



Overview

- Motivation: Trees and leaves
- Methodology
 - Model estimation
 - Tests for parameter instability
 - Segmentation
 - Pruning
- Applications
 - Costly journals
 - Beautiful professors
 - Choosy students
- Software

Model-Based Regression Trees in Economics and the Social Sciences

Achim Zeileis

<http://eeecon.uibk.ac.at/~zeileis/>

Motivation: Trees

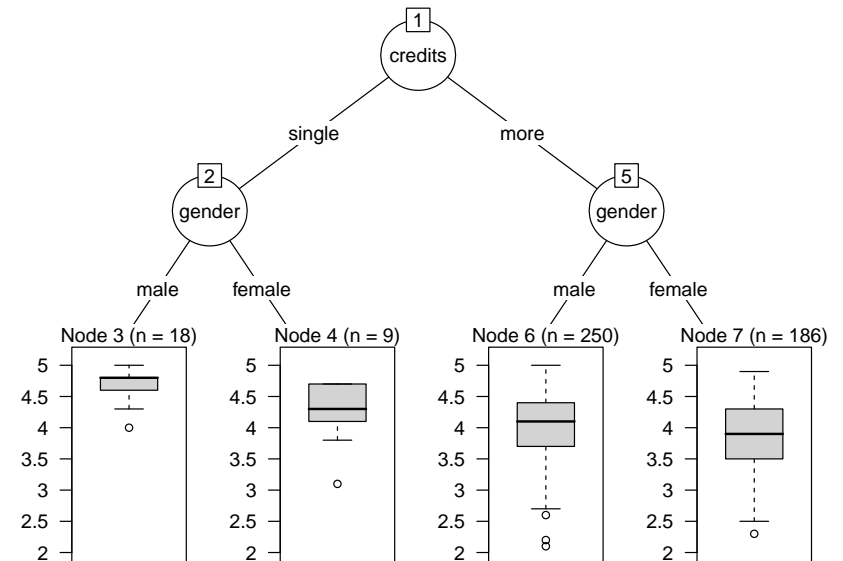
Breiman (2001, *Statistical Science*) distinguishes two cultures of statistical modeling.

- **Data models:** Stochastic models, typically parametric.
 - Predominant strategy in economics and social sciences. Regression models are workhorse for empirical analyses.
- **Algorithmic models:** Flexible Models, data-generating process unknown. → Few applications in social sciences and especially in economics.

Examples: Recursive partitioning, decision trees. Class of flexible methods for classification and regression.

Illustration: Average evaluation of professors (on a scale 1–5) by gender and type of course.

Motivation: Trees



Motivation: Trees

Formally: Modeling of dependent variable Y by “learning” a recursive partition w.r.t explanatory variables Z_1, \dots, Z_l .

Key features:

- 1 Predictive power in nonlinear regression relationships.
- 2 Interpretability (enhanced by visualization), i.e., no “black box” methods.

Examples: CART and C4.5 in statistical and machine learning, respectively.

Recursive partitioning

Base algorithm:

- 1 Fit model for Y .
- 2 Assess association of Y and each Z_j .
- 3 Split sample along the Z_{j^*} with strongest association: Choose breakpoint with highest improvement of the model fit.
- 4 Repeat steps 1–3 recursively in the subsamples until some stopping criterion is met.

Here: Segmentation (3) of parametric models (1) with additive objective function using parameter instability tests (2) and associated statistical significance (4).

Motivation: Leaves

Typically: Simple models for univariate Y , e.g., mean.

Idea: More complex models for multivariate Y , e.g., multivariate normal model, regression models, etc.

Here: Synthesis of parametric data models and algorithmic tree models.

Goal: Fitting local models by partitioning of the sample space.

1. Model estimation

Models: $\mathcal{M}(Y, \theta)$ with (potentially) multivariate observations $Y \in \mathcal{Y}$ and k -dimensional parameter vector $\theta \in \Theta$.

Parameter estimation: $\hat{\theta}$ by optimization of objective function $\Psi(Y, \theta)$ for n observations Y_i ($i = 1, \dots, n$):

$$\hat{\theta} = \operatorname{argmin}_{\theta \in \Theta} \sum_{i=1}^n \Psi(Y_i, \theta).$$

Special cases: Maximum likelihood (ML), weighted and ordinary least squares (OLS and WLS), quasi-ML, and other M-estimators.

1. Model estimation

Estimating function: $\hat{\theta}$ can also be defined in terms of

$$\sum_{i=1}^n \psi(Y_i, \hat{\theta}) = 0,$$

where $\psi(Y, \theta) = \partial\Psi(Y, \theta)/\partial\theta$.

