regressions, and compare findings across species and plots with different management histories. By working through 19 conditional quantiles (0.05–0.95), students can see where an average relationship leaves important variation unexplained. The teaching sequence includes preparatory exercises, collaborative analyses, and a management-oriented briefing, supported by an assessment rubric. We distinguish results mentioned in the project materials from estimates that students would need to verify with the original observations. In particular, differences between thinned and untreated plots cannot, by themselves, establish an effect of thinning. This article sets out a teaching design rather than reporting an intervention trial; classroom learning outcomes remain to be evaluated. The case illustrates how statistical computation can be taught alongside ecological interpretation, reproducibility, and caution in drawing conclusions.
Keywords: Applied statistics; Quantile regression; Forestry education; Case-based learning; Tree height; Teaching reform

1 INTRODUCTION

A familiar difficulty in applied-statistics teaching appears when students leave a worked example and encounter an unfamiliar dataset. They may know the commands for fitting a regression but still be uncertain about which observations belong in the analysis, what a coefficient represents, or whether the model addresses the question posed. Forestry makes these decisions particularly concrete. Tree dimensions vary with species, stand conditions, competition, management history, and the way sample plots have been selected. A numerical estimate cannot be interpreted well without first identifying the trees being compared and the assumptions behind that comparison.

Height and diameter at breast height are routinely recorded in forest inventories. A conventional regression can describe mean height at a given diameter, which is often a useful starting point. Yet trees with similar diameters do not necessarily occupy similar positions in a stand's height distribution. Quantile regression provides a way to examine relationships at different parts of that conditional distribution. Forestry examples include studies of height–diameter relationships, branch characteristics, and basal-area increment [1-4]. Koenker's review discusses the broader rationale for looking beyond the conditional mean [5]. In the classroom, this creates a meaningful choice rather than a competition between two techniques: students must explain what the mean model answers and what, if anything, a quantile analysis adds.

Our teaching material draws on the project Individual-Tree Growth Based on Quantile Regression, intended for a Forestry Statistical Cases component of a professional master's course. The Chinese source describes a case and a proposed teaching approach; it does not report a controlled evaluation of student learning. We therefore treat the material as the basis for an instructional design, retaining its stated forestry setting and measurements without presenting proposed activities as completed experiments.

The course revolves around three practical questions. What must a student establish before turning a forest-management concern into a statistical analysis? How should OLS and quantile-regression results be compared without mistaking software output for interpretation? And what evidence should accompany a recommendation to a forestry practitioner? The teaching plan moves from understanding the site and measurements to model comparison, group critique, and a brief report that states both findings and their limits.

2 CASE CONTEXT AND TEACHING CHALLENGES

The project concerns a mixed larch–birch forest at Cuigang Forest Farm in the Greater Khingan region. According to the source description, four sample plots yielded 677 valid tree records. Variables include diameter at breast height, total tree height, height to the live crown, and crown width measured in four directions. Some plots had undergone thinning, whereas others were described as untreated. These details offer enough context to formulate realistic statistical

© By the Author(s) 2026, under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0).

questions. They do not provide the tree-level dataset needed to reproduce model estimates, so the case must be taught with that distinction made explicit.

Several features make the setting useful for a graduate class. Species differ, plot histories differ, and tree height may respond to diameter in ways that are not adequately summarized by one line. The source proposes estimates at 19 quantiles, from 0.05 through 0.95 in increments of 0.05, and presents selected curves at 0.10, 0.25, 0.50, 0.75, and 0.90. Students can use these comparisons to ask whether the estimated diameter association changes across the conditional distribution. They should not interpret a low conditional quantile as a particular tree age or an upper-quantile slope as a stand-wide average.

The same case also exposes errors that are easy to miss in a routine modeling assignment. A strong fitted association does not show that one tree characteristic causes another. A short tree at a given diameter need not be young. Pooling different species and plots can hide differences in biological relationships and can overstate the amount of independent information in the sample. Finally, an appealing coefficient graph may draw attention away from uncertain estimates, influential observations, extrapolation, or differences between plots that existed before thinning. Rather than treating these points as a final checklist, instructors can bring them into discussions whenever teams propose a model or defend a result.

