Information processing algorithms in the brain

Philippe G. Schyns[1], Frédéric Gosselin[2] and Marie L. Smith[1]

[1] Centre for Cognitive Neuroimaging, Department of Psychology, University of Glasgow, 58 Hillhead Street, Glasgow G12 8QB, UK
[2] Département de Psychologie, Université de Montréal, CP 6128, succ. Centre-ville, Montréal H3C 3J7, Canada

If the brain is a machine that processes information, then its cognitive activity can be interpreted as a set of information processing states linking stimulus to response (i.e. as a mechanism or an algorithm). The cornerstone of this research agenda is the existence of a method to translate the measurable states of brain activity into the information processing states of a cognitive theory. Here, we contend that reverse correlation methods can provide this translation and we frame the transitions between information processing states in the context of automata theory. We illustrate, using examples from visual cognition, how this novel framework can be applied to understand the information processing algorithms of the brain in cognitive neuroscience.

Algorithms in the brain

Consider your favourite email program running on your desktop computer. As you depress 'get new mail', an electronic storm of microstates occurs in several regions of the motherboard to implement this function. Depressing 'store mail' elicits a different storm. Detailed measurements of the motherboard electronics would reveal to a computer engineer that the microstates implementing 'get new mail' differ from those implementing 'store mail'. These measurements, aggregated over time, could also reveal the subsets of motherboard components mostly activated to perform these functions. However, as Marr [1] famously pointed out, to reverse engineer the computational steps of 'getting' or 'storing new mail' (i.e. to understand 'how' separate algorithms perform the functions) additional knowledge is required. First, the engineer must understand the abstract goals of the computation (i.e. 'what' information is manipulated, such as sender, receiver, subject and message fields) and they must also understand the hardware implementing the algorithm (i.e. 'where') such as a desktop computer connected to a mail server.

Cognitive neuroscientists face the similar problem of reverse engineering the algorithms of cognitive functions from measurements of brain activity. Of crucial importance is identifying the information processed and how it is distributed and transferred throughout the different brain areas involved. Careful control of cognitive parameters during experimentation and use of advanced analysis methods have provided considerable insight into what

information is represented by the brain [2]. However, although subtle experimental conditions (e.g. a happy versus a fearful face) might reveal which brain areas respond to a particular visual stimulus, they cannot inform us about the information processing that subsumes these crucial categorizations. For example, what information is being extracted from the fearful face? How is it extracted? What are the networks involved in this extraction? Is their information processing content changing over time?

We believe that cognitive neuroscience must now embrace such detailed questions of information processing to provide a fuller account of cognition from brain activity measured with a millisecond time resolution. Here, we present a novel approach that uses a method called 'Bubbles' [3,4] to characterize the information goals of cognitive tasks (i.e. 'what information is processed?'). We then model 'how' brain algorithms process this information between stimulus onset and behavioural response.

Diagnostic use of Information

Consider an observer asked to perform a cognitive task, for example, judging the expression of the face presented on a computer screen. By using reverse correlation methods (i.e. Bubbles technique [3,4]), we can precisely determine the facial information that the observer uses to make this judgement (e.g. reveal that the wide-opened eyes predict 'fear' or that the smiling mouth predicts 'happy'. See Figure 1; The Bubbles Procedure). Therefore, we know that the brain of this observer must process at least these specific diagnostic features to reach these particular judgments.

This is an obvious point, but its implications are important and often neglected. When the brain processes the features diagnostic of facial expressions, then it must go through information processing (i.e. functional) states such as 'F = process the eyes' or 'G = process the mouth' before categorizing 'fearful' or 'happy', respectively. Diagnostic information, therefore, provides a lower bound on the information processing states that the brain must go through. Although we cannot search the brain directly for these states, we can search for correlated states of brain activity in the brain measurements m[F] and m[G] – henceforth called 'microstates m[F] and m[G]' – which refer to the brain measurements m[X] of information processing state X. Knowledge of what the diagnostic features are then constrains what we are searching the brain for, that is, the microstates correlated with the processing of the diagnostic features. The processing of these diagnostic

