Supplemental Instructions for the AMEs for GAMs Program

This document provides additional information for using the program to calculate average marginal effects for generalized additive models. I describe the fundamentals of doing so in my book.

I assume you are familiar with how to specify models for the Generalized Additive Models program. Be sure to read the material in the 'details' link for it and watch the video for the program. I will repeat selected portions of that material here for convenience.

For specifying your model, when you enter the equation of interest, you have two choices for a given predictor, (1) treat it as a linear predictor of Y, much like you would in traditional regression, or (2) treat it as a non-linear predictor whose relationship with Y is modeled using a smoother. A predictor is treated with a smoother if it is enclosed in s() when you specify the equation, such as

Y ~ s(X)

You can expand the notation within s() for a given predictor to take advantage of options provided by the GAM function from the R package mgcv that I make use of. The program offers a variety of regression based spline smoothers for modeling the relationship between X and Y. Examples include P-splines, B-splines, thin plate splines, and tensors. The default smoother is a penalized thin plate regression spline but you can force the program to use a different spline type for any given predictor. For example, to use a P spline for predictor X, I would specify

Y ~ s(X,bs='ps')

where bs is the keyword to signify the type of spline and 'ps' stands for a P spline. Options to specify after bs include 'tp' for thin plate regression splines, 'ds' for Duchon splines, 'cr' for cubic regression splines ('cs' specifies a shrinkage version of 'cr' and 'cc' specifies a cyclic cubic regression spline), 'sos' for splines on the sphere, 're' for random effects, 'mrf' for Markov random effects 'gp' for Gaussian process smooths, 'so' for soap film smooths, 'ad' for adaptive smoothers, and 'fs' for factor smooth interactions. The strengths and weaknesses of each are discussed in Wood (2017). I tested the average marginal effect program using the default smoother and have not explored it with the others. However, I do not a priori see any problems with using other smoothers.

GAMs can explore non-linear interactions between continuous predictors. The most common approach is to use either a full tensor product smooth (called 'te' in the mgcv package) or a tensor product interaction (called 'ti'), with the latter used in the spirit of the classic philosophy of adjusting for or removing main effect influences from the product term. For predictors X and Z, the equation would be specified as

Y ~ ti(X) + ti(Z) + ti(X,Z)

For details, see Wood (2017) and my book. If Z is a nominal moderator, then the interaction often is modeled using the default penalized thin plate regression spline smoother in conjunction with the 'by' command, like this:

Y ~ Z + s(X) + s(X, by = Z)

See my book for details. The average marginal effects program can handle statements with the 'by' subcommand but you must specify the 'by' variable in the average marginal effect program as a factor variable even if it is only binary. **THIS IS IMPORTANT**. When using the 'by' command, the marginal effects program I use reports only one marginal effect that includes both the ‘main effect’ for a variable, X, in the interaction term and the ‘interaction’ per se simultaneously. That is, the coefficient for X tells you the amount of change that a one unit change in X produces taking into account both the main effect and interaction term. This is standard practice for average marginal effects programs. See the details for the other average marginal effects programs on my website.

GAMS can be difficult to interpret, so I like to pursue profile analyses within them or pursue the average marginal effects. The latter is a bit controversial for models with penalty functions but can be viable in many cases.