3 REFORM OBJECTIVES AND FRAMEWORK

By the end of the case, students should be able to move between a forestry question and a defensible statistical answer. This requires more than recalling definitions. Students need to explain how a conditional mean differs from a conditional quantile and why the quantile loss function leads to a different fitted relationship. They should be able to assess measurement quality, motivate baseline and alternative specifications, examine uncertainty, and describe the scope of a conclusion. A further objective is to write for a forest manager without turning an observed association into a claim about what thinning caused.

The activities are organized around the decisions an analyst would face. Students begin with a forest-inventory question and identify the variables and comparisons required to investigate it. They prepare a data dictionary and analysis plan before estimating OLS and quantile models. Peer review then asks each team to defend its choices, including why a seemingly reasonable alternative was rejected. The final product is a short technical briefing in which the model output supports an argument; the coefficients are not treated as an end in themselves.

We propose three stages: orientation, laboratory analysis, and a synthesis seminar. During orientation, students learn the necessary forestry terminology and try both estimators on a small practice dataset. In the laboratory, they work with the case records if access is authorized, maintaining a log of their cleaning and modeling decisions. At the seminar, they compare interpretations, discuss limitations, and present recommendations in language accessible to practitioners. This structure draws on the contextual reasoning emphasized in case-method teaching [6], but its effectiveness in this particular course has not yet been tested.

4 STATISTICAL WORKFLOW AS A TEACHING CASE

Teams begin by writing down the observational unit, outcome, predictors, and groups they intend to compare. Total height H and diameter D offer a straightforward first model; live-crown height C provides a second response for discussion. Before estimating anything, students check measurement units, missing values, repeated identifiers, implausible ranges, and unusual observations. Plots of the raw measurements should be examined separately for each species and management category. Crown width may be explored as another predictor, but a decision to include or exclude it should follow the scientific question rather than a significance test conducted after looking at the data.

The mean-regression baseline expresses tree height H as a function of diameter D and a grouping indicator G. Quantile regression instead describes the corresponding association at a chosen conditional quantile, with parameters estimated by minimizing asymmetric check loss. Equations (1)–(4) give the models and objectives. Depending on the number of observations within each category, teams can also consider species-specific fits or interactions. The purpose of this comparison is to connect the estimators' different objective functions to the different interpretations of their coefficients [5].

$$E(H|D,G) = \alpha + \beta D + \gamma G \tag{1}$$

$$Q_\tau(H|X) = X^{\mathrm{T}}\boldsymbol{\beta}(\tau), \quad 0 < \tau < 1 \tag{2}$$

$$\widehat{\boldsymbol{\beta}}(\tau) = \operatorname{argmin}_{\mathbf{b}} \sum_{i=1}^{n} \rho_\tau\left(H_i - X_i^{\mathrm{T}}\mathbf{b}\right) \tag{3}$$

$$\rho_\tau(u) = u[\tau - \mathbb{I}(u < 0)] \tag{4}$$

After fitting the OLS model, teams first examine quantile models at 0.10, 0.25, 0.50, 0.75, and 0.90. They then review the full set of 19 quantile-specific estimates. Tables and figures must identify predictor units, sample sizes, observed diameter ranges, and uncertainty where it can be estimated appropriately. An important discussion concerns what the slopes do not mean. A steeper upper-quantile relationship does not imply that older trees grow more rapidly, nor that every large tree is on a faster growth trajectory. If all measurements were collected at one time, the analysis describes relationships among dimensions, not observed changes in those dimensions over time. This qualification matters because the project title uses the word 'growth'.

One preliminary example in the source materials gives a diameter coefficient of roughly 89.6 cm in fitted height per additional centimetre of diameter for an upper-quantile birch subgroup in untreated plots. Students should identify the exact subgroup, units, and model before discussing this figure. Without tree-level data, we cannot check the value or

attach a new confidence interval or significance test to it. The source also describes relatively shallow diameter relationships for some thinned larch groups and flatter live-crown-height relationships at middle and higher quantiles. These observations invite ecological hypotheses, but they cannot establish that thinning produced the differences.