Central limit theorem: If there is a true parameter θ_0 and given certain weak regularity conditions:

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, V(\theta_0)),$$

where $V(\theta_0) = \{A(\theta_0)\}^{-1}B(\theta_0)\{A(\theta_0)\}^{-1}$. A and B are the expectation of the derivative of ψ and the variance of ψ , respectively.

2. Tests for parameter instability

Generalized M-fluctuation tests capture instabilities in $\hat{\theta}$ for an ordering w.r.t Z_j .

Basis: Empirical fluctuation process of cumulative deviations w.r.t. to an ordering $\sigma(Z_{ij})$.

$$W_j(t, \hat{\theta}) = \hat{B}^{-1/2} n^{-1/2} \sum_{i=1}^{\lfloor nt \rfloor} \psi(Y_{\sigma(Z_{ij})}, \hat{\theta}) \quad (0 \leq t \leq 1)$$


Functional central limit theorem: Under parameter stability $W_j(\cdot) \xrightarrow{d} W^0(\cdot)$, where W^0 is a k -dimensional Brownian bridge.

1. Model estimation

Idea: In many situations, a single global model $\mathcal{M}(Y, \theta)$ that fits **all** n observations cannot be found. But it might be possible to find a partition w.r.t. the variables $Z = (Z_1, \dots, Z_l)$ so that a well-fitting model can be found locally in each cell of the partition.

Tool: Assess parameter instability w.r.t to partitioning variables $Z_j \in \mathcal{Z}_j$ ($j = 1, \dots, l$).

2. Tests for parameter instability



2. Tests for parameter instability

Test statistics: Scalar functional $\lambda(W_j)$ that captures deviations from zero.

Null distribution: Asymptotic distribution of $\lambda(W^0)$.

Special cases: Class of test encompasses many well-known tests for different classes of models. Certain functionals λ are particularly intuitive for numeric and categorical Z_j , respectively.

Advantage: Model $\mathcal{M}(Y, \hat{\theta})$ just has to be estimated once. Empirical estimating functions $\psi(Y_i, \hat{\theta})$ just have to be re-ordered and aggregated for each Z_j .

2. Tests for parameter instability

Splitting categorical variables: Assess instability using χ^2 statistics.

$$\lambda_{\chi^2}(W_j) = \sum_{c=1}^C \frac{n}{|I_c|} \left\| \Delta_{I_c} W_j \left(\frac{i}{n} \right) \right\|_2^2$$

Feature: Invariant for re-ordering of the C categories and the observations within each category.

Interpretation: Captures instability for split-up into C categories.

Limiting distribution: χ^2 with $k \cdot (C - 1)$ degrees of freedom.

2. Tests for parameter instability

Splitting numeric variables: Assess instability using supLM statistics.

$$\lambda_{\text{supLM}}(W_j) = \max_{i=\underline{l}, \dots, \bar{l}} \left(\frac{i}{n} \cdot \frac{n-i}{n} \right)^{-1} \left\| W_j \left(\frac{i}{n} \right) \right\|_2^2.$$

Interpretation: Maximization of single shift LM statistics for all conceivable breakpoints in $[\underline{l}, \bar{l}]$.

Limiting distribution: Supremum of a squared, k -dimensional tied-down Bessel process.

3. Segmentation

Goal: Split model into $b = 1, \dots, B$ segments along the partitioning variable Z_j associated with the highest parameter instability. Local optimization of

$$\sum_b \sum_{i \in I_b} \Psi(Y_i, \theta_b).$$

$B = 2$: Exhaustive search of order $O(n)$.

$B > 2$: Exhaustive search is of order $O(n^{B-1})$, but can be replaced by dynamic programming of order $O(n^2)$. Different methods (e.g., information criteria) can choose B adaptively.

Here: Binary partitioning.

4. Pruning


Pruning: Avoid overfitting.

Pre-pruning: Internal stopping criterium. Stop splitting when there is no significant parameter instability.

Post-pruning: Grow large tree and prune splits that do not improve the model fit (e.g., via cross-validation or information criteria).

Here: Pre-pruning based on Bonferroni-corrected p values of the fluctuation tests.

Costly journals



Costly journals

Task: Price elasticity of demand for economics journals.

Source: Bergstrom (2001, *Journal of Economic Perspectives*) "Free Labor for Costly Journals?", used in Stock & Watson (2007), *Introduction to Econometrics*.

Model: Linear regression via OLS.

- Demand: Number of US library subscriptions.
- Price: Average price per citation.
- Log-log-specification: Demand explained by price.
- Further variables without obvious relationship: Age (in years), number of characters per page, society (factor).

