Contrast weights in flexible factorial design with multiple groups of subjects

Jan Gläscher[1] and Darren Gitelman[2]

[1] Div. of Humanities and Social Science, Caltech
[2] Dept. of Neurology and Radiology, Northwestern

May 30, 2008

contact: glascher@hss.caltech.edu

Contrast weights for factorial designs that include a group factor (e.g. 2 different groups of subjects) are often a cause of confusion because the contrast weights that have to be entered in SPM diverge from the simple scheme for a single group. The reason for this is that the group sizes have to be taken into account when deriving contrast weights in order to yield an estimable contrast.

The situation is further complicated because the definition of factors in an experiment is dissociated from the configuration of the design matrix in SPM5. While the former are defined under the "Factors" section of the flexible factorial design, the actual regressors of the design matrix are configured under "Main Effects and Interactions". SPM5 does not impose any restriction on which main effect or interaction to include in the design matrix, but the decision affects the necessary contrast weights dramatically[1]. Although termed "Main Effects and Interactions", this part only configures the design matrix, but does not provide the contrast weights to test for these factorial effects. Understanding how to enter the contrast weights is the goal of this tutorial.

This tutorial provides algorithms for computing the contrast weights for the main effects and interactions in experiments that includes a group factor (e. g. patients and control subjects). We will demonstrate these for various design matrices that include regressors for different combinations of main effects and interactions (as configured in the above mentioned section). Finally, we will provide some preliminary advice regarding which effects should be included in the design matrix.

We begin this tutorial with some general comment about configuring a flexible factorial design in SPM5. We will then review contrast weights for a single group of subjects before moving ahead to the more complex multi-groups. Contrast vectors are typeset in `monospaced Courier Font`, but we will omit the brackets (`[]`) from the contrast vector notation. Factor names are typeset in **bold typeface.**

---

[1] Note that SPM5 allows up to 3 factors and 2-way interactions to be estimated using the flexible factorial interface.
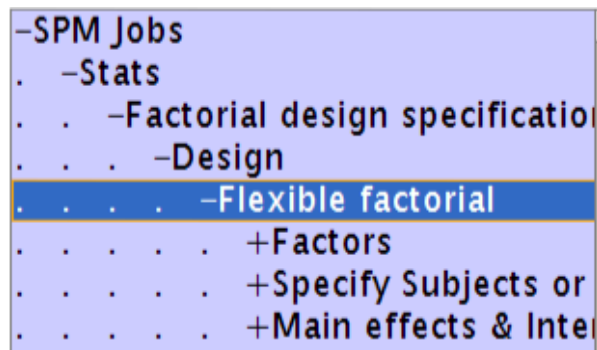
Figure 1. Configuration of a flexible factorial design

General comments about configuring a Flexible Factorial Design

There are 3 essential parts in configuring a flexible factorial design in SPM5: Factors, Subjects, and Main effects & Interactions (see Figure 1):

- Factors: Different experimental effects are configured in this section (e.g. **subject**, **condition**, **group**). The options on "Independence" and "Variance" can be used to formulate expected violations of the sphericity assumption (i.e. whether the errors of the model are expected to be identically and independently distributed (i.i.d.) or whether SPM5 should configure a set of variance components in order to model the expected deviation from i.i.d.) (see below for common settings for "Independence" and "Variance").
- Subjects: This section provides all images for each subject as well as a matrix that maps a subject's images on the different factor levels.
- Main Effects & Interactions: These options affect the regressors that will be included in the design matrix (but it does not provide the contrast weights to test these factorial effects).

Factors

Each factor is identified by a name (e.g. **group**, **condition**). The number of factors determines the size of the factor matrix in the Subjects section. The "factor matrix" specifies the mapping between scans and factors, and is configured under "Conditions" in the Subjects section.

The factor **subject** is a special name for SPM5 in the sense that its factor levels can be left out of the factor matrix.

