Meta-Analysis of Functional Imaging Data Using Replicator Dynamics

Jane Neumann,* Gabriele Lohmann, Jan Derrfuss,
and D. Yves von Cramon

Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

Abstract: Despite the rapidly growing number of meta-analyses in functional neuroimaging, the field lacks formal mathematical tools for the quantitative and qualitative evaluation of meta-analytic data. We propose to use replicator dynamics in the meta-analysis of functional imaging data to address an important aspect of neuroimaging research, the search for functional networks of cortical areas that underlie a specific cognitive task. The replicator process requires as input only a list of activation locations, and it results in a network of locations that jointly show significant activation in most studies included in the meta-analysis. These locations are likely to play a critical role in solving the investigated cognitive task. Our method was applied to a meta-analysis of the Stroop interference task using data provided by the publicly accessible database BrainMap DBJ. *Hum Brain Mapp* 25:165–173, 2005. © 2005 Wiley-Liss, Inc.

Key words: functional neuroimaging; fMRI; mathematical analysis; replicator dynamics; meta-analysis; BrainMap DBJ

INTRODUCTION

Functional neuroimaging provides a powerful tool for the detection of task-related neural activity in the human brain. Functional magnetic resonance imaging (fMRI) in particular, with its unsurpassed spatial resolution, allows us to determine the location of activation foci corresponding to a cognitive task in question. While a subject is presented with a stimulus or performing a particular task in the MRI scanner, changes in the blood oxygenation level, which are supposed to be linked to neural activation, are recorded. Results of the statistical evaluation of these measurements are then mapped onto the subject's anatomical image.

Results from single imaging sessions have to be treated with caution, as they can vary considerably not only between subjects, but even between imaging sessions of the

same subject [Maitra et al., 2002; McGonigle et al., 2000; Neumann et al., 2003; Noll et al., 1997]. Consequently, generalizations about the neural correlates underlying a specific cognitive task can be drawn only from repeated measurements of many subjects and from comparisons of different experiments to exclude any bias introduced by a particular scanning hardware, analysis procedure, or experimental setup. Sophisticated spatial normalization and registration methods [Crivello et al., 2002; Crum et al., 2003; Friston et al., 1995], a standardized coordinate system [Talairach and Tournoux, 1988], and the increasing availability of experimental results in publicly accessible databases such as BrainMap DBJ [Fox and Lancaster, 2002] have opened up the exciting possibility to compare and integrate results from many studies carried out independently in laboratories around the globe. An excellent overview of such meta-analyses in the context of functional neuroimaging can be found in [Fox et al. 1998]. More recent publications include Farah and Aguirre [1999]; Indefrey and Levelt [2000]; Phan et al. [2002], and Ullsperger and von Cramon [2004].

Despite the rapidly growing number of meta-analyses in functional imaging, the field lacks formal mathematical tools for the qualitative and quantitative evaluation of meta-analytic data. In the most simple and still most often reported

*Correspondence to: Dr. Jane Neumann, Max Planck Institute for Human Cognitive and Brain Sciences, Stephanstrasse 1a, D-04103 Leipzig, Germany. E-mail: neumann@cbs.mpg.de

Received for publication 16 March 2004; Accepted 27 August 2004
DOI: 10.1002/hbm.20133

Published online in Wiley InterScience (www.interscience.wiley.com).

case, a meta-analysis consists of plotting activation foci from a number of experiments into a single figure. Visual inspection is then used to detect anatomical regions that show significant activation in all or most of the included studies. Unfortunately, this methodology introduces a high degree of subjectivity into the interpretation of the findings. More objective results can be obtained by calculating activation probability distributions from reported activation foci and testing them against some given threshold or against randomly generated activation foci representing noise [Chein et al., 2002; Turkeltaub et al., 2002; Wager et al., 2003]. This might not only reveal clusters of activation that are hard to detect by visual inspection; these methods also take into account the unavoidable but considerable intersubject and interlaboratory variability of activation locations and allow to correct for variable sample sizes of different studies [Fox et al., 1999].

An important aspect of functional imaging research augmenting the exact localization of activation foci is the search for functional networks of cortical regions related to a specific cognitive task. We are interested in finding cortical areas that jointly show significant activation for a particular cognitive task and therefore can be thought of as belonging to a functional network involved in solving the task in question. Finding such networks of co-occurring activation foci in meta-analytic data by means of the existing methodology is problematic. The number of co-occurrences of two or three activation foci in the investigated studies might be detectable upon visual inspection, for example by color coding or labeling of activation foci. However, as the potential number of activation foci within the same functional network increases, the process of finding co-occurrences quickly becomes too complex to manage manually.

We propose to use replicator dynamics as a tool for finding networks of cortical areas that jointly show significant activation across a number of studies. The proposed method is easy to implement and requires as input only a list of activation maxima from the individual experiments included in the analysis. It can therefore be applied to different imaging techniques with good spatial resolution such as fMRI or positron emission tomography (PET).

Here, we introduce the principle of replicator dynamics and its application to the meta-analysis of neuroimaging data. We then present results from a meta-analysis of fMRI experiments addressing the well-known Stroop interference task [Stroop, 1935].

