Contents

APPENDIX A

The ANOVA decomposition of $[0,1]^d$

The ***analysis of variance*** (ANOVA) is a statistical model for analyzing experimental data. Given a rectangular table of data it quantifies how important the rows are relative to the columns and also captures the non-additivity under the term 'interaction'. The ANOVA can be applied to any number of independent variables and the variables do not have to be at discrete levels such as row and column names. The ANOVA on $[0,1]^d$ that we emphasize here is sometimes called the ***functional ANOVA***.

For Monte Carlo and quasi-Monte Carlo methods, the ANOVA provides a convenient way to quantify the importance of input variables to a function, through the related notions of effective dimension, Sobol' indices, and mean dimension.

A.1 ANOVA for tabular data

The ANOVA originated in agriculture. Suppose that we plant seeds of types $i = 1, \dots, I$ and apply fertilizers $j = 1, \dots, J$, and then measure the resulting crop yield Y_{ij} for all $I \times J$ combinations. We may then want to know which seed type is best, which fertilizer is best, the relative importance of these two variables and also the extent to which the best fertilizer varies with the type of seed and vice versa. As a toy example, suppose we have the following yields

$$\begin{array}{c} Y_{ij} \\ i=1 \\ i=2 \\ i=3 \end{array} \begin{array}{c} \begin{array}{cc} j=1 & j=2 \end{array} \\ \begin{pmatrix} 25 & 9 \\ 20 & 28 \\ 27 & 11 \end{pmatrix} \end{array}.$$

By inspection, we can see that column 1 has a higher average yield than colunn 2 and that row 2 has the highest row average. The average yield in row i is $\bar{Y}_{i\bullet} = (1/J)\sum_{j=1}^J Y_{ij}$. The average yield overall is $\bar{Y}_{\bullet\bullet} = (1/(IJ))\sum_{i=1}^I\sum_{j=1}^J Y_{ij}$. Taking this average as a baseline we can attribute an incremental yield of $\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}$ to seed type i, and an incremental yield of $\bar{Y}_{\bullet j} - \bar{Y}_{\bullet\bullet}$ to fertilizer j where $\bar{Y}_{\bullet j} = (1/I)\sum_{i=1}^I Y_{ij}$. If yields were additive, then Y_{ij} would be the baseline plus an increment for row i and an increment for column j. That is, we would find that $Y_{ij} = \bar{Y}_{\bullet\bullet} + (\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}) + (\bar{Y}_{\bullet j} - \bar{Y}_{\bullet\bullet})$. Subtracting this additive approximation from Y_{ij} yields the ***interaction*** term

$$Y_{ij} - \bar{Y}_{\bullet\bullet} - (\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}) - (\bar{Y}_{\bullet j} - \bar{Y}_{\bullet\bullet}) = Y_{ij} - \bar{Y}_{i\bullet} - \bar{Y}_{\bullet j} + \bar{Y}_{\bullet\bullet}.$$

The baseline $\bar{Y}_{\bullet\bullet}$ is usually called the 'grand mean' while $\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}$ for $i = 1, \dots, I$ is the ***main effect*** of the row variable and $\bar{Y}_{\bullet j} - \bar{Y}_{\bullet\bullet}$ for $j = 1, \dots, J$ is the main effect of the column variable.

We can display this decomposition as

$$\underbrace{\begin{pmatrix} 25 & 9 \\ 20 & 28 \\ 27 & 11 \end{pmatrix}}_{Y_{ij}} = \underbrace{\begin{pmatrix} 20 & 20 \\ 20 & 20 \\ 20 & 20 \end{pmatrix}}_{\bar{Y}_{\bullet\bullet}} + \underbrace{\begin{pmatrix} -3 & -3 \\ 4 & 4 \\ -1 & -1 \end{pmatrix}}_{\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}} + \underbrace{\begin{pmatrix} 4 & -4 \\ 4 & -4 \\ 4 & -4 \end{pmatrix}}_{\bar{Y}_{\bullet j} - \bar{Y}_{\bullet\bullet}} + \underbrace{\begin{pmatrix} 4 & -4 \\ -8 & 8 \\ 4 & -4 \end{pmatrix}}_{Y_{ij} - \bar{Y}_{i\bullet} - \bar{Y}_{\bullet j} + \bar{Y}_{\bullet\bullet}}.$$

Notice that the row effects $\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}$ average to zero over i within all columns $j = 1, \dots, J$ while the colunn effects average to zero over columns for each row. This is a consequence of the way we centered the data. The final interaction term averages to zero within each row and also within each column. In this made up example, the benefit of combining row 2 and column 2 is so strong that the best yield actually came from the worst column.

Interaction terms are differences of differences. An interaction effect from two factors can be written in these two ways

$$(Y_{ij} - \bar{Y}_{i\bullet}) - (\bar{Y}_{\bullet j} - \bar{Y}_{\bullet\bullet}) \quad \text{or} \quad (Y_{ij} - \bar{Y}_{\bullet j}) - (\bar{Y}_{i\bullet} - \bar{Y}_{\bullet\bullet}).$$

When there are more than 2 factors, then interactions of order $k > 2$ are k-fold differences of differences.

The importance of rows, columns and interactions can be measured by their sums of squares $\sum_{ij}(Y_{i\bullet} - Y_{\bullet\bullet})^2$, $\sum_{ij}(Y_{\bullet j} - Y_{\bullet\bullet})^2$, and $\sum_{ij}(Y_{ij} - Y_{i\bullet} - Y_{\bullet j} + Y_{\bullet\bullet})^2$. In the above example, these are 52, 96 and 192 respectively.

Much of the complexity of statistical experimental design, outside the scope of this text, arises because the yields Y_{ij} are themselves averages of noisy data. They have statistical uncertainty and then so do the estimates of the grand mean, main effects and interactions. In the toy example above, there were two factors, seed and fertilizer, while in applications there can be many more than two factors, so there are higher order interactions than two. Also, one must plan how to gather the data. See Box et al. (2005) and Wu and Hamada (2011) for more.

A.2 The functional ANOVA

The example ANOVA in §A.1 had two factors, one for rows and one for columns, and our main tool was averaging. We can replace averages over a finite number of levels by averages over some other distribution. Here we present the functional ANOVA for real-valued functions $f(\boldsymbol{x})$ where $\boldsymbol{x} = (x_1, \dots, x_d) \sim \mathbf{U}[0,1]^d$. We let $\mu = \mathbb{E}(f(\boldsymbol{x}))$ and $\sigma^2 = \mathrm{Var}(f(\boldsymbol{x}))$, and assume that $\sigma^2 < \infty$. It is not critical that $x_j \sim \mathbf{U}[0,1]$. The ANOVA can be defined for other distributions of x_j. Independence of all components of $\boldsymbol{x}$ is however critically important and finite variance is necessary for the most important results.

The functional ANOVA that we develop here writes $f(\boldsymbol{x})$ as a sum of functions that may each depend on some but not all x_j and it apportions the variance of $f(\boldsymbol{x})$ over all $2^d - 1$ non-empty subsets of the variables $x_1, \dots, x_d$.

The variable indices are in the set $\{1, \dots, d\}$ that we abbreviate to $1{:}d$. We use $|u|$ for the cardinality of each $u \subseteq 1{:}d$. If $u = \{j_1, j_2, \dots, j_{|u|}\}$ then we write $\boldsymbol{x}_u$ for $(x_{j_1}, \dots, x_{j_{|u|}}) = (x_j)_{j \in u}$. The complementary set $1{:}d \setminus u$ is denoted by $-u$. For singleton sets $u = \{j\}$ it is notationally convenient in a few places, especially subscripts, to replace $\{j\}$ by j. Some further shorthands are introduced as needed.