"Independence" and "Variance" have to be configured for each factor separately. These settings are used to express the expected distribution of the errors (whether they are i.i.d. or not). "Independence" refers to whether the errors are independent between the different factor levels; "Variance" refers to whether the error variances are equal or unequal between the factor levels. Typical settings are:

- Within-subjects factor: independence – no, variance – equal.
  This is typically a factor of an experimental condition (e.g. different stimulus intensities). All subjects are presented with every level of this factor. With "Independence – no" we expect that the different factor levels are correlated within a subject (SPM2 calls this option "correlated repeated measures"). "Variance – equal" is a reasonable choice, because we usually do not expect different variances between the levels of a within-subject factor.

- Between-subjects factor: independence – yes, variance – equal.
  Typically, this is the "subject" factor, which is used to model the subject constants (see below). Usually, we expect that the errors between different subject are not correlated (they are different individuals after all), but that the error variances are equal, because the subjects are sampled from the same underlying population. In general, it is good practice to configure the factor **subject**, even if it is not included in the design matrix, because it provides the possibility to change the configurations for "independence" and "variance" over all subjects if necessary.
- Between-group factors: independence – yes, variance – unequal.
  We usually expect that the errors between different groups are independent (two different underlying populations). In addition, we usually also expect differences in error variances, especially between patients and control subjects for the same reason.

The effects of these settings (i.e. which variance components are configured by SPM5) can be seen under Review →Design→Explore→Covariance Structure

In order to avoid confusion, the slowest rotating factor should be configured first, then the next slowest, and so on. If the factor **subject** is included, then this should be factor 1 (because it is a special factor), followed by a potential factor **group** as factor 2, and then followed by various experimental factors sorted ascendingly from slowest to fastest rotating.

Subjects

In the subject section, the input images for each subject have to be specified along with a factor matrix that maps the images onto the different factor levels. This matrix has the dimensions nImages $\times$ nFactors. Factor levels for the factor **subject** can be left out, and then the factor matrix has the dimension nImages $\times$ (nFactors-1). Each factor (and its levels) is configured in a single <u>column</u> (!) of the factor matrix.

For example, in a 2x3 factorial design with factors **A** (2 levels) and **B** (3 levels), six images have to specified for each subject. If the order of the images corresponds to A1 A2 A3 B1 B2 B3, then each subject's factor matrix should look like:

```
1 1
1 2
1 3
2 1
2 2
2 3
```

This would be entered under Conditions using this command:
```
1 1; 1 2; 1 3; 2 1; 2 2; 2 3
```

Main Effects & Interactions

The settings in this section affect the regressors that are included in the design matrix. The different effects of these settings will be explored in greater detail below. In general, for each main effect, SPM will include as many regressors as the factor has levels; for each interaction, SPM will includes as many regressors as the product of the number of the levels of both factors (e.g. in a 2×3 design, the number of regressors for an interaction is 6). At this point, it is only possible to configure 2-way interactions.

One group of subjects

We will consider a 2×3 factorial design with the (within-subject) factor **A** (2 levels) and **B** (3 levels) in a sample of 11 subjects. The main effect of **A** is given by a comparison of level 1 > level 2, whereas the main effect of **B** is a linear increase across the 3 levels (level 1 < level 2 < level 3). Because there are 3 levels in Factor **B**, this main effect amounts to a differential contrast B3–B1 (B2 receives a weight of 0 in this contrast). For the more complex design with main effects and interactions included in the design matrix, we will use a coloring scheme to visualize the different weights of the contrast vector that pertain to the different partitions of the design matrix configured by the main effects and interactions.

We provide contrast weights for each main effect and interaction for all the designs below. In addition, we also compute contrast weights to test for one or two individual regressors within the ANOVA design. Although these may be more appropriately tested in simpler random effects design (e.g. one sample t-tests), a contrast for a single regressor adds a lot of flexibility to exploring the results within a single 2nd level analysis. These contrasts for one or two regressors are just examples that can be easily generalized to multiple regressors.