MATERIALS AND METHODS

Replicator Equations

Consider the following dynamical system:

$$\frac{d}{dt}x_i(t) = x_i(t)[(\mathbf{C}\mathbf{x}(t))_i - \mathbf{x}(t)^T\mathbf{C}\mathbf{x}(t)], \quad i = 1 \dots n, \quad (1)$$

or the corresponding discrete-time version,

$$x_i(t+1) = x_i(t) \frac{(\mathbf{C}\mathbf{x}(t))_i}{\mathbf{x}(t)^T\mathbf{C}\mathbf{x}(t)}. \quad (2)$$

These equations are known from theoretical biology as replicator equations [Hofbauer and Sigmund, 1988b; Schuster and Sigmund, 1983], as they can be used to model evolution of self-replicating interacting entities over time. Here, replicators $i = 1 \dots n$ are entities in the evolutionary process, each of which is associated with a fitness, encoded in the matrix $\mathbf{C}$, and with a proportion x_i indicating the percentage of the population that is identical to i . The term $\mathbf{x}^T\mathbf{C}\mathbf{x}$ denotes the mean fitness of the population.

The dynamics of such replicator systems are described by the fundamental theorem of natural selection, which goes back in its original formulation to Fisher [1930] and Kimura [1958]. It states that if $\mathbf{C}$ is a nonnegative, real-valued, symmetric matrix, then the function $\mathbf{x}(t)^T\mathbf{C}\mathbf{x}(t)$ is strictly increasing with increasing t along any nonstationary trajectory $\mathbf{x}(t)$ under both continuous-time (1) and discrete-time (2) replicator dynamics. Moreover, any such trajectory converges toward a stationary point $\bar{\mathbf{x}}$, and a vector $\bar{\mathbf{x}} \in \mathbb{R}^n$ with $\bar{x}_i \geq 0$, $i = 1 \dots n$, and $\sum \bar{x}_i = 1$ is asymptotically stable, if and only if it is a strict local maximizer of $\mathbf{x}^T\mathbf{C}\mathbf{x}$ [Hofbauer and Sigmund, 1988b]. In other words, starting from an initial vector $\mathbf{x}$, the replicator process evolves toward a vector $\bar{\mathbf{x}}$, which maximizes $\mathbf{x}^T\mathbf{C}\mathbf{x}$. During this process, replicators with a particularly good fitness will be assigned a high proportion, whereas the proportion of replicators with poor fitness will decrease. Eventually, the vector $\mathbf{x}$ becomes stable, i.e., the proportions x_i no longer change from iteration to iteration. The replicators with a particularly high proportion then form a dominant group or network, whereas the proportion of other replicators will have decreased and in most cases become negligible.

Membership in the group of dominant replicators is a fuzzy concept, because the proportion of a replicator can take on any value between 0 and 1. However, we can define a membership function that classifies each replicator as belonging to this group or not, based on the replicator's proportion after reaching stationarity. One choice for this function is to eliminate all replicators from the network whose proportion drops close to zero during the replicator process. Such a test is difficult to implement without introducing an arbitrary threshold defining the vicinity of zero. A second, numerically more stable function was suggested by Lohmann and Bohn [2002]. We can declare a replicator to be a member of the dominant network, if its proportion exceeds the average proportion of a replicator in the population, i.e., replicator i is a member of the network, if $x_i > 1/n$. As will be described in the following section, in our application the vector $\mathbf{x}$ will be initialized as $x_i = 1/n$ for all $i = 1 \dots n$. The membership function therefore determines those replicators whose proportions improve during the process above the initial state, the equality of all replicators.

Application to Functional Imaging Data

The application of replicator equations is not restricted to modeling natural evolutionary processes. Rather, they de-

scribe a wide class of dynamical systems and have been applied successfully, for example, in game theory [Hofbauer and Sigmund, 1988a; Taylor and Jonker, 1978] and optimization [Menon et al., 1995; Pelillo, 2001], to complex problems such as matching hierarchical structures and finding maximum cliques in graphs [Bomze et al., 2000; Pelillo et al., 1999]. In the context of functional neuroimaging replicator dynamics have been used recently to analyze fMRI data from individual subjects [Lohmann and Bohn, 2002]. Here, voxels of a functional image were represented by replicators whose fitness was derived from pairwise correlations of individual time series.

We propose to apply replicator dynamics to meta-analytic data where lists of coordinates of activation maxima are often the only source of information available. First, activation foci are calculated from the lists of activation maxima. Each activation focus is then represented by a replicator. The fitness of each focus is derived from the number of co-occurrences with other activation foci in all experiments included in the analysis. The rationale behind this choice of fitness function is as follows.

We are interested in finding the underlying neural correlates for a particular cognitive task, which will usually include a number of different activation foci. It is safe to assume that cortical areas that play a critical role in solving this task will be observed as jointly activated in most experiments. Additional activation foci might appear in some experiments, given the differences in experimental setup, stimuli presentation, scanning hardware, and statistical evaluation of the considered experiments. Consequently, in a meta-analysis we want to find those cortical areas that are most often jointly activated and separate them from foci whose activation is the result of a particular experimental design or measurement noise. The fitness assigned to each activation focus should therefore be determined by the number or relative frequency of co-occurrences with other activation foci. This information can be captured in a co-occurrence matrix.