There is also an issue of replication. The case includes 677 trees but only four plots; trees within a plot may share environmental conditions and management history. Treating every tree as an independent replicate of a plot-level treatment would therefore be misleading. With only four clusters, conventional cluster-robust standard errors may be unreliable. Teams can examine each plot separately, compare the stability of results when plots are omitted, and explain why broad claims about thinning are not justified. For predictive exercises, a held-out test should reflect the intended use: if predictions are to be transferred to another plot, separating training and validation observations by plot is more informative than a random split of individual trees. In-sample fit alone is not enough.

More advanced analyses may consider nonlinear diameter terms, crossings between fitted quantile curves, and changes in results under alternative specifications. Students should avoid predictions outside the observed predictor range and examine whether species composition or initial stand differences could account for apparent management contrasts. Where suitable validation data exist, quantile-specific pinball loss and interval calibration provide ways to compare predictions. Teams interested in mixed-effects or hierarchical models can investigate the relevant literature [2-3] while recognizing that four plots leave little information for reliable plot-level inference. These extensions should be chosen to answer a stated question, not simply to make the model more elaborate.

5 CLASSROOM IMPLEMENTATION AND ASSESSMENT

The opening stage occupies two class periods. A short forest-inventory briefing introduces measurements, species-specific height–diameter relationships, and the difference between association and a causal claim. Students then estimate OLS and median regression on an instructor-supplied practice dataset and compare what the fitted lines summarize. A short diagnostic task includes incorrect statements—for example, the claim that the 0.90 conditional quantile identifies 90-year-old trees. Students explain the error in their own words before moving to the larger case.

Three or four laboratory periods are allocated to the main analysis. When permission allows, teams receive de-identified tree records, variable definitions, contextual notes, and a starter script. Each team checks the data, keeps a record of exclusions, estimates comparable models, and prepares a coefficient-profile figure with understandable labels. The instructor does not supply a single 'correct' cleaned file or a fixed sequence of commands. Instead, checkpoints require students to explain choices as they are made. Peer review pairs a team presenting its analysis with another team asked to identify hidden assumptions and unclear explanations.

At the final seminar, teams present a two-page briefing to a hypothetical forest manager and respond to questions from classmates. The briefing distinguishes what the measurements show, what the models might predict, and which ecological explanations remain speculative. Discussion can focus on whether the thinned and untreated plots were comparable, whether tree age was measured, and whether a cross-sectional coefficient can legitimately be described as a growth rate. Students also identify what additional data would be needed before making a causal recommendation. The instructor closes by connecting these specific decisions to statistical work in other applied fields.

We propose assessing the work students produce rather than relying on expressions of satisfaction with the course. The rubric assigns 20% to problem definition and domain knowledge, 20% to data checking and reproducibility, 25% to model development and diagnosis, 20% to interpretation and discussion of uncertainty, and 15% to presentation and reflection. Four performance levels distinguish an incomplete attempt from a well-supported, independently defensible analysis. For example, recognizing that plots differ is an initial step; checking how those differences affect estimated relationships and explaining the consequences is stronger evidence of learning. Feedback is provided after the initial quiz, data audit, and draft figures so teams can revise their work. Each team must also describe at least one model it decided not to use and justify that decision. The instructor can then trace the final recommendation back through the figures, estimates, variable choices, and original question. These are proposed classroom procedures, not observations of learning gains already achieved.

A later implementation could be evaluated with comparable transfer tasks administered before and after instruction, blind scoring of students' work on unseen datasets, reproducibility checks, and reflective accounts of modeling decisions. A parallel class or repeated cohorts might offer additional context, although differences in student selection would remain a concern. Reports should include effect sizes and uncertainty rather than relying on p-values alone. Because the project documentation contains no completed classroom outcome assessment, the present article makes no claim that the reform has already raised achievement.

Table 1 Case-Based Learning and Assessment Blueprint

Stage	Core student output	Assessment evidence
Preparation	Research question; baseline model; diagnostic quiz	Conceptual accuracy and unit interpretation
Integrated laboratory	Data audit; OLS/quantile code; coefficient profiles	Reproducibility, diagnostics, and uncertainty
Synthesis seminar	Management memo; peer critique; oral defense	Ecological reasoning and defensible claims
Future evaluation	Pre/post transfer tasks; blinded scoring	Measured change with transparent limitations

Table 1 brings together the tasks and evidence that could be collected at each stage. Its final row describes a possible later evaluation, not a completed classroom study.