For all the following designs, these factors are configured:
Factor 1: **subject**
Factor 2: **A** (2 levels)
Factor 3: **B** (3 levels)

The ordering of each subject's images is: A1 A2 A3 B1 B2 B3

Accordingly, the factor matrix for each subject is (1st column = **A**, 2nd column = **B**):
```
1 1
1 2
1 3
2 1
2 2
2 3
```



Design 1

Main Effects and Interaction:
1. Interaction: 2  3  (**A** x **B**)

Contrast Weights
1. Main effect of **A**:                      1   1   1  -1  -1  -1
2. Main effect of **B** (level 3-level 1):   -1   0   1  -1   0   1
3. Interaction **A** × **B** (level 3-level 1): -1   0   1   1   0  -1
4. Test for a single regressor (A1B2):        0   1   0   0   0   0

4

**Design 2**

Input Images

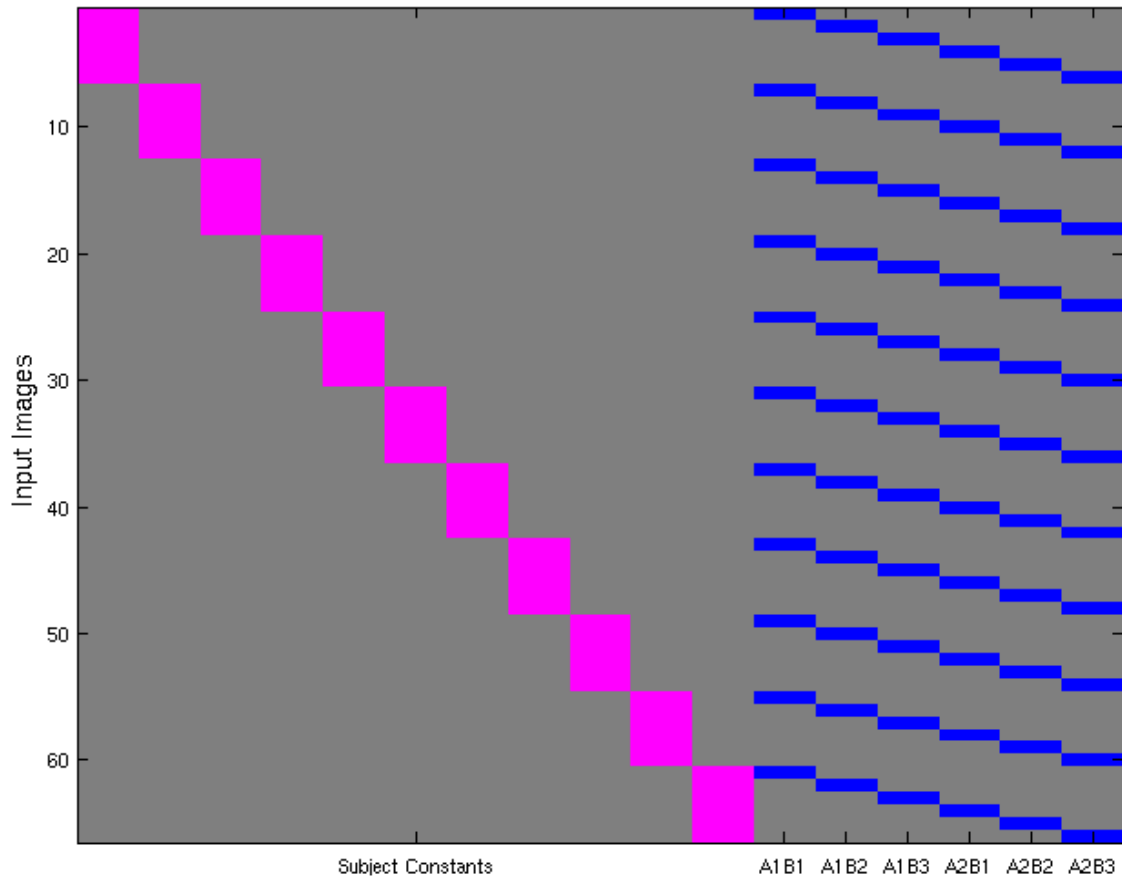A1  A2  B1  B2  B3  A1B1  A1B2  A1B3  A2B1  A2B2  A2B3