A co-occurrence matrix is a quadratic matrix with as many rows and columns as activation foci included in the replicator process. The individual entries report how often any pair of activation foci was found co-activated in the same study. To be more precise, for two activation foci i and j , the entry c_{ij} in the co-occurrence matrix

$$\mathbf{C} = \begin{pmatrix} c_{11} & c_{12} & \cdots & c_{1j} & \cdots & c_{1n} \\ c_{21} & c_{22} & \cdots & c_{2j} & \cdots & c_{2n} \\ \vdots & \vdots & & & & \\ c_{i1} & c_{i2} & \cdots & c_{ij} & \cdots & c_{in} \\ \vdots & \vdots & & & & \\ c_{n1} & c_{n2} & \cdots & c_{nj} & \cdots & c_{nn} \end{pmatrix} \quad (3)$$

represents the number of individual studies for which foci i and j have both been observed as activated. The resulting matrix is nonnegative, real-valued, and symmetric and thus meets the requirements for the theorem of natural selection.

The choice of values for the diagonal elements in a fitness matrix depends on the problem to be modeled. When modeling an evolutionary process, for example, an entry c_{ij} in the fitness matrix reflects the influence species i has on the growth of species j . Consequently, a diagonal element c_{ii} should reflect the contribution of species i to its own growth [Bouzat and Zanette, 2003]. In our application, one conceivable choice is the number of occurrences for each activation focus in all included studies. However, note that the number of times a region is activated is always at least as big as the number of times it is co-activated with other areas. The diagonal elements would therefore very much dominate the fitness matrix, and the fitness of a replicator would reflect mostly its overall occurrence rather than its co-occurrence with other replicators. The replicator process is then likely to favor activation foci that simply occur most often in the list of activation maxima.

We are interested in finding a network of cortical areas that are activated most often together. We therefore set $c_{ii} = 0$ for all $i = 1 \dots n$. This allows us to emphasize the co-activation of cortical areas rather than their mere individual activation. The number of occurrences of single foci still enters the replicator process as it is encoded implicitly in the number of co-occurrences with other activation foci.

In our context, the proportion of a replicator evolving during the replicator process can be interpreted as the “importance” of an activation focus for performing the experimental task. To avoid any bias in the input data, the proportion vector $\mathbf{x}$ should be initialized as $\mathbf{x}(0) = (1/n \ 1/n \dots 1/n)$ where n denotes the number of activation foci included in the analysis.

What happens now when a replicator process is applied to a number of activation foci? First, for each focus the number of joint activations with other foci is represented in the co-occurrence matrix $\mathbf{C}$. Equation (2) is then iteratively applied, where the new proportion of a replicator is determined from its individual fitness, its proportion, and the mean fitness of the population. More precisely, if for a focus i the term $(\mathbf{C}\mathbf{x}(t))_i$ exceeds the mean fitness $\mathbf{x}(t)^T \mathbf{C}\mathbf{x}(t)$, its proportion will increase; if it does not exceed the mean fitness, the proportion will decrease.

We can see that in the first step of the iteration the numerator in equation (2) only depends on the matrix $\mathbf{C}$ and on the overall number of replicators, because the initial proportion x_i is the same for all activation foci:

$$(\mathbf{C}\mathbf{x}(0))_i = \sum_j \mathbf{C}_{ij} x_j(0) = \sum_j \mathbf{C}_{ij} 1/n = 1/n \sum_j \mathbf{C}_{ij} \quad (4)$$

In other words, only the total number of joint activations with other foci determines the new proportion for a particular replicator. Consequently, two replicators i and j with the same number of co-occurrences with other activation foci will be assigned the same proportion in the first step of the iteration: $x_i(1) = x_j(1)$. For foci that relatively often are jointly activated with other activation foci, the ratio of individual and mean fitness is high. Such foci will be assigned a

high new proportion in the population. In contrast, foci with only a few co-occurrences will be assigned a lower proportion.

In the following steps of the iteration, the term $(Cx(t))_i$ no longer only depends on the total number of joint activations for a focus, it now also becomes important with which other foci it occurs. A focus i that co-occurs with other “important” foci, i.e., with foci that were assigned a high proportion in the previous step, will increase its own proportion more than would a focus j that is equally often activated jointly but with less “important” other foci.

As the theorem of natural selection states, the mean fitness $x(t)^T Cx(t)$ is maximized as the process goes on. Therefore, only the fitness and the proportion of the most important foci will be high enough to exceed the mean fitness in later stages of the replicator process and, as a result, a network of the most important activation foci will emerge. The replicator process becomes stationary when $x^T Cx$ reaches a maximum. The changes in the value of $x^T Cx$ can be used as criterion to stop the iterative process. The process can be terminated when the changes in the value of $x^T Cx$ between successive iterations become negligible, i.e., $x(t+1)^T Cx(t+1) - x(t)^T Cx(t) < \epsilon$ for small ϵ .