Sometimes we have to make up a new point by putting together components from two other points. If $\boldsymbol{x}, \boldsymbol{z} \in [0,1]^d$ and $u \subseteq 1{:}d$, then the hybrid point $\boldsymbol{y} = \boldsymbol{x}_u{:}\boldsymbol{z}_{-u}$ is the one with $y_j = x_j$ for $j \in u$ and $y_j = z_j$ for $j \notin u$.

The ANOVA of the unit cube develops in a way that parallels the ANOVA of tabular data. When we are done, we will be able to write

$$f(\boldsymbol{x}) = \sum_{u \subseteq 1:d} f_u(\boldsymbol{x}) \tag{A.1}$$

where the function $f_u(\cdot)$ depends on its argument $\boldsymbol{x}$ only through $\boldsymbol{x}_u$. For $u = \varnothing$, the function $f_\varnothing(\boldsymbol{x})$ does not depend on any components x_j of $\boldsymbol{x}$; it is a constant function that will be equal to the grand mean. Indexing the terms in (A.1) by subsets is convenient, because it replaces ungainly expressions like

$$f(\boldsymbol{x}) = f_\varnothing + \sum_{r=1}^{d} \sum_{1 \leqslant j_1 < j_2 < \cdots < j_r \leqslant d} f_{j_1,\dots,j_r}(x_{j_1}, \dots, x_{j_r})$$

that become difficult to manipulate.

We begin the functional ANOVA by generalizing the grand mean to

$$f_\varnothing(\boldsymbol{x}) = \int_{[0,1]^d} f(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x} \equiv \mu, \tag{A.2}$$

for all $\boldsymbol{x}$. Next, for $j = 1, \dots, d$, the main effects are

$$f_{\{j\}}(\boldsymbol{x}) = \int_{[0,1]^{d-1}} (f(\boldsymbol{x}) - \mu) \,\mathrm{d}\boldsymbol{x}_{-j} \tag{A.3}$$

which depends on $\boldsymbol{x}$ only through x_j as all components of $\boldsymbol{x}_{-j}$ have been integrated out.

The general expression for a set $u \subseteq \{1,\dots,d\}$ is

$$f_u(\boldsymbol{x}) = \int_{[0,1]^{d-|u|}} \Bigl(f(\boldsymbol{x}) - \sum_{v \subsetneq u} f_v(\boldsymbol{x})\Bigr)\,\mathrm{d}\boldsymbol{x}_{-u}. \tag{A.4}$$

We don't want to attribute anything to $\boldsymbol{x}_u$ that can be explained by $\boldsymbol{x}_v$ for strict subsets $v \subsetneq u$ so we subtract the corresponding $f_v(\boldsymbol{x})$. Then we average the difference over all the other variables not in u. The definition of $f_{1:d}(\boldsymbol{x})$ ensures that the functions defined in (A.4) satisfy (A.1).

There are many ways to make a decomposition of the form (A.1). Indeed an arbitrary choice of f_u for all $|u| < d$ can be accomodated by taking $f_{1:d}$ to be f minus all the other terms. The anchored decomposition of §A.7 is an important alternative to the ANOVA.

The effects f_u can also be written

$$f_u(\boldsymbol{x}) = \int_{[0,1]^{d-|u|}} f(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}_{-u} - \sum_{v \subsetneq u} f_v(\boldsymbol{x}), \tag{A.5}$$

because f_v does not depend on any component of $\boldsymbol{x}_{-u}$ when $v \subsetneq u$.

A.3 Orthogonality of ANOVA terms

In this section we show that ANOVA terms are mutually orthogonal. We saw that ordinary ANOVA terms average to zero over any of their indices. Similarly, we will show that

$$\int_0^1 f_u(\boldsymbol{x})\,\mathrm{d}x_j = 0, \quad \text{for} \quad j \in u.$$

Lemma A.1 proves this result which we then use to show orthogonality of ANOVA components.

Lemma A.1. *Let the function f be defined on $[0,1]^d$ with $\int f(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} < \infty$. For $u \subseteq \{1,\dots,d\}$, let f_u be the ANOVA effect defined by* (A.4)*. If $j \in u$, then*

$$\int_0^1 f_u(\boldsymbol{x}_{-j}{:}x_j)\,\mathrm{d}x_j = 0 \tag{A.6}$$

holds for all $\boldsymbol{x}_{-j} \in [0,1]^{d-1}$.

Proof. The proof is by induction on $|u|$. The statement of the lemma implies that $1 \leqslant |u| \leqslant d$. For $|u| = 1$, let $u = \{j\}$. Then by (A.3)

$$\int_0^1 f_{\{j\}}(\boldsymbol{x})\,\mathrm{d}x_j = \int_0^1 \int_{[0,1]^{-\{j\}}} (f(\boldsymbol{x}) - \mu) \prod_{k \neq j} \mathrm{d}x_k\,\mathrm{d}x_j$$

$$
\begin{aligned}
&= \int_{[0,1]^d} (f(\boldsymbol{x}) - \mu)\,\mathrm{d}\boldsymbol{x} \\
&= 0.
\end{aligned}
$$

Now suppose that $\int_0^1 f_v(\boldsymbol{x})\,\mathrm{d}x_j = 0$ for $j \in v$ whenever $1 \leqslant |v| \leqslant r < d$. Choose u with $|u| = r+1$, pick $j \in u$, and let $-u+j$ be a shorthand for $\{j\} \cup -u$. To complete the induction,

$$
\begin{aligned}
\int_0^1 f_u(\boldsymbol{x})\,\mathrm{d}x_j &= \int_{[0,1]^{d-|u|+1}} \Bigl(f(\boldsymbol{x}) - \sum_{v \subsetneq u} f_v(\boldsymbol{x})\Bigr)\,\mathrm{d}\boldsymbol{x}_{-u+j} \\
&= \int_{[0,1]^{d-|u|+1}} \Bigl(f(\boldsymbol{x}) - \sum_{v \subsetneq u,\, j \notin v} f_v(\boldsymbol{x})\Bigr)\,\mathrm{d}\boldsymbol{x}_{-u+j} \\
&= \int_{[0,1]^{d-|u|+1}} \Bigl(f(\boldsymbol{x}) - \sum_{v \subseteq u-\{j\}} f_v(\boldsymbol{x})\Bigr)\,\mathrm{d}\boldsymbol{x}_{-u+j} \\
&= \int_{[0,1]^{d-|u|+1}} \Bigl(f(\boldsymbol{x}) - \sum_{v \subsetneq u-\{j\}} f_v(\boldsymbol{x})\Bigr)\,\mathrm{d}\boldsymbol{x}_{-u+j} + f_{u-\{j\}}(\boldsymbol{x}) \\
&= f_{u-\{j\}}(\boldsymbol{x}) - f_{u-\{j\}}(\boldsymbol{x}) \\
&= 0. \qquad \square
\end{aligned}
$$

Now consider the product $f_u(\boldsymbol{x})f_v(\boldsymbol{x})$. If $u \neq v$ then there is some j that is in u but not v, or vice versa. Integrating $f_u f_v$ over x_j then yields zero and the orthogonality we want. Using this argument involves Fubini's theorem and to get a sufficient condition for Fubini's theorem, we need to establish a technical point first. If f is square integrable, then so are all the ANOVA effects f_u.