Main Effects and Interaction:
1. Main effect: 2 (**A**)
2. Main effect: 3 (**B**)
3. Interaction: 2  3  (**A** $\times$ **B**)

Contrast Weights
1. Main effect of **A**:        1  -1   0   0   0   ones(1,3)/3  -ones(1,3)/3
2. Main effect of **B**:        0   0  -1   0   1   [-1 0 1]*[1/2]  [-1 0 1]*[1/2]
3. Interaction **A** $\times$ **B**:   0   0   0   0   0   -1   0   1   1   0  -1
4. Test for a single regressor in main effect of **A** (e.g. A1)
                                1   0   ones(1,3)/3  ones(1,3)/3  zeros(1,3)
5. Test for a single regressor in main effect of **B** (e.g. B2)
                                0.5 0.5 0 1 0 0 0.5 0 0 0.5 0
6. Test for a single regressor in interaction **A** x **B** (e.g. A1B3)
                                1 0 0 0 1 0 0 1 0 0 0
                                -

5

Main Effects and Interaction:
1. Main effect: `1` (**subject**)
2. Interaction: `2  3` (**A** × **B**)

Contrast Weights
1. Main effect of **A**:   `zeros(1,11)  1   1   1  -1  -1  -1`
2. Main effect of **B**:   `zeros(1,11) -1   0   1  -1   0   1`
3. Interaction **A** × **B**:   `zeros(1,11) -1   0   1   1   0  -1`
4. Test for a single regressor (e.g. A2B2)
   `ones(1,11)*[1/11]  0   0   0   0   1   0`
5. Test for a two single regressor (e.g. A1B2 and A2B2)
   `ones(1,11)*[2/11]  0   1   0   0   1   0`

# Two groups of subjects

In the case of two or more groups, a factor **group** has to be included as a factor. In this case, the contrast weights for main effects are always influenced by the number of subjects in each group. For this part of the tutorial, we consider a 2×3 factorial design with the (between-group) factor **group** (2 levels) and the (within-subject) factor **condition** (2 levels). There are 6 subjects in the first group and 5 subjects in the second group.[2] There are 3 images per subject corresponding to the 3 levels of the factor **condition**. For the more complex design with main effect and interactions we will again use a coloring scheme to visualize the contrast weight for the different partitions of the design matrix.

For all the designs below, the following factors were configured:
Factor 1: **subject**      independence – yes, variance – equal
Factor 2: **group**       independence – yes, variance – unequal
Factor 3: **condition**   independence – no, variance – equal

The factor matrices are (1st column = **group**, 2nd column = **condition**):
- subjects in group 1
    1  1
    1  2
    1  3
- subjects in group 2
    2  1
    2  2
    2  3

Let's define some MATLAB variables to achieve a general algorithm for the contrast weights below:

n1 = 6;  (number of subjects in group 1)
n2 = 5;  (number of subjects in group 2)
nc = 3;  (number of levels in **condition** factor)
ng = 2;  (number of groups)

MEc = [1:nc]-mean(1:nc);  (main effect of condition, here: [-1 0 1])
MEg = [1 -1];  (main effect of group: Group 1 > Group 2)

---

2   In contrast to the previous section, we only have one experimental factor (**condition**) here. However, the contrast weights below can be easily generalized to 2 or more experimental factors. However, due to the current limitation to 3 factors in SPM5, one would have to drop the factor **subject** from these designs.

