9

Machine learning for functional brain mapping

Malin Björnsdotter
Institute of Neuroscience and Physiology
University of Gothenburg
Sweden

1. Introduction

Human brain activity – the product of multiple neural networks encoding a multitude of concurrent cognitive states and responses to various stimuli – is a prime example of highly complex data. Irrespective of measuring technique, acquired brain signals are invariably exceedingly noisy, multivariate and high-dimensional. These inherent properties produce analysis difficulties traditionally overcome by utilization of descriptive statistical methods averaging across numerous time-fixed events, and conclusions about brain function are thus typically based on mean signal variations in single measuring points. It is, however, commonly believed that neural network activity responsible for cognitive function is distributed across time, frequency and space – aspects that traditional descriptive statistics inevitably fail to capture. Machine learning approaches do, in contrast, supply tools for detection of single, subtle signal patterns, rather than characteristic univariate, average activity. The field of pattern recognition harbors great potential in answering research questions fundamentally relevant for understanding the functional organization of the human brain where conventional methods fail.

Machine learning techniques are, consequently, being increasingly used by the neuroimaging community for a variety of data analyses. In functional magnetic resonance imaging (fMRI; described in section 2), where neural activity is estimated through measuring changes in blood flow and oxygenation (collectively called 'hemodynamics') at an excellent spatial resolution, for example, brain activation magnitudes are minute – at a magnetic field of 1.5 Tesla, for example, the hemodynamic change in response to a stimulation is typically only about 2% of the total amplitude Buxton (2002). This effect, nonetheless, provides researchers a non-invasive window into the human brain, and can be utilized to localize regions activated by various conditions in an attempt to understand the functional organization of the cortex. Conventionally, regional activity is estimated by fitting a spatially invariant model of the expected hemodynamic change at each individual measuring point (voxel) and, subsequently, testing for differences in signal levels between two experimental conditions Friston et al. (1994a). This massively univariate (voxel-by-voxel) analysis produces statistical maps highlighting brain locations that are "activated" by the condition in an appealing fashion (see figure 1E). Although this analytical scheme has proven tremendously useful in functional brain mapping, it is limited to revealing average stimuli - single location relationships and cannot reflect individual events or effects in distributed activation patterns.

In recent years, state-of-the-art machine learning methods allowing identification of subtle, spatially distributed single-trial fMRI signal effects not detectable using conventional analysis

have instead won substantial ground Haynes & Rees (2006); Norman et al. (2006). Typically, a classifier model is trained to identify signal patterns related to a given experimental condition by integrating information across multiple voxels, thus allowing detection of information which only produce weak single-voxel effects. Subsequently, the trained classifier is applied to new data samples and attempts to predict the experimental condition to which each sample belongs. This approach, termed multivoxel pattern analysis (MVPA; Norman et al. (2006); or more seductively 'brain reading'), was initiated with a breakthrough fMRI study of the vision pathway Haxby et al. (2001). Volunteers were shown a number of visual stimuli from various object categories (cats, faces, houses, etc.), and the researchers found that fMRI responses evoked by each object category was associated with a distinct spatial signal pattern which could be identified ('decoded') by a classification scheme. The authors concluded that information is encoded across spatially wide and distributed patterns of fMRI responses, undetectable by conventional approach designed to detect voxel-by-voxel statistically significant activations. Since then, numerous studies have utilized multivoxel pattern analysis for exploring human brain function with outstanding results (see e.g. Björnsdotter et al. (2009); Cox & Savoy (2003); Formisano et al. (2008); Haynes & Rees (2005a;b); Howard et al. (2009); Kamitani & Tong (2005)).

The variety of machine learning approaches utilized in fMRI spans a vast range: from classical linear discriminant analysis Carlson et al. (2003) to state-of-the-art particle swarm optimization Niiniskorpi et al. (2009). This chapter introduces machine learning concepts in the context of functional brain imaging, with focus on classification based fMRI brain mapping. First, the fMRI technique for measuring brain activity is presented, including technical background and a number of considerations for subsequent analysis. Second, the essential steps required for effective machine learning classification analysis are described, and, ultimately, an example of our work utilizing evolutionary algorithms for mapping human brain responses to pleasant touch is presented.

2. Functional magnetic resonance imaging

Functional magnetic resonance imaging (fMRI) derives from magnetic resonance imaging (MRI), which is utilized for obtaining structural (as opposed to functional) images of tissue (see figure 1A and B). Magnetic resonance imaging is made possible by intrinsic physical properties of atoms in matter, including magnetism and nuclear spin. Measurable effects in such inherent properties due to local blood oxygenation changes in response to neural activity provide an effective, albeit indirect, entirely non-invasive measure of brain activation. Below follows a description of the technical MRI and fMRI background as well as the conventional approach to fMRI signal processing and analysis (see also e.g. Buxton (2002) or Norris (2006) for more details).

2.1 Magnetic resonance imaging

In magnetic resonance imaging (MRI), the magnetism and nuclear spin of atoms are utilized to obtain information about the environment in which the they are contained (e.g. brain tissue; see Buxton (2002) for more technical details on MRI). The spin refers to the inherent angular momentum possessed by all atomic nuclei with an odd number of protons or neutrons, and, fundamentally important for brain imaging, only one substance abundantly found in organic tissue possesses such a spin – the hydrogen atom. Due to the spin, the hydrogen atom also has a magnetic dipole moment and, therefore, behaves like a small magnet.

When there is no external magnetic field, the magnetic dipole moments are randomly orientated and there is a net magnetization of approximately zero. When an external magnetic field is applied, a majority of the magnetic moments will gradually align with the magnetic field (a process referred to as longitudinal relaxation) with a time constant T_1 (typically around 1 s), and a net magnetization in the direction of the field is produced. Also, the nuclei will change the direction of the axis of rotation around the field axis in a process called *precession*. The frequency of precession, v_0, is called the *Larmor* (or nuclear magnetic resonance) frequency. The Larmor frequency is proportional to the strength of the external magnetic field: $v_0 = \gamma B_0$, where B_0 (measure in Tesla, T) is the external magnetic field strength and γ (units: MHz/T) is the gyromagnetic ratio. The latter is a physical property of the given element, and for hydrogen the ratio is 42.58 MHz/T. Magnetic fields used for human fMRI are typically in the range of 1.5 T (approximately 30 000 times the earth's magnetic field) to 9 T.

The precession of the nuclei could potentially induce a current in a receiver coil placed nearby, and thus be measured. However, although the nuclei precess with the same frequency, their phase is different and the net magnetism is zero – thus no current would be induced in the coil. The phase of the precessing nuclei must therefore be coordinated. This is achieved by applying a rapidly oscillating magnetic radio frequency (RF) pulse at the precession frequency (v_0) to the nuclei. The rapid oscillations of the RF pulse gradually nudge the magnetic moments away from their initial axis of precession. The RF pulse can be applied to rotate all the magnetic moments 90 degrees, and thus change the net magnetization from being aligned with the external magnetic field to pointing perpendicular to the field – still, however, precessing around the field. As a result, the magnetic moments are in phase, producing a precessing net magnetization which can be registered by a receiver coil. The induced current, alternating with the Larmor frequency, is informative about the physical environment, such as the number of nuclei (spins) in the sample.

Notably, the current induced in the coil decays over time (a process called *relaxation*). This decay is partly due to thermal motion on the molecular level, realigning the net magnetic moments with the external magnetic field (T_1 relaxation), partly to that the random interactions of nuclei result in a loss of coherence of the precession which reduces the net magnetization (T_2 relaxation). Moreover, inhomogeneities in the magnetic field cause dephasing, since the precession frequency of the nuclei is proportional to the strength of B_0. This effect in combination with the random nuclei interactions is referred to as T_2^* relaxation.

Fundamentally important for functional and structural imaging, the relaxation times differ between various tissues (such as muscle and bone, gray and white brain matter, and so on). Using an MRI scanner, structural images (akin to X-rays) can therefore be reconstructed from the acquired relaxation signals (figure 1A).

2.2 BOLD functional imaging

In addition to structural images, magnetic resonance techniques can also be utilized for acquiring *functional* data, that is, signals related to some sort of active function of the brain. Functional magnetic resonance imaging (fMRI) is based on the presence of hemoglobin (the molecule in red blood cells which contains oxygen) combined with various hemodynamic changes (such as blood flow, blood volume, oxygen consumption etc). Hemoglobin is diamagnetic when oxygenated and paramagnetic when deoxygenated and, therefore, possesses different magnetic characteristics depending on oxygenation state. This phenomenon, in combination with the measured T_2^* (transversal) relaxation (see the section on Magnetic resonance imaging above), is used in fMRI to detect magnetic differences between oxygenated and de-

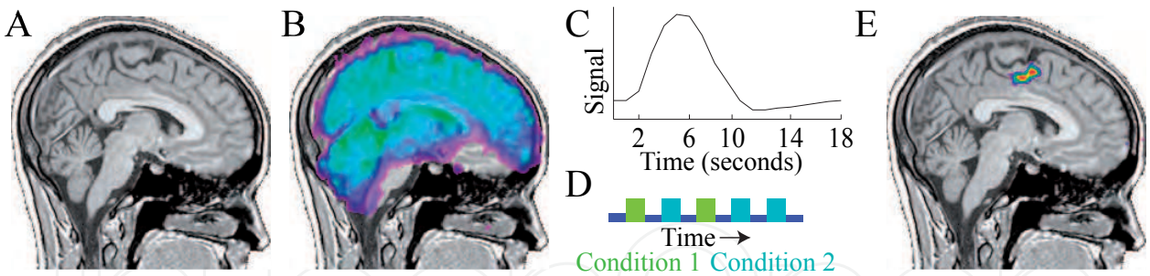

Fig. 1. **A**. MRI structural image. **B**. fMRI functional image. **C**. Hemodynamic response function (in arbitrary units). **D**. Schematic of experimental paradigm with two conditions. **E**. General linear model (GLM) statistical activation map.

oxygenated blood in the brain. Specifically, blood-oxygen-level dependent (BOLD) fMRI is the technique used to identify temporal and spatial variations in the proportion of oxygenated to deoxygenated blood, which, in turn, is an indication of blood flow changes Ogawa et al. (1990). A relative increase in blood flow results in a positive signal, and vice versa.
For brain function related studies, the BOLD signal is acquired volume by volume, and each volume element is referred to as a 'voxel' (see figure 1B). The minimum time required for one whole brain volume is typically in the range of 2-3 seconds, and each voxel is in the range of 2-4 mm per side. Thus, fMRI has comparatively poor temporal resolution and excellent spatial resolution.