Lemma A.2. *Let f be a real-valued function on $[0,1]^d$ with $\int_{[0,1]^d} f(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} < \infty$. Then $\int_{[0,1]^d} f_u(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} < \infty$ for all $u \subseteq \{1,\dots,d\}$.*

Proof. We will proceed by induction on $|u|$. If $|u| = 0$, then $f_u(\boldsymbol{x})$ is a constant function and it is then square integrable. For $|u| > 0$

$$
f_u(\boldsymbol{x}) = \int_{[0,1]^{d-|u|}} f(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}_{-u} - \sum_{v \subsetneq u} f_v(\boldsymbol{x}), \tag{A.7}
$$

and all of the f_v on the right hand side of (A.7) are square integrable, so it is enough to show that $f_{\bar{u}}(\boldsymbol{x}) \equiv \int_{[0,1]^{d-|u|}} f(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}_{-u}$ is square integrable. Now $f_{\bar{u}}(\boldsymbol{x}) = \mathbb{E}(f(\boldsymbol{x}) \mid \boldsymbol{x}_u)$ for $\boldsymbol{x} \sim \mathbf{U}[0,1]^d$ and so

$$
\int f_{\bar{u}}(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} \leqslant \int f(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} < \infty. \quad \square
$$

The following Lemma is a very general orthogonality result for ANOVA.

Lemma A.3. *Let f and g be real-valued functions on $[0,1]^d$ with $\int_{[0,1]^d} f(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x} < \infty$ and $\int_{[0,1]^d} g(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x} < \infty$. Let $u, v \subseteq \{1, \dots, d\}$. If $u \neq v$, then*

$$\int_{[0,1]^d} f_u(\boldsymbol{x}) g_v(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x} = 0.$$

Proof. Since $u \neq v$, there either exists $j \in u$ with $j \notin v$, or $j \in v$ with $j \notin u$. Without loss of generality suppose that $j \in u$ and $j \notin v$. Next by Lemma A.2, both $\int f_u(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x}$ and $\int g_v(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x}$ are finite. Therefore $\int |f_u(\boldsymbol{x}) g_v(\boldsymbol{x})| \,\mathrm{d}\boldsymbol{x} < \infty$ by Cauchy-Schwarz. It follows that we may use Fubini's theorem to integrate x_j out of $f_u g_v$ first as follows:

$$\begin{aligned}\int_{[0,1]^d} f_u(\boldsymbol{x}) g_v(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x} &= \int_{[0,1]^{d-1}} \int_0^1 f_u(\boldsymbol{x}) g_v(\boldsymbol{x}) \,\mathrm{d}x_j \,\mathrm{d}\boldsymbol{x}_{-j} \\ &= \int_{[0,1]^{d-1}} \int_0^1 f_u(\boldsymbol{x}) \,\mathrm{d}x_j \, g_v(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x}_{-j} = 0,\end{aligned}$$

using Lemma A.1 on the inner integral. □

Corollary A.1. *Let f be a real-valued function on $[0,1]^d$ with $\int f(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x} < \infty$. If $u \neq v$ are subsets of $\{1, \dots, d\}$, then*

$$\int f_u(\boldsymbol{x}) f_v(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x} = 0.$$

Proof. Take $f = g$ in Lemma A.3. □

Now we can explain the name ANOVA by decomposing (analyzing) the variance of f. The variance of $f_u(\boldsymbol{x})$ for $\boldsymbol{x} \sim \mathbf{U}[0,1]^d$ is

$$\sigma_u^2 \equiv \begin{cases} \int f_u(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x}, & |u| > 0 \\ 0, & u = \varnothing. \end{cases} \tag{A.8}$$

Lemma A.4. *Let f be a real-valued function on $[0,1]^d$ with $\mu = \int f(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x}$ and $\sigma^2 = \int (f(\boldsymbol{x}) - \mu)^2 \,\mathrm{d}\boldsymbol{x} < \infty$. Then*

$$\sigma^2 = \sum_{|u|>0} \sigma_u^2. \tag{A.9}$$

Proof. From the definition of σ^2,

$$\int (f(\boldsymbol{x}) - \mu)^2 \,\mathrm{d}\boldsymbol{x} = \int \sum_{|u|>0} \sum_{|v|>0} f_u(\boldsymbol{x}) f_v(\boldsymbol{x}) \,\mathrm{d}\boldsymbol{x} = \sum_{|u|>0} \int f_u(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x}$$

using Corollary A.1. The result follows by the definition of σ_u^2 at (A.8). □

ANOVA is an acronym for analysis of variance. Equation (A.9) shows that the variance of f decomposes into a sum of variances of ANOVA effects. The quantity σ_u^2 is a ***variance component***. We may also write $\sigma^2 = \sum_u \sigma_u^2$, summing over all 2^d subsets, because $\sigma_\varnothing^2 = 0$. Similarly $\int f(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x} = \sum_u \int f_u(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x} = \mu^2 + \sum_{|u|>0} \sigma_u^2$.

A.4 Best approximations via ANOVA

We can use the ANOVA decomposition to define some best approximations to a function f. We suppose that f has domain $[0,1]^d$ and that $\int f(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} < \infty$ and best means minimizing the squared error. We begin with the best additive approximation to f.

Definition A.1. The function $g : [0,1]^d \to \mathbb{R}$ is ***additive*** if $g(\boldsymbol{x}) = \widetilde{g}_0 + \sum_{j=1}^d \widetilde{g}_j(x_j)$ where $\widetilde{g}_j$ are real-valued functions on $[0,1]$ and $\widetilde{g}_0 \in \mathbb{R}$ is a constant.

It is convenient to rewrite g in its ANOVA decomposition. If g is additive, then the ANOVA decomposition of g is

$$g(\boldsymbol{x}) = g_{\varnothing}(\boldsymbol{x}) + \sum_{j=1}^d g_{\{j\}}(\boldsymbol{x})$$

where $g_{\varnothing}(\boldsymbol{x}) = \widetilde{g}_0 + \sum_{j=1}^d \int_0^1 \widetilde{g}_j(x)\,\mathrm{d}x$ and $g_{\{j\}}(\boldsymbol{x}) = \widetilde{g}_j(x_j) - \int_0^1 \widetilde{g}_j(x)\,\mathrm{d}x$.

Definition A.2. Given a function $f \in L^2[0,1]^d$, the ***additive part*** of f is

$$f_{\mathrm{add}}(\boldsymbol{x}) = f_{\varnothing}(\boldsymbol{x}) + \sum_{j=1}^d f_{\{j\}}(\boldsymbol{x}). \tag{A.10}$$

It is sometimes convenient to simplify f_{add} to $\mu + \sum_{j=1}^d f_j(x_j)$ where $f_j(x_j) = f_{\{j\}}(x_j{:}\boldsymbol{x}_{-j})$. Evidently f_{add} is additive. It may be concisely written as

$$f_{\mathrm{add}}(\boldsymbol{x}) = \sum_{|u|\leqslant 1} f_u(\boldsymbol{x}).$$

The next lemma shows an optimality property of f_{add}.