Costly journals

Recursive partitioning:

	Regressors		Partitioning variables				
	(Const.)	log(Pr./Cit.)	Price	Cit.	Age	Chars	Society
1	4.766	-0.533	3.280	5.261	42.198	7.436	6.562
	< 0.001	< 0.001	0.660	0.988	< 0.001	0.830	0.922
2	4.353	-0.605	0.650	3.726	5.613	1.751	3.342
	< 0.001	< 0.001	0.998	0.998	0.935	1.000	1.000
3	5.011	-0.403	0.608	6.839	5.987	2.782	3.370
	< 0.001	< 0.001	0.999	0.894	0.960	1.000	1.000

(Wald tests for regressors, parameter instability tests for partitioning variables.)

Beautiful professors

Task: Correlation of beauty and teaching evaluations for professors.

Source: Hamermesh & Parker (2005, *Economics of Education Review*). "Beauty in the Classroom: Instructors' Pulchritude and Putative Pedagogical Productivity."

Model: Linear regression via WLS.

- Response: Average teaching evaluation per course (on scale 1–5).
- Explanatory variables: Standardized measure of beauty and factors gender, minority, tenure, etc.
- Weights: Number of students per course.

Beautiful professors

	All	Men	Women
(Constant)	4.216	4.101	4.027
Beauty	0.283	0.383	0.133
Gender (= w)	−0.213		
Minority	−0.327	−0.014	−0.279
Native speaker	−0.217	−0.388	−0.288
Tenure track	−0.132	−0.053	−0.064
Lower division	−0.050	0.004	−0.244
R^2	0.271	0.316	

(Remark: Only courses with more than a single credit point.)

Beautiful professors

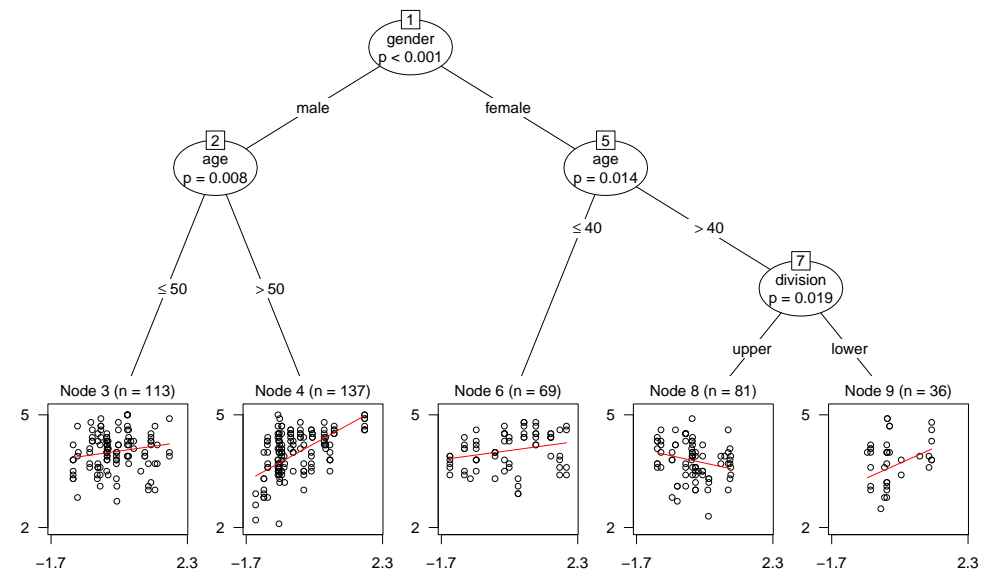
Hamermesh & Parker:

- Model with all factors (main effects).
- Improvement for separate models by gender.
- No association with age (linear or quadratic).

Here:

- Model for evaluation explained by beauty.
- Other variables as partitioning variables.
- Adaptive incorporation of correlations and interactions.

Beautiful professors



Beautiful professors

Recursive partitioning:

	(Const.)	Beauty
3	3.997	0.129
4	4.086	0.503
6	4.014	0.122
8	3.775	-0.198
9	3.590	0.403


Model comparison:

Model	R^2	Parameters
full sample	0.271	7
nested by gender	0.316	12
recursively partitioned	0.382	10 + 4

Beautiful professors

Single credit courses:

- Different type of courses: Yoga, aerobic, etc.
- Associated with second strongest instability (after gender).
- Subsamples too small for separated models: 18 (m), 9 (f).