From the theorem of natural selection we know that when reaching this maximum, the proportion vector x is asymptotically stable, i.e., the membership of replicators in the dominant network does not change anymore. This information could also be used to determine the stationarity of the process. The iterative process can be terminated if the membership status of all replicators remains unchanged for a predefined number of iterations. In all our experiments, membership remained the same after as few as 20 iterations.

An Example

The following example is designed to illustrate the behavior of a replicator process for a number of replicators with different fitness values. Although the example is not derived from real data, it is constructed to demonstrate some of the mechanisms described above. Results from applications to real meta-analytic data will be presented in the next section.

Suppose that within 14 experiments investigating the same cognitive task, we have observed the activation of six different foci as shown in Table I. The co-occurrence of the activation foci can be pictured conveniently in a graph

TABLE I. Distribution of 6 replicators in 14 observations

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
A	x	x		x	x		x	x	x			x	x	x
B	x	x	x	x			x	x		x			x	
C			x			x	x		x			x		x
D			x		x	x				x	x			
E					x						x			
F										x				

Activation for 6 foci (A–F) as they could have been observed in 14 experiments.

Figure 1.

Co-occurrences of activation foci from Table I represented as a graph. Nodes represent activation foci and weighted connections represent the number of co-occurrences of two foci.

where nodes represent the activation foci and connections between nodes represent joint activations. The weight of a connection is the number of joint activations found in all experiments. The graph corresponding to the data in Table I is depicted in Figure 1.

As graph and table show, activation foci A and B have the most co-occurrences with other activation foci, 12 and 11, respectively. Focus C and focus D have the same total number of co-occurrences, but co-occur partly with different other activation foci. E and F appear only rarely. We can derive the co-occurrence matrix for our example as follows. The order of rows and columns in the matrix C corresponds to the order of the foci in Table I:

$$C = \begin{pmatrix} c_{AA} & c_{AB} & \cdots & \cdots & c_{AF} \\ c_{BA} & c_{BB} & \cdots & \cdots & c_{BF} \\ \vdots & \vdots & & & \\ c_{FA} & c_{FB} & \cdots & \cdots & c_{FF} \end{pmatrix} \quad (5)$$

The individual entries correspond to the connections between two nodes in the graph, e.g., $c_{AB} = c_{BA} = 6$. As explained in the previous section, all diagonal elements are set to zero. This results in the co-occurrence matrix:

$$C = \begin{pmatrix} 0 & 6 & 4 & 1 & 1 & 0 \\ 6 & 0 & 2 & 2 & 0 & 1 \\ 4 & 2 & 0 & 2 & 0 & 0 \\ 1 & 2 & 2 & 0 & 2 & 1 \\ 1 & 0 & 0 & 2 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 \end{pmatrix} \quad (6)$$

The co-occurrence matrix can also be calculated very efficiently by viewing the transpose of Table I as an occurrence matrix

$$Z = \begin{pmatrix} z_{1A} & z_{1B} & \cdots & \cdots & z_{1F} \\ z_{2A} & z_{2B} & \cdots & \cdots & z_{2F} \\ \vdots & \vdots & & & \\ z_{14A} & z_{14B} & \cdots & \cdots & z_{14F} \end{pmatrix} \quad (7)$$

with elements $z_{ij} = 1$ if focus j was activated in study i and $z_{ij} = 0$ otherwise. The co-occurrence matrix is then simply $\mathbf{C} = \mathbf{Z}^T \mathbf{Z}$ with the diagonal elements set to zero.

Because six activation foci are included in the analysis, the proportion vector $\mathbf{x}$ is initialized as

$$\mathbf{x}(0) = (1/6 \ 1/6 \ 1/6 \ 1/6 \ 1/6 \ 1/6)^T. \quad (8)$$

After the first iteration through equation (2) we get the new proportion vector

$$\mathbf{x}(1) = (0.27 \ 0.25 \ 0.18 \ 0.18 \ 0.07 \ 0.04)^T. \quad (9)$$

As we would expect from the distribution of co-occurrences, foci A and B were assigned the largest proportion; C and D with the same total number of co-occurrences also have the same proportion after the first iteration. The proportions of E and F have decreased from the initial values, because $(\mathbf{C}\mathbf{x}(0))_E = 0.50$ and $(\mathbf{C}\mathbf{x}(0))_F = 0.33$ do not exceed the mean fitness $\mathbf{x}(0)^T \mathbf{C}\mathbf{x}(0) = 1.22$. After applying the membership function, we see that the proportions of A, B, C, and D exceed the mean proportion $1/6$. These foci are therefore members of the dominant network.

In the second iteration, the proportion vector changes to

$$\mathbf{x}(2) = (0.35 \ 0.31 \ 0.18 \ 0.12 \ 0.02 \ 0.01)^T. \quad (10)$$

The proportions of A and B have increased again; the proportions of E and F have decreased further. The proportion of C now differs from that of D. Although the proportion of C stayed approximately the same in this iteration, the proportion of D has decreased, because C and D co-occur with different activation foci. C is jointly activated most often with focus A and focus B, which in turn have the two highest proportions among all activation foci after the previous iteration. In contrast, focus D occurs more often with E and F, which have relatively low proportions after the previous step. Consequently, $(\mathbf{C}\mathbf{x}(1))_C = 1.95$ is higher than $(\mathbf{C}\mathbf{x}(1))_D = 1.32$. In fact, $(\mathbf{C}\mathbf{x}(1))_D$ no longer exceeds the new mean fitness $\mathbf{x}(1)^T \mathbf{C}\mathbf{x}(1) = 1.93$. The proportion of D thus decreases, and after applying the membership function again, only foci A, B, and C remain in the dominant network.