Lemma A.5. *Let $f \in L^2[0,1]^d$ and let $f_{\mathrm{add}}(x)$ be defined at* (A.10)*. If $g(\boldsymbol{x})$ is an additive function, then*

$$\int (f(\boldsymbol{x}) - g(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x} \geqslant \int (f(\boldsymbol{x}) - f_{\mathrm{add}}(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x}.$$

Proof. In this proof summations over u are over all $u \subseteq 1{:}d$ unless otherwise indicated. If $\int g(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x} = \infty$ then the conclusion follows easily, so we assume $g \in L^2[0,1]^d$ as well, so g has an orthogonal ANOVA decomposition and hence so does $f - g$.

Orthogonality of ANOVA terms yields

$$\int (f(\boldsymbol{x}) - g(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x} = \sum_u \int (f_u(\boldsymbol{x}) - g_u(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x}.$$

The ANOVA effects of f_{add} for $|u|>1$ are $f_{\mathrm{add},u}(\boldsymbol{x})=g_u(\boldsymbol{x})=0$ while for $|u|\leqslant 1$ they are $f_{\mathrm{add},u}(\boldsymbol{x})=f_u(\boldsymbol{x})$. Therefore

$$\begin{aligned}
\int (f(\boldsymbol{x})-g(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x} &= \sum_u \int (f_u(\boldsymbol{x})-g_u(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x}\\
&= \sum_u \int (f_u(\boldsymbol{x})-f_{\mathrm{add},u}(\boldsymbol{x})+f_{\mathrm{add},u}(\boldsymbol{x})-g_u(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x}\\
&= \sum_{|u|\leqslant 1} \int (f_{\mathrm{add},u}(\boldsymbol{x})-g_u(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x} + \sum_{|u|>1}\int f_u(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x}\\
&\geqslant \sum_{|u|>1}\int f_u(\boldsymbol{x})^2\,\mathrm{d}\boldsymbol{x}\\
&= \int (f(\boldsymbol{x})-f_{\mathrm{add}}(\boldsymbol{x}))^2\,\mathrm{d}\boldsymbol{x}. \qquad \Box
\end{aligned}$$

We can get the best additive approximation to f by simply removing from the ANOVA decomposition all terms f_u with $|u|>1$. The same argument shows that the best approximation (in mean square) having interactions of order at most 2 is

$$f_{\mathrm{two}}(\boldsymbol{x}) \equiv \sum_{|u|\leqslant 2} f_u(\boldsymbol{x}).$$

More generally, the best approximation with interactions up to order $k\leqslant d$ is

$$f_{\mathrm{order}\ k}(\boldsymbol{x}) \equiv \sum_{|u|\leqslant k} f_u(\boldsymbol{x}).$$

Now suppose that we want the best approximation to $f(\boldsymbol{x})$ on $[0,1]^d$ that can be obtained using only $\boldsymbol{x}_u$. Modifying the argument that we used to identify the best additive approximation to f we find that

$$f_{\bar{u}}(\boldsymbol{x}) \equiv \sum_{v\subseteq u} f_v(\boldsymbol{x}) = \int_{[0,1]^{d-|u|}} f(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}_{-u} = \mathbb{E}(f(\boldsymbol{x})\mid \boldsymbol{x}_u) \tag{A.11}$$

is that best approximation. This function appeared earlier in the proof of the technical Lemma A.2.

Equation (A.11) expresses each of 2^d cumulative effects $f_{\bar{u}}$ as a sum of original ANOVA effects f_v. There is also an inverse relationship (Exercise A.1)

$$f_u(\boldsymbol{x}) = \sum_{v\subseteq u} (-1)^{|u-v|} f_{\bar{v}}(\boldsymbol{x}). \tag{A.12}$$

Equation (A.12) is an example of the Möbius inversion formula.

A.5 Effective dimension

It is often observed empirically that a function f defined on $[0,1]^d$ is very nearly equal to the sum of its interactions of order up to $s \ll d$. When this happens we consider the function to have an ***effective dimension*** much lower than its nominal dimension d.

The benefit of low effective dimension comes up in quadrature formulas. Let $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n \in [0,1]^d$. Then

$$\hat{\mu} \equiv \frac{1}{n}\sum_{i=1}^{n} f(\boldsymbol{x}_i) = \sum_{u \subseteq \{1,\dots,d\}} \frac{1}{n}\sum_{i=1}^{n} f_u(\boldsymbol{x}_i), \tag{A.13}$$

with a similar formula holding in the case of unequally weighted quadrature rules. The error is

$$\hat{\mu} - \mu = \sum_{|u|>0} \hat{\mu}_u$$

where $\hat{\mu}_u = (1/n)\sum_{i=1}^n f_u(\boldsymbol{x}_i)$.

Let us split the error into high dimensional contributions $\hat{\mu}_u$ for $|u| > k$ and low dimensional ones for $1 \leqslant |u| \leqslant k$. If all of the k dimensional projections of $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$ have good equidistribution properties then unless f_u is particularly awkward, we should expect a small error $\hat{\mu}_u$. Similarly, if all the high dimensional components f_u are nearly zero, then we expect a small error $\hat{\mu}_u$ for them.

If 99% of σ^2 can be attributed to ANOVA effects u with $|u| \leqslant s$ then we can approach a 100 fold variance reduction if we can find a Monte Carlo method with $\mathbb{E}(\hat{\mu}_u^2) = o(1/n)$ (for $|u| \leqslant s$) while $\mathbb{E}(\hat{\mu}_u^2) \doteq \sigma_u^2/n$ for $|u| > s$. Some randomized quasi-Monte Carlo methods, presented in Chapter 17 behave this way.

Definition A.3. Let f be a square integrable function on $[0,1]^d$. The ***effective dimension of f in the superposition sense*** is the smallest integer s such that $\sum_{|u|\leqslant s} \sigma_u^2 \geqslant 0.99\sigma^2$.

Another notion of effective dimension is that only a small number s of the input variables are important. In such cases we might treat those variables differently, but to do that we need to know which ones they are. Without loss of generality, we suppose that the first s variables are most important.

Definition A.4. Let f be a square integrable function on $[0,1]^d$. The ***effective dimension of f in the truncation sense*** is the smallest integer s such that $\sum_{u\subseteq\{1,\dots,s\}} \sigma_u^2 \geqslant 0.99\sigma^2$.

The value 0.99 is a consequence of the somewhat arbitrary target of a 100-fold variance reduction. A different threshold might be more suitable for some problems.

A.6 Sobol' indices and mean dimension

Given a black box function of independent variables we might want to measure and compare the importance of those variables. If $\sigma^2_{\{j\}} > \sigma^2_{\{k\}}$, then other things being equal, we would consider x_j to be more important than x_k. The other things that might not be equal, include the extent to which those variables contribute to interactions. Additionally, we might want to quantify the importance of $\boldsymbol{x}_u$ for a set u of more than one of the variables.

Sobol's indices are

$$\underline{\tau}_u^2 = \sum_{v\subseteq u} \sigma_v^2 \quad \text{and} \quad \overline{\tau}_u^2 = \sum_{v:v\cap u\neq\varnothing} \sigma_v^2.$$

The reader should verify that $\overline{\tau}_u^2 = \sigma^2 - \underline{\tau}_{-u}^2$. The lower index $\underline{\tau}_u^2$ measures the importance of $\boldsymbol{x}_u$ through all main effects and interactions in u. The upper index $\overline{\tau}_u^2$ includes any interaction to which one or more of the components of $\boldsymbol{x}_u$ contribute. These indices are usually expressed as normalized forms $\underline{\tau}_u^2/\sigma^2$ and $\overline{\tau}_u^2/\sigma^2$, where they then quantify the proportion of variance of f attributable to subsets of u, and subsets intersecting u, respectively. The ***closed sensitivity index*** is $\underline{\tau}_u^2/\sigma^2$ and the ***total sensitivity index*** is $\overline{\tau}_u^2/\sigma^2$.