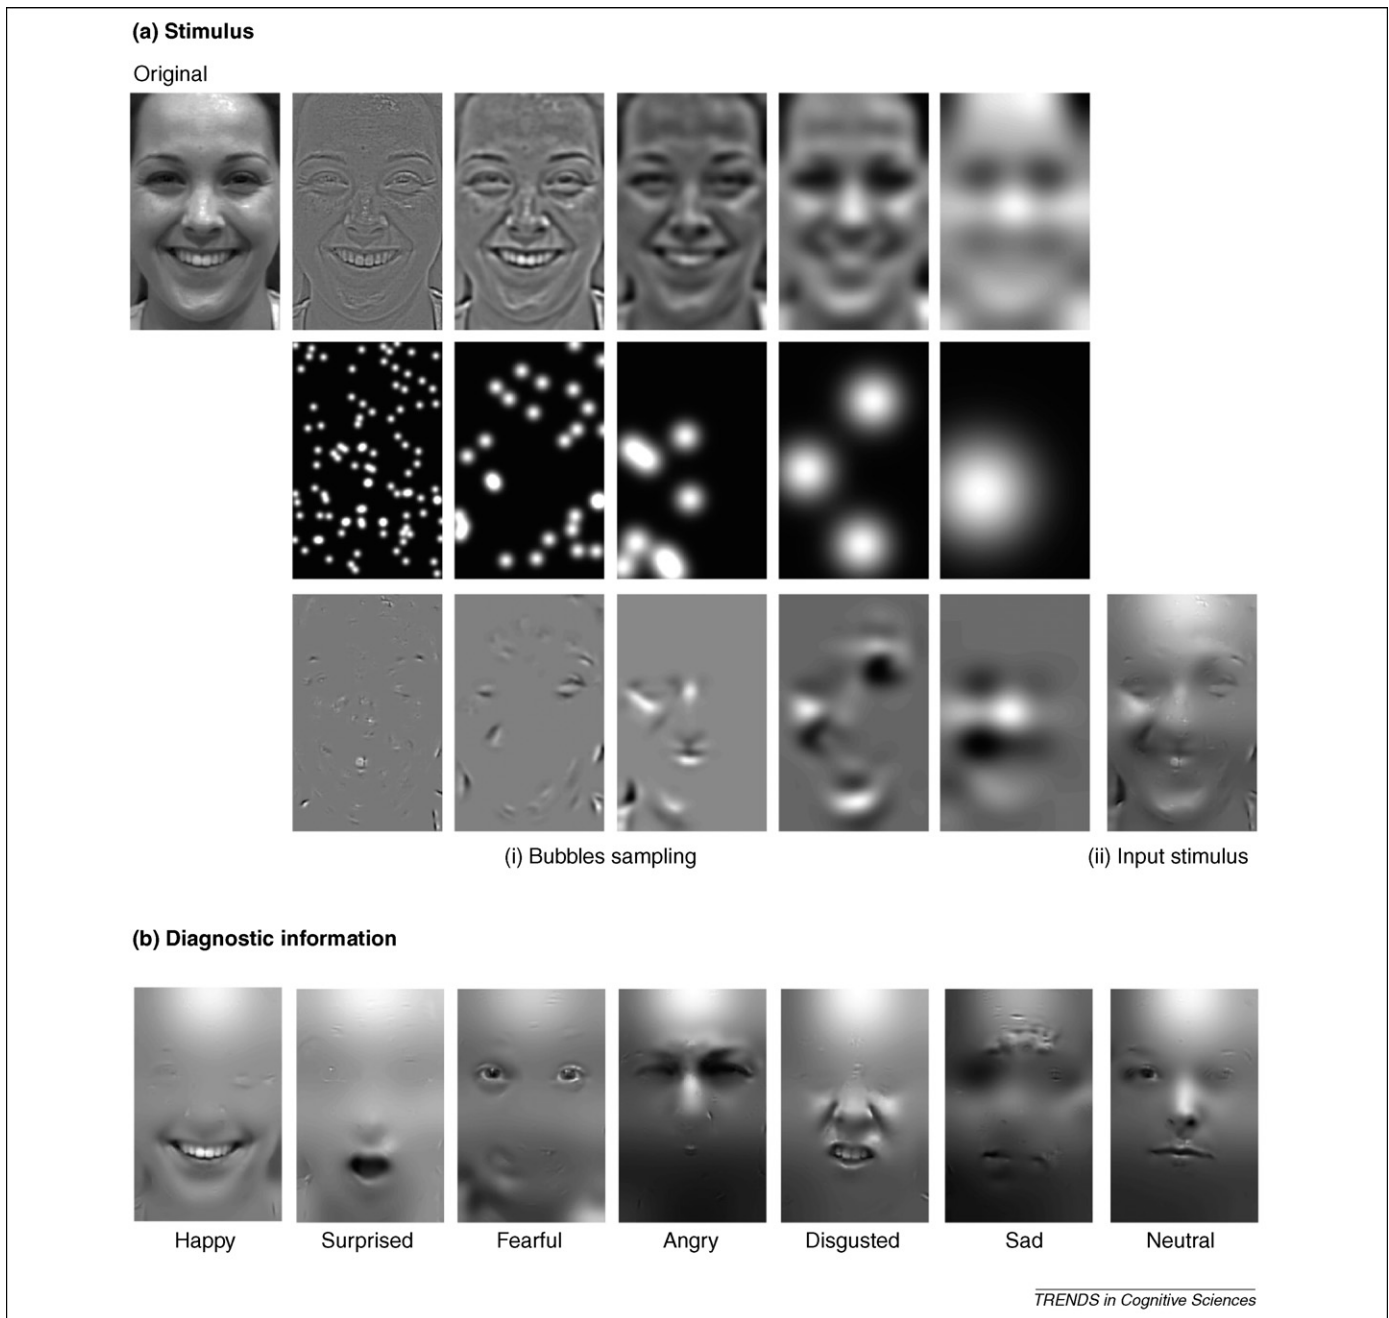Corresponding author: Schyns, P.G. (philippe@psy.gla.ac.uk).

Figure 1. The Bubbles Procedure. **(a)** Stimulus. (i) Bubbles sampling. In this application of Bubbles, an original face picture is initially filtered into five independent spatial frequency bands. On each sampling trial, Gaussian apertures are randomly allocated across spatial frequency bands and image locations to sample and reveal different portions of the input face. (ii) Input stimulus. On each trial, the input face will present different samples of facial information. As the location of the bubbles randomly change across trials, the entire face will be sampled throughout the experiment. **(b)** Diagnostic information. After many trials, a multiple linear regression associates correct and incorrect categorization responses (here, for different facial expressions) with the sampled facial features. Diagnostic information, therefore, represents the facial information that the observer's brain must process to correctly perform the task. In the example, different features are processed for correctly categorizing different facial expressions of emotion [3,22,23]. Adapted from Ref. [22].

features can be construed as the 'information goals' of the brain during a particular visual categorization task. Note that diagnostic features are found in any sensory modality (e.g. in audition, to distinguish the vocalizations of different bird species; in taste, to classify wines; in touch, to categorize smooth from rough surfaces; in smell, to avoid noxious substances, and so forth). So, the approach of identifying the states of brain activity correlated with the processing of diagnostic features is a potentially pervasive approach to cognitive neuroimaging in many perceptual and cognitive tasks [5,6].