The third and fourth iteration result in

$$\mathbf{x}(3) = (0.39 \ 0.35 \ 0.17 \ 0.07 \ 0.00 \ 0.00)^T \quad (11)$$

and

$$\mathbf{x}(4) = (0.42 \ 0.37 \ 0.15 \ 0.03 \ 0.00 \ 0.00)^T. \quad (12)$$

The proportion of C decreases for the first time in the third iteration, because $(\mathbf{C}\mathbf{x}(2))_C = 2.26$ is smaller than the new mean fitness $\mathbf{x}(2)^T \mathbf{C}\mathbf{x}(2) = 2.41$. After the third iteration, the proportion of C remains large enough to keep C as a member of the dominant network, whereas after the fourth iteration it drops below the threshold. The membership of only

A and B in the dominant network now remains unchanged for all following iterations. As the result of the replicator process, foci A and B are determined as dominant network in this example. Note that using the number of occurrences of each focus as diagonal elements of the fitness matrix results in a dominant network only containing focus A. A is the focus that is activated more often than all other foci. This result would not reflect the strong connection between A and B.

Choice of Activation Foci

The remaining task is to determine the replicators, i.e., activation foci included in the replicator process, from the lists of activation maxima in the considered experiments. Given anatomical differences between subjects and differences in scanning procedures, hardware, and spatial normalization and registration procedures, it is very unlikely that reported Talairach coordinates of activation maxima are identical in all experiments. We therefore have to decide whether two coordinates should be regarded as belonging to the same activation focus or not. Visual inspection by experts might often be possible but introduces subjectivity into the data and can be a tedious and labor-intensive task. We applied a more automated procedure described, for example, by Turkeltaub et al. [2002] and Chein et al. [2002]. First, activation likelihood estimates (ALE) for each voxel are calculated. Activation maxima are modeled by 3D Gaussian probability distributions centered at their Talairach coordinates. The probability that a given activation maximum lies within a particular voxel is then

$$p = \frac{1}{(2\pi)^{3/2} \sigma^3} \exp \left[\frac{-d^2}{2\sigma^2} \right], \quad (13)$$

where σ is the standard deviation of the distribution and d is the Euclidean distance of the voxel to the activation maximum. For each voxel, the union of these probabilities calculated for all activation maxima gives the ALE. In regions with a relatively high density of reported activation maxima, voxels will be assigned a high ALE, in contrast to regions where few and widely spaced activation maxima have been reported. We can thus search in a second step for clusters of voxels exceeding a certain activation likelihood threshold. Each such region forms an activation focus for the subsequent analysis. The occurrence of an activation focus in a particular study is recorded if at least one of the activation maxima reported for this study lies within this focus. The resulting occurrence matrix can be transformed finally into the co-occurrence matrix $\mathbf{C}$, which serves as input to the replicator process.

To summarize, the process of finding a dominant network of activation foci given a list of Talairach coordinates from different studies consists of a series of processing steps:

1. Form activation foci from the given coordinates using, for example, activation likelihood estimates.

2. Record for each activation focus the occurrence in the investigated studies.
3. Calculate the co-occurrence matrix of activation foci.
4. Apply the replicator process to determine the dominant network.

RESULTS

The method was applied to fMRI experiments addressing the well-known Stroop interference task [Stroop, 1935]. This task requires a subject to respond to a particular stimulus dimension while a competing stimulus dimension has to be suppressed. In the color–word Stroop task, for example, subjects are presented with a color name, such as “blue.” In the congruent condition, this word is printed in blue letters, whereas a different color is used in the incongruent condition. In a neutral condition, only patterns or noncolor words are presented. Other variants of the Stroop interference task include the spatial–word Stroop task (the word “above” is written below a horizontal line), the counting Stroop task (the word “two” appears three times on the screen), and the object–color Stroop task (an object is presented in an atypical color, e.g., a blue lemon).

For our meta-analysis, different forms of the task were considered. Only selection criteria for the included studies were the imaging modality fMRI and that some form of Stroop interference was caused by the presented stimulus such that an incongruent condition could be contrasted with a congruent or neutral one. Data were obtained from the publicly accessible database BrainMap DBJ [Fox and Lancaster, 2002]. This online database provides Talairach coordinates of activation maxima from published functional neuroimaging experiments covering a variety of experimental paradigms and imaging modalities. Searching the database for fMRI experiments investigating variants of the Stroop interference task yielded experimental results reported in 16 peer-reviewed journal publications. One publication was a meta-analysis and therefore excluded from further processing. Within the remaining 15 studies, 17 data sets provided Talairach coordinates for the contrasts incongruent versus congruent and incongruent versus neutral obtained for groups of subjects. These data sets, containing 239 activation maxima in total, were included in our analysis. Plots of all coordinates projected onto a single axial, sagittal, and coronal slice are shown in the top row of Figure 2. Coordinates reported from different studies are represented by different colors. As can be seen, activation maxima are distributed over large parts of the cortex, and visual inspection does not reveal any obvious clustering of activation foci.