These indices are interpreted as follows. If $\underline{\tau}_u^2$ is large, then $\boldsymbol{x}_u$ is important. If $\overline{\tau}_u^2$ is small, then $\boldsymbol{x}_u$ is unimportant, because even with all interactions included, it does not make much difference. One can then freeze $\boldsymbol{x}_u$ at a default value, call it $\boldsymbol{c}_u$, and devote more attention to studying $f(\boldsymbol{c}_u{:}\boldsymbol{x}_{-u})$ as a function of $\boldsymbol{x}_{-u} \in [0,1]^{d-|u|}$. Freezing $\boldsymbol{x}_u$ this way requires a hidden but often very reasonable assumption that f is well enough behaved, that unimportance of $\boldsymbol{x}_u$ in our mean square sense is enough for our application. For instance, if there are points $\boldsymbol{x}$ for which $|f(\boldsymbol{x}) - f(\boldsymbol{c}_u{:}\boldsymbol{x}_{-u})|$ is not small, then those points could pose a problem when freezing $\boldsymbol{x}_u$ at $\boldsymbol{c}_u$.

The Sobol' indices can be estimated by ***pick-freeze*** methods described next. We do not have to estimate any of the ANOVA effects f_u. Instead, for $\underline{\tau}_u^2$, we may use the identity

$$\int_{[0,1]^{2d-|u|}} f(\boldsymbol{x})f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z}_{-u} = \underline{\tau}_u^2 + \mu^2. \tag{A.14}$$

In (A.14), we sample $\boldsymbol{x}$ to get $f(\boldsymbol{x})$, then freeze the selection $\boldsymbol{x}_u$ and pick new values $\boldsymbol{z}_{-u}$ independently of $\boldsymbol{x}$ and take the expected value of the product $f(\boldsymbol{x})f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})$ over the distribution of $\boldsymbol{x}$ and $\boldsymbol{z}_{-u}$. To prove (A.14), we use the ANOVA decomposition $f(\boldsymbol{x}) = \sum_{v\subseteq 1:d} f_v(\boldsymbol{x})$. Then

$$\begin{aligned}\int_{[0,1]^{2d-|u|}} f(\boldsymbol{x})f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z}_{-u} &= \int_{[0,1]^{2d-|u|}} \sum_v f_v(\boldsymbol{x})f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z}_{-u}\\ &= \int_{[0,1]^{2d-|u|}} \sum_{v\subseteq u} f_v(\boldsymbol{x})f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z}_{-u}\end{aligned}$$

because if v has an element $j \notin u$, then

$$\int_0^1 f_v(\boldsymbol{x})f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})\,\mathrm{d}x_j = f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})\int_0^1 f_v(\boldsymbol{x})\,\mathrm{d}x_j = 0.$$

Next, by orthogonality of ANOVA terms, we find for $v \subseteq u$, that

$$\begin{aligned}\int_{[0,1]^{2d-|u|}} f_v(\boldsymbol{x}) f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u}) \,\mathrm{d}\boldsymbol{x} \,\mathrm{d}\boldsymbol{z}_{-u} &= \int_{[0,1]^{2d-|u|}} f_v(\boldsymbol{x}) f_v(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u}) \,\mathrm{d}\boldsymbol{x} \,\mathrm{d}\boldsymbol{z}_{-u} \\ &= \int_{[0,1]^{|v|}} f_v(\boldsymbol{x})^2 \,\mathrm{d}\boldsymbol{x}_v \\ &= \begin{cases} \sigma_v^2, & |v| > 0, \\ \mu^2, & v = \varnothing. \end{cases}\end{aligned}$$

Summing over $v \subseteq u$ completes the proof of (A.16).

Taking $\boldsymbol{x}_i \sim \mathbf{U}[0,1]^d$ and $\boldsymbol{z}_i \sim \mathbf{U}[0,1]^d$ all independently, we may form the estimate

$$\hat{\underline{\tau}}_u^2 = \frac{1}{n} \sum_{i=1}^n f(\boldsymbol{x}_i) f(\boldsymbol{x}_{i,u}{:}\boldsymbol{z}_{i,-u}) - \hat{\mu}^2 \tag{A.15}$$

where $\hat{\mu} = (1/n) \sum_{i=1}^n (f(\boldsymbol{x}_i) + f(\boldsymbol{x}_{i,u}{:}\boldsymbol{z}_{i,-u}))/2$.

For the upper index, we may use the identity

$$\frac{1}{2} \int_{[0,1]^{d+|u|}} \bigl(f(\boldsymbol{x}) - f(\boldsymbol{x}_{-u}{:}\boldsymbol{z}_u)\bigr)^2 \,\mathrm{d}\boldsymbol{x} \,\mathrm{d}\boldsymbol{z}_u = \overline{\tau}_u^2. \tag{A.16}$$

See Exercise A.2. We can estimate $\overline{\tau}_u^2$ by Monte Carlo, via

$$\hat{\overline{\tau}}_u^2 = \frac{1}{2n} \sum_{i=1}^n \bigl(f(\boldsymbol{x}_i) - f(\boldsymbol{x}_{i,-u}{:}\boldsymbol{z}_{i,u})\bigr)^2,$$

and, because we don't need to estimate $\hat{\mu}^2$, we have $\mathbb{E}(\hat{\overline{\tau}}_u^2) = \overline{\tau}_u^2$.

The most important Sobol' indices are the ones for singletons $\{j\}$. Then $\underline{\tau}_j^2$ is the mean square of the main effect for x_j, while $\overline{\tau}_j^2$ includes all interaction mean squares that x_j contributes to. Now

$$\sum_{j=1}^d \overline{\tau}_j^2 = \sum_{j=1}^d \sum_{u \subseteq 1:d} \mathbb{1}\{j \in u\} \sigma_u^2 = \sum_{u \subseteq 1:d} \sum_{j=1}^d \mathbb{1}_{\{j \in u\}} \sigma_u^2 = \sum_{u \subseteq 1:d} |u| \sigma_u^2. \tag{A.17}$$

We can use this cardinality weighted sum of variance components to define the ***mean dimension*** of f. If $\sigma^2 \neq 0$, then the mean dimension of f is

$$\nu(f) = \frac{1}{\sigma^2} \sum_u |u| \sigma_u^2.$$

The mean dimension is easier to estimate than the effective dimension. We only need to compute d averages, one for each of the $\overline{\tau}_j^2$.

The above $\nu(f)$ is a mean dimension in the superposition sense. We can also define a mean dimension in the truncation sense. For non-empty u, define $\lceil u \rceil = \max\{j \mid j \in u\}$ and set $\lceil \varnothing \rceil = 0$. Then we can use

$$\nu_{\mathrm{trunc}}(f) = \frac{1}{\sigma^2} \sum_u \lceil u \rceil \sigma_u^2$$

as the mean dimension of f in the truncation sense.