The previous paragraph highlighted a natural relationship between diagnostic features in the stimulus, their processing in the brain and the observer's categorization behaviour. However, is this all there is? For example, when the visual stimulus is an expressive face, does the brain process other features of the face that are not diagnostic for this task? How could we know what these features are if they do not underlie behaviour? For example, perhaps the information processing state 'F = process the eyes' systematically precedes the state 'G = process the mouth', when behaviour only requires the presence of the smiling mouth.
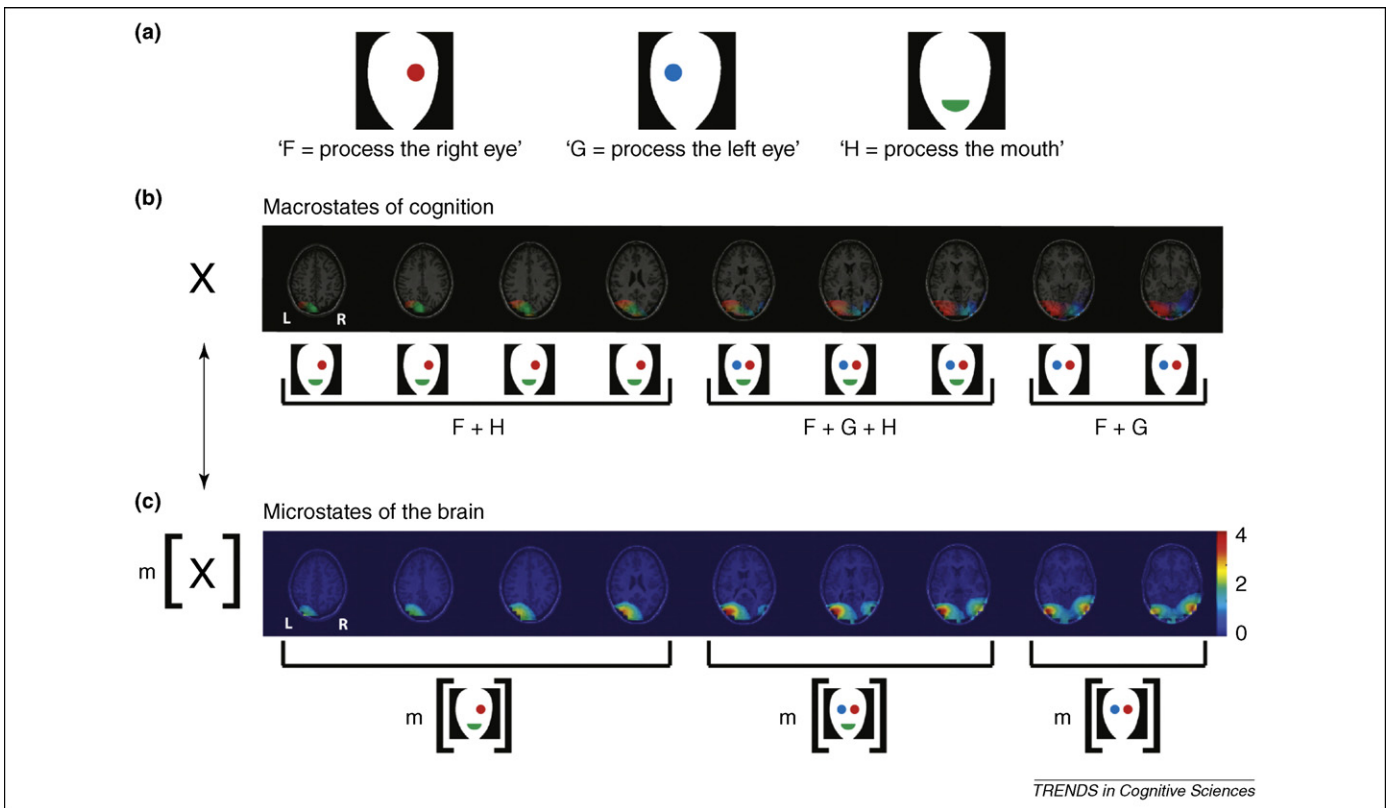
Figure 2. Translation of macrostates of cognition X into microstates of the brain m[X] (part 1). Imagine that the task of an observer is to determine whether the input face is smiling or not. Suppose we know that the observer must use information from the mouth to respond correctly. Suppose further that we measure the observer's brain with an MEG scanner while she performs this task. **(a)** We find that different brain regions process different combinations of three basic facial features: 'F = process the right eye', 'G = process the left eye' and 'H = process the mouth'. The data report an experiment carried out in our laboratory using Bubbles. Figure 3 explains how Bubbles derives these estimates of information processing from behavioural [3] and brain data [4]. **(b)** and **(c)** present 2D slices of the observer's 3D brain (going from the top of the head on the left to the bottom on the right) measured with MEG between 165 ms to 175 ms after stimulus onset. (b) Macrostates of Cognition X. We translate the activity of the brain microstates m[X] into the macrostates X of an information processing function. (c) Microstates of the Brain m[X]. Occipito-temporal hotspots correspond to the typical sources of the left and right M170 (the MEG version of the N170 ERP). The arrow between X and m[X] illustrates the correspondence that our method establishes between information processing state X (e.g. 'F = process the left eye') and its brain measurement m[X] (e.g. a MEG 'hotspot' in the right occipito-temporal cortex between 165 ms to 175 ms following stimulus onset).

How would we know that the brain correlate of F did, indeed, correspond to processing of the eyes, and not to processing of the nose or forehead? This presents a difficulty: the processing of diagnostic features in the brain is necessary, but it might not be sufficient to understand the full process. The next sections develop these themes.