**Design 1**



Main effects & interactions:
1. main effect: `1` (**subject**)
2. interaction: `2 3` (**group** $\times$ **condition**)

Contrast weights
1. Main effect of **group**:
   `ones(1,n1)/n1 -ones(1,n2)/n2 ones(1,nc)/nc -ones(1,nc)/nc`
2. Main effect of **condition**:
   `zeros(1,n1+n2) MEc MEc`
3. Interaction **group** $\times$ **condition**:
   `zeros(1,n1+n2) MEc -MEc`
4. Test for single regressor in one group (e.g. G1C2)
   `ones(1,n1)/n1 zeros(1,n2) 0 1 0 0 0 0`
5. Test for two single regressors in both groups (e.g. G1C2 and G2C2)
   `ones(1,n1)/n1 ones(1,n2)/n2 0 1 0 0 1 0`
6. Average effect of **condition**:
   `ones(1/n1+n2)/(n1+n2) ones(1,3)*n1/((n1+n2)*nc)`
   `ones(1,3)*n2/((n1+n2)*nc)`

**Design 2**



Main effects and interactions:

1. Main effect: 2 (**group**)
2. Main effect: 3 (**condition**)
3. Interaction: 2  3 (**group** × **condition**)

Contrast weights:

1. Main effect of **group**:
   `MEg zeros(1,nc) ones(1,nc)/nc -ones(1,nc)/nc`
2. Main effect of **condition**:
   `zeros(1,ng) MEc MEc*[n1/(n1+n2)] MEc*[n2/(n1+n2)]`
3. Interaction **group** × **condition**:
   `zeros(1,ng) zeros(1,nc) MEc -MEc`
4. Test for a single regressor in main effect **group** (e.g. G1)
   `1 0 ones(1,nc)/nc ones(1,nc)/nc zeros(1,nc)`
5. Test for a single regressor in main effect **condition** (e.g. C1):
   `n1/(n1+n2) n2/(n1+n2) 1 0 0 n1/(n1+n2) 0 0 n1/(n1+n2) 0 0`
6. Test for a single regressor in interaction **group** × **condition** (e.g. G2C2)
   `0 1 0 1 0 0 0 0 0 1 0`
7. Test for two single regressors in interaction **group** × **condition** (e.g. G1C1 and G2C1)
   `n1/(n1+n2) n2/(n1+n2) 1 0 0 n1/(n1+n2) 0 0 n2/(n1+n2) 0 0`

Contrasts 5 & 7 are actually identical which is an effect of the over-parameterized design. When looking at 2 conditions in an interaction, one has to weight the same condition in the main effect and vice versa. Therefore, as in this example, these two contrasts are indistinguishable.

**Design 3**



Main effect & interactions:
1. Main effect: 1 (**subject**)
2. Main effect: 2 (**group**)
3. Main effect: 3 (**condition**)
4. Interaction: 2 3 (**group** × **condition**)

Contrast weights:
1. Main effect of **group**:
   `ones(1,n1)/n1 -ones(1,n2)/n2 MEg zeros(1,nc) ones(1,nc)/nc -ones(1,nc)/nc`
2. Main effect of **condition**
   `zeros(1,n1+n2) zeros(1,ng) MEc MEc*[n1/(n1+n2)] MEc*[n2/(n1+n2)]`
3. Interaction **group** x **condition**:
   `zeros(1,n1+n2) zeros(1,ng) zeros(1,nc) MEc -MEc`
4. Test for a single regressor in main effect of **group** (e.g. G1)
   `ones(1,n1)/n1 zeros(1,n2) 1 0 ones(1,nc)/nc ones(1,nc)/nc zeros(1,nc)`
5. Test for a single regressor in main effect of **condition** (e.g. C2):
   `ones(1,n1+n2)/(n1+n2) n1/(n1+n2) n2/(n1+n2) 0 1 0 0 n1/(n1+n2) 0 0 n2/(n1+n2) 0`
6. Test for a single regressor in the interaction **group** × **condition** (e.g. G1C2)
   `ones(1,n1)/n1 zeros(1,n2) 1 0 0 1 0 0 1 0 0 0 0`
7. Test for two single regressor in the interaction **group** × **condition** (e.g. G1C1 and G2C1)
   `ones(1,n1+n2)/(n1+n2) n1/(n1+n2) n2/(n1+n2) 1 0 0 n1/(n1+n2) 0 0 n2/(n1+n2) 0 0`

# Which main effects/interactions should be included in the design matrix?

In theory, the inclusion of various main effects should increase the sensitivity for an interaction effect because the main effects regressors model, and hence remove, that part of the variance that is not crucial to the interaction effect. For instance, in the case of a **group** $\times$ **condition** interaction, the inclusion of the main effect of **group** models the general differences between patients and control subjects that are *not* specific to the experimental task. Similarly, the main effect of **condition** models differences between the levels of the **condition** factor that are common among patients and control subject. By "removing" these two main effects from the interaction effect, the sensitivity to detect an interaction should be increased (which is, of course, encoded in the regressors for the **group** $\times$ **condition** interaction).