There are additional ways to estimate Sobol' indices. For instance, using

$$\underline{\tau}_u^2 = \iint f(\boldsymbol{x})(f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u}) - f(\boldsymbol{z}))\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z} \tag{A.18}$$

we don't have to subtract an estimate of μ^2. Another choice is

$$\underline{\tau}_u^2 = \iiint (f(\boldsymbol{x}) - f(\boldsymbol{y}_u{:}\boldsymbol{x}_{-u}))(f(\boldsymbol{x}_u{:}\boldsymbol{z}_{-u})) - f(\boldsymbol{z}))\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{y}\,\mathrm{d}\boldsymbol{z} \tag{A.19}$$

for independent $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \sim \mathbf{U}[0,1]^d$. Either of these can be the basis of a Monte Carlo algorithm to estimate $\underline{\tau}_u^2$. There is some discussion in the end notes.

A.7 Anchored decompositions

The ***anchored decomposition*** is another way to write f as a sum of 2^d functions each one depending only on $\boldsymbol{x}_u$ for one set $u \subseteq \{1, \dots, d\}$. The functions are defined with respect to a special point $\boldsymbol{c} \in [0,1]^d$, called the anchor. We will obtain the decomposition

$$f(\boldsymbol{x}) = \sum_u f_{u,\boldsymbol{c}}(\boldsymbol{x}) \tag{A.20}$$

where, after the anchor has been chosen, $f_{u,\boldsymbol{c}}$ depends on $\boldsymbol{x}$ only through $\boldsymbol{x}_u$.

Before giving a general expression, we show an example for the case with $d = 3$ and $\boldsymbol{c} = \mathbf{0}$. In this case, the constant term is $f_{\varnothing,\mathbf{0}}(\boldsymbol{x}) = f(0,0,0)$. The main effect for x_1 is $f_{\{1\},\mathbf{0}}(\boldsymbol{x}) = f(x_1,0,0) - f(0,0,0)$ and those of x_2 and x_3 are similarly defined. Instead of using $\mathbb{E}(f(\boldsymbol{x}))$ as the baseline we subtract $f(\boldsymbol{c}) = f(0,0,0)$. If we use only the terms with $|u| = 0$ or 1, we get an additive approximation

$$f(x_1,0,0) + f(0,x_2,0) + f(0,0,x_3) - 2f(0,0,0).$$

This approximation is not generally the closest additive function to f in mean square. We can however compute it at any $\boldsymbol{x}$ that we like, unlike $f_{\mathrm{add}}(\boldsymbol{x})$.

The term for $u = \{1,2\}$ is

$$\begin{aligned} f_{\{1,2\},\mathbf{0}}(\boldsymbol{x}) &= f(x_1,x_2,0) - f_{\varnothing,\mathbf{0}}(\boldsymbol{x}) - f_{\{1\},\mathbf{0}}(\boldsymbol{x}) - f_{\{2\},\mathbf{0}}(\boldsymbol{x}) \\ &= f(x_1,x_2,0) - f(x_1,0,0) - f(0,x_2,0) + f(0,0,0), \end{aligned}$$

after simplification. The terms for $u = \{1,3\}$ and $u = \{2,3\}$ are similar. The term for $u = \{1,2,3\}$ is

$$\begin{aligned} f_{\{1,2,3\},\mathbf{0}}(\boldsymbol{x}) &= f(x_1,x_2,x_3) - f_{\varnothing,\mathbf{0}}(\boldsymbol{x}) - f_{\{1\},\mathbf{0}}(\boldsymbol{x}) - f_{\{2\},\mathbf{0}}(\boldsymbol{x}) - f_{\{3\},\mathbf{0}}(\boldsymbol{x}) \\ &\quad - f_{\{1,2\},\mathbf{0}}(\boldsymbol{x}) - f_{\{2,3\},\mathbf{0}}(\boldsymbol{x}) - f_{\{1,3\},\mathbf{0}}(\boldsymbol{x}) \\ &= f(x_1,x_2,x_3) - f(0,0,0) + f(x_1,0,0) + f(0,x_2,0) + f(0,0,x_3) \end{aligned}$$

$$- f(0,x_2,x_3) - f(x_1,0,x_3) - f(0,x_2,x_3),$$

after some algebra. The alternating signs above generalize to higher dimensions. Just as in the ANOVA, the terms are k-fold differences of differences.

The anchored decomposition is constructed just like the ANOVA but at each step instead of subtracting an average over one of the x_j we subtract the value we get in f at $x_j = c_j$. To begin with, we take

$$f_{\varnothing,\boldsymbol{c}}(\boldsymbol{x}) = f(\boldsymbol{c}).$$

Then for non-empty $u \subseteq \{1,\dots,d\}$, we define

$$f_{u,\boldsymbol{c}}(\boldsymbol{x}) = f(\boldsymbol{x}_u{:}\boldsymbol{c}_{-u}) - \sum_{v \subsetneq u} f_{v,\boldsymbol{c}}(\boldsymbol{x}).$$

The counterpart to the Möbius equality (A.12) that the ANOVA satisfied is

$$f_{u,\boldsymbol{c}}(\boldsymbol{x}) = \sum_{v \subseteq u} (-1)^{|u-v|} f(\boldsymbol{x}_v{:}\boldsymbol{c}_{-v}) \tag{A.21}$$

In the ANOVA decomposition, the term $f_u(\boldsymbol{x})$ integrated to 0 over x_j for any $j \in u$. Here $f_{u,\boldsymbol{c}}(\boldsymbol{x}) = 0$ if $x_j = c_j$ for any one of $j \in u$. We can prove that using (A.21) to write

$$\begin{aligned} f_{u,\boldsymbol{c}}(\boldsymbol{x}) &= \sum_{v \subseteq u} (-1)^{|u-v|} f(\boldsymbol{x}_v{:}\boldsymbol{c}_{-v}) \\ &= \sum_{v \subseteq u-j} (-1)^{|u-v|} \bigl(f(\boldsymbol{x}_v{:}\boldsymbol{c}_{-v}) - f(\boldsymbol{x}_{v+j}{:}\boldsymbol{c}_{-v-j}) \bigr). \end{aligned}$$

If j is in u but not v and $c_j = x_j$, then $\boldsymbol{x}_v{:}\boldsymbol{c}_{-v}$ and $\boldsymbol{x}_{v+j}{:}\boldsymbol{c}_{-v-j}$ are the same point and then each term in the sum above is zero, making $f_{u,\boldsymbol{c}}(\boldsymbol{x}) = 0$.

For any anchor $\boldsymbol{c}$ that we choose, we can compute all terms of the anchored decomposition at any point $\boldsymbol{x}$ that we choose. We may have to evaluate f up to 2^d times to do so. If possible, we should choose $\boldsymbol{c}$ to be a point where the computation of $f(\boldsymbol{x})$ is simpler when some of the $x_j = c_j$. That could be $\mathbf{0}$ or $\mathbf{1}$ or $(1/2,\dots,1/2)$.

Appendix end notes

The ANOVA was introduced by Fisher and Mackenzie (1923). The frequent occurence of physical phenomena well explained by a small number of low order interactions among experimental variables, known as ***factor sparsity***, has often been remarked on by G. E. P. Box. The book of Box et al. (2005) presents experimental designs geared to exploiting factor sparsity.

The extension of ANOVA to the continuum was made by Hoeffding (1948) in his study of U–statistics. Sobol' (1969) introduced it independently to study multidimensional integration problems. It was used by Efron and Stein (1981) in the study of the jackknife. In his study of Latin hypercube sampling, Stein (1987) used the ANOVA to find the best additive approximation to a given function on $[0, 1]^d$. Owen (1998) defines an ANOVA for the case $d = \infty$.