Choosy students

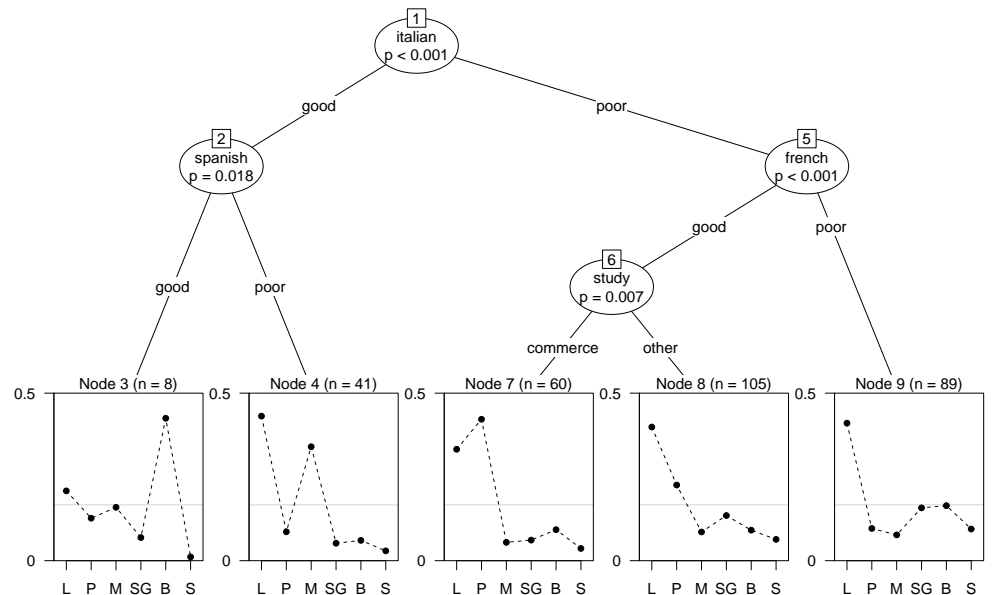
Task: Choice of university in student exchange programmes.

Source: Dittrich, Hatzinger, Katzenbeisser (1998, *Journal of the Royal Statistical Society C*). "Modelling the Effect of Subject-Specific Covariates in Paired Comparison Studies with an Application to University Rankings."

Model: Paired comparison via Bradley-Terry(-Luce).

- Ranking of six european management schools: London (LSE), Paris (HEC), Milano (Luigi Bocconi), St. Gallen (HSG), Barcelona (ESADE), Stockholm (HHS).
- Interviews with about 300 students from WU Wien.
- Additional information: Gender, studies, foreign language skills.

Choosy students



Choosy students

Recursive partitioning:

	London	Paris	Milano	St. Gallen	Barcelona	Stockholm
3	0.21	0.13	0.16	0.07	0.43	0.01
4	0.43	0.09	0.34	0.05	0.06	0.03
7	0.33	0.42	0.05	0.06	0.09	0.04
8	0.40	0.23	0.09	0.13	0.09	0.06
9	0.41	0.10	0.08	0.16	0.16	0.09

(Standardized ranking from Bradley-Terry model.)

Summary

Model-based recursive partitioning:

- Synthesis of classical parametric data models and algorithmic tree models.
- Based on modern class of parameter instability tests.
- Aims to minimize clearly defined objective function by greedy forward search.
- Can be applied general class of parametric models.
- Alternative to traditional means of model specification, especially for variables with unknown association.
- Object-oriented implementation freely available: Extension for new models requires some coding but not too extensive if interfaced model is well designed.

Software

All methods are implemented in the R system for statistical computing and graphics. Freely available under the GPL (General Public License) from the Comprehensive R Archive Network:

- Trees/recursive partitioning: **party**,
- Structural change inference: **strucchange**,
- Bradley-Terry regression/tree: **psychotree**.

<http://www.R-project.org/>
<http://CRAN.R-project.org/>

References

- Zeileis A, Hornik K (2007). "Generalized M-Fluctuation Tests for Parameter Instability." *Statistica Neerlandica*, **61**(4), 488–508. doi:10.1111/j.1467-9574.2007.00371.x
- Zeileis A, Hothorn T, Hornik K (2008). "Model-Based Recursive Partitioning." *Journal of Computational and Graphical Statistics*, **17**(2), 492–514. doi:10.1198/106186008X319331
- Kleiber C, Zeileis A (2008). *Applied Econometrics with R*. Springer-Verlag, New York. URL <http://CRAN.R-project.org/package=AER>
- Strobl C, Wickelmaier F, Zeileis A (2011). "Accounting for Individual Differences in Bradley-Terry Models by Means of Recursive Partitioning." *Journal of Educational and Behavioral Statistics*, Forthcoming. URL <http://statmath.wu.ac.at/~zeileis/papers/Strobl+Wickelmaier+Zeileis-2010.pdf>