The following is paraphrased from one of Darren's emails (S=main effect **subject**, G= main effect **group**, C= main effect **condition**, I= **group** x **condition** interaction):

> I ran a number of designs that modeled the individual factors and their various combinations (S, G, C, S+G, etc.). It seems that including the subject effects is critical and dramatically improves the sensitivity of the model. By this I mean t-scores that are 4-9 units greater (but one's mileage may vary depending on the design, subjects, etc.). Adding other factors only slightly affects the various results. So, for example, estimating a group main effect was very similar for the models (S, S+G, S+G+I, and S+G+C+I). There were some differences in t-scores and p-values but these were generally < 0.2. The more elaborated models (S+G+I and S+G+C+I) generally had slightly lower t-scores and higher p-values, both because of lower degrees of freedom (-2) and a greater number of resels in the analysis. The greater number of resels was because of a lower number of voxels per resel.
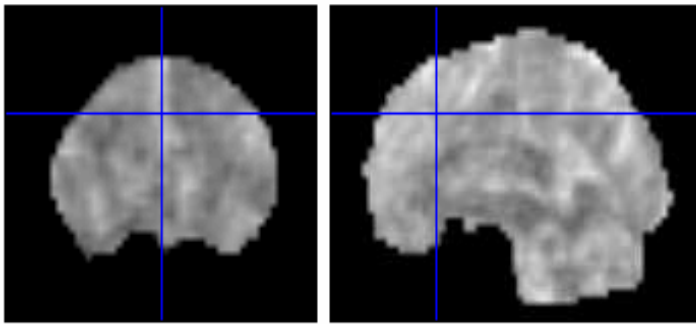
Tom Nichols provided some insight into this behavior (see this thread on the SPM mailing list: http://www.jiscmail.ac.uk/cgi-bin/webadmin?A2=ind0803&L=spm&O=D&F=&S=&P=9865).

> Variation that is not modeled in the design ends up in the residuals, and the smoothness is estimated from the residuals. All other things being equal, more large-scale structure in the residuals will make the residuals "smoother." This will increase the estimate of the FWHM and decrease the RESEL count. (If the data weren't very smooth, and the unmodeled effects were really 'sharp', it is conceivable that such structure could deflate the smoothness instead, but that's the less likely scenario). You might want to look at the image of local smoothness for model 3, i.e. compute FWHM = RPV^(-1/3) with ImCalc. I bet you'll see areas of greater smoothness that correspond to areas in model 1 where you find a significant interaction (or other) effects.

These effects can be seen by looking at the FWHM images, which is calculated from the RPV (resels per voxel) image using ImCalc and the formula (i1.^(-1/3)). Notice that the image at the top has much smoother parts of the image (whiter) than the one at the bottom.
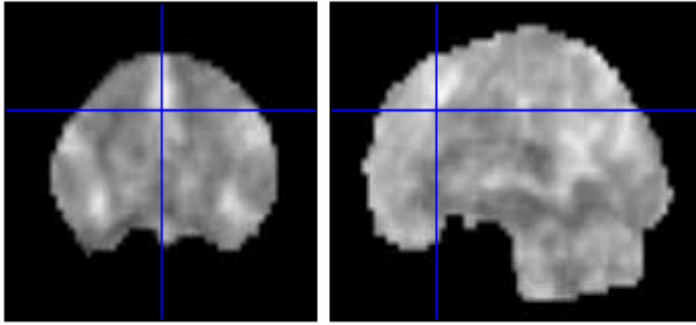
Simpler model



More complex model