Translating brain microstates into functional states of information processing

Coming back to the example of facial expressions, here, we show that a method exists to translate the microstates of brain measurements m[X] into the functional states of information processing X. In Figure 2, the brain is measured with a magnetoencephalographic (MEG) scanner when an observer is instructed to categorize faces appearing one at a time on the screen according to whether they are expressive or not. Because MEG brain measurements are finely resolved in space and time, one can measure 'where' and 'when' activity happens in the brain. To simplify matters, consider a single snapshot of brain activity measured at 170 ms after stimulus onset. Figure 2 presents this activity as different 2D slices of the 3D brain (from the top to the bottom of the head), with hotspots of activity (indicated in red) in the left and right occipito-temporal regions corresponding to microstates m[X]. Our

translation function reveals that activity in the right hotspot corresponds to 'G = process the left eye' (encoded with a blue colour). Activity in the left hotspot corresponds to 'F = process the right eye' (encoded with a red colour) and 'H = process the mouth' (encoded with a green colour, in the occipital region). The main point of this article is that reverse correlation can be extended from behavioural measurements to translate brain microstate m[X] into information processing macrostate X. Figure 3 explains how this translation can be accomplished using the Bubbles technique [3,4], although other forms of translation could be developed, possibly leading to other macrostates.

Pursuing the example, we now consider the brain as a dynamic machine and so we generalize from Figure 2 and repeat the translation exercise for different time slices between stimulus onset, when information is presented, and response, when the observer acts on this information. Figure 4 presents the generalization of Figure 2, every 10 ms, between 125 ms until 185 ms after stimulus onset. The point here is to note the sequence of transitions of information processing states in the brain. In the example, transitions start from processing the contra-lateral right eye in the left brain, to later process the contra-lateral right eye and then the mouth (the latter in parietal regions), to
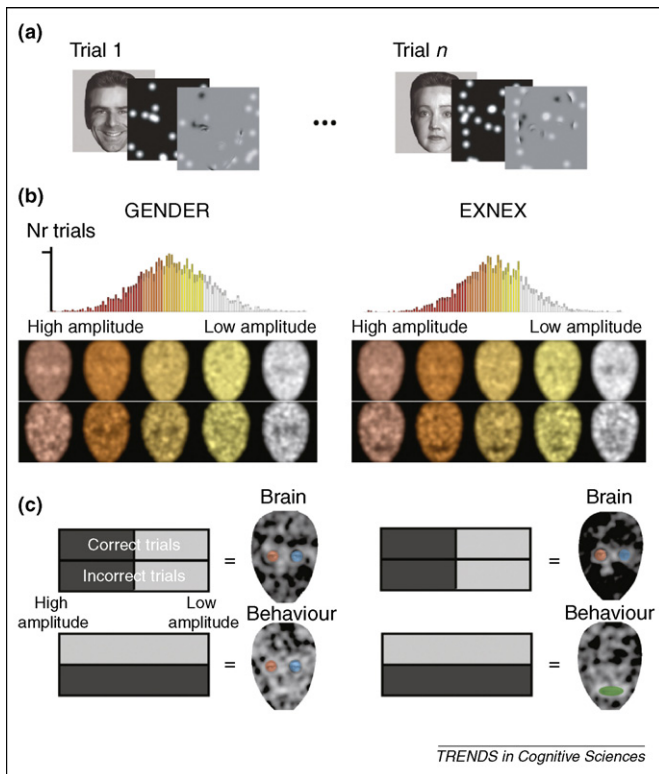
Figure 3. Translating brain microstate m[X] into functional macrostate X. **(a)** Stimulus information. Trials labelled 1 to *n* illustrate that a mask punctured with randomly located Gaussian apertures (indicated as white apertures against a black background) samples information from the input face on each trial. In GENDER, the observer categorizes the gender of the faces; in EXNEX, the same observer categorizes whether the same faces are expressive or not. In both conditions, categorization accuracy and the observers' brain are measured on each trial (here EEG voltages. But, the same method can be applied to MEG source activations, fMRI blood oxygen level-dependent responses or any other brain response for which single trial measurements can be obtained). **(b)** Information bins: from brain amplitudes to stimulus information. In GENDER and EXNEX, a histogram represents the distribution of brain responses (e.g. amplitudes) across trials. We split the histogram into five colour-coded bins of equal trial numbers. For each trial, we substitute the measured amplitude with the corresponding stimulus mask, to derive bins of stimulus masks corresponding to bins of amplitudes (top row, correct trials; bottom row, incorrect trials). Finally, for each amplitude bin, we average the stimulus masks to compute the corresponding stimulus information. For example, in GENDER, for correct trials, the bin of highest brain amplitudes is associated with the presence of two eyes in the input stimulus. **(c)** Computation of brain and behaviour classification images. With two simple subtractions schematized as (dark information bins and light information bins) we derive two classification images from the stimulus information bins. The 'Brain' classification images depict the statistically significant features discriminating high from low amplitudes ($p < .01$, colour-coded for different features), irrespective of correct versus incorrect behaviour. The 'Behaviour' classification images represent the features discriminating correct from incorrect face categorizations ($p < .01$, colour-coded for different features), irrespective of low versus high brain amplitudes. This approach enables a direct comparison between functional brain states (here, 'F + G = process the left and right eyes' in both tasks) and behaviour (requiring the eyes in GENDER and the mouth in EXNEX) (see Refs [4,15,16,24,25] for EEG and Ref. [26] for fMRI). Adapted from Ref. [4].

finally process the left eye, the right eye and the mouth in different brain regions. The right hand side of Figure 4 illustrates these transitions in macrostates of information processing.