Activation likelihood estimates were calculated for all voxels using three-dimensional Gaussian distributions with variance $\sigma = 3$ mm. The subsequent search for clusters exceeding an activation likelihood threshold of 0.05 yielded the 15 activation foci depicted in the bottom row of Figure 2. The foci were again projected onto a single axial, sagittal and coronal slice. A co-occurrence matrix was calculated from all coordinates laying within these 15 activation foci and the

Figure 2.

Top: 239 activation maxima obtained for fMRI experiments addressing the Stroop interference task. Bottom: 15 activation foci derived from ALE of the above coordinates. Activation maxima and foci are shown in the glass brain view, i.e., projected onto three orthogonal single coronal, sagittal, and axial slices.

replicator process was applied until changes in the mean fitness became negligible.

During the replicator process, five activation foci were selected as members of the dominant network: the pre-supplementary motor area (pre-SMA), the inferior frontal sulcus (IFS) extending onto the middle frontal gyrus, the anterior cingulate cortex (ACC) of both hemispheres extending into BA8, and the inferior frontal junction area (IFJ). These five activation foci are depicted in Figure 3. The fitness of the remaining 10 activation foci dropped to values close to zero. Volume, mean ALE values, the number of joint activations, and the fitness of the five selected regions are listed in Table II.

The highest fitness was observed for the ACC of the right hemisphere, which was the second largest activation focus and also showed the second largest number of joint activations. The largest number of joint activations was found for IFS. Although this is the smallest activation focus of the dominant network, it was assigned the second highest fitness by the replicator process. These results confirm our earlier observation that the fitness of a replicator does not merely reflect the total number of joint activations with other activation foci, nor does it only reflect its anatomical size. Rather, while the size and ALE of a focus are derived from the number and exact location of reported activation maxima within this cortical region, the result of the replicator process is determined by the relationship between different activation foci.

DISCUSSION

The application of replicator dynamics in meta-analytic studies addresses an important aspect of functional imaging

research, the search for functional networks of cortical areas underlying a specific cognitive task. It is based on the assumption that cortical areas involved in solving the same cognitive task are likely to be found jointly activated in experiments addressing this task. Consequently, the method searches for cortical areas that are jointly activated in most experiments considered in the meta-analysis. It is important to be clear that this search goes beyond the usually performed simple counting of activation maxima in specific cortical regions. Counting the number of studies that activated a cortical region can provide useful information about the role this single region might play in performing a particular task; however, it leaves any information about the relationship between cortical areas unexplored. Our new method draws on both sources of information. The number of studies that activated a particular region and the distribution of activation over anatomical regions first influence the choice of activation foci that enter the replicator process. The co-occurrences of activation foci then determine the members of the dominant network.

The results of any meta-analysis depend on the completeness of the included data. This can be problematic, particularly if experimental data extracted from the literature are used. Often very specific aspects of an experimental paradigm or a particular cortical region are investigated in a functional imaging experiment and thus not all significant activations found for a stimulus might be reported in its publication. Data included in subsequent meta-analyses are, therefore, likely to be biased toward particular cortical areas, a problem referred to previously as “literature or publication bias.” When applying a replicator process to such data, we cannot guarantee that all cortical areas relevant to the task in question are included in the resulting dominant network.

Figure 3.

The five selected members of the dominant network. They include the IFJ, IFS extending into middle frontal gyrus, pre-SMA, and ACC of both hemispheres extending into BA8. **a, b:** Left lateral cortex at $x = -44$ and $x = -3$, respectively. **c:** Right lateral cortex at $x = 7$.

TABLE II. Five members of the dominant network resulting from the replicator process

Region	Volume (mm ³)	Mean ALE	Joint activations (n)	Fitness
IFJ	3,012	0.09	38	0.17
IFS	1,912	0.08	46	0.27
Pre-SMA	2,023	0.08	30	0.14
ACC left	2,134	0.07	34	0.12
ACC right	2,500	0.08	40	0.29

Volume, mean activation likelihood estimate (ALE), number of joint activations, and fitness of the five activation foci selected by the replicator process.

Because this article is aimed primarily at introducing a new analysis tool and not at answering a specific neurocognitive question, we chose a well-studied paradigm for the presentation of our methodology and, to include a reasonable amount of data, placed only minimal restrictions on the selected studies. Although the results of the replicator process agree well with results reported in many other Stroop studies, they might be biased at least to some degree. Some of the included studies, for example, were aimed at investigating the possible functionality of the anterior cingulate cortex. It is thus not surprising to find a high proportion of reported activations in this particular area, whereas other activation locations might have been omitted in the publication of these studies. This might also explain the absence of any parietal areas in the result of the replicator process. Some activation foci were derived from parietal activation maxima reported in the considered experiments (see Fig. 2, bottom). However, their co-occurrences with other foci were outnumbered by those of frontal regions that were of particular interest in some of the experiments. Like in all meta-analytic approaches, the careful selection of studies included in the analysis is therefore the most important prerequisite for the successful application of the replicator process and the meaningful interpretation of its results.