The definitions of effective dimension are from Caflisch et al. (1997). The first notion of effective dimension appeared in Richtmyer (1952). Numerous ways to define effective dimension have been used. Owen (2019) includes a historical survey. See Wasilkowski (2019) for an emphasis on information based complexity.

Sobol' indices, the identities (A.14) and (A.16), and the idea of freezing unimportant variables are from Sobol' (1990), which was translated into English as Sobol' (1993). The index $\underline{\tau}_j^2/\sigma^2$ appears independently in Ishigami and Homma (1990). These indices are the cornerstone of global sensitivity analysis, as distinct from a local sensitivity analysis that just uses small perturbations of the variables. See Saltelli et al. (2008) for more about global sensitivity analysis. Oakley and O'Hagan (2004) study a Bayesian approach to estimating global sensitivity indices.

Janon et al. (2014) study the estimator (A.15) of $\underline{\tau}_u^2$. Mauntz (2002) and Saltelli (2002) independently propose the unbiased estimator (A.18) of $\underline{\tau}_u^2$. The estimator (A.19) of $\underline{\tau}_u^2$ that uses $\boldsymbol{x}$, $\boldsymbol{y}$, and $\boldsymbol{z}$ is from Owen (2013a). There is no universally best estimator of $\underline{\tau}_u^2$ among all of these and other possible choices. The one in (A.19) has an advantage when the corresponding upper index $\overline{\tau}_u^2$ is small.

The mean dimension (A.17) is from Liu and Owen (2006) who also consider mean square dimensions. Many other interesting and potentially useful quantities can be estimated using the pick-freeze ideas. We can easily estimate $\sum_{|u|=1} \sigma_u^2 = \sum_{j=1}^d \underline{\tau}_j^2$. It is possible to estimate $\sum_{|u|=2} \sigma_u^2$ using an integral with only $2d + 2$ different evaluations of f (see Exercise A.4b) despite it being a sum of $d(d-1)/2$ variance components. Hooker (2004) considers $\Upsilon_u^2 = \sum_{v \supseteq u} \sigma_v^2$. This superset importance measure quantifies the effect of dropping all of the effects involving $\boldsymbol{x}_u$ and possibly more x_j from a formula. Fruth et al. (2014) compare methods of estimating Υ_u^2 from samples. See Owen (2013b) for these and other examples of things to estimate.

The efficiency with which an ANOVA derived quantity can be estimated is hard to predict because the variance of these estimators depends on some fourth moments. Those are expectations of products of f evaluated at up to four pick-freeze locations. Further complicating the problem is that estimators of these quantities may differ in the number of function evaluations that they consume, and when we have a large list of sensitivity indices and related quantities to estimate, then some function evaluations can be reused in multiple estimates. Then the cost of estimating a set of Sobol' index quantities can be less than the sum of their individual costs. See Saltelli (2002), Owen (2013a), Tissot and Prieur (2015) and Gilquin et al. (2019) for some of those issues.

Sobol' indices provide a global sensitivity analysis, while derivatives are used for a local sensitivity analysis. Sobol' and Kucherenko (2009) show that

$$\underline{\tau}_j^2 \leqslant \frac{1}{\pi^2}\int_{[0,1]^d}\Bigl(\frac{\partial f}{\partial x_j}\Bigr)^2 \,\mathrm{d}\boldsymbol{x}$$

connecting the two notions. Kucherenko and Iooss (2017) have more results of this type including ones for Gaussian random variables.

There are philosophical reasons to prefer the Shapley value from economics and game theory (Shapley, 1953) to Sobol' indices as a way to measure the importance of independent inputs x_j to the function $f(\boldsymbol{x})$. In this context, the Shapley value for x_j is

$$\phi_j = \sum_{u:j\in u}\frac{\sigma_u^2}{|u|}.$$

It has $\sum_{j=1}^d \phi_j = \sigma^2$. The Shapley value shares σ_u^2 equally over all $j \in u$. By comparison, $\underline{\tau}_j^2$ has a zero coefficient on σ_u^2 if $|u| \geqslant 2$, while $\overline{\tau}_j^2$ counts all of σ_u^2 if $j \in u$. We easily find that $\underline{\tau}_j^2 \leqslant \phi_j \leqslant \overline{\tau}_j^2$. There are no simple identities that let us efficiently estimate the Shapley value, though we can efficiently estimate both $\underline{\tau}_j^2$ and $\overline{\tau}_j^2$ by Sobol' identities and they bracket ϕ_j. For more about how Shapley value relates to Sobol' indices, see Owen (2014), Song et al. (2016), Owen and Prieur (2017) and Iooss and Prieur (2017).

The anchored decomposition goes back at least to Sobol' (1969). The Möbius relation (A.21) for it is from Kuo et al. (2010). They consider very general types of decompositions with the ANOVA and the anchored decomposition as just two examples.

Throughout this appendix, the inputs to f have been independent random variables. Many real problems involve dependent variables but it is exceedingly challenging to make the ANOVA work in such cases. This problem has been considered by Chastaing et al. (2012), Hooker (2007) and Stone (1994) among many others.

Exercises

A.1. Prove equation (A.12).

A.2. Prove equation (A.16).

A.3. Suppose that $f(\boldsymbol{x})$ is a constant function on $[0,1]^d$ for $d \geqslant 1$. What then is its effective dimension according to Definition A.3?

A.4. Let $\boldsymbol{x}$ and $\boldsymbol{z}$ be independent $\mathbf{U}[0,1]^d$ random vectors.

a) If $d \geqslant 2$ and $j \neq k$, show that

$$\iint f(\boldsymbol{x}_j{:}\boldsymbol{z}_{-j})f(\boldsymbol{x}_{-k}{:}\boldsymbol{z}_k)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z} = \mu^2 + \underline{\tau}_{\{j,k\}}^2.$$

This is from Saltelli (2002).

b) Define

$$\Omega \equiv \frac{1}{2}\iint \Biggl(df(\boldsymbol{z}) - \sum_{j=1}^{d} f(\boldsymbol{x}_j{:}\boldsymbol{z}_{-j})\Biggr)\Biggl((d-2)f(\boldsymbol{x}) - \sum_{k=1}^{d} f(\boldsymbol{x}_{-k}{:}\boldsymbol{z}_k)\Biggr)\,\mathrm{d}\boldsymbol{x}\,\mathrm{d}\boldsymbol{z}.$$

Show that

$$\Omega = \sum_{|u|=2} \sigma_u^2.$$

Note that $df(\boldsymbol{z})$ in the integrand for Ω is $d \times f(\boldsymbol{z})$, i.e., the d there is the dimension, not a differential.

c) For $d = 1$ there are no two factor interactions and so Ω should be 0. Show that the integrand in the definition of Ω reduces to zero for $d = 1$.