2.3 Neural correlates of BOLD

It was observed as early as in the 1890s that nerve cell activation level is positively correlated with blood flow Roy & Sherrington (1890). The temporal pattern of blood flow changes in response to activated nerve cells is called the hemodynamic response function (HRF; see figure 1C), and different brain areas respond differently Leoni et al. (2008). The full course of the blood flow response to a briefly presented stimulus is about 20 s and a maximum is obtained at approximately 6 s. As a result, the temporal resolution of fMRI is limited due to the inherent delay in the hemodynamic response. Capillaries, small arteries and veins, as well as large arteries and veins all contribute to the registered BOLD signal.
Although it is generally assumed that changes in blood flow (and supply of oxygenated blood) is prompted by increased oxygen consumption by activated nerve cells, the exact relationship between neural activity and the BOLD signal is not fully understood. In fact, the measurable (highest in amplitude) portion of the HRF appears to be a substantial over-compensation (supplying more blood than is required by metabolic demands) and the mechanisms for this are unknown. Moreover, the BOLD signal is an indirect measure of brain activity, and is susceptible to influence by physical parameters of non-neural nature, and, can, in fact, potentially represent increased blood flow into an area despite no local neural activity (see e.g. Sirotin & Das (2008)). Concurrent intracortical recordings of neural signals and fMRI responses in the visual cortex in monkeys have shown, however, that local field potentials are significantly correlated to the hemodynamic response (Logothetis et al. (2001); see also Goense & Logothetis (2008) for a review).

3. Conventional brain mapping analysis

The aim of conventional analysis is typically to identify regions of the brain which are activated by the processing of a given stimulation or condition. In order to achieve effective analysis, a number of steps are required. First, the experimental paradigm must be carefully designed in order ensure that the actual effect of interest is analyzed. Second, the acquired data must be pre-processed, followed by a statistical analysis which estimates significantly activated regions. A variety of softwares exist for both pre-processing and statistical analysis, including the freely available Neurolens (neurolens.org), SPM (fil.ion.ucl.ac.uk/spm; Frackowiak et al. (1997)) and ANFI (afni.nimh.nih.gov), as well as commercial software such as BrainVoyager (brainvoyager.com). The required analysis steps are described in detail below.

3.1 Experimental paradigm

Careful attention needs to be paid to the type and organization of conditions presented during the experiment (the experimental paradigm) in order to isolate the effect of interest (as opposed to noise or other cognitive process). Typically, paradigms involve a number of stimulus conditions which are contrasted in subsequent analysis to remove confounding variables (such as attentional effects). During the scanning session, the conditions are presented in a predetermined fashion in one of a number of ways. Influenced by positron emission tomography (PET) imaging where extended stimulation periods are required in order to produce stable activations Muehllehner & Karp (2006), fMRI studies often utilize experimental paradigms which alternate extended periods of stimuli being 'on' or 'off' (see figure 1D; Turner et al. (1998)). This so called *block design* is appealing due to ease of presentation and analysis, as well as to the relatively high signal-to-noise ratios achieved. Brief stimuli can, however, produce a measurable BOLD response (e.g. 34 ms; Rosen et al. (1998)), and such are applied in *event-related designs*. Various types of conditions unsuitable for study with block designs, such as the oddball paradigm McCarthy et al. (1997), are instead made possible by event-related studies. Also, more dynamic responses, suitable in situations where habituation is a concern, are produced, and, given similar scanning times, more stimulus repetitions can be applied (see e.g. Kriegeskorte et al. (2008) for a study utilizing numerous stimuli applications). A draw-back of even-related paradigms is, however, the lower functional signal-to-noise ratio than in block design paradigms Bandettini & Cox (2000).

3.2 Pre-processing

In order to reduce noise, the acquired fMRI data is subjected to a series of pre-processing operations. The following steps are typically applied, although all are not necessarily required and further steps can be included to improve the analysis (see e.g. Friston et al. (2007) or Henson (2003) for more details, and note that virtually all fMRI analysis software include functions for these corrections).

- *Slice-time correction*: The acquisition of an entire brain volume generally takes in the order of 2-3 seconds (depending on MRI scanner parameters), during which slices of brain tissue are scanned consecutively. The resulting shift in acquisition time between slices is typically corrected by resampling the time courses with linear interpolation such that all voxels in a given volume represent the signal at the same point in time.
- *Motion correction*: The excellent spatial resolution of fMRI means that slight movements of the head can affect the signal analysis substantially, and head movement effects must

therefore be corrected. A variety of more or less sophisticated algorithms are available in any of the software packages mentioned above.

- *Signal filtering*: Temporal drifts which can significantly affect the results are typically removed using temporal high-pass filtering. Noise can further be reduced by temporal low pass filtering.
- *Spatial smoothing*: In order to reflect some spatial integration, spatial smoothing is typically applied on the volume time series using a Gaussian kernel with the parameter FWHM (full width at half maximum) in the range of 3-12 millimeters. Spatial smoothing increases subsequent mapping sensitivity when two conditions differ in terms of their regional mean activation levels. In these cases, local signal differences coherently point in the same direction and are enhanced by spatial averaging. However, if two conditions differ in terms of their fine-grained spatial activation patterns, spatial smoothing has a destructive effect and cancels out the discriminative information, which can be detected by pattern recognition methods (see the section on Machine learning brain mapping analysis below).
- *Spatial normalization*: Individual brains are highly anatomically and functionally variable. Thus, for group analysis and comparison with brain atlases, the fMRI data must be projected into a standard brain space such as Talairach Talairach & Tournoux (1988) or so called MNI (Montreal Neurological Institute; Evans et al. (1993)) space. This is performed with a number of different algorithms (see e.g. Collins et al. (1994), and see Crinion et al. (2007) for issues with spatial normalization).

3.3 General linear modelling and activation detection

Numerous variations of fMRI analysis techniques are widely used, and the field is under active research. The most lucrative approach thus far, however, includes statistical analysis to produce images (statistical parametric maps) which identify brain regions that show significant signal changes in response to the conditions present during scanning (see e.g. Henson (2003)). Typically, a spatially invariant model of the expected blood oxygenation level dependent (BOLD) response is fitted independently at each voxel's time course and the differences between estimated activation levels during two or more experimental conditions are tested Friston et al. (1994a). Parametric tests, assuming that the observations are drawn from normal populations, are typically applied. Most such parametric modeling techniques are versions of the general linear model (GLM). The GLM aims to explain the variation of the time course $y_1 ... y_i ... y_n$, in terms of a linear combination of explanatory variables and an error term. For a simple model with only one explanatory variable $x_1 ... x_i ... x_n$, the GLM can be written as:

$$y_i = x_i \beta + \epsilon_i \tag{1}$$

where β is the scaling (slope) parameter, and ϵ_i is an error term. The model is also often written in matrix form when containing more variables:

$$Y = X\beta + \epsilon \tag{2}$$

where Y is the vector of observed voxel values, β is the vector of parameters and ϵ is the vector of error terms. The matrix X is termed the design matrix, containing one row per time point and one column per explanatory variable in the model (e.g. representing the presence or absence of a specific condition). In order to detect activations, the magnitude of the parameter in

β corresponding to these vectors are computed. β can be determined by solving the following equations:

$$X^T Y = (X^T X)\hat{\beta} \tag{3}$$

where $\hat{\beta}$ corresponds to the best linear estimate of β. Given that $X^T X$ is invertible, $\hat{\beta}$ can be estimated as:

$$\hat{\beta} = (X^T X)^{-1} X^T Y \tag{4}$$

These parameter estimates are normally distributed, and since the error term can be determined, statistical inference can be made as to whether the β parameter corresponding to the model of an activation response is significantly different from the null hypothesis. A number of additional parameters (regressors) can be included in the GLM analysis, such as drift, respiration, motion correction parameters or other information of interest.
Importantly, the massively univariate testing results in one statistic per voxel, and thus produces a classical problem of multiple comparisons which requires correction Friston et al. (1994a). Depending on the number of voxels included for analysis, the threshold for which a voxel can be considered significant varies – for smaller regions the threshold is lower. While whole-brain analyses might not yield any significant results, directed searches in carefully, a priori identified regions of interests can potentially yield significantly activated voxels. Thus, in combination with methods for mitigating the multiple comparisons problem, this massively univariate analysis produces statistical maps of response differences, highlighting brain locations that are activated by a certain condition dimension (see figure 1E; Friston et al. (1994a)). As such, univariate activation detection maps average single-voxel responses to experimental conditions, and does not reflect direct relationships between distributed patterns of brain processing and single condition instances.