The research programme: macrostate transitions, automata and functional theories
The ability to translate brain activity into a set of distinct information processing states and their transitions is important: it projects cognitive neuroscience into the familiar territory of the formal study of mechanisms. Auto-

mata theory is a branch of mathematics, pre-dating computer science, which provides a generic definition of a computing machine, also called a 'mechanism', a 'finite automaton' or an 'algorithm'. A finite automaton comprises a finite set of states, a finite set of possible inputs and a table specifying how successive transitions between the states implement a computation. When sufficiently powerful (as in the most general sense of a Turing machine), such an automaton is believed to be capable of computing any 'intuitively computable function'. With this in mind, cognitive neuroscience can be framed in terms of the so-called 'Church-Turing thesis' [7]: cognitive functions of the brain, such as determining whether a face is expressive or not, or its gender, can be implemented as cognitive algorithms. That is, any cognitive function can be implemented as a (finite) set of information processing states and a transition table. The transition table specifies how the computation should progress from state to state to produce behavioural responses from sensory inputs.

Such mechanistic conception was very much prevalent at the inception of cognitive science, based on Alan Turing's work [8–10], but its meaning has become somewhat diluted in modern cognitive neuroimaging. Although the term 'mechanism' appears in many studies, few genuinely aim at unravelling brain mechanisms as suggested earlier. This probably arises from the success of measurements of brain activity without a sufficient time resolution to identify brain microstates at a precise timescale (such as functional magnetic resonance imaging [fMRI]), precluding the identification of a precisely timed sequence of cognitive macrostates, with millisecond resolution, without which the study of dynamic cognitive mechanisms is not possible.

When brain microstates are measured at a sufficiently precise timescale (by using electroencephalography [EEG] [11,12] and MEG), their transitions have been successfully modelled (e.g. Refs [13,14]). However, these approaches do not reveal the information that each brain microstate processes (if any) and how transitions between the information processing states implement a cognitive algorithm. Our approach seeks to do this by translating brain measurements m[X] into information processing states X. Once this translation is performed, the evolution of brain activity over time can be framed as a succession of states that process information (i.e. brain activity can be framed as an information processing algorithm). In summary, by translating brain activity into information processing states, our approach based on reverse correlation enables the neuroimaging of the flow of cognitive information, not just the flow of brain activity.

To illustrate, suppose an observer who must determine whether the input face is smiling or not (Figure 2 and Figure 4), or whether the face is male or female. Figure 5 illustrates how automata can be derived from the output of the reverse correlation methods (described in Figure 1 and Figure 3) to model state transitions at an abstract, information processing, functional level of interpretation. Such automata detail the sequences of states in the occipito-temporal regions that is linking the input face to behavioural responses via a categorization-specific sequence of information processing states. For example, in the gender
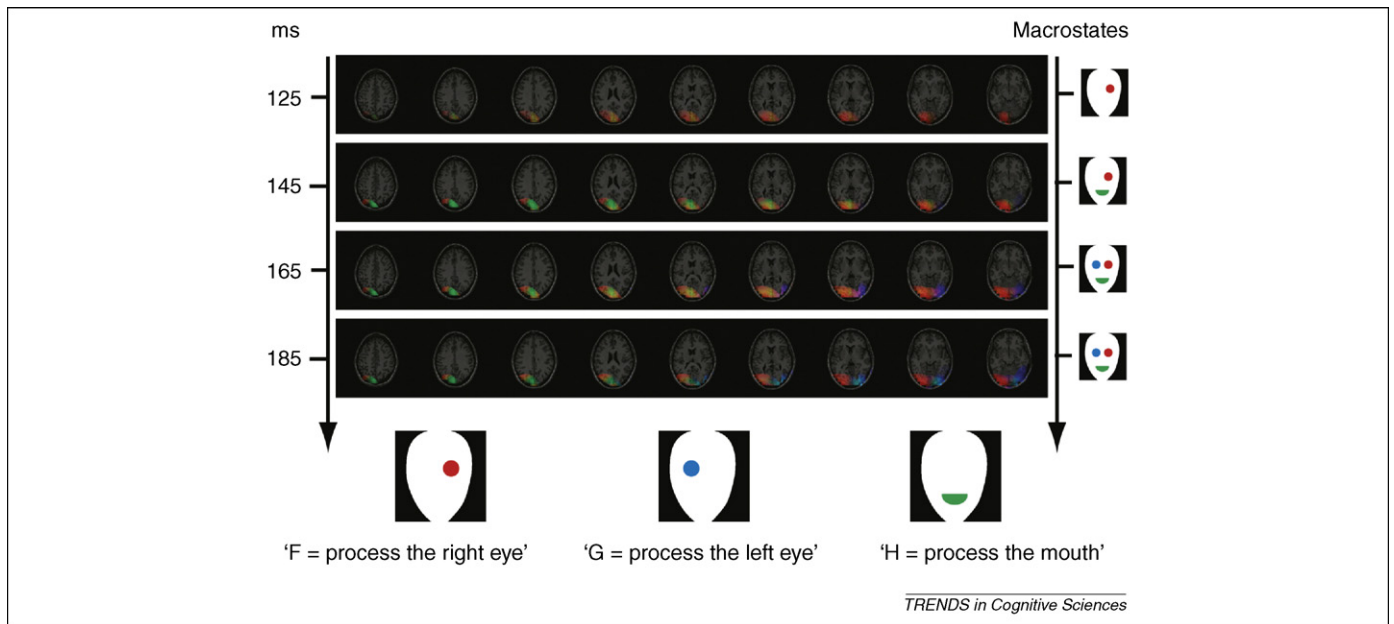
Figure 4. Translation of macrostates of cognition X into microstates of the brain m[X] (part 2). This Figure repeats the computations illustrated in Figure 2 and Figure 3 for four different time slices between 125 ms and 185 ms after stimulus onset. The schematized faces on the right hand side represent the macrostates of information (corresponding to processing different combinations of the basic features) that brain regions process across time. The schematized faces also illustrate the transitions in macrostates that the brain goes through over the time slices considered. That is, processing starts with the right eye, followed by the right eye and the mouth, followed by two eyes and the mouth.

task, the brain goes through the states of first encoding one eye and then two eyes. To judge whether the same faces are expressive or not, it will instead encode the eye and then the mouth. Thus, the automata represent the cognitive algorithms used by the brain to resolve the gender and expression of a face.