6 DISCUSSION

The teaching design is distinctive chiefly in the questions it asks students to confront. They must determine what a dataset permits them to say before interpreting numerical output. They also see the difference between a conditional average and other parts of the height distribution in a setting where that difference has an ecological meaning. Finally, they must translate a reproducible analysis into advice for someone who does not work with regression models. In combination, these tasks put judgment alongside computational competence.

The aim is not to persuade students that quantile regression is always preferable to OLS. A manager interested in expected height may have good reasons to choose a mean model. Questions about unusually short or tall trees at a given diameter may call for a model that represents more of the conditional distribution. Evidence from published forest applications likewise suggests that model choice varies with species, competitive environment, response, and prediction goal [1-4, 7-8]. Students should learn to explain their estimator in relation to the decision at hand rather than its technical complexity.

Several limitations shape how this article should be read. The source project supplies descriptions and figures, but the present materials do not contain machine-readable tree observations for independent re-estimation. Nor do they document random allocation to thinning, repeated observations of growth, measured ages, or assessed learning outcomes from an implemented class. The four plots also provide a narrow basis for generalization about management effects. For these reasons, this is an instructional proposal supported by an exploratory forestry example, not a demonstration of treatment effects or educational effectiveness. Making these boundaries visible is itself part of the proposed lesson: some questions remain unanswered because the necessary evidence has not been collected.

7 CONCLUSION

The reported measurements for 677 trees provide a concrete setting for teaching model selection and interpretation to applied-statistics graduate students. In the proposed sequence, students compare OLS with quantile regression, inspect species and plot differences, document their analytical choices, and prepare a briefing grounded in the evidence available. The main educational purpose is to help students recognize heterogeneity and uncertainty without overstating what cross-sectional data can establish. Implementation in a future course should preserve the analysis files and assessment records and use prespecified learning measures so the value of the approach can be evaluated empirically.

COMPETING INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

FUNDING

Supported by the Professional Degree Graduate Teaching Case Development Project, 'Individual-Tree Growth Based on Quantile Regression'.

REFERENCES

[1] Ning J, Wang G, Xu N, et al. Study on growth of basal area increment of individual trees in broad-leaved secondary forest based on quantile regression model. Journal of Central South University of Forestry & Technology, 2023, 43(12): 24-34. DOI: 10.14067/j.cnki.1673-923x.2023.12.003.

[2] Miao Z, Widagdo F R A, Dong L, et al. Prediction of branch growth using quantile regression and mixed-effects models: An example with planted Larix olgensis Henry trees in Northeast China. Forest Ecology and Management, 2021, 496: 119407. DOI: 10.1016/j.foreco.2021.119407.

[3] Miao Z, Li F, Zhao X, et al. Utilizing quantile regressions to predict vertical distribution of branch size in Larix olgensis Henry. Forest Ecosystems, 2025, 13(1): 100300. DOI: 10.1016/j.fecs.2025.100300.

[4] Wang Y, Zhao W, Zhang Q, et al. Predicting tree height-diameter relationship from relative competition levels using quantile regression models for Chinese fir (Cunninghamia lanceolata) in Fujian Province, China. Forests, 2020, 11(2): 183. DOI: 10.3390/f11020183.

[5] Koenker R. Quantile regression: 40 years on. Annual Review of Economics, 2017, 9: 155-176. DOI: 10.1146/annurev-economics-063016-103651.

[6] Rosier G. The case method evaluated in terms of higher education research: A pilot study. The International Journal of Management Education, 2022, 20(3): 100660. Article identifier: S1472811722000623. https://www.sciencedirect.com/science/article/pii/S1472811722000623.

[7] Zea-Camano J D, Soto J R, Arce J E, et al. Improving the modeling of the height-diameter relationship of tree species with high growth variability: Robust regression analysis of Ochroma pyramidale. Forests, 2020, 11(3): 313. DOI: 10.3390/f11030313.

[8] Chiou C R, Cheng C P, Yang S I. Comparison of parametric and nonparametric methods for estimating size-density relationships in old-growth Japanese cedar plantations. Forests, 2020, 11(6): 625. DOI: 10.3390/f11060625.