4. Machine learning brain mapping analysis

Contrary to conventional univariate analysis (described above), where average voxel-by-voxel signal increases or decreases are estimated using statistical techniques, machine learning approaches utilize information distributed across multiple voxels. Specifically, classifier-based machine learning algorithms attempt to identify and distinguish the specific spatial activity patterns produced by single experimental conditions.
To this end, multivoxel fMRI activity patterns (samples) can be represented as points in a multidimensional space where the number of dimensions equals that of voxels. In the simplified situation of a two-voxel volume, each pattern can be considered as a point in a plane corresponding to the magnitude measured in each voxel (see figure 2). The aim of a classifier is to distinguish the conditions, that is, to separate the points belonging the each of the condition classes. As shown in the figure, the method of doing so depends on the structure of the data – if the conditions are sufficiently different (figure 2A) this can be done on a single voxel level (with conventional univariate statistics), whereas if the voxel distributions overlap (figure 2B and C) multiple voxels must be taken into account to distinguish the conditions.
After initial pre-processing and estimation of single condition responses, application of classifier-based machine learning techniques to fMRI generally entails a number of steps (see figure 3). The data is partitioned into datasets – one for training the classifier (that is, estimating classifier parameters), and one exclusively used in conjunction with the trained classifier to evaluate the classification performance. Voxel selection, aiming at reducing the complexity of

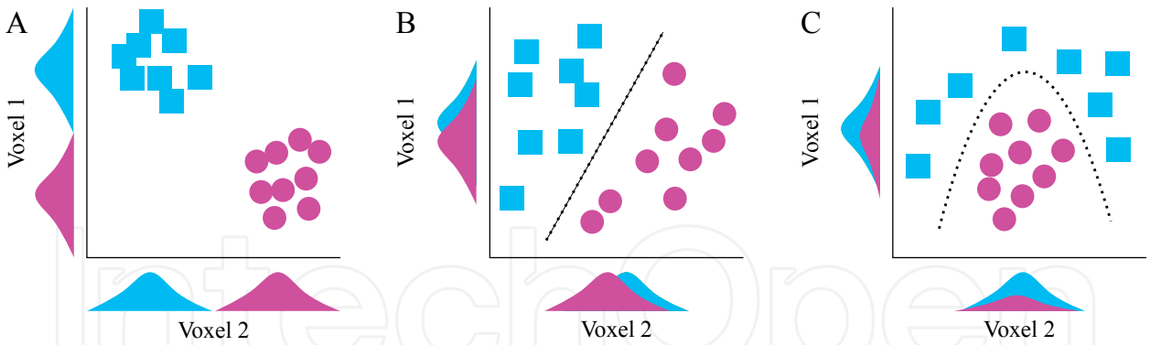

Fig. 2. Two-voxel illustration of the multivoxel analysis approach, where blue and purple represent two different experimental conditions. In **A**, the response distributions to the conditions (the Gaussian curves) are separable in each single voxel and a univariate statistical approach is feasible to distinguish the conditions. In **B**, however, the two conditions can not be separated in each individual voxels due to the overlap of the distributions, and a univariate measure would fail in distinguishing the conditions. A linear decision boundary (dotted line) can, however, separate the conditions. Similarly, in **C**, the conditions can be separated but a non-linear decision boundary is required.

the dataset and improving classification performance, is performed, often in intimate conjunction with classifier training. During classifier training, several processing and voxel selection cycles may therefore be explored. After voxel selection and classifier training, discriminative maps are produced indicating regions encoding information regarding the conditions. Finally, the capability of the trained classifier to accurately discriminate the experimental conditions when presented with data from the hitherto unused partition is tested to assess classifier generalization ability.

4.1 Experimental paradigm

The experimental paradigm considerations for machine learning based analysis are typically the same as for conventional approaches, including those for block-design (see e.g. Björnsdotter et al. (2009)) and event-related (see e.g. Beauchamp et al. (2009); Burke et al. (2004)) paradigms. For machine learning analysis, event-related designs have the benefit of producing more independent datapoints, which, in turn, yields less contaminated estimations of the spatial pattern related to each condition. In theory, this can improve the machine learning algorithm's sensitivity in detecting information contained in the spatial patterns. However, rapid-event related designs risk temporal overlap of hemodynamic responses, although various techniques can be applied to reduce this effect (see e.g. Beauchamp et al. (2009)).

4.2 Pre-processing

The pre-processing steps required for machine learning analysis are essentially identical to conventional analysis pre-processing, with the notable exception of spatial smoothing – if the conditions differ in terms of their fine-grained spatial activation patterns, spatial smoothing will reduce the discriminative information content. Moreover, without smoothing the highest possible spatial resolution offered by the fMRI scanner is preserved, and small differences in location can be maximally resolved. Smoothing may, nonetheless, have a beneficial impact on classification performance (LaConte et al., 2003)

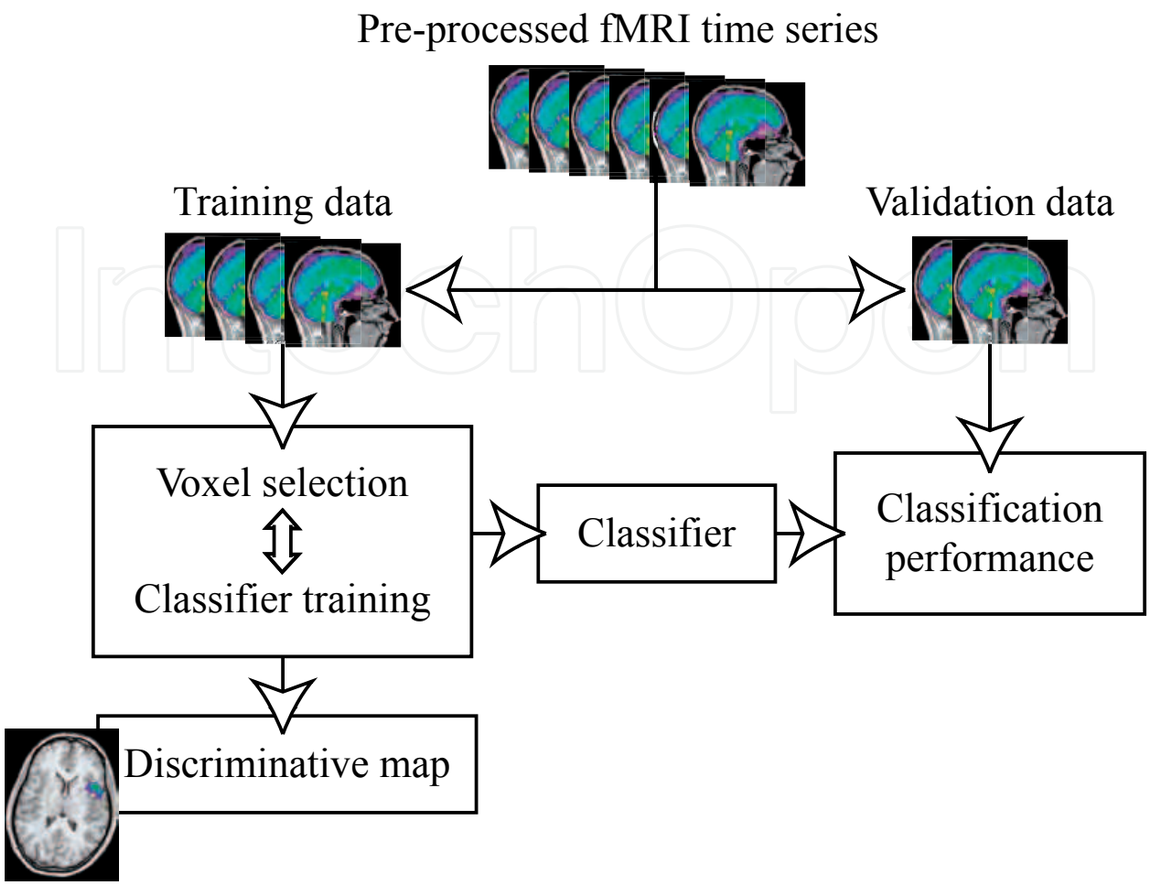

Fig. 3. General multivoxel pattern analysis (MVPA) workflow.

4.3 Condition response estimation

The continuous fMRI brain activity signal, consisting of a series of intensity values for each voxel across the scanning time course, must be re-represented as single condition responses per time unit for subsequent analysis. Generally, one time-measure (volume) is used as a sample, but as outlined in the section on BOLD fMRI above, the hemodynamic response (with delays of approximately 6 s after onset of stimulation) must be accounted for when estimating temporal single-trial responses. A number of activation representations have been used in multivoxel fMRI studies, including the following:

1. *Single-volume intensities*: The intensity at a single acquisition volume can be used to represent a condition Haynes & Rees (2005a); Mourão-Miranda et al. (2005), given appropriate hemodynamic lag compensation. A simple approach to compensating for hemodynamic lag is to simply shift the data labels the corresponding amount of acquisition time points (typically around 6 s).
2. *Volume-average intensities*: When the conditions are applied during multiple volumes (e.g. in a block design study), the average intensity across the volumes Björnsdotter et al. (2009); Kamitani & Tong (2005); Mourão-Miranda & Reynaud (2006) can be used. Typically the first few volumes are also discarded. This approach has the added benefit of increased signal-to-noise ratio.

3. *Single-trial GLM fitting*: A third option is to directly estimate the single-trial response based on the hemodynamic response function De Martino et al. (2008); Formisano et al. (2008); Staeren et al. (2009). Here, a trial corresponds to one application of the stimulus, or, in the block design case, the duration of the 'on' block. A trial estimate of the voxel-wise response is obtained by fitting a general linear model (GLM) with one predictor coding for the trial response and one linear predictor accounting for a within-trial linear trend to each single trial. The trial-response predictor can be obtained by convolution of a boxcar with a double-gamma hemodynamic response function (HRF; Friston et al. (1998)). At every voxel, the corresponding regressor coefficient (beta) is taken to represent the trial response.