Segmentation of brain microstates using functional macrostates

Automata are not only important for the fundamental relationship that they offer between computability and the understanding of brain functions. They also represent a temporal segmentation of the function of brain activity that can enhance the understanding of how a cognitive function is implemented in the brain. To illustrate, consider again Figure 5a. Starting with 'process the left eye' on the right side of the brain, a crucial transition occurs between 141 ms and 161 ms to 'process the two eyes'. Consider further that 'process the right eye' is typically the first state on the left side of the brain. The transition to the macrostate 'process the two eyes' is crucial because it indicates either a transition of information across the two hemispheres or a transition of processing mode from contra-lateral to bi-lateral in each hemisphere. From this observation, one could then examine brain activity around this particular time window for evidence of this inter-hemispheric transmission, or for evidence of a transition from uni- to bi-lateral processing within a hemisphere. Another important transition occurs when the diagnostic features of behaviour become reflected in the states of information processing (e.g. the mouth in the 'expressive or not' task in Figure 5). At this time point of its dynamics, the brain has encoded the information that is required for behavioural decision – we know this by examining the classification image derived from behaviour in this task.

These examples illustrate that an automaton segments the macrostates of the brain into several epochs with possibly different functional interpretations. One can then use these epochs and transitions between epochs as temporal guidelines to further examine the crucial microstates of the brain implementing these transitions. We now use two examples from our own research to illustrate this point. In Ref. [15], we sought to understand the information processing function of the N170 event related potential (ERP) – the N170 is a negative deflection measured in the occipital region ~170 ms after stimulus onset. We showed that the N170 curvature reflects a process that integrates visual information specific to each expression according to a pattern. Specifically, starting 50 ms before the N170 peak, facial information tends to be integrated from the eyes downward on the face. The integration stops and the N170 peaks when the information diagnostic for judging a particular expression has been integrated (e.g. the corners of the nose in 'disgust', the smiling mouth in 'happy'). So, we explicitly related two microstates of brain activity (the curvature and peak of the N170) to two states of information processing (integration of information from the eyes downward on the face and the integration of diagnostic, expression specific information). In Ref. [16], we analyzed the oscillatory activity of the brain when observers consciously perceived an ambiguous figure (Dali's painting Slave Market with the *Disappearing Bust of Voltaire*). We found that the theta band (4–8 Hz) encoded the features related to one conscious perception (i.e. 'Voltaire') whereas the β band (12–24 Hz) encoded the features related to the other conscious perception (i.e. 'The Nuns'). Again, we related microstates of brain activity (energy in the theta and β oscillatory bands) with two states of information pro-
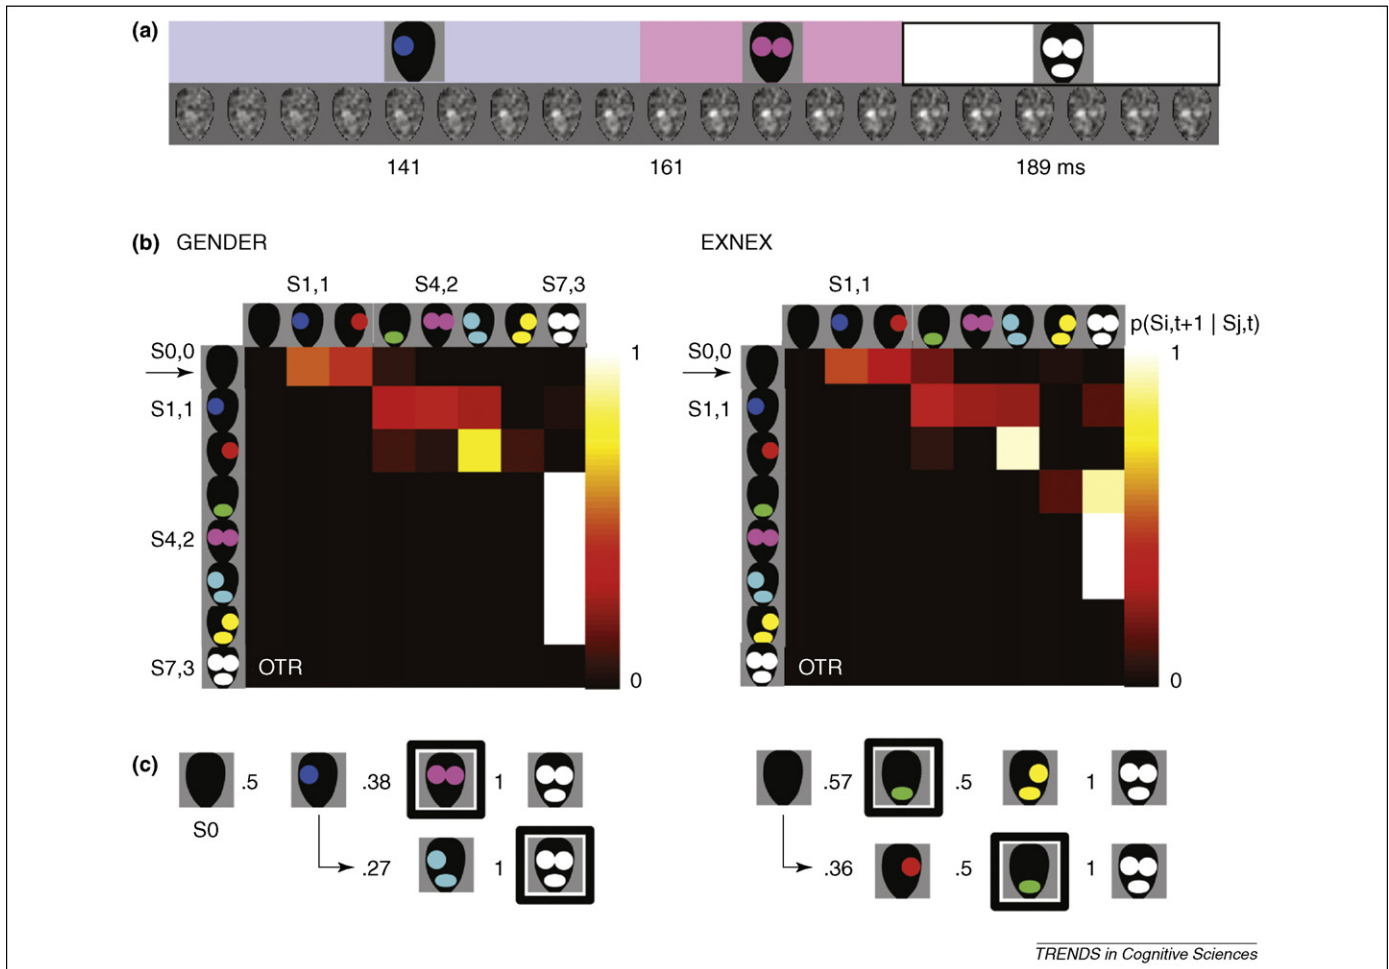
Figure 5. From transitions of information processing states to mechanisms. **(a)** Identifying information processing macrostates. Classification images computed every 4 ms between stimulus onset and behavioural response are first sorted for the feature combination that they encode. The blue, purple and white set of classification images correspond to colour-coded macrostates 'G = processing the left eye', 'F + G = processing the right and left eyes' and 'F + G + H = processing the left and right eyes and the mouth'. **(b)** Characterizing the mechanisms of state transitions. For the GENDER and EXNEX categorizations on right occipito-temporal electrode (OTR), transitions of macrostates are then reported in a table. The table reports the conditional probabilities of transitions between state Si (e.g. processing the left eye) and state Sj (e.g. processing the left and right eyes) computed by pooling transition data across observers. To illustrate, starting at time 0 in GENDER with a state of no information processing $S_{0,0}$ (see arrow) the first row describes the conditional probabilities of a transition to any other possible state $p(Si = 1...7,1| S0,0)$. The most likely transition at time = 1 is S_1, the contra-lateral left eye (encoded in blue), with a .5 probability. From the left eye in the left column (i.e. $S_{1,1}$), the most likely conditional transition at time 2 is S_4, the two eyes (encoded in purple) with a .38 probability. From the two eyes at time 2 (i.e. $S_{4,2}$), the most likely transition at time 3 is S_7 (encoding all three features), with a probability of 1. Halting states (i.e. S_7) are those for which there is no transition to any other state. **(c)** Most likely state transitions. Across categorizations (compare GENDER and EXNEX), diagnostic information is reached within two transitions and is represented by schematic faces surrounded by a black box. Note also the progression in the content of the information states: the transition to a feature combination (e.g. one eye and the mouth) can only be reached from one of the component features, never from $S_{0,0}$. Similarly, the probability of a transition from $S_{0,0}$ directly to S_7 (comprising all three features) is 0. Adapted, with permission, from Ref. [25].