One way of alleviating the problems introduced by a publication bias is to include many studies in a meta-analysis. Assuming that not all studies are biased toward the same aspect, the influence of some biased data on the overall result of a meta-analysis should decrease as more studies are included. This in turn requires efficient analysis tools. The method presented here is very fast and easy to implement. After extracting activation maxima from the database, the results of our example analysis were calculated within a few seconds using a MATLAB implementation and an AMD Athlon XP1800 processor. Our method can, therefore, be applied easily to large-scale studies.

The analysis presented here is related to a method introduced recently by Lohmann and Bohn [2002], who used replicator dynamics to detect functionally coherent networks in single-subject data. In contrast to meta-analyses where only lists of activation coordinates are available, this method draws on the information encoded in the fMRI time

series of individual voxels. The replicator process takes as input a fitness matrix that is derived from pairwise correlations of these time series and determines those voxels as dominant that show a strong correlation to all other network members. Group analyses can be carried out by averaging fitness matrices of individual subjects.

An interesting point to note is the close relation of replicator dynamics to multivariate data analysis techniques such as multidimensional scaling and principal component analysis (PCA). For example, given some data $\mathbf{z}$, the first principal component of $\mathbf{z}$ is the linear combination $y = \mathbf{x}^T \mathbf{z}$ with maximum variance. This component is found by choosing the vector $\mathbf{x}$ that maximizes the objective function $\mathbf{x}^T \mathbf{C}_z \mathbf{x}$ under the constraint $\sum x_i = 1$. Here, $\mathbf{C}_z$ is the covariance matrix of the data, and it turns out that the first eigenvector of $\mathbf{C}_z$ is the solution for $\mathbf{x}$. In a similar manner, the replicator process finds a vector $\mathbf{x}$ that maximizes the objective function $\mathbf{x}^T \mathbf{C} \mathbf{x}$. $\mathbf{C}$ is now the co-occurrence matrix for our observations or, in case of single-subject analyses, the correlation matrix for the observed fMRI time series. The only difference in the two methods lies in the stronger constraint $\bar{x}_i \geq 0$, $i = 1 \dots n$ and $\sum \bar{x}_i = 1$ on the solution of the replicator process.

Techniques such as multidimensional scaling and PCA have been applied already to functional imaging data [Friston et al., 1996; Sugiura et al., 2004; Welchew et al., 2002]. Moreover, Lohmann and Bohn [2002] proposed the combination of replicator dynamics and PCA for the extraction of a series of functionally coherent networks from single-subject data, making use of the strong similarity between the two techniques.

For reasons of simplicity, the activation foci included in the replicator process in our example study were determined by thresholding the ALE at 0.05. The threshold was chosen to provide a reasonable number of clusters that agree in size with activated areas typically reported in the fMRI literature. Lower thresholds resulted in implausibly large clusters; much higher thresholds excluded large parts of the data from the subsequent analysis. However, activation likelihood estimates strongly depend on the overall number of activation maxima and the parameter σ of the Gaussian probability distributions, and suitable thresholds might vary considerably from study to study. One could conceive of applying somewhat more objective measures, as was done for example by Turkeltaub et al. [2002]. Here, the distribution of observed activation maxima was compared to randomly generated sets of activation maxima. It is important to note that the results of such comparisons also depend on the choice of a threshold for the statistical test. The exact influence of these thresholds on the result of the subsequent replicator process is a topic of our current research.

CONCLUSIONS

We have presented an efficient new method for the meta-analysis of functional imaging data. The method requires as input only lists of comparable activation locations such as coordinates in the Talairach stereotactic space. It can be

applied to experiments employing different hardware configurations and imaging modalities and thus enables the comparison and integration of results from imaging studies carried out in different laboratories. With the increasing availability of such experimental data in searchable indexed databases, we would expect meta-analyses to become a viable source of new information derived from functional imaging that further our understanding of how the human brain works.

ACKNOWLEDGMENTS

We thank the BrainMap DBJ development team for providing access to the database and for very helpful technical support.