4.4 Data partitioning

In order to avoid classifier overfitting and unsound prediction accuracies (see e.g. Kriegeskorte et al., 2009 for a review of this problem in functional brain imaging), care is required when partitioning the data samples into training (used for any aspect of establishing the classifier, including classifier parameter estimation and voxel selection) and validation (exclusively used in the validation of the already established classifier) data.

Potential dependencies between datasets must be carefully avoided – the inherent temporal sluggishness of the hemodynamic response producing temporal dependencies is of particularly concern when fMRI is considered. Thus, any randomization of training and validation samples must be preceded by a condition response estimation (see section above) ensuring no temporal dependencies between samples (or the prediction accuracies will be biased towards the higher end of the spectrum). Another possibility is to select a temporally independent validation sample from samples collected towards the end of the scanning session.

4.5 Classifiers

Classifiers employed for multivoxel pattern analysis of fMRI data range from various versions of linear discriminant analysis Carlson et al. (2003); Haynes & Rees (2005a;b); Kriegeskorte et al. (2006); O'toole et al. (2005), correlation-based classification Haxby et al. (2001); Spiridon & Kanwisher (2002) linear Cox & Savoy (2003); De Martino et al. (2008); Formisano et al. (2008); Kamitani & Tong (2005); LaConte et al. (2005); Mitchell et al. (2004); Mourão-Miranda et al. (2005); Mourão-Miranda & Reynaud (2006); Staeren et al. (2009) and non-linear Cox & Savoy (2003); Davatzikos et al. (2005) support vector machine (SVM), artificial neural networks (ANNs; Hanson et al. (2004); Polyn et al. (2005)) and gaussian naive bayes (GNB) classifiers Mitchell et al. (2004). See e.g. Duda et al. (2000) for details on these classifiers.

Despite the theoretical superiority of non-linear classifiers, linear classifiers have by far been most popular in fMRI multivoxel research (see citations above). A highly appealing advantage of linear classifiers is the direct relation between classifier weights and voxels, providing a means to understand which regions of the brain are multivariately informative De Martino et al. (2008); Mourão-Miranda et al. (2005). Although linear discriminant analysis and linear support vector machines have dominated the field, there appears to be little practical difference between linear classifiers Ku et al. (2008).

4.6 Performance metrics

The performance metrics indicates the ability of the classifier to predict the condition categories to which previously unseen samples belong. Typically, the performance metric is expressed in terms of classification performance (e.g. proportion correctly or incorrectly labelled

instances of the validation data, or area under the receiver operating characteristic curve). Due to the limited number of samples available in fMRI studies, various data partitioning schemes such as cross-validation are utilized in order to obtain a good estimate of the performance Duda et al. (2000).

4.7 Voxel selection

As is well-known from other areas of machine learning, classification performance degrades as the number of irrelevant features increases (party due to the curse of dimensionality; Bellman (1961); Guyon & Elisseeff (2003)). Given the excessive number of voxels in a typical brain volume (in the order of tens to hundreds of thousands) compared to the limited number of available volumes (typically in the range of tens to hundreds), voxel selection is an acute issue (as also pointed out by Norman et al. (2006)). Selection of an adequate subset of voxels is not only of critical importance in order to obtain classifiers with good generalization performance, but also to provide insight into what brain regions encode information relevant to the conditions under investigation ("brain mapping"). Moreover, fMRI voxel selection deviates from conventional feature selection in the sense that the smallest possible subset of voxels which is *sufficient* for maximal classification is not necessarily desired. In fact, *all* voxels which potentially contain relevant information are interesting, as is the relative degree of information content (see Sato et al. (2009) for a discussion on this topic).

In early stages of fMRI multivoxel analysis, the feature selection problem was resolved by region-of-interest (ROI) based methods where classifiers were applied to voxels in anatomically or functionally pre-defined areas Cox & Savoy (2003); Haynes & Rees (2005a); Kamitani & Tong (2005). Coarse brain maps can be obtained, given that some ROIs produce higher classification results than others. Although this approach can be of high utility provided previously determined ROIs, only testing of a highly limited set of spatial hypotheses is possible and no information regarding which number and combination of voxels form a discriminative pattern can be obtained.

Another popular early method is voxel ranking and selection according to various univariate measures. These include estimations of activation magnitude due to any condition, as measured using an F-test (activation-based voxel selection) or the ability to differentiate the conditions, as quantified by parametric (t) or non-parametric (Wilcoxon) statistical tests (discrimination-based voxel selection; Haynes & Rees (2005a); Mitchell et al. (2004); Mourão-Miranda & Reynaud (2006)). The univariate ranking is either used directly (selecting a number of the highest ranked voxels for classification), or for initial, fast but coarse ranking for improved speed and accuracy in subsequent multivariate voxel selection (see e.g. De Martino et al. (2008); Niiniskorpi et al. (2009)). Such activation- and discrimination-based voxel selection, however, disregards any distributed aspects of the brain processing and is thus suboptimal in the processing pipe-line of multivariate analysis.

A second generation of voxel selection methods which utilize the multivariate nature of fMRI data can be categories into two distinct classes – locally multivariate analysis, where information is integrated across multiple voxels in a small neighborhood of adjacent spatial locations (see for example Kriegeskorte et al. (2006)) and globally multivariate, where voxels are jointly analyzed in spatially remote regions or across the entire brain volume (see e.g. De Martino et al. (2008); Mourão-Miranda et al. (2005)).

Although not strictly a voxel selection approach, the locally multivariate method termed "the searchlight" introduced by Kriegeskorte and colleagues (2006) scans the brain with a fixed-size sphere to identify regions which encode information regarding the experimental condi-

tions. This method relies on the assumption that the discriminative information is encoded in neighboring voxels within the sphere. Such locally-distributed analysis might, however, be suboptimal when no hypothesis is available on the size of the neighborhood and might fail to detect discriminative patterns jointly encoded by distant regions (e.g. bilateral activation patterns). Other locally multivariate, fixed-size search approaches have followed suit Björnsdotter & Wessberg (2009). The evolutionary algorithm described in more detail below belongs to the class of locally multivoxel methods. However, the evolutionary algorithm is fundamentally different in that it optimizes voxel cluster size, shape and location in a more traditional feature selection sense Björnsdotter Åberg & Wessberg (2008). The evolutionary algorithm thus produces highly sensitive maps specifically tailored to the spatial extent of the informative region and is suitable in studies where low contrast-to-noise ratio, single optima are expected (such as in the somatotopy study described below; Björnsdotter et al. (2009)).
Global multivoxel selection schemes taking any number of spatially separate regions into account is represented by recursive feature elimination (RFE; De Martino et al. (2008); Formisano et al. (2008); Staeren et al. (2009)) which is initiated by a massively multivariate, whole-brain classification approach Mourão-Miranda et al. (2005). This method requires the use of a linear classifier, where the contribution of each voxel to the classification can be estimated by the classifier weights. The ranking obtained from the classifier weights is subsequently used for iterative elimination of voxels until a voxel subset which maximally discriminate the conditions is obtained. This approach is appropriate when the discrimination of the experimental conditions is reflected by widely distributed activation patterns that extend and include a number of separate brain regions.
In summary, the various voxel selection schemes used in multivoxel pattern analysis for brain mapping differ in scope and sensitivity, and care is required when choosing a voxel selection scheme suitable for the given fMRI study. In all voxel selection cases, group maps describing informative voxels across a number of individuals can be formed using various statistical techniques (see Wang et al. (2007) for details).

5. Brain mapping using evolutionary algorithms

As mentioned above, a key issue in the analysis of fMRI data is the identification of brain areas involved in the processing of given conditions, and, consequently, selection of voxels which can be used to effectively classify the conditions. Given a typical brain volume of $64 \times 64 \times 25$ voxels and the combinatorial explosion of possible voxel combinations, however, voxel selection is a daunting task, and the excessive number of possible voxel combinations renders any exhaustive search virtually impossible (but see the discussion of locally-multivariate mapping above). To address this issue, we developed a machine learning optimization method based on evolutionary algorithms Reeves & Rowe (2002), that extracts voxels yielding optimal brain state classification results in an intelligent and efficient fashion. In the following section we illustrate the evolutionary algorithm in detail, and subsequently demonstrate the utility of machine learning in general and our evolutionary algorithm in particular for highly sensitive mapping of an authentic physiological problem – that of the body-map representation of pleasant touch in the posterior insular cortex.