Box 1. Outstanding questions

- What is the relationship between the distinct neural microstates observed in EEG, reflecting stable topography of the brain electric landscapes [11,12] and macrostates of information processing?
- Can the reverse-correlation based, reverse engineering approach be extended from single or a small number of EEG electrodes to the many cortical sources of oscillatory cortical activity as measured with MEG [27] to identify independent information processing networks in the brain?
- How are the different oscillatory networks reported in MEG data [27] related to the processing of different sorts of features (e.g. fine scale versus coarse scale [16]) and/or different cognitive functions (e.g. attention and memory)?
- Reverse correlation is a multiple linear regression. It therefore assumes a linear mapping between visual information and brain signals. This mapping will determine the estimates of information states of the brain. When should nonlinear regressions be used?

cessing (processing the features subtending the perception of 'Voltaire' versus 'The Nuns' in an ambiguous figure).

Conclusions

The approach presented here uses reverse correlation to understand cognitive functions from brain measurements and behaviour. Starting with the information goals of a cognitive task ('What information is computed?' – i.e. Diagnostic information), we apply reverse correlation to time-resolved brain measurements to model an algorithm of the cognitive function resolving these information goals. That is, a model showing the transitions of the information processing states of the brain between stimulus onset and behaviour. This approach attempts to address the question of 'How is information computed in the brain?' Here, we discuss a number of questions for future research (see also Box 1).

One important question concerns the meaning of the computed information states. They reveal that the brain 'does something' to specific visual features at different time points. What is unclear is what this 'doing something' is. Revealing the information over which an information processing function is performed (e.g. the contra-lateral eye) is not sufficient. We must also understand what this information processing function is (e.g. 'expecting the contra-lateral eye', 'attending to it', 'encoding it', 'recognizing it', 'memorizing it' and so forth). Drawing definitive inferences about functional information from brain imaging data is a generic difficulty, given that it logically represents a 'reverse inference' (Ref. [17]; see also Ref. [18] for discussions of the 'sufficient' neural systems that perform a particular cognitive function; see Ref. [19] for discussions of how different coordinations of large scale cortical networks can implement different cognitive functions).