REFERENCES

- Bomze IM, Pelillo M, Stix V (2000): Approximating the maximum weight clique using replicator dynamics. *IEEE Trans Neural Netw* 11:1228–1241.
- Bouzat S, Zanette DH (2003): Survival and extinction in the replicator model: dynamics and statistics. *Adv Complex Syst* 6:223–240.
- Chein JM, Fissell K, Jacobs S, Fiez JA (2002): Functional heterogeneity within Broca's area during verbal working memory. *Physiol Behav* 77:635–639.
- Crivello F, Schormann T, Tzourio-Mazoyer N, Roland PE, Zilles K, Mazoyer BM (2002): Comparison of spatial normalization procedures and their impact on functional maps. *Hum Brain Mapp* 16:228–250.
- Crum WR, Griffina LD, Hilla DL, Hawkesa DJ (2003): Zen and the art of medical image registration: correspondence, homology, and quality. *Neuroimage* 20:1425–1437.
- Farah MJ, Aguirre GK (1999): Imaging visual recognition: PET and fMRI studies of the functional anatomy of human visual recognition. *Trends Cogn Sci* 3:179–186.
- Fisher RA (1930): *The genetical theory of natural selection*. London, UK: Oxford University Press.
- Fox PT, Lancaster JL (2002): Mapping context and content: the BrainMap model. *Nat Rev Neurosci* 3:319–321.
- Fox PT, Lancaster JL, Parson LM, Xiong JH, Zarnarripa F (1999): Functional volumes modeling: theory and preliminary assessment. *Hum Brain Mapp* 5:306–311.
- Fox PT, Parsons LM, Lancaster JL (1998): Beyond the single study: functional/location metanalysis in cognitive neuroimaging. *Curr Opin Neurobiol* 8:178–187.
- Friston KJ, Ashburner J, Frith CD, Poline JB, Heather JD, Frackowiak RS (1995): Spatial registration and normalization of images. *Hum Brain Mapp* 2:165–189.
- Friston KJ, Frith CD, Fletcher P, Liddle PF, Frackowiak RS (1996): Functional topography: multidimensional scaling and functional connectivity in the brain. *Cereb Cortex* 6:165–164.
- Hofbauer J, Sigmund K (1988a): *Evolutionary games and population dynamics*. Cambridge: Cambridge University Press.
- Hofbauer J, Sigmund K (1988b): *The theory of evolution and dynamical systems*. Cambridge: Cambridge University Press.
- Indefrey P, Levelt W (2000): The neural correlates of language production. In: Gazzaniga M, editor. *The new cognitive neurosciences* (2nd ed.). Cambridge, MA: MIT Press. p 845–865.
- Kimura M (1958): On the change of population fitness by natural selection. *Heredity* 12:145–167.

- Lohmann G, Bohn S (2002): Using replicator dynamics for analyzing fMRI data of the human brain. *IEEE Trans Med Imaging* 21:485–492.
- Maitra R, Roys SR, Gullapalli RP (2002): Test–retest reliability estimation of functional MRI data. *Magn Reson Med* 48:62–70.
- McGonigle DJ, Howseman AM, Athwal BS, Friston KJ, Frackowiak RS, Holmes AP (2000): Variability in fMRI: an examination of intersession differences. *Neuroimage* 11:708–734.
- Menon A, Mehrotra K, Mohan CK, Rank S (1995): Optimization using replicators. In: Eshelman L, editor. *Proceedings of the Sixth International Conference on Genetic Algorithms*. San Francisco, CA: Morgan Kaufmann Publishers. p 209–216.
- Neumann J, Lohmann G, Zysset S, von Cramon DY (2003): Within-subject variability of BOLD response dynamics. *Neuroimage* 19:784–796.
- Noll DC, Genovese CR, Nystrom L, Forman S, Eddy WF, Cohen J (1997): Estimating test–retest reliability in fMRI II: application to sensory-motor and cognitive activation. *Magn Reson Med* 38: 508–517.
- Pelillo M (2001): Replicator dynamics in combinatorial optimization. In: Floudas CA, Pardalos PM, editors. *Encyclopedia of optimization*, Vol. 5. Boston: Kluwer Academic Publishers. p 23–35.
- Pelillo M, Siddiqi K, Zucker SW (1999): Matching hierarchical structures using association graphs. *IEEE Trans Pattern Anal Mach Intell* 21:1105–1120.
- Phan K, Wagner T, Taylor SF, Liberzon I (2002): Functional neuro-anatomy of emotion: a meta-analysis of emotion activation studies in PET and fMRI. *Neuroimage* 16:331–348.
- Schuster P, Sigmund K (1983): Replicator dynamics. *J Theor Biol* 100:533–538.
- Stroop JR (1935): Studies of inference in serial verbal reactions. *J Exp Psychol* 18:643–662.
- Sugiura M, Watanabe J, Maeda Y, Matsue Y, Fukuda H, Kawashima R (2004): Different roles of the frontal and parietal regions in memory-guided saccade: a PCA approach on time course of BOLD signal changes. *Hum Brain Mapp* 23:129–180.
- Talairach J, Tournoux P (1988): *Co-planar stereotaxic atlas of the human brain*. Stuttgart: Thieme. 122 p.
- Taylor P, Jonker L (1978): Evolutionary stable strategies and game dynamics. *Math Biosci* 40:145–156.
- Turkeltaub PE, Eden GF, Jones KM, Zeffiro TA (2002): Meta-analysis of the functional neuroanatomy of single-word reading: method and validation. *Neuroimage* 16:765–780.
- Ullsperger M, von Cramon DY (2004): Neuroimaging of performance monitoring: error detection and beyond. *Cortex* 40:593–604.
- Wager TD, Phan KL, Liberzon I, Taylor SF (2003): Valence, gender, and lateralization of functional brain anatomy in emotion: a meta-analysis of findings from neuroimaging. *Neuroimage* 19:513–531.
- Welchew DE, Honey GD, Sharma T, Robbins TW, Bullmore ET (2002): Multidimensional scaling of integrated neurocognitive function and schizophrenia as a disconnection disorder. *Neuroimage* 17:1227–1239.