5.1 Evolutionary algorithms

An evolutionary algorithm is an optimization scheme inspired by Darwinian evolution, where potential problem solutions are encoded as individuals in a population Reeves & Rowe (2002). Various genetic schemes, including mutation operators, selection and reproduction (sexual or

asexual) are subsequently applied in order to improve the over-all fitness of the population and eventually find an individual which fulfills the required fitness. Importantly, biologically interpretable maps require spatial consistency - that is, where a few single, distributed voxels might be sufficient for good classification results (see e.g. Åberg et al. (2008)), from a neuroimaging perspective it is more interesting to identify *all* informative voxels in a given region (again, see Sato et al. (2009) for a discussion on this topic). Our evolutionary algorithm was therefore designed to identify a spatially coherent voxel cluster of unrestricted size which optimally differentiated the conditions.
Below follows details on the implementation of the algorithm, and pseudo-code for the algorithm is presented in figure 4.
Representation: One chromosome representing one voxel cluster was encoded per individual. Due to the exceedingly high dimensionality of fMRI data (in the order of tens to hundreds of thousands of features), we chose to encode the clusters sparsely as indexed lists (and not as typical binary strings).
Initialization: The population of individuals was initialized in a stochastic fashion, where, for each individual, one seed voxel was randomly selected. The voxel cluster was then constructed by the addition of random voxels which neighbor the seed voxel, or, subsequently, any voxel already in the cluster. We have also obtained excellent results from covering the entire brain volume with randomized clusters and subsequently selecting a population of individuals from the best-performing random clusters.
Mutation operations: The following mutation operations were implemented in the algorithm: the addition of a number of voxels, the deletion of a number of voxels, and the substitution of a voxel with another voxel. All voxel additions and substitutions were performed on neighboring voxels, that is, voxels within the 26 voxel cube surrounding any voxel already contained in the cluster. Also, deletions or substitutions resulting in voxels disconnecting from the clusters were disallowed. The frequency of mutation was regulated by a constant mutation rate parameter for each mutation operation. Also, a voxel cluster in the population was occasionally substituted for a new, randomly generated cluster to add fresh genetic material and aid in escaping potential local maxima.
Selection and reproduction: A standard tournament scheme was used for parent selection. In order to retain a variety of the genetic material and maintain searches in widespread regions of the brain, the proportion of parents to discarded individuals was set high. It should be noted, however, that the suitable proportion of parents depends on the expected signal-to-noise-ratio of the data as well as the number of regions of interest. Since all individuals in the population represent different locations and crossover thus would destroy the spatial integrity of the voxel clusters, reproduction was asexual and the new generation was formed by cloning the parents.
Fitness computation: The fitness value, that is, the condition classification success, of each individual cluster was computed using a classifier. Any classifier can be applied (including non-linear schemes; see the discussion on Classifiers above), and we used linear support vector machines Suykens et al. (2002). To ensure high generalization capability, the algorithm is supplied with three datasets. The first was used in classifier training (training data, 35% of the total volumes) while the second was used for fitness estimation (testing data, 45%). The third dataset was exclusively used with the already trained and optimized classifier and voxel cluster (validation data, 20%). Any fitness measure indicative of classification performance can be used, including proportion correctly classified instances or the area under the receiver operating characteristic curve (see the discussion on Performance metrics above).

Termination criterium: The algorithm was run for either a pre-determined maximum number of generations or until a cluster yielding testing data classification rates above a given threshold was obtained. The algorithm can, however, overtrain when allowed to run the full course. We therefore used the cluster with the best result on the mean of the training and testing data performance for validation classification.

```
BEGIN
  Initialize population;
  While (termination criteria not met);
        For (each individual);
              Apply mutation operations;
                    1. Add n_a random voxels;
                    2. Remove n_r random voxels;
                    3. Substitute n_s random voxels;
        End For
        Select parents;
        Reproduce;
        generation = generation + 1;
  End While
END
```

Fig. 4. Pseudo-code for evolutionary voxel selection (Åberg & Wessberg, 2008).

5.2 Example: Somatotopic organization of the insular cortex due to pleasant touch

Machine learning in general and the evolutionary algorithm described above in particular have been of high utility in our research on how the brain processes pleasant touch. In the study described here we specifically used the advantages of the superior sensitivities of multi-voxel pattern analysis to demonstrate that a specific region of the cortex is activated differently depending on what body part is stimulated.

The brain receives information about touch through two fundamentally different systems – one network of thick, myelinated fibers, termed Aβ afferents, which transmit discriminative aspects of touch (e.g. what type of texture am I touching?), and a parallel network of thin, unmyelinated so called C-tactile fibers, which are primarily activated by gentle types of touch and signal hedonic, pleasant features of the tactile sensation Löken et al. (2009); Olausson et al. (2002); Vallbo et al. (1999; 1993). The C-tactile system was recently discovered in humans, and is currently under intense research Edin (2001); Löken et al. (2007); Nordin (1990); Vallbo et al. (1999).

The C-tactile system is difficult to study in relation to brain processing, mainly since it cannot be activated selectively – any mechanical stimulation of the skin invariably co-activates thick, myelinated fibers in healthy subjects. C-tactile physiology has, nevertheless, been successfully explored in a patient (GL) with neuronopathy syndrome Sterman et al. (1980), who lacks large myelinated Aβ afferents but whose C fibers are intact Olausson, Cole, Rylander, McGlone, Lamarre, Wallin, Krämer, Wessberg, Elam, Bushnell & Vallbo (2008); Olausson et al. (2002). Due to the lack of Aβ fibers, GL has no sensation of touch. We demonstrated, however, that

she can detect the light stroking of a brush and reports a pleasant sensation in response to the stimuli Olausson et al. (2002). Surprisingly, we recently found that she can not only detect the stimulus, but also distinguish the body quadrant to which it was applied at an accuracy of 72% Olausson, Cole, Rylander, McGlone, Lamarre, Wallin, Krämer, Wessberg, Elam, Bushnell & Vallbo (2008). GL's performance suggests that the C-tactile system projects some, albeit crude, information about stimulus location.

The C-tactile system belongs to a class of fibers (C afferents) which projects various information about the physiological condition of the body, including temperature and pain, from the periphery to the brain (see Craig (2002) for a review of this system). As opposed to Aβ fibers, which project directly to brain regions specifically processing tactile information (the somatosensory cortices), C afferents have been shown to connect from the thalamus to a part of the brain called the insular cortex on the opposite side of the brain to which the stimulation was applied Craig et al. (2000; 1994); Hua le et al. (2005). Similarly, functional imaging in GL and a similarly deafferented subject (IW) revealed that C-tactile stimulation also activates the insular cortex Olausson, Cole, Vallbo, McGlone, Elam, Krämer, Rylander, Wessberg & Bushnell (2008); Olausson et al. (2002). Moreover, for pain and temperature stimulation, the posterior part of the insular cortex has been shown to be organized in a particular fashion – upper body stimulation stimulations project anterior (closer to the nose) than lower body part stimulations Brooks et al. (2005); Henderson et al. (2007); Hua le et al. (2005). This type of organization is referred to as somatotopic and is observed in several brain regions where the physical location of the stimulation is important, such as the primary somatosensory and motor cortices.

A corresponding somatotopic organization of the insular cortex due to pleasant touch could potentially explain GL's surprisingly good localization performance. However, as opposed for example painful stimuli, C-tactile activations are typically weak and often difficult to identify using conventional statistical methods. In order to investigate whether posterior insular cortex projections due to pleasant touch are organized in a body-map similar to that of painful and temperature stimulation, we therefore applied the evolutionary algorithm to data acquired as described below.

6. Data acquisition

Participants: Informed consent was obtained from six healthy subjects, as well as one subject (GL) with sensory neuronopathy syndrome Sterman et al. (1980). At the age of 31, GL suffered permanent loss specifically of thick myelinated afferents (the so called $A\beta$ fibers), leaving unmyelinated and small-diameter myelinated afferents intact Forget & Lamarre (1995). She can detect temperature and pain normally Olausson, Cole, Rylander, McGlone, Lamarre, Wallin, Krämer, Wessberg, Elam, Bushnell & Vallbo (2008); Olausson et al. (2002), but denies any ability to sense touch below the level of the nose Forget & Lamarre (1995). In a forced choice task she could, however, perceive light touch, and she failed to detect vibratory stimuli (which poorly excite C-tactile afferents; Olausson, Cole, Rylander, McGlone, Lamarre, Wallin, Krämer, Wessberg, Elam, Bushnell & Vallbo (2008)). Moreover, in a four-alternative forced choice procedure, she identified 72% of soft brush stimuli to the correct extremity (at chance level of 25%). Healthy subjects, on the other hand, detect light touch as well as vibration without fail, and can localize point indentation on the skin to an accuracy in the range of two centimeters Norrsell & Olausson (1994). The Ethical Review Board at the University of Gothenburg approved the study, and the experiments were performed in accordance with the Declaration of Helsinki.