In Figure 5, state transitions occurred over one or two electrodes. Extrapolating from the MEG data presented in Figure 2 and Figure 4, state transitions could potentially be measured across voxels in the 3D microstates of the brain, to describe the dynamics of information processing networks. Notwithstanding the sheer complexity of the combinatorial explosion arising from centimetre spatial resolution and millisecond temporal resolution, resulting in thousands of separate precise measurements, a crucial problem is that of determining functional causality in the sense of 'F causes G'. Ideally, one might want to track the 'information flow' of macrostates, starting from the expectation of task relevant information (e.g. the mouth in 'happy') in frontal or pre-frontal cortex, to the encoding of this information in occipito and occipito-temporal cortex, to categorization or perceptual decision in parietal cortex (all of this of course with possible mutual interactions between the listed cortical regions). However, in doing so, one needs to allow distal transitions to be modelled, in which distal happens both in time (because there is *a priori* no reason why two different states contiguous in time are causally related) and space (because there is *a priori* no reason why two different states proximal in space are causally related). If causality were resolved, one would derive a probabilistic automaton forming a network, in which the nodes represent functional regions of the brain going through different information processing states over the time course of processing. The automata presented Figure 3 are, in this sense, over-simplified. They are derived from a single electrode, eliminating the problem of transitions in the 3D space of the brain and they assume that temporal contiguity between states warrants causality. One approach would be to consider brain regions as pairs of variables and examine their causal relationships (e.g. with Granger Causality [20] or Dynamic Causal Modelling [21]). Still, it will require considerable conceptual and formal developments to expand from 'single electrode automata' to 'networks of nodes automata'. Nevertheless, the main point remains that we can start relating stimulus to behaviour in the brain via the dynamics of a set of information processing states and their transitions. This is one important step towards Turing's original agenda of understanding the brain as a machine that processes information.

References
1 Marr, D. (1982) *Vision.* Henry Colt & Co
2 Norman, K.A. *et al.* (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* 10, 424–430
3 Gosselin, F. and Schyns, P.G. (2001) Bubbles: a technique to reveal the use of information in recognition. *Vision Res.* 41, 2261–2271
4 Schyns, P.G. *et al.* (2003) A principled method for determining the functionality of ERP components. *Neuroreport* 14, 1665–1669
5 Schyns, P.G. (1998) Diagnostic recognition: task constraints, object information and their interactions. *Cognition* 67, 147–179
6 Gosselin, F. and Schyns, P.G. (2002) RAP: a new framework for visual categorization. *Trends Cogn. Sci.* 6, 70–77
7 Kleene, S.C. (1967) *Mathematical Logic.* Wiley
8 Turing, A.M. (1936) On Computable Numbers, with an application to the Entscheidungsproblem. *Proc. Lond. Math. Soc* 42, 230–265
9 Turing, A.M. (1950) Computing machinery and intelligence. *Mind* 49, 433–460
10 Chomsky, N. (1957) *Syntactic Structures.* Mouton
11 Strik, W.K. *et al.* (1997) Psychiatry research. *Neuroimaging* 75, 183–191
12 Lehmann, D. *et al.* (1987) EEG alpha map series: brain microstates by space oriented adaptive segmentation. *Electroencephalogr. Clin. Neurophysiol.* 67, 271–288
13 Pascual-Marqui, R.D. *et al.* (1995) Segmentation of brain activity: model estimation and validation. *IEEE T. Biomed. Eng.* 42, 658–665
14 Rabiner, L.R. (1989) A tutorial on Hidden Markov Models and selected applications in speech recognition. *Proc. IEEE* 77, 257–286
15 Schyns, P.G. *et al.* (2007) Dynamics of visual information integration in the brain to categorize facial expressions. *Curr. Biol.* 17, 1580–1585
16 Smith, M.L. *et al.* (2006) Perceptual moments of conscious visual experience. *Proc. Natl. Acad. Sci. U. S. A.* 103, 5626–5631
17 Poldrack, R.A. (2006) Can cognitive processes be inferred from neuroimaging data? *Trends Cogn. Sci.* 10, 59–63
18 Price, C.J. and Friston, K.J. (2002) Degeneracy and cognitive anatomy. *Trends Cogn. Sci.* 6, 416–421
19 Bressler, S.L. and Tognoli, E. (2006) Operational principles of neurocognitive networks. *Int. J. Psychophysiol.* 60, 139–148
20 Roebroeck, A. *et al.* (2005) Mapping directed influence over the brain using Granger causality and fMRI. *Neuroimage* 25, 230–242
21 Friston, K.J. *et al.* (2003) Dynamic causal modelling. *Neuroimage* 19, 1273–1302
22 Smith, M.L. *et al.* (2005) Transmitting and decoding facial expressions. *Psychol. Sci.* 16, 184–189
23 Schyns, P.G. *et al.* (2002) Show me the features. Understanding recognition from the use of visual information. *Psychol. Sci.* 13, 402–409
24 Smith, M. *et al.* (2004) Receptive fields for flexible face categorizations. *Psychol. Sci.* 15, 753–761
25 Smith, M.L. *et al.* (2007) From a face to its category via a few information processing states in the brain. *Neuroimage* 37, 974–984
26 Smith, F.W. *et al.* (2008) Classification images reveal the information sensitivity of brain voxels in fMRI. *Neuroimage* 40, 1643–1654
27 Gross, J. *et al.* (2001) Dynamic imaging of coherent sources: studying neural interactions in the human brain. *Proc. Natl. Acad. Sci. U. S. A.* 98, 694–699