A.5. The Sobol' g function on $[0,1]^d$ is

$$g(\boldsymbol{x}) = \prod_{j=1}^{d} g_j(x_j), \quad \text{for} \quad g_j(x) = \frac{|4x-2| + a_j}{1 + a_j}$$

where $a_j \neq -1$. Typically $a_j \geqslant 0$ with larger values of a_j making x_j less important. Note that $\int_0^1 g_j(x)\,\mathrm{d}x = 1$.

a) Find a closed form expression for the ANOVA terms g_u of this function.

b) Find a closed form expression for $\sigma_u^2(g) = \mathrm{Var}(g_u(\boldsymbol{x}))$.

c) For $d = 10$ and $a_j = (j-1)^2$ find the true value of $\underline{\tau}_j^2(g) = \sigma_j^2(g)$ for $j = 1, \dots, 10$.

d) Compute plain Monte Carlo estimates for $\underline{\tau}_j^2$ for $j = 1, \dots, 10$ using (A.15), using (A.18) and using (A.19). For each $j = 1, \dots, 10$ determine which estimator you think is best, and explain how you decided. For sake of simplicity, pretend that the only cost to the user is the number of times that they must evaluate g and that the goal is to estimate $\underline{\tau}_j^2$ with a small squared error.

Bibliography

Box, G. E. P., Hunter, W. G., and Hunter, J. S. (2005). *Statistics for Experimenters: Design, Innovation, and Discovery.* Wiley, New York, 2nd edition.

Caflisch, R. E., Morokoff, W., and Owen, A. B. (1997). Valuation of mortgage backed securities using Brownian bridges to reduce effective dimension. *Journal of Computational Finance*, 1(1):27–46.

Chastaing, G., Gamboa, F., and Prieur, C. (2012). Generalized Hoeffding-Sobol' decomposition for dependent variables-application to sensitivity analysis. *Electronic Journal of Statistics*, 6:2420–2448.

Efron, B. and Stein, C. (1981). The jackknife estimate of variance. *Annals of Statistics*, 9(3):586–596.

Fisher, R. A. and Mackenzie, W. A. (1923). The manurial response of different potato varieties. *Journal of Agricultural Science*, 13(3):311–320.

Fruth, J., Roustant, O., and Kuhnt, S. (2014). Total interaction index: A variance-based sensitivity index for second-order interaction screening. *Journal of Statistical Planning and Inference*, 147:212–223.

Gilquin, L., Arnaud, E., Prieur, C., and Janon, A. (2019). Making the best use of permutations to compute sensitivity indices with replicated orthogonal arrays. *Reliability Engineering & System Safety*, 187:28–39.

Hoeffding, W. (1948). A class of statistics with asymptotically normal distribution. *Annals of Mathematical Statistics*, 19(3):293–325.

Hooker, G. (2004). Discovering additive structure in black box functions. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '04, pages 575–580, New York. ACM.

Hooker, G. (2007). Generalized functional ANOVA diagnostics for high-dimensional functions of dependent variables. *Journal of Computational and Graphical Statistics*, 16(3):709–732.

Iooss, B. and Prieur, C. (2017). Shapley effects for sensitivity analysis with dependent inputs: comparisons with Sobol indices, numerical estimation and applications. Technical Report arXiv:1707.01334, Université Grenoble Alpes. To appear in International Journal for Uncertainty Quantification.

Ishigami, T. and Homma, T. (1990). An importance quantification technique in uncertainty analysis for computer models. In *Proceedings of the First International Symposium on Uncertainty Modeling and Analysis*, pages 398–403.

Janon, A., Klein, T., Lagnoux, A., Nodet, M., and Prieur, C. (2014). Asymptotic normality and efficiency of two Sobol' index estimators. *ESAIM: Probability and Statistics*, 18:342–364.

Kucherenko, S. and Iooss, B. (2017). Derivative-based global sensitivity measures. In Ghanem, R., Higdon, D., and Owhadi, H., editors, *Handbook of uncertainty quantification*, pages 1241–1263. Springer, Cham, Switzerland.

Kuo, F., Sloan, I., Wasilkowski, G., and Woźniakowski, H. (2010). On decompositions of multivariate functions. *Mathematics of computation*, 79(270):953–966.

Liu, R. and Owen, A. B. (2006). Estimating mean dimensionality of analysis of variance decompositions. *Journal of the American Statistical Association*, 101(474):712–721.

Mauntz, W. (2002). Global sensitivity analysis of general nonlinear systems. Master's thesis, Imperial College. Supervisors: C. Pantelides and S. Kucherenko.

Oakley, J. E. and O'Hagan, A. (2004). Probabilistic sensitivity analysis of complex models: a Bayesian approach. *Journal of the Royal Statistical Society, Series B*, 66(3):751–769.

Owen, A. B. (1998). Latin supercube sampling for very high dimensional simulations. *ACM Transactions on Modeling and Computer Simulation*, 8(2):71–102.

Owen, A. B. (2013a). Better estimation of small Sobol' sensitivity indices. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 23(2):11.

Owen, A. B. (2013b). Variance components and generalized Sobol' indices. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):19–41.

Owen, A. B. (2014). Sobol' indices and Shapley value. *Journal on Uncertainty Quantification*, 2:245–251.

Owen, A. B. (2019). Effective dimension of some weighted pre-Sobolev spaces with dominating mixed partial derivatives. *SIAM Journal on Numerical Analysis*, 57(2):547–562.

Owen, A. B. and Prieur, C. (2017). On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):986–1002.

Richtmyer, R. D. (1952). The evaluation of definite integrals, and a quasi-Monte Carlo method based on the properties of algebraic numbers. Technical Report LA-1342, University of California.

Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145:280–297.

Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., and Tarantola, S. (2008). *Global Sensitivity Analysis. The Primer.* John Wiley & Sons, Ltd, New York.

Shapley, L. S. (1953). A value for n-person games. In Kuhn, H. W. and Tucker, A. W., editors, *Contribution to the Theory of Games II (Annals of Mathematics Studies 28)*, pages 307–317. Princeton University Press, Princeton, NJ.

Sobol', I. M. (1969). *Multidimensional Quadrature Formulas and Haar Functions.* Nauka, Moscow. (In Russian).

Sobol', I. M. (1990). On sensitivity estimation for nonlinear mathematical models. *Matematicheskoe Modelirovanie*, 2(1):112–118. (In Russian).

Sobol', I. M. (1993). Sensitivity estimates for nonlinear mathematical models. *Mathematical Modeling and Computational Experiment*, 1:407–414.

Sobol', I. M. and Kucherenko, S. (2009). Derivative based global sensitivity measures and their link with global sensitivity indices. *Mathematics and Computers in Simulation*, 10:3009–3017.

Song, E., Nelson, B. L., and Staum, J. (2016). Shapley effects for global sensitivity analysis: Theory and computation. *SIAM/ASA Journal on Uncertainty Quantification*, 4(1):1060–1083.

Stein, M. (1987). Large sample properties of simulations using Latin hypercube sampling. *Technometrics*, 29(2):143–51.

Stone, C. J. (1994). The use of polynomial splines and their tensor products in multivariate function estimation. *The Annals of Statistics*, 22:118–171.

Tissot, J.-Y. and Prieur, C. (2015). A randomized orthogonal array-based procedure for the estimation of first-and second-order Sobol' indices. *Journal of Statistical Computation and Simulation*, 85(7):1358–1381.

Wasilkowski, G. (2019). ε-superposition and truncation dimensions and multivariate decomposition method for ∞-variate linear problems. In Hickernell, F. J. and Kritzer, P., editors, *Multivariate Algorithms and Information-Based Complexity*, Berlin/Boston. De Gruyter.

Wu, C. F. J. and Hamada, M. S. (2011). *Experiments: planning, analysis, and optimization*, volume 552. John Wiley & Sons.