Stimulation: Light stimulation, known to vigorously activate C-tactile afferents in humans as well as in other species Bessou et al. (1971); Douglas & Ritchie (1957); Edin (2001); Kumazawa & Perl (1977); Nordin (1990); Vallbo et al. (1999; 1993); Zotterman (1939), was delivered using a seven centimeter wide soft brush with an indentation force 0.8 N. The experimenter manually stroked the brush in a proximal to distal direction on the right forearm or thigh.
Scanning protocol: The experimenter applied the tactile stimulation according to timing cues from the scanner, and all subjects were instructed to focus on the stimulus throughout the fMRI scanning session. In the healthy subjects, the distance covered was 16 centimeters for a duration of three seconds, whereas in the case of GL the distance was 30 centimeters and the duration four seconds. Three-volume blocks of forearm brushing, thigh brushing or no brushing (rest) were alternated in a pseudo-random order with equal numbers of each of the three conditions. The condition order remained fixed throughout each scan and across participants, and the scanning session consisted of one anatomical and six functional scans. During each functional scan, 13 blocks were obtained in the healthy subjects and 10 in GL, totaling in 78 and 60 three-volume blocks per condition respectively. A 1.5 T fMRI scanner (healthy subjects: Philips Intera; GL: Siemens Sonata) with a SENSE head coil (acceleration factor 1) was used to collect whole brain scans.
Pre-processing: The standard pre-processing steps described previously were applied to the data. Motion correction was performed using the Neurolens software package (www.neurolens.org; developed at the Neurovascular Imaging Lab, UNF Montreal, Canada), whereas the remaining pre-processing steps were performed using custom-coded scripts in Matlab (The Mathworks, Natick, MA). To offset hemodynamic delay and minimize within-trial variability, the first volume in each block was discarded and an average over the remaining two volumes was obtained (leaving a total of 78 volumes per stimulus for the healthy subjects and 60 for GL). The posterior contralateral (left) insula, containing a subject mean of 222 (range 177-261) voxels for the healthy participants and 97 voxels for GL, was subsequently identified using an anatomical reference Naidich et al. (2004) and previous multivoxel analysis (see Björnsdotter et al. (2009) for further details on the identification of the ROI). All analysis was performed in original individual space, and the resulting maps were transformed into MNI (Montreal Neurological Institute) standard stereotactic space Evans et al. (1993) using SPM5 Friston et al. (1994b) with the supplied EPI brain as template. For data visualization, the programs MRIcron (by Chris Rorden, `www.sph.sc.edu/comd/rorden/mricron/`) and Cartool (by Denis Brunet, `http://brainmapping.unige.ch/Cartool.htm`) were used.
Conventional analysis: A general linear model (GLM) whole-brain analysis was performed on smoothed data (Gaussian filter FWHM 6 mm). A fixed effect model was used to generalize healthy subject activations to the group level. The resulting activation maps were thresholded to a false discovery rate (FDR) of < 0.01).
Machine learning analysis: In order to compare the forearm and thigh brushing projections, the evolutionary clustering scheme was applied to the forearm/rest and thigh/rest datasets separately within the region of interest (ROI) in the posterior insular cortex. The number of voxels allowed in the cluster was fixed in order to obtain directly comparable classification performances within and between individuals. In GL, seven voxels was empirically determined to be a suitable cluster size for high classification performance. The corresponding volume in the healthy subjects, due to the higher spatial sampling frequency of the functional data, was 20 voxels. The algorithm was iterated 200 times, and the clusters which maximally differentiated the forearm and thigh stimuli from rest were identified. The magnitude of condition separabil-

ity (classification score) was quantified by the area under the receiver operating characteristic (ROC) curve (AUC).

7. Results

Conventional analysis: The whole-brain general linear model approach identified a variety of expected activated regions in the healthy subjects, including the contralateral postcentral gyrus (the primary somatosensory cortex; T-value peak at MNI [X, Y, Z] coordinates 5.2 [-26, -42, 64] for forearm stimulation and 6.9 [-16, -48, 70] for thigh stimulation) and bilateral parietal operculum (secondary somatosensory cortex). No significant activations were identified in the insular cortex, however, and for the neuronopathy patient GL no significantly activated voxels were identified in any region of the brain.
Machine learning analysis: The evolutionary multivoxel pattern recognition approach was substantially more successful, however, and the condition discrimination performance of the identified clusters were highly significant in GL as well as for the healthy volunteers. Forearm and thigh tactile stimulation were found to project to distinctly separate locations in GL, with a substantial euclidean distance between cluster centroids of 8.9 mm (figure 5A). The forearm cluster centroid was located at MNI (X, Y, Z) coordinates (-34, -10, 4), and the thigh cluster was found at (-34, -18, 0). The distance between clusters was thus maximal in the anterior-posterior (Y) plane at 8 mm, whereas the location differences in the remaining planes were non-existent or small (X: 0 mm, Z: 4 mm). Validating the pattern observed in GL, the healthy subject insular projections were also arranged in a clear somatotopic fashion. The individual forearm and thigh cluster centroid locations are shown in figure 5B, illustrating the consistency in activation pattern across all individuals including GL. The difference in location was significant only in the Y-plane (anterior-posterior; two-tailed paired t-test, $p < 0.05$). The subject mean euclidean distance between the cluster centroids equaled that of GL at 9.3 (range 6.6-12) mm.
Our evolutionary machine learning algorithm was able to identify the subtle fMRI pattern produced by pleasant touch, invisible to conventional univariate analysis. Moreover, we demonstrated that such pleasant touch projects to different regions of the insular cortex depending on what part of the body was stimulated.
The observation of the body-map topology has important physiological consequences, mainly in strengthening the suggestion that C-tactile afferents are organized in a fashion similar to that of the pain and temperature systems. Also, our study suggests that, although the Aβ system is clearly dominant for touch discrimination and localization, there is some localization functionality to the C-tactile system. It appears improbable that the C-tactile system plays a significant role in discriminative touch, yet it can be presumed that the general stimulus location significantly modulates *affective* sensations which are intimately related to C-tactile activity Löken et al. (2009). Propagation of such affective information is of fundamental importance in preparing appropriate actions in response to emotionally relevant stimuli. Thus, we hypothesize that the crude localization capacity of the C-tactile system serves an affective function, where, for example, a gentle stroke on the cheek evokes a different motivational and hedonic response than that on the leg, thus signaling various emotional aspects with corresponding social consequences.

8. Concluding remarks

In contrast to conventional statistical techniques based on average voxel-by-voxel activations, machine learning-based multivoxel pattern analysis utilizes the inherent multivariate nature

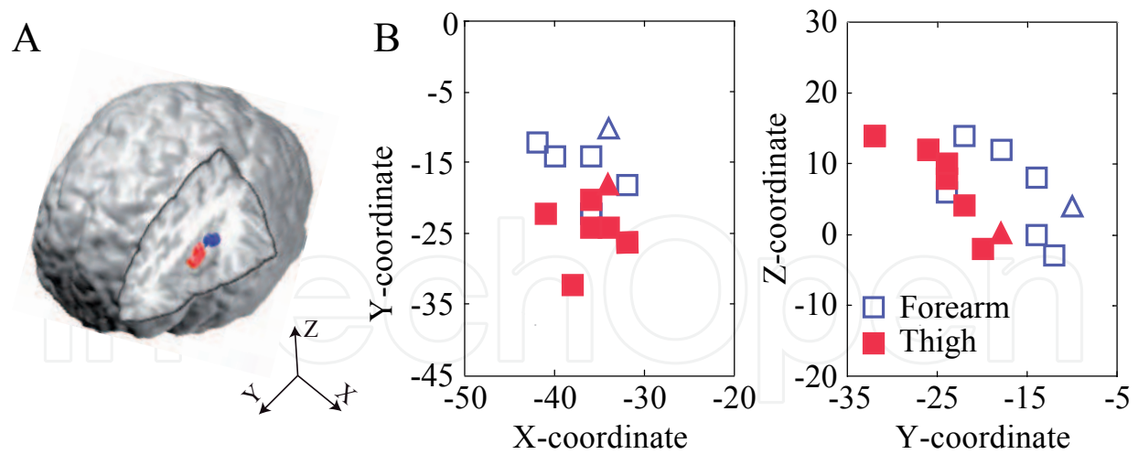

Fig. 5. Insular somatotopy of the contra-lateral posterior insular cortex due to pleasant touch identified using the evolutionary voxel selection scheme. **A**, The voxel clusters maximally differentiating forearm (red) and thigh (blue) stimulation from rest in the neuronopathy syndrome patient GL, reflecting the projection of pleasant touch afferents. A similar somatotopic organization was consistently identified also in neurologically intact subjects, as demonstrated in panel **B**, showing the forearm (red/filled) and thigh (blue/empty) cluster centroid MNI coordinates for each of the healthy subjects (□) and the neuronopathy syndrome patient GL (△). There was a significant difference between forearm and thigh cluster centroid location in the Y-plane only (two-tailed paired t-test, $p < 0.05$).

of brain activity and highlight informative spatial patterns. As such, these state-of-the-art analysis methods outperform conventional techniques in terms of brain mapping sensitivity, but also provides a direct link between brain state and brain activity patterns. Moreover, brain responses to conditions are treated as independent - as opposed to average - evoked activation patterns, allowing interesting and highly novel studies of functional representation Kriegeskorte et al. (2008). Also, as the spatial resolution is continuously improved with advances in high field imaging Yacoub et al. (2008), issues such as reduced signal-to-noise ratio and escalating multiple comparison problem will render univariate analysis less feasible and pave the way for multivoxel analyses Kriegeskorte & Bandettini (2007).
The utility of machine learning in neuroimaging is evidenced by the recent surge of studies taking advantage of the appealing benefits of multivoxel analysis – despite the fact that appropriate application of machine learning concepts to fMRI analysis currently requires not only an understanding of brain physiology, but also solid technical and mathematical knowledge. Further interdisciplinary research aiming to refine, develop and integrate machine learning techniques for standard fMRI analysis promises exciting possibilities for improved insight into the inner workings of the human brain.

9. Acknowledgments

This research was supported by the Swedish Research Council (grant K2007-63X-3548, Dr. J. Wessberg) and the Sahlgrenska University Hospital (grant ALFGBG 3161). The fMRI data was acquired in collaboration with Line Löken, Karin Rylander, Linda Lundblad, Dr. Håkan

Olausson and Dr. Catherine Bushnell. I am highly grateful to GL for her invaluable contribution, and to Dr. Johan Wessberg and Dr. Federico De Martino for continuous support and expert comments on the manuscript.

10. References

Åberg, M., Löken, L. & Wessberg, J. (2008). An evolutionary approach to multivariate feature selection for fMRI pattern analysis, *Proceedings of the International Conference on Bio-inspired Systems and Signal Processing* .

Bandettini, P. A. & Cox, R. W. (2000). Event-related fMRI contrast when using constant interstimulus interval: theory and experiment., *Magnetic Resonance in Medicine* **43**(4): 540–548. **URL:** *http://view.ncbi.nlm.nih.gov/pubmed/10748429*

Beauchamp, M., Laconte, S. & Yasar, N. (2009). Distributed representation of single touches in somatosensory and visual cortex., *Human Brain Mapping* **(in press)**.

Bellman, R. E. (1961). *Adaptive Control Processes*, Princeton University Press, Princeton, NJ.

Bessou, P., Burgess, P. R., Perl, E. R. & Taylor, C. (1971). Dynamic properties of mechanoreceptors with unmyelinated (c) fibers, *Journal of neurophysiology* **34**(1): 116–31.

Björnsdotter Åberg, M. & Wessberg, J. (2008). An evolutionary approach to the identification of informative voxel clusters for brain state discrimination, *IEEE Journal of Selected Topics in Signal Processing* **2**(6): 919–928.

Björnsdotter, M., L. S. Löken, H. Olausson, A. B. Vallbo, and J. Wessberg (2009). Somatotopic organization of gentle touch processing in the posterior insular cortex. Journal of Neuroscience 29(29): 9314-9320.

Björnsdotter, M. & Wessberg, J. (2009). Particle swarm voxel clustering for multivariate fmri mapping, *Organization for Human Brain Mapping Annual Meeting, San Fransisco, USA* .

Brooks, J. C. W., Zambreanu, L., Godinez, A., Craig, A. B. & Tracey, I. (2005). Somatotopic organisation of the human insula to painful heat studied with high resolution functional imaging, *Neuroimage* **27**: 201–209.

Burke, D., Murphy, K., Garavan, H. & Reilly, R. (2004). Pattern recognition approach to the detection of single-trial event-related functional magnetic resonance images., *Medical and Biological Engineering and Computing* **42**(5): 604–9.

Buxton, R. B. (2002). *An introduction to functional magnetic resonance imaging*, Cambridge University Press, Cambridge, United Kingdom.

Carlson, T. A., Schrater, P. & He, S. (2003). Patterns of activity in the categorical representations of objects., *Journal of Cognitive Neuroscience* **15**(5): 704–717.
URL: *http://www.mitpressjournals.org/doi/abs/10.1162/jocn.2003.15.5.704*

Collins, D. L., Neelin, P., Peters, T. M. & Evan, A. C. (1994). Automatic 3d intersubject registration of MR volumetric data in standardized Talairach space., *Journal of Computer Assisted Tomography* **18**(2): 192–205.

Cox, D. D. & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) 'brain reading': detecting and classifying distributed patterns of fMRI activity in human visual cortex., *Neuroimage* **19**(2 Pt 1): 261–270.

Craig, A., Chen, K., Bandy, D. & Reiman, E. (2000). Thermosensory activation of insular cortex., *Nature Neuroscience* **3**(2): 184–190.

Craig, A. D. (2002). How do you feel? interoception: the sense ofthe physiological condition of the body, *Nature Reviews Neuroscience* **3**(8): 655–666.

Craig, A. D., Bushnell, M., Zhang, E. & Blomqvist, A. (1994). A thalamic nucleus specific for pain and temperature sensation., *Nature* **372**(6508): 770–3.

Crinion, J., Ashburner, J., Leff, A., Brett, M., Price, C. & Friston, K. (2007). Spatial normalization of lesioned brains: Performance evaluation and impact on fMRI analyses, *Neuroimage* **37**(3): 866–875.

Davatzikos, C., Ruparel, K., Fan, Y., Shen, D., Acharyya, M., Loughead, J., Gur, R. & Langleben, D. (2005). Classifying spatial patterns of brain activity with machine learning methods: Application to lie detection., *Neuroimage* **28**: 663–668.

De Martino, F., Valente, G., Staeren, N., Ashburner, J., Goebel, R. & Formisano, E. (2008). Combining multivariate voxel selection and support vector machines for mapping and classification of fMRI spatial patterns, *Neuroimage* **43**(1): 44–58.
URL: *http://dx.doi.org/10.1016/j.neuroimage.2008.06.037*

Douglas, W. W. & Ritchie, J. M. (1957). Nonmedullated fibres in the saphenous nerve which signal touch., *Journal of Physiology* **139**(3): 385–99.

Duda, R. O., Hart, P. E. & Stork, D. G. (2000). *Pattern Classification*, Wiley-Interscience Publication.

Edin, B. (2001). Cutaneous afferents provide information about knee joint movements in humans., *Journal of Physiology* **531**(Pt 1): 289–97.

Evans, A. C., Collins, D. L., Mills, S. R., Brown, E. D., Kelly, R. L. & Peters, T. M. (1993). 3d statistical neuroanatomical models from 305 MRI volumes, *Proceedings of the IEEE-Nuclear Science Symposium and Medical Imaging Conference* pp. 1813–1817.

Forget, R. & Lamarre, Y. (1995). Postural adjustments associated with different unloadings of the forearm: effects of proprioceptive and cutaneous afferent deprivation., *Canadian Journal of Physiology and Pharmacology* **73**(2): 285–94.

Formisano, E., De Martino, F., Bonte, M. & Goebel, R. (2008). "who" is saying "what"? brain-based decoding of human voice and speech., *Science* **5903**(322): 970–3.

Frackowiak, R., Friston, K., Frith, C., Dolan, R. & Mazziotta, J. (1997). *Human Brain Function.*, Academic Press.

Friston, K., Ashburner, J., Kiebel, S., Nichols, T. & Penny, W. (eds) (2007). *Statistical Parametric Mapping: The Analysis of Functional Brain Images*, Academic Press.
URL: *http://books.elsevier.com/neuro/?isbn=9780123725608&srccode=89660*

Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M. D. & Turner, R. (1998). Event-related fMRI: Characterizing differential responses, *Neuroimage* **7**(1): 30–40.
URL: *http://dx.doi.org/10.1006/nimg.1997.0306*

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D. & Frackowiak, R. S. J. (1994a). Statistical parametric maps in functional imaging: A general linear approach, *Human Brain Mapping* **2**(4): 189–210.
URL: *http://dx.doi.org/10.1002/hbm.460020402*

Friston, K. J., Holmes, A. P., Worsley, K. J., Poline, J. P., Frith, C. D. & Frackowiak, R. S. J. (1994b). Statistical parametric maps in functional imaging: A general linear approach, *Human Brain Mapping* **2**(4): 189–210.

Goense, J. B. & Logothetis, N. K. (2008). Neurophysiology of the BOLD fMRI signal in awake monkeys, *Current Biology* **18**(9): 631–640.

Guyon, I. & Elisseeff, A. (2003). An introduction to variable and feature selection, *Journal of machine learning research* **3**(1): 1157–1182.

Hanson, S. J., Matsuka, T. & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: is there a "face" area?, *Neuroimage* **23**(1): 156–66.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L. & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex., *Science* **293**: 2425–2430.

Haynes, J. & Rees, G. (2005a). Predicting the orientation of invisible stimuli from activity in human primary visual cortex, *Nature Neuroscience* **8**(5): 686–691.
URL: *http://www.nature.com/neuro/journal/v8/n5/abs/nn1445.html*

Haynes, J. & Rees, G. (2005b). Predicting the stream of consciousness from activity in human visual cortex., *Current Biology* **15**: 1301–1307.

Haynes, J. & Rees, G. (2006). Decoding mental states from brain activity in humans, *Nature Reviews Neuroscience* **7**(7): 523–534.
URL: *http://dx.doi.org/10.1038/nrn1931*

Henderson, L., Gandevia, S. & Macefield, V. (2007). Somatotopic organization of the processing of muscle and cutaneous pain in the left and right insula cortex: A single-trial fMRI study, *Pain* **128**: 20–30.

Henson, R. (2003). Analysis of fMRI time series, *in* R. Frackowiak, K. Friston, C. Frith, R. Dolan, K. Friston, C. Price, S. Zeki, J. Ashburner & W. Penny (eds), *Human Brain Function*, 2nd edn, Academic Press.

Howard, J., Plailly, J., Grueschow, M., Haynes, J. & Gottfried, J. (2009). Odor quality coding and categorization in human posterior piriform cortex., *Nature Neuroscience* **(In press)**.

Hua le, H., Strigo, I. A., Baxter, L. C., Johnson, S. C. & Craig, A. D. (2005). Anteroposterior somatotopy of innocuous cooling activation focus in human dorsal posterior insular cortex., *American Journal of Physiology - Regulatory, Integrative, and Comparative Physiology* **289**(2): 319–325.

Kamitani, Y. & Tong, F. (2005). Decoding the visual and subjective contents of the human brain, *Nature Neuroscience* **8**(5): 679–685.
URL: *http://www.nature.com/neuro/journal/v8/n5/abs/nn1444.html*

Kriegeskorte, N. & Bandettini, P. (2007). Combining the tools: activation- and information-based fmri analysis, *NeuroImage* **38**: 666–668.

Kriegeskorte, N., Goebel, R. & Bandettini, P. (2006). Information-based functional brain mapping, *PNAS* **103**: 3863–3868.

Kriegeskorte, N., Mur, M., Ruff, D. A. A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K. & Bandettini, P. A. A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey., *Neuron* **60**(6): 1126–1141.
URL: *http://dx.doi.org/10.1016/j.neuron.2008.10.043*

Ku, S. P., Gretton, A., Macke, J. & Logothetis, N. K. (2008). Comparison of pattern recognition methods in classifying high-resolution bold signals obtained at high magnetic field in monkeys., *Magnetic resonance imaging* **26**(7): 1007–1014.
URL: *http://dx.doi.org/10.1016/j.mri.2008.02.016*

Kumazawa, T. & Perl, E. (1977). Primate cutaneous sensory units with unmyelinated (C) afferent fibers., *Journal of Neurophysiology* **40**(6): 1325–38.

LaConte, S., Strother, S., Cherkassky, V., Anderson, J. & Hu, X. (2005). Support vector machines for temporal classification of block design fMRI data., *Neuroimage* **26**(2): 317–29.

Leoni, R., Mazzeto-Betti, K., Andrade, K. & de Araujo, D. (2008). Quantitative evaluation of hemodynamic response after hypercapnia among different brain territories by fMRI., *Neuroimage* **41**(4): 1192–8.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T. & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fmri signal., *Nature* **412**(6843): 150–157.
URL: *http://dx.doi.org/10.1038/35084005*

Löken, L., Wessberg, J., Morrison, I., McGlone, F. & Olausson, H. (2009). Coding of pleasant touch by unmyelinated afferents in humans., *Nature Neuroscience* **12**(5): 547–8.

Löken, L., Wessberg, J. & Olausson, H. W. (2007). Unmyelinated tactile (CT) afferents are present in the human peroneal and radial nerves, *Society for Neuroscience 38th Annual Meeting, San Diego, USA* (827.2).

McCarthy, G., Luby, M., Gore, J. & Goldman-Rakic, P. (1997). Infrequent events transiently activate human prefrontal and parietal cortex as measured by functional mri., *Journal of Neurophysiology* **77**(3): 1630–1634.

Mitchell, T. M., Hutchinson, R., Niculescu, R. S., Pereira, F., Wang, X., Just, M. & Newman, S. (2004). Learning to decode cognitive states from brain images, *Machine Learning* **57**(1-2): 145–175.

Mourão-Miranda, J., Bokde, A. L., Born, C., Hampel, H. & Stetter, M. (2005). Classifying brain states and determining the discriminating activation patterns: Support vector machine on functional MRI data, *Neuroimage* **28**(4): 980–95.

Mourão-Miranda, J. & Reynaud, E., M. F. C. G. B. M. (2006). The impact of temporal compression and space selection on SVM analysis of single-subject and multi-subject fMRI data., *Neuroimage* **33**(4): 1055–65.

Muehllehner, G. & Karp, J. S. (2006). Positron emission tomography, *Physics in Medicine and Biology* **51**(13): R117–R137.
URL: *http://dx.doi.org/10.1088/0031-9155/51/13/R08*

Naidich, T., Kang, E., Fatterpekar, G., Delman, B., Gultekin, S., Wolfe, D., Ortiz, O., Yousry, I., Weismann, M. & Yousry, T. (2004). The insula: anatomic study and MR imaging display at 1.5 T., *American Journal of Neuroradiology* **25**(2): 222–32.

Niiniskorpi, T., Björnsdotter Åberg, M. & Wessberg, J. (2009). Particle swarm feature selection for fmri pattern classification, Proceedings of the International Conference on Bio-inspired Systems and Signal Processing, Porto, Portugal.

Nordin, M. (1990). Low-threshold mechanorecepetive and nociceptive units with unmyelinated C fibres in the human supraorbital nerve, *Journal of Physiology* **426**(310): 229–240.

Norman, K. A., Polyn, S. M., Detre, G. J. & Haxby, J. V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data, *Trends in Cognitive Sciences* **10**(9): 424–430.

Norris, D. G. (2006). Principles of magnetic resonance assessment of brain function, *Journal of Magnetic Resonance Imaging* **23**(6): 794–807.
URL: *http://dx.doi.org/10.1002/jmri.20587*

Norrsell, U. & Olausson, H. (1994). Spatial cues serving the tactile directional sensibility of the human forearm., *Journal of Physiology* **478**(Pt 3): 533–40.

Ogawa, S., Lee, T. M., Kay, A. R. & Tank, D. W. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation., *PNAS* **87**(24): 9868–9872.
URL: *http://dx.doi.org/10.1073/pnas.87.24.9868*

Olausson, H., Cole, J., Rylander, K., McGlone, F., Lamarre, Y., Wallin, B., Krämer, H., Wessberg, J., Elam, M., Bushnell, M. & Vallbo, Å. (2008). Functional role of unmyelinated tactile afferents in human hairy skin: sympathetic response and perceptual localization., *Experimental Brain Research* **184**(1): 135–40.

Olausson, H., Cole, J., Vallbo, Å., McGlone, F., Elam, M., Krämer, H., Rylander, K., Wessberg, J. & Bushnell, M. (2008). Unmyelinated tactile afferents have opposite effects on insular and somatosensory cortical processing., *Neuroscience Letters* **436**(2): 128–32.

Olausson, H., Lamarre, Y., Backlund, H., Morin, C., Wallin, B. G., Starck, G., Ekholm, S., Strigo, I., Worsley, K., Vallbo, Å. B. & Bushnell, M. C. (2002). Unmyelinated tactile afferents signal touch and project to insular cortex., *Nature Neuroscience* **5**(9): 900–904.

O'toole, A. J., Jiang, F., Abdi, H. & Haxby, J. V. (2005). Partially distributed representations of objects and faces in ventral temporal cortex, *Journal of Cognitive Neuroscience* **17**(4): 580–590.

Polyn, S. M., Natu, V. S., Cohen, J. D. & Norman, K. A. (2005). Category-specific cortical activity precedes retrieval during memory search, *Science* **310**(5756): 1963–6.

Reeves, C. & Rowe, J. E. (2002). *Genetic Algorithms - Principles and Perspectives: A Guide to GA Theory*, Kluwer Academic Publishers, Norwell, MA, USA.

Rosen, B. R., Buckner, R. L. & Dale, A. M. (1998). Event-related functional MRI: Past, present, and future, *PNAS* **95**(3): 773–780.
URL: *http://dx.doi.org/10.1073/pnas.95.3.773*

Roy, C. S. & Sherrington, C. S. (1890). On the regulation of the blood-supply of the brain., *Journal of Physiology* **11**(1-2).
URL: *http://view.ncbi.nlm.nih.gov/pubmed/16991945*

Sato, J., Fujita, A., Thomaz, C., Martin Mda, G., Mourão-Miranda, J., Brammer, M. & Amaro Junior, E. (2009). Evaluating SVM and MLDA in the extraction of discriminant regions for mental state prediction., *Neuroimage* **46**(1): 105–14.

Sirotin, Y. B. & Das, A. (2008). Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity, *Nature* **457**(7228): 475–479.
URL: *http://dx.doi.org/10.1038/nature07664*

Spiridon, M. & Kanwisher, N. (2002). How distributed is visual category information in human occipito-temporal cortex? an fMRI study., *Neuron* **35**: 1157–1165.

Staeren, N., Renvall, H., F., D., Goebel, R. & Formisano, E. (2009). Sound categories are represented as distributed patterns in the human auditory cortex., *Current Biology* **19**(6): 498–502.

Sterman, A. B., Schaumburg, H. H. & Asbury, A. K. (1980). The acute sensory neuronopathy syndrome: a distinct clinical entity., *Annals of Neurology* **7**(4): 354–8.

Suykens, J., Gestel, T. V., Brabanter, J. D., Moor, B. D. & Vandewalle, J. (2002). *Least Squares Support Vector Machines*, World Scientific.

Talairach, J. & Tournoux, P. (1988). *Co-Planar Stereotaxic Atlas of the Human Brain: 3-Dimensional Proportional System : An Approach to Cerebral Imaging*, Thieme Medical Publishers.
URL: *http://www.amazon.co.uk/exec/obidos/ASIN/0865772932/citeulike-21*

Turner, R., Howseman, A. & Friston, K. (1998). Functional magnetic resonance imaging of the human brain: Data acquisition and analysis., *Experimental Brain Research* **123**: 5.

Vallbo, Å., Olausson, H. & Wessberg, J. (1999). Unmyelinated afferents constitute a second system coding tactile stimuli of the human hairy skin, *Journal of Neurophysiology* **81**(310): 2753–2763.

Vallbo, Å., Olausson, H., Wessberg, J. & Norrsell, U. (1993). A system of unmyelinated afferents for innocuous mechanoreception in the human skin, *Brain Research* **628**(310): 301–304.

Wang, Z., Childress, A., Wang, J. & Detre, J. (2007). Support vector machine learning-based fmri data group analysis., *Neuroimage* **15**(36): 1139–51.

Yacoub, E., Harel, N. & Ugurbil, K. (2008). High-field fMRI unveils orientation columns in humans., *PNAS* **105**(30): 10607–12.
Zotterman, Y. (1939). Touch, pain and tickling: an electro-physiological investigation on cutaneous sensory nerves., *Journal of Physiology* **95**(1): 1–28.

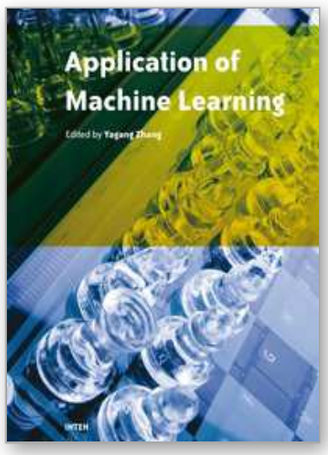

Application of Machine Learning
Edited by Yagang Zhang

ISBN 978-953-307-035-3
Hard cover, 280 pages
Publisher InTech
Published online 01, February, 2010
Published in print edition February, 2010

The goal of this book is to present the latest applications of machine learning, which mainly include: speech recognition, traffic and fault classification, surface quality prediction in laser machining, network security and bioinformatics, enterprise credit risk evaluation, and so on. This book will be of interest to industrial engineers and scientists as well as academics who wish to pursue machine learning. The book is intended for both graduate and postgraduate students in fields such as computer science, cybernetics, system sciences, engineering, statistics, and social sciences, and as a reference for software professionals and practitioners. The wide scope of the book provides them with a good introduction to many application researches of machine learning, and it is also the source of useful bibliographical information.

InTech Europe
University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China
Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